Consumption Wedges: Measuring and Diagnosing Distortions *

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

February 11, 2026

Abstract

Ample empirical evidence documents deviations from the canonical consumption-savings model; yet, it remains difficult to assess the roles of different underlying distortions, such as financial constraints and behavioral preferences. We develop a sufficient-statistics approach that measures individual-level wedges between observed and counterfactual "frictionless" consumption. Since different distortions imply different wedge properties, wedges provide a diagnostic to distinguish between models. We measure wedges using administrative transactions data linked to surveyed expectations for a population of middle-income, low-liquidity US consumers. The expectations data allow us to distinguish wedges attributable to frictions and behavioral preferences from wedges driven by deviations from full-information rational expectations (FIRE). We find that consumption wedges are large and heterogeneous: the median wedge is 40% of frictionless consumption in absolute value, with 49% of consumers under-consuming (negative wedges) and 51% over-consuming (positive wedges). Borrowing constraints cannot rationalize this pattern because they only generate negative wedges. Models combining present bias with borrowing constraints, or featuring consumption adjustment costs, best account for the wedge properties we document.

*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 Andy Abel, David Berger, John Beshears, Gideon Bornstein, Taha Choukhmane, Francesco D'Acunto, Tim de Silva, Marty Eichenbaum, Ben Enke, Joao Gomes, Yuriy Gorodnichenko, Fatih Guvenen, Erik Hurst, Greg Kaplan, Ben Keys, Dirk Krueger, Karen Lewis, Igor Livshits, Pierre Mabille, Peter Maxted, Emi Nakamura, Gisle Natvik, Matt Notowidigdo, Guillermo Ordoñez, Jonathan Parker, Michael Roberts, Nick Roussanov, Rob Stambaugh, Jón Steinsson, Luke Taylor, Jessica Wachter, and Steve Zeldes. 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 EarnIn and the Wharton School's Jacobs Levy and Rodney L. White Centers. The surveys were approved under Penn and Berkeley's IRBs (#851816 and #2024-03-17263, respectively) and analysis of the transactions data was covered by Harvard IRB (#IRB22-1281).

†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

The seminal consumption-savings model going back to Friedman (1957) continues to be the backbone of much theoretical and empirical work in macroeconomics. Yet extensive empirical evidence documents 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 (Johnson, Parker and Souleles, 2006).[1] The dominant explanation for these deviations are borrowing constraints, which play a central role in theories of consumption-savings behavior in macroeconomics and household finance (e.g., Kaplan and Violante, 2014). However, alternative distortions such as present bias (Attanasio, Kovacs and Moran, 2024; Maxted, Laibson and Moll, 2025), consumption adjustment costs (Berger and Vavra, 2015), and bounded rationality (Ilut and Valchev, 2023) can also account for these empirical patterns.

Distinguishing between these explanations is important because they imply different distributional and aggregate consequences of fiscal policy, monetary policy, and business cycle fluctuations. For example, Lee and Maxted (2023) and Maxted et al. (2025) show that adding present bias to an otherwise standard heterogeneous agent New Keynesian model (with borrowing constraints and an illiquid asset) leads to significant amplification of both fiscal and monetary policy. Moreover, at the micro level, it is necessary to understand which forces shape consumer behavior in order to identify and ameliorate the causes of financial distress. Empirical evidence on which distortions best explain the consumption-savings behavior of low-liquidity consumers is therefore crucial to guide theory and, ultimately, policies that affect their well-being and the broader macroeconomy.

In this paper, we develop a new sufficient statistics approach to assess the explanatory role of alternative models. We measure consumer-level "wedges" between observed consumption and the hypothetical "frictionless" consumption implied by an Euler equation and budget constraint. 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). We show that consumption wedges can be used as a diagnostic to discriminate between models of consumer behavior since different distortions have different predictions for the properties of consumption wedges, such as their sign, correlates, and responses to shocks.

These inferences rely on two innovations of our approach relative to existing research that

[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).

measures the effect of distortions as 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. This innovation is similar in spirit to Hsieh and Klenow (2009), who estimate firm-level productivity wedges (misallocation). This proves to be essential in revealing distortions: while the aggregate (median) wedge is less than 1% of frictionless consumption, the median *absolute value* of consumer-level wedges is 40%, with similar shares of over-consumption (positive wedges) and under-consumption (negative wedges).

Second, we use subjective beliefs data instead of assuming full-information rational expectations (FIRE). Any consumption choice can be rationalized by some set of beliefs, so assuming FIRE risks conflating the effects of deviations from FIRE with that of frictions and behavioral preferences.[2] Using subjective beliefs data, we can isolate the influence of distortions from that of deviations from FIRE. This distinction proves quantitatively important: frictionless consumption explains 82% of the variation in the cross-section of consumption when measured using subjective beliefs, compared to 57% under FIRE. Distortions explain the remaining 18%, suggesting that both deviations from FIRE and distortions are first-order determinants of variation in consumption.

We estimate consumption wedges using a sample of 5,028 consumers for whom we obtain both de-identified administrative banking data and linked survey expectations. The consumers are customers of EarnIn, an American financial technology company that offers users access to their wages prior to their 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 covers a significant swathe of the income distribution, with over-representation between the 30th and 80th percentiles and under-representation 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 in the 2022 Survey of Consumer Finances. This combination of middle-of-the-distribution earnings and low liquid wealth means our respondents are precisely the low-liquidity, high-MPC households that heterogeneous-agent macroeconomic models emphasize: hand-to-mouth behavior among such consumers plays a central role in shaping the macroeconomic predictions and welfare implications of workhorse models like Kaplan et al. (2018). Our sample of wage workers with low liquid wealth is also of particular policy relevance. Many of the largest tax-expenditure and social-insurance programs in the United States—

[2] There is abundant evidence suggesting that consumer expectations deviate from FIRE. For example, inflation expectations are excessively influenced by grocery prices (D'Acunto et al., 2021), while D'Acunto et al. (2024) finds evidence of extrapolative income expectations.

including the Earned Income Tax Credit, the Child Tax Credit, health insurance premium subsidies, and unemployment insurance—explicitly target low- and moderate-income wage workers with limited liquid assets. This same group of households also relies increasingly on technology-enabled financial products such as earned wage access to manage short-term liquidity shortfalls, which are the subject of ongoing regulatory attention (Marek et al., 2025).

We measure consumption wedges as the gap between each individual's observed consumption and the hypothetical frictionless level they would choose absent distortions. We formulate a frictionless benchmark model in which a consumer chooses consumption and saving 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.[3] Frictionless consumption is characterized by an Euler equation and a budget constraint.

The purpose of the frictionless benchmark is not to explain observed consumption, nor does it describe what consumers should do. Rather, it enables us to measure the hypothetical consumption that would arise in the absence of distortions (frictions and behavioral preferences). The wedge between observed and frictionless consumption quantifies the total net impact of *all* distortions on consumption. Quantitatively, the frictionless benchmark proves to be a useful modeling foundation; it is able to account for 82% of the cross-sectional variation in consumption for the EarnIn sample when calculated using consumer's subjective expectations.

We then show that our wedge measurement approach generalizes beyond our highly stylized benchmark to a large class of 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. 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, such as labor supply or a richer asset environment (including the case of complete markets).

Empirically, 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. Expressed in dollars, the median absolute value wedge is $13,301 per year, which amounts to 34% of respondents' median annual post-tax income of $39,615. The large magnitude of the wedges implies that distortions are significant determinants of the consumption among low-liquidity consumers. In contrast, the mean and median

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

wedges are much smaller, at 15% and 0.9%, respectively. This discrepancy highlights the importance of measuring micro-level wedges, as aggregate wedges mask underlying cross-sectional heterogeneity. This applies especially when heterogeneity in consumer behavior matters for a model's predictions (e.g., Kaplan et al., 2018; Patterson, 2023), but less so for questions where a representative agent model is sufficient.

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 simultaneous existence of a large share of under- and over-consumers implies that uni-directional distortions are insufficient to rationalize the consumption choices in our sample. Borrowing constraints, in particular, cannot be the dominant distortion for the majority of our sample as they only generate under-consumption. We identify two candidate modeling approaches that can account for these patterns. The first is to augment models featuring borrowing constraints to include frictions that produce positive wedges, such as present bias (e.g., Maxted, 2025; Attanasio, Kovacs and Moran, 2024). The second approach features frictions that generate both positive and negative wedges, such as consumption adjustment costs (e.g., Berger and Vavra, 2015; Fuster et al., 2021; Beraja and Zorzi, 2024) or bounded rationality (e.g., Ilut and Valchev, 2023). Both can produce consumer inertia, which can create both positive and negative wedges by limiting the consumption response to shocks.

We also evaluate the importance of using subjective beliefs data in for measuring micro-level wedges and for explaining the cross-section of consumption. We decompose observed consumption into three terms: frictionless consumption under FIRE, a wedge induced by subjective beliefs deviating from FIRE, and the remaining wedge induced by distortions. For the first term, we impute FIRE over income, inflation, and returns for each person using their realizations. The sum of the second and third term corresponds to the wedge we would have measured if we did not have subjective beliefs data. This combined wedge exhibits higher rates of over-consumption (67% versus 51%) and a smaller median absolute value of 27% (versus 40%), indicating that subjective beliefs are also quantitatively important for measuring the distribution of wedges. A covariance decomposition further reveals that deviations from FIRE explain 25% of the variation in observed consumption and distortions explain 18%, suggesting the two are similarly important determinants of the cross-section of consumption.

To assess the robustness of our results, we provide evidence on their sensitivity to two sets of assumptions: the modeling of preferences and measurement error. As for the preference specification, we find little sensitivity of our results to the choice of the inverse IES (γ). The discount factor β matters more; however, even for a very low annual factor (0.80), we estimate a sizable (20%)

share of over-consumers. We also consider heterogeneity in preferences, since it could manifest as a wedge. To allow for a plausible degree of heterogeneity, we set preferences to match one of the three types in Aguiar et al. (2025) such that they minimize wedges. We obtain similar results, with a median absolute value wedge of 37.2% 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 Falk et al. (2018) Global Economic Preferences Survey and Andreoni and Sprenger (2012), respectively. We find that these measures of preferences are neither statistically nor economically significant predictors of wedges. These (and additional) exercises suggest that our evidence on the role of distortions is unlikely to be confounded by unobserved preference heterogeneity.

Regarding measurement error, we perform four sets of robustness exercises. In the first, we drop respondents who are likely to exhibit more measurement error, such as users who have low financial literacy, spent less than 6 minutes on the survey, or rounded inflation expectations to a multiple of five percentage points. 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. Here, we recalculate 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 aggregate the wedge inputs for each cluster to "average out" idiosyncratic measurement error and use these to calculate a wedge for each cluster. Fourth, 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 median absolute value wedge. Hence, measurement error—whether in consumption, income, or beliefs data—is unlikely to significantly distort our findings.

Finally, we demonstrate how wedges can serve as a diagnostic tool in both reduced-form and structural applications. Our reduced-form analyses validate the wedges by examining their correlations with other variables of economic interest, which also provide diagnostic evidence on the source of distortions. We begin by correlating the wedges with survey responses about how consumers would adjust savings if they expected higher inflation. We find that over-consumers are more likely to report saving less under high inflation than under-consumers, revealing an internal consistency between consumers' anticipated behavior and their wedge-determined consumption type. Turning to the diagnostic aspect, we correlate wedges with MPCs as well as with proxies for consumption commitments, financial distress, and homeownership. We generally find strong correlations with signs inconsistent with borrowing constraints being the dominant distortion (but consistent with either present bias or inertia). For example, MPCs (as measured out of

the 2021 stimulus checks) are positively correlated with wedges—the opposite of what borrowing constraints would predict if they were the main driver of high MPCs. Similarly, consumption commitments, measured as the ratio of spending on housing and childcare to income, are positively correlated with wedges. This is consistent with consumption commitments acting as a dominant distortion. Over-consumption is also strongly related to financial distress (measured as perceptions of financial anxiety, unmanageable debt, and difficulty borrowing). These patterns fit present-biased behavior and inertia but are at odds with borrowing constraints. Finally, we find that over-consumption is rarer among consumers with a mortgage. This suggests that borrowing constraints may be a more important distortion for consumers with substantial illiquid assets, resembling the "wealthy hand-to-mouth" of Kaplan and Violante (2014), while distortions like present bias and inertia dominate for those without such assets.

In our structural application, we study a series of simple quantitative heterogeneous-agent incomplete-markets models that feature at least one of three distortions: borrowing constraints, present bias, and consumption adjustment costs. We calibrate the distortion parameters to target the 51% over-consumer share as well as the 40% median absolute value wedge found in the empirical estimations. We find that present bias combined with a borrowing constraint is best able to match these empirical moments. This model can account for 78% of the observed over-consumer share and 66% of the median absolute value wedge. We also discuss the potential role of non-rational expectations in order to fully account for consumer behavior. An important for role non-rational expectations is also consistent with our covariance decomposition finding that deviations from FIRE explain a large share of the cross-sectional variation in consumption.

Taken together, these findings imply that distortions beyond borrowing constraints—such as present bias and consumer inertia—are important drivers of high MPCs, financial distress, and overall consumption.

Related Literature. Our paper contributes to several literatures. A first strand is the macroeconomics literature using wedges to quantify frictions and distortions. We most directly build on Zeldes (1989), which estimates average wedges in the Euler equations of high- and low-wealth households. Zeldes focuses on testing for borrowing constraints, assumes FIRE, and uses data on food consumption from the PSID. We also build on the business cycle accounting methodology of Chari, Kehoe and McGrattan (2007), which popularized studying wedges between actual and frictionless values of aggregate variables. Subsequent work in this tradition has focused on quantifying the importance of misallocation across firms (Hsieh and Klenow, 2009; Baqaee and Farhi, 2020) and risk-sharing across households (Berger, Bocola and Dovis, 2023).

We take this approach in a new direction by both measuring wedges at the individual level and using subjective beliefs. By using subjective expectations, we can identify wedges due to frictions and behavioral preferences separately from those due to deviations from FIRE. Moreover, the ability to measure micro-level consumer wedges enables us to show that the typical magnitude of the wedge is much larger than what the mean or median of its distribution suggests.

Second, we contribute to the empirical macroeconomics 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 research documenting large MPCs, especially among consumers with low liquidity (Johnson, Parker and Souleles, 2006; Baker, 2018; Ganong and Noel, 2019; 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 findings have motivated a growing literature proposing behavioral explanations, such as bounded rationality and present bias (Ilut and Valchev, 2023; Lian, 2023; Attanasio, Kovacs and Moran, 2024; Ganong, Greig, Noel, Sullivan and Vavra, 2024; Maxted, 2025; Boutros, 2026).

We contribute to this ongoing debate by providing new evidence and tools. Our findings point to a prominent role for distortions like inertia or present bias, which can generate overconsumption. Methodologically, the micro-level consumption wedge that we introduce here is a diagnostic tool that can be applied in other settings. An appealing feature is that, unlike MPCs, measuring consumption wedges does not require quasi-experimental variation. And while transactions data is ideal for calculating consumption, wedges can also be computed using survey data alone.

Third, we build on work in empirical macroeconomics that emphasizes the central role of consumer beliefs, including departures from FIRE, in explaining consumer behavior (see reviews by Weber et al., 2022; D'Acunto et al., 2023b). Recent papers have linked beliefs to consumption decisions using surveys (Coibion et al., 2023; D'Acunto et al., 2023a), grocery purchases (Weber et al., 2023), German bank data (Hackethal et al., 2023), and credit card data (Kanz et al., 2020). Consumer beliefs systematically deviate from FIRE by, for example, over-weighting grocery

prices in inflation expectations (D'Acunto et al., 2021) and forming extrapolative income forecasts (D'Acunto et al., 2024). These findings motivate our use of subjective beliefs data.

Additionally, we provide new evidence on the relative importance of beliefs versus distortions for consumption. Our covariance decomposition indicates that both deviations from FIRE and distortions explain similarly large fractions of the cross-section of consumption. Hence both are important for understanding empirical consumption.

Outline. We present the frictionless benchmark 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, Section 5 explores the robustness of the results, and Section 6 provides additional diagnostic evidence. Section 7 concludes.

2 Theory: Measuring Consumption Wedges

We derive a formula for micro-level consumption wedges. We first solve a stylized consumption-savings model to characterize frictionless consumption. The wedge is the difference between frictionless and observed consumption. We then show that the formula generalizes from the stylized benchmark model to much richer environments. 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.

Frictionless Benchmark Model. Consumer i lives for T periods. Every period t, 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 (gross) nominal rate of return. A negative value of $A_{i,t}$ corresponds to borrowing. The price level in period t is P_t. The consumer has "standard preferences," which we take to mean time consistent, time separable, homothetic, strictly increasing, strictly concave, and continuously differentiable.

Differently from much of the prior work on wedges in the tradition of Chari et al. (2007), we allow consumer i's beliefs to flexibly depart from FIRE. As such, we do not assume that her subjective conditional expectation operator follows Bayes' rule, nor that it uses valid probability distributions.

At time t, consumer i chooses nominal consumption $C_{i,t}$ and savings $A_{i,t+1}$ to maximize her

expected utility subject to a budget constraint, solving:

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

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

where the operator $\widetilde{E}_{i,t}(\cdot)$ denotes i's subjective expectation conditional on her information set at time t. Optimal consumption $C^{\star}_{i,t}$ in the frictionless benchmark satisfies 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 model has three key features. First, there are no economic frictions (borrowing constraints, adjustment costs, etc.). Second, the model assumes standard preferences and thus precludes behavioral preferences like present bias or habit formation, which can result in "as-if" constrained behavior. The benchmark *intentionally* omits both of these types of distortions so that we can measure their impact as a wedge between frictionless and observed consumption. In other words, the purpose of the benchmark is not to provide a realistic model of observed consumer behavior, but to facilitate measurement. 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 is a important distinction, as it enables us to isolate the impact of distortions separately from the impact of deviations from FIRE.

Frictionless Consumption. We characterize 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. Appendix A.1 derives the following expression:

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

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 $R_{i,t+j}$. The preference parameters are the inverse in-

tertemporal 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. Frictionless consumption represents a counterfactual level of consumption. It holds constant the consumer's wealth and the prices they expect to face, but it removes any distortions. Therefore, we can quantify the total impact of all distortions omitted from the frictionless benchmark on observed consumption (holding wealth and prices constant)[4] as a wedge

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

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. Because different distortions imply different properties of wedges, they can be used as a diagnostic tool. We illustrate this for some commonly used distortions. First, consider borrowing constraints, which can only reduce consumption. Therefore, a testable implication of borrowing constraints is that they always generate negative consumption wedges. A second example is present bias, which increases the preference for consuming more in the present. Therefore, it produces positive wedges.[5] Third, consider distortions producing inertia such as consumption commitments (adjustment costs), habit formation (reference-dependent preferences), or some forms of bounded rationality (e.g., where cognition is costly and limits consumption adjustments). By limiting the reaction of consumption to shocks, either over- or under-consumption can emerge depending on the nature of the shock. For example, over-consumption can arise following a negative income shock if the consumer fails to adjust her consumption downward. Therefore, distortions inducing inertia can generate either positive or negative wedges. See Appendix A.2 for further details on these examples.

[4] In general equilibrium, "turning off" distortions will generally imply a different wealth distribution and prices. Hence the counterfactual we study (holding wealth and prices constant) is partial equilibrium. In other words, the wedge describes how distortions shift the consumer's consumption policy function.

[5] Note that, when measuring the consumption wedge (Equation 5) using the approximate frictionless consumption formula (Equation 4), the measured wedge includes some negative approximation bias, which we measure and discuss in more detail in Section 5.5. Therefore, while present bias produces only positive wedges (when measured exactly), both positive and negative wedges are possible when measured approximately.

Robustness to Model Extensions. The frictionless benchmark abstracts away from additional choice variables, such as labor supply, multiple assets, and durable goods. In Appendix A.3, we show that this is without loss of generality. Intuitively, adding choice variables introduces further necessary conditions that characterize the overall optimum (i.e., *all* optimal choices). But since we need only characterize optimal consumption, our formula applies to the entire class of models where an Euler equation and budget constraint are necessary conditions for optimality, even if they are not sufficient conditions for overall optimality. This still holds if other arguments of utility (e.g., leisure) enter utility in a non-separable way, as our first-order approximation of the Euler equation remains unchanged (see Lemma 1 in Appendix A.1). Appendix A.3 discusses extensions featuring additional choices, assets, and durable goods.

3 Data and Survey Design

3.1 EarnIn Administrative Data

We utilize 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. The firm had over 2.5 million active users as of the start of the sample period.[6] 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 key variables: nondurable spending and income. Following prior literature (Lusardi, 1996; Ganong and Noel, 2019), we focus on nondurable spending to derive total consumption, since the relationship between spending and consumption of durables depends on factors such as their rate of depreciation and utilization of the durable. To calculate nondurable spending, we classify outflow transactions based on their assigned transaction category (e.g., "Grocery stores"). We map from the more than 500 transaction categories to more aggregated spending categories building on Ganong and Noel (2019).

We obtain total or "notional" consumption (i.e., the argument of the utility function) by dividing each respondent's nondurable consumption by the typical expenditure share of nondurable goods (79.37%).[7] We obtain this nondurable expenditure share from Beraja and Zorzi (2024), which calculates it using Consumer Expenditure Survey data.

We measure income as the sum of post-tax earnings and unemployment insurance (UI) benefits. EarnIn maintains a separate database of earnings used to determine eligibility for EWA. We

[6] "More than 2.5 million active users" is reported by EarnIn as of November 2, 2021 (link).

[7] This calculation yields notional consumption under the assumption that notional consumption is a Cobb-Douglas aggregate of durable and nondurable good consumption flows, as detailed in Appendix A.3.3.

classify inflow transactions as either earnings (observed post-tax), UI benefits, or "other" using the earnings data together with the transaction category, memo line, and periodicity of the transaction. Our wedge analysis measures income and consumption using data from the 12-month period preceding each survey (e.g., annual nondurable spending and income between October 2021 to September 2022 for wave 1).[8]

Appendix B provides additional details on the EarnIn data, sample construction, and variable definitions.

3.2 Survey Data

We analyze 5,028 consumers using EarnIn administrative data linked with survey responses covering demographics, personal finances, and subjective economic expectations. Our dataset connects expectations with earnings, spending, and savings data that provide a near-comprehensive picture of a consumer's economic activity. Datasets linking subjective economic expectations and bank account transactions are rare, and we are the first (to our knowledge) to collect this data for US consumers.[9]

EarnIn fielded three survey waves between 2022 and 2024. 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. All respondents received an Amazon gift card after completing the survey ($5 for waves 1 and 3, $10 for wave 2).

EarnIn users received an email invitation to complete the survey if they met minimum thresholds for activity in their linked bank account over the prior 12 months (detailed in Appendix B.2) and had not opted out of receiving emails from EarnIn. Wave 2 was a follow-up survey of wave 1 respondents, whereas waves 1 and 3 surveyed the cross-section of eligible EarnIn users. Across all three surveys, EarnIn sent 463,267 survey invitations and received 15,853 responses, yielding an aggregate response rate of 3.3%. 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, other than gender: women make up slightly more than half of

[8] Conceptually, we want to measure time t consumption wedges using time t consumption and time t beliefs about time $t+k$ variables. We verify that our wedge analysis obtains similar results when aggregating over fewer than 12 pre-survey months (see Appendix Figure E.4).

[9] D'Acunto et al. (2021) link economic expectations to grocery spending and Kanz et al. (2020) 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.

the sample, but two thirds of the responses (see Appendix Table B.2).

Each survey wave elicited expectations for inflation, income growth, and interest rates. 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 personal income growth over the next 12 months, the annual percentage yield they would expect on additional savings over the next 12 months, and the annual percentage rate they would expect to pay on new borrowing over the same period. We also asked respondents how a hypothetical increase in inflation would affect their saving and why, and elicited measures of financial distress, use of alternative financial products, financial literacy, demographics, and total household income, savings, and debt (across categories).[10]

In waves 2 and 3, we added survey questions to expand our wedge measurement and analysis. We elicited 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 also added measures of committed consumption by asking respondents whether they own or rent their home, monthly rent/mortgage payments, and childcare spending.

Sample Restrictions for Analysis. Measuring consumption wedges requires high-fidelity data on consumption, income, and beliefs. Thus, we restrict our analysis to the sample of consumers with reliable survey responses and comprehensive transactions data that can be cleanly categorized as consumption or income. All sample restrictions are detailed in Appendix B.3, and we summarize the main steps here.

We begin by excluding low-effort and inattentive survey responses: we require a minimum time spent on the survey of 3.5 minutes, internally consistent responses, and economic expectations within reasonable bounds (in the spirit of the best practices for ex-post survey data quality checks outlined in Stantcheva, 2023). We then restrict our analysis sample to users whose transactions data provide a comprehensive and reliable measure of consumption and income. The most important restriction is that we require users to consume primarily through the bank account linked to EarnIn, as proxied for by having at least 20 outflow transactions in all months throughout the analysis period. We apply further restrictions (shown in Appendix Table B.3) which exclude users with missing transactions memos or account balances and trim outliers for key financial variables (nondurables spending, income, average propensity to consume, wealth-

[10] The survey instruments are available online: http://bit.ly/3VUPf4o for wave 1, https://bit.ly/3KqKTzq for wave 2, and http://bit.ly/46IfpNW for wave 3.

to-income, and expected levered return). Our final analysis sample includes 5,263 responses from 5,028 unique users.[11]

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, we leverage our detailed economic, financial, and demographic data and train a machine learning model (XGBoost) to predict illiquid assets, liquid assets, and debt (where the latter two are 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 C.1.

3.3 Summary Statistics

Table 1 presents summary statistics for our analysis sample. Compared to the US population, our sample skews younger (with an inter-quartile age range of 31-43), female (69%), and non-white (42%). Figure 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 over-representation between the 30th and 80th percentiles and fewer workers in the tails (Panel A). However, our sample is predominantly low-liquidity relative to respondents to the Survey of Consumer Finances (Panel B). The median respondent has $39,615 annual post-tax income, $17,541 in total debt, and $20,585 in total assets—but only $250 in liquid assets. Only 18% of respondents have a mortgage.

Panel C of Figure 1 presents the distribution of inflation expectations in our survey relative to MSC and SCE respondents from the same time period. Inflation expectations in our sample are broadly consistent with those of nationally representative samples, particularly the SCE. 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 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 sensible, with medians of 2% and 15%, respectively.

[11] There are 235 users who responded in both wave 1 and the follow-up survey in wave 2, giving us a total of 5,263 linked survey responses.

Table 1. Summary Statistics

	Mean (1)	SD (2)	P25 (3)	P50 (4)	P75 (5)	N (6)
Panel A: Demographics						
Female (%)	69	.	.	.	.	5,263
White (%)	54	.	.	.	.	5,175
Age	37	9	31	36	43	5,245
Panel B: Household finances						
Liquid assets ($)	1,520	3,683	236	250	750	5,263
Total assets ($)	47,719	72,833	8,303	20,585	46,438	5,263
Total debt ($)	43,843	57,683	7,408	17,541	41,239	5,263
Has mortgage (%)	18	.	.	.	.	5,263
Total net worth ($)	3,876	46,791	-13,970	484	16,337	5,263
Nondurables spending ($)	29,989	13,589	19,997	26,941	37,299	5,263
Income ($)	43,655	19,346	29,934	39,615	53,212	5,263
Total net worth to income (%)	7	98	-37	1	41	5,263
Nondurables C/Y (%)	72	26	54	68	85	5,263
Panel C: Economic expectations						
E(Inflation, 1Y) (%)	7	7	3	5	10	5,263
E(Inflation, 3Y) (%)	4	8	-2	4	8	5,263
E(Income growth) (%)	5	9	2	3	5	5,263
E(Rate on savings) (%)	3	3	1	2	4	5,263
E(Rate on borrowing) (%)	15	10	7	15	24	5,263
E(Levered return) (%)	16	10	7	15	25	5,263
E(Spending growth) (%)	6	66	-2	5	10	2,347

Notes: The table presents summary statistics for our analysis sample, which consists of 5,263 EarnIn users. See Appendix B.4 for variable definitions. The smaller observation counts for race and age are due to missing survey responses. The count is smaller for expected spending growth because this question was only included in waves 2 and 3.

Figure 1. Sample Benchmarks

Panel A. Distribution of Income Percentiles

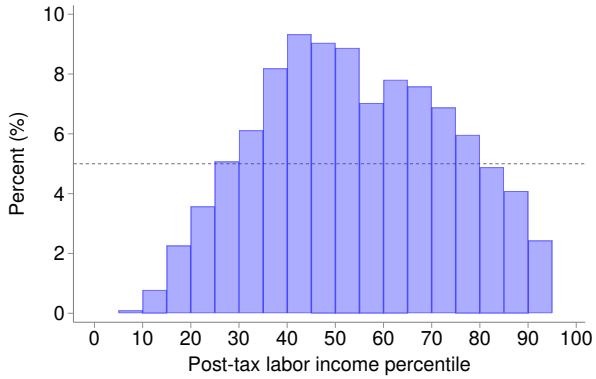

Panel B. Distribution of Liquid Wealth

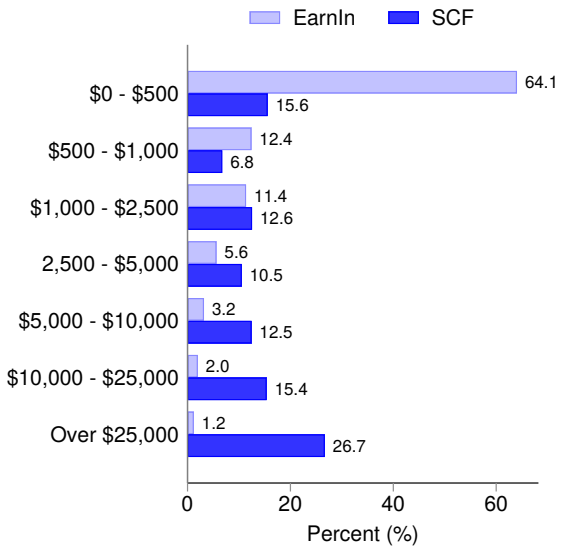

Panel C. Distribution of Inflation Expectations

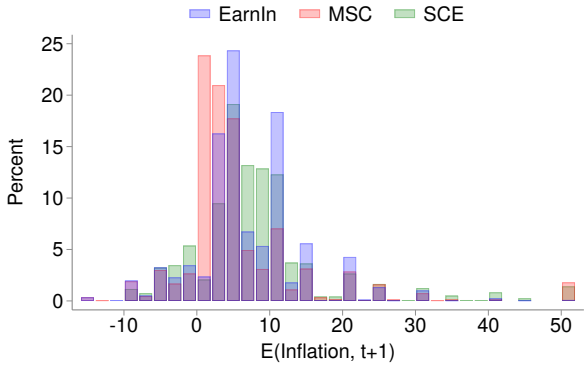

Notes: Panel A presents the distribution of post-tax labor income percentiles among our EarnIn analysis sample, estimated using CPS ASEC data (United States Census Bureau, 2025) and the CPS ASEC Tax Model (Lin, 2022). The CPS ASEC sample restricts to adults age 18 or older who received wage or salary earnings during the reference year. The 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. Panel C shows the distribution of one-year-ahead inflation expectations for our sample compared to those of 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). For each of the three survey months, we reweight MSC and SCE data from the corresponding month so that the share of respondents from each survey month matches the distribution from the EarnIn sample.

3.4 Quantifying the Frictionless Benchmark

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

Preference Parameters. Our baseline parameterization uses standard values for the annual discount factor ($\beta = 0.92$) and the inverse IES ($\gamma = 2$).[12] In Section 5, we perform sensitivity analyses confirming 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 $\widetilde{E}_{i,t}R_{i,t+1}$ using expected returns and 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 $\widetilde{E}_{i,t}R^D_{i,t}$ (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. In a robustness analysis, we also allow the expected return to reflect beliefs about default (i.e., nonpayment), and find little impact on the properties of the distribution of wedges (see Appendix Figure E.1).

Term Structure of Beliefs. The frictionless benchmark is a function of the full term structure of beliefs. Since we only observe one- and three-year-ahead beliefs for inflation and one-year-ahead beliefs for nominal income growth and interest rates, we impute the remaining term structure. For inflation expectations, we use data on one to thirty-year-ahead expected inflation (measured via inflation swaps). For interest rates, we assume the expected gross levered rate is constant (i.e., a flat term structure). Lastly, we impute expected income growth expectations using a procedure that allows for two key empirical properties: (1) large but temporary variation in expected growth rates of income in the short-run due to anticipated shocks (e.g., job changes) and (2) lifecycle dynamics. The SCE reveals significant reversion in income expectations over a year, which we use to discipline the relationship between one-year- versus two-year-ahead expected income growth. We use the MSC to estimate a lifecycle profile of income growth expectations, which we apply to our

[12] 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).

sample for three-year-ahead (and later) expectations. We detail all of our imputation procedures in Appendix C. In Section 5, we document robustness to other term-structure assumptions as well an alternative wedge formula that requires different data but no term-structure assumptions.

4 Results: Consumption Wedges

This section reports the empirical distribution of consumption wedges and discusses how their properties relate to those implied by various models of distortions. We then use a covariance decomposition to examine the importance of subjective beliefs data in both explaining the cross-section of consumption and in affecting the measured wedge distribution.

4.1 The Distribution of Consumption Wedges

Our first main finding is that consumption wedges are large and heterogeneous. Figure 2 Panel A shows the distribution of consumption wedges as a percent of frictionless consumption, and Panel B plots the distribution of their absolute value. The mean wedge is 15%, implying that observed consumption is 15.2% above frictionless consumption on average, while the median absolute value wedge is 40.1%. In terms of dollars, the median wedge is $259 while the median absolute value wedge is $13,301. (See Appendix Figure D.1.) This indicates that distortions to consumption are significant relative to the median post-tax income of $39,615 in our sample.

Figure 2. Distribution of Consumption Wedges

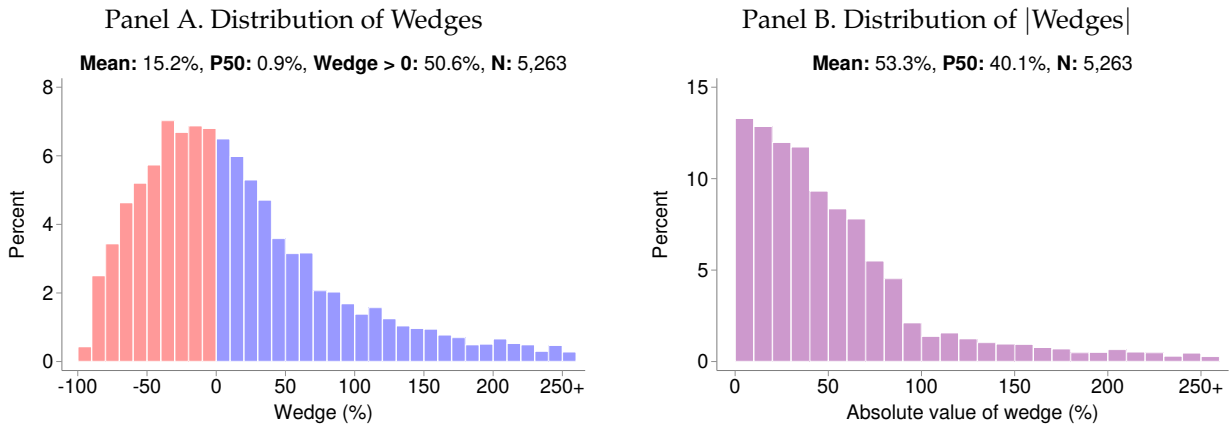

Notes: The graphs show the distribution of consumption wedges (left) and their absolute values (right). Wedges are reported as the percent deviation of observed consumption from frictionless consumption.

The substantial dispersion of wedges highlights the value of studying micro-level wedges. Mean and median wedges are relatively small at 15% and 1%, respectively, but are less than half

the magnitude of the median absolute value wedge (40%). Without observing the distribution, one could significantly underestimate the importance of distortions in the consumption decisions of low-liquidity consumers. For phenomena where heterogeneity in consumer behavior drives aggregate outcomes, like fiscal and monetary policy transmission (Kaplan et al., 2018) or business cycles (Patterson, 2023), these patterns imply that heterogeneity in distortions are important to incorporate into the modeling of consumers. For phenomena where representative agent models of consumers suffice, the small median and average wedges suggest that incorporating distortions may be less critical.

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 view that borrowing constraints are the key friction shaping the consumption choices of low-liquidity consumers and can be relied on as the dominant modeling paradigm in household finance and macroeconomics. As discussed in Section 2, borrowing constraints only generate negative wedges. Thus, borrowing constraints could explain the 49% of consumers with negative wedges, but cannot be the dominant friction for the other 51%.

We propose and explore two promising approaches to account for the behavior of all consumers. A first solution is including both present bias and borrowing 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).

4.2 The Importance of Subjective Expectations

Our unique combination of survey and transactions data allows us not only to disentangle distortions from other, belief-based deviations from FIRE, but also to assess the relative importance of subjective beliefs data. We quantify the role of subjective beliefs both for measuring wedges and, through a novel exercise, as a determinant of consumption.

We begin by calculating each respondent's "frictionless FIRE consumption" via the same benchmark formula (Equation 4) but using imputed FIRE rather than subjective expectations. We impute FIRE by setting beliefs equal to their realized values and applying the term structure modeling described in Section 3.4. We measure income realizations using the transactions data, use the CPI for inflation, and the typical interest rates on deposits and credit cards.[13] We restrict this

[13] For the savings rate, we use the average prevailing deposit APY across banks and credit unions (as reported in the Federal Deposit Insurance Corporation's "National Rates and Rate Caps" report). We obtain average credit card rates from the Federal Reserve Board's *Consumer Credit*, G.19 statistical release. Credit card interest rates reflect the APR averaged across all credit card accounts at all reporting banks for August and November 2023 in the G.19

analysis to wave 1 because we have realized income only for this wave.

Based on these calculations, we can decompose observed consumption into frictionless FIRE consumption ($C_{i,t}^{\star,\text{FIRE}}$), a "subjective expectations (SE) wedge" ($\eta_{i,t}^{\text{SE}}$), and a distortions wedge ($\eta_{i,t}^{\text{distortions}}$):

$$C_{i,t} = \underbrace{C_{i,t}^{\star,\text{FIRE}} + \eta_{i,t}^{\text{SE}}}_{\equiv C_{i,t}^{\star}} + \eta_{i,t}^{\text{distortions}}. \tag{7}$$

We calculate the SE wedge by taking the difference between frictionless consumption calculated under subjective expectations ($C_{i,t}^{\star}$ from Equation 4) and frictionless *FIRE* consumption ($C_{i,t}^{\star,\text{FIRE}}$). The SE wedge measures the impact of deviations in subjective expectations from FIRE on consumption. The distortion wedge ($\eta_{i,t}^{\text{distortions}}$) is the same wedge from Equation (5) that has been our focus, but we add the "distortions" superscript to distinguish it from the SE wedge.

We first assess the importance of using of subjective expectations data in calculating wedges. We compare the distribution of wedges we saw in Figure 2 with the distribution we would have obtained if we did not have such data and instead assumed consumers form FIRE. Appendix Figure D.2 reports the results, displaying the distribution of the sum of the SE and distortion wedges as a percentage of frictionless FIRE consumption. Without incorporating subjective expectations into the benchmark, we would have found a higher share of over-consumers (67% versus 51%) and smaller median absolute value distortions (27% versus 40%).

Next, we use the decomposition to quantify the relative importance of belief-based deviations from FIRE versus distortions (frictions and behavioral preferences) in explaining the cross-section of consumption. We calculate a covariance decomposition of Equation (7) to assess how much of the variation in observed consumption is explained by its three components. Starting with the first component, we find that frictionless FIRE consumption ($C_{i,t}^{\star,\text{FIRE}}$) accounts for 57% of the cross-sectional variation in observed consumption. The sizable proportion is to be expected, given that income and contemporaneous consumption tend to be highly correlated (Carroll, 1994), and $C_{i,t}^{\star,\text{FIRE}}$ is proportional to contemporaneous income (see Equation 4).

More strikingly, we find that subjective beliefs significantly improve the explanatory power of the frictionless benchmark, with the subjective expectations wedge accounting for more than half of the residual variation in consumption (25%). Thus, the frictionless benchmark measured using subjective expectations explains 82% of the cross-section of consumption. The role of distortions is similar in magnitude to deviations from FIRE, accounting for the remaining 18% of the variation

release. Savings interest rates reflect the average rate paid by all depository institutions for which data is available in the FDIC's National Rates and Rate Cap September 2022 dataset, based on a $2,500 balance and weighted by the institution's share of domestic deposits.

in the cross-section of consumption.

Taking stock, we find that distortions are both an important determinant of the level of consumption (as evidenced by the large typical distortion wedge) and explain a significant amount of cross-sectional variation in consumption. Additionally, the large share of variation in consumption explained by subjective beliefs indicates that differences in how consumers' beliefs deviate from FIRE is important for understanding the differences in their consumption decisions. This result underscores the importance of studying and modeling non-rational expectations for understanding consumption-savings decisions.

5 Robustness

In this section, we examine the robustness of our two main findings: (1) a significant share (51%) of over-consumers and (2) a median absolute value wedge of 40%. We focus on four classes of concerns: sensitivity to parameter choices, preference heterogeneity, measurement error, and approximation bias. Finally, we discuss external validity.

5.1 Sensitivity to Parameter Choices

We vary the value of each input parameter (discount factor β, inverse IES γ, and the nondurable share of expenditure $\frac{C}{Y}$) and recalculate the distribution of wedges to assess the sensitivity of our results to these parameter choices. We summarize the resulting share of over-consumers and the median absolute value wedge in Table 2 and the relationship between the results and a fuller range of inputs in Appendix Figure E.2.

Table 2. Sensitivity Analysis: Impact of Alternative Calibration Choices

	Parameter Range			Over-consumer (%)		Med. |Wedge| (%)	
	Baseline (1)	Min (2)	Max (3)	Min (4)	Max (5)	Min (6)	Max (7)
β	0.92	0.80	0.98	17.8	73.1	40.1	53.4
γ	2.00	1.00	5.00	49.1	51.0	37.4	60.1
Nondurable share of $\frac{C}{Y}$	0.7937	0.72	0.90	42.0	56.7	39.4	43.3

Notes: Table presents the sensitivity of our two main results to our parameter calibration: (1) the percent of users who over-consume and (2) the median absolute value wedge. We report the baseline parameters in column (1), and columns (2) and (3) indicate the range of parameter values considered in the analysis. Columns (4) and (5) report the smallest and largest observed over-consumer share across the considered parameter values (respectively). Columns (6) and (7) report the same for the median absolute value wedge. The sensitivity analysis varies one parameter at a time.

Discount Factor (β). Wedges depend on the discount factor (β): if a consumer discounts future consumption more strongly, their frictionless consumption is higher, leading to lower measured wedges. Relative to our baseline value of $\beta = 0.92$, a β of 0.80 reduces the share of over-consumers from 51% to 17.8% while a β of 0.98 increases the share of over-consumers to 73.1%. The median absolute value of the wedge is less sensitive because, in either case, wedges on one side of the distribution become larger while those on the other side become smaller. We conclude that, even when assuming extremely low discount rates, there remain both a meaningful share of over-consumers and large wedges.

Inverse Intertemporal Elasticity of Substitution (γ). The wedges are less sensitive to the inverse IES (γ). As we vary γ from one to five, the share of over-consumers remains close to 51%. The median absolute value wedge is stable around 40% for values of γ ranging from 2 to 5, but increases to 60% for γ equal to 1 (see Appendix Figure E.2 Panel B for the full range).

Nondurable Expenditure Share. In contrast to β and γ, the nondurable expenditure share does not affect the frictionless benchmark but it does affect the measurement of total observed consumption. Recall that we divide observed nondurable spending by the nondurable expenditure share to infer total consumption; therefore, assuming a lower nondurable expenditure share translates to higher consumption and more over-consumers. Varying the nondurable expenditure share from 72% to 90% monotonically decreases the share of over-consumers from 56.7% to 42.0%.[14] As above, the median absolute value wedge is broadly insensitive to the calibration choice.

Term Structure of Beliefs. We assess the sensitivity of the wedge estimates to the imputed term structure of beliefs by comparing our preferred imputations to the alternative assumption of constant beliefs in Appendix Figure E.3. Allowing one-year beliefs to persist indefinitely results in more outliers and larger wedges in absolute value, but the shape of the distribution is otherwise similar to those resulting from our baseline extrapolated expectations.

Robustness to Consumption Time Horizon. Our baseline approach uses 12 months of transactions data preceding the survey to calculate the wedges. Including a full year of data can smooth out the influence of typical expense shocks and seasonal spending patterns. Appendix Figure E.4 reports the share of over-consumers and the median absolute value wedge re-estimated using 1,

[14] The range of plausible expenditure shares is motivated by estimates in the existing literature. Our baseline value of 79.37% is from (Beraja and Zorzi, 2024). 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.

3, 6, and 9 months of consumption data. The median absolute value wedge is notably consistent across time horizons, declining modestly with the inclusion of more data from 46% to 40%. The share of over-consumers ranges from 43% to 51% and is highest when including all 12 months.

5.2 Preference Heterogeneity

Our baseline parameterization assumes homogeneous preferences. Below, we discuss why we restrict preference heterogeneity, report wedges allowing for some preference heterogeneity, and use proxies for risk aversion and discount rates to show that preference heterogeneity is unlikely to explain the cross-section of wedges.

Why restrict preference heterogeneity? A frictionless benchmark with unrestricted preference heterogeneity cannot generate consumption wedges; there exists some set of preferences to rationalize any feasible consumption choice. In contrast, a large "Euler equation estimation" literature has sought to infer preferences from consumption data by assuming that consumers face no distortions (e.g., Hansen and Singleton, 1982; Attanasio and Weber, 1995). We instead impose restrictions on preference heterogeneity in order to infer distortions. Our approach reflects the goal to inform the design of a large and important class of models that rely primarily on distortions, rather than flexible preference heterogeneity. This class has proven useful in terms of its ability to explain empirical patterns like the wealth distribution and high MPCs, and is a widely-used tool for studying policy, business cycles, and inequality (e.g., Kaplan et al., 2018; Maxted, 2025; Attanasio et al., 2024). Moreover, there is generally an isomorphism between any distortion and some (possibly state- and/or time-varying) preferences in terms of their implied consumer behavior.[15] Given this, our analysis of the wedges does not aim to distinguish between "true" distortions and preferences that result in "as-if" distorted behavior. Nevertheless, we conduct two exercises to explore the extent to which heterogeneity in preferences could explain the distribution of wedges we estimate.

Wedges with Heterogeneous Preference Types. We modify our baseline results to allow for some heterogeneity in preferences, using the three preference types estimated in Aguiar et al. (2025). For each person in our sample, we assign them the discount factor and risk aversion type that *minimizes* the absolute value of their consumption wedge. As such, this exercise is conservative in that it produces a lower bound on the typical wedge size under these preference types.

[15] For example, Gul-Pesendorfer temptation preferences are isomorphic to a particular wealth-dependent discount factor (Kaplan and Violante, 2022).

The median absolute value wedge declines only modestly from 40% to 37%, and the share of over-consumers falls from 51% to 39% (see Appendix Tables E.1 and Appendix Figure E.5).

Survey-Elicited Preferences. In survey waves 2 and 3, we elicited measures of risk aversion and discount rates using questions based on the Falk et al. (2018) Global Economic Preferences Survey and the Andreoni and Sprenger (2012) convex budgets approach. These questions require an assumed "background" level of consumption to convert responses into a β and γ and thus most likely provide an ordinal rather than cardinal measure of preferences. We begin by comparing these elicited preferences with "zero-wedge" implied parameter values, which are calculated as the $\beta_{i,t}$ and $\gamma_{i,t}$ that imply a wedge of zero.[16] If the wedges primarily reflect preference heterogeneity, and these survey elicitations are reasonable reflections of underlying preferences, then we would expect to see a strong relationship between elicited preferences and the zero-wedge preferences. Appendix Figure E.6 shows that the survey measures have very weak relationships with the zero-wedge preferences. We generally find weak, economically insignificant relationships between these measures of preferences, at odds with the wedges reflecting persistent differences in preferences.

5.3 Measurement Error

Subgroups with Milder Measurement Error. We test the sensitivity of the wedge estimates to measurement error in the inputs (beliefs, income, and consumption) by dropping respondents who are likely to exhibit larger measurement error. For beliefs, we sequentially drop subgroups who are less likely to effectively communicate their economic expectations: users who finish the survey in less than six minutes, who report round numbers for inflation expectations (divisible by 5), and who incorrectly respond to at least one of two financial literacy questions. For income, we drop users with any months of zero income or any months of UI income. For consumption, we implement are more involved approach: We re-estimate the wedges sequentially excluding users with any of peer-to-peer transfers, payments to external accounts (e.g., unobserved credit cards), durables purchases, or cash payments exceeding 25% of their nondurable spending.

Appendix Figure E.7 reports the median absolute value wedge and share of over-consumers when wedges are re-estimated excluding these groups that may be subject to higher measurement error. The median absolute value wedge varies from 37–41% and the share of over-consumers

[16] We perform a joint grid search starting at the values of $\beta_{i,t} = 0.92$ and $\gamma_{i,t} = 2.0$ and proceed in intervals of 0.01 and 0.1 respectively, to find the pair of preference parameters that minimizes the Euclidean distance conditional on setting the wedge to zero.

range from 40–54%. The insensitivity of our wedge statistics to the exclusion of these higher measurement error groups suggests that input noise is not significantly influencing our findings.

Repeated Measurement. The second survey wave re-surveyed the users from wave one, which allows us to test whether the distortions a consumer faces tend to be persistent over time and our wedge estimates are capturing their impacts, rather than idiosyncratic measurement error. While distortions are likely to evolve somewhat over time, we expect that consumers facing borrowing constrained or exhibiting present bias in 2022 are more likely to be subject to those same distortions in 2024. In contrast, if the consumption wedges primarily reflect idiosyncratic measurement error, we would not expect wedges estimated for the same consumers to be correlated.

Panel A of Appendix Figure E.8 shows that the wedges are highly correlated across waves ($p < 0.001$) and demonstrate strong test-retest reliability when the inputs are re-measured independently in a different time period.

***k*-Prototype Clustering.** We next attempt to reduce measurement error by aggregating over wedge inputs for clusters of similar respondents. We group respondents using k-prototype clustering based on their similarity across demographics, income, assets, expectations, and other financial characteristics.[17] Under the assumptions that (1) respondents within a cluster have the same data generating process for their wedge inputs (wealth, expectations, etc.) and (2) the noise in measurement of inputs is zero-mean, averaging over these inputs yields a consistent (cluster-specific) estimate of each.

We construct 500 clusters, resulting in 10 observations per cluster on average. Appendix Figure E.9 displays re-estimated histograms. The share of over-consumers increases from 51% to 52%. The median absolute value wedge falls from 40% to 34%. This provides further reassurance that measurement error exerts limited influence on our main results.

Simulating Additional Measurement Error. As a final check, we simulate adding noise to the wedge inputs (uncorrelated or correlated across inputs). The purpose of this exercise is to learn how much mean-zero noise is required to significantly alter the wedge estimates. Wedges are a nonlinear object and it is not obvious if such noise will result in attenuation nor even bias that is monotonic in the variance of the measurement error. We conduct several versions of these

[17] Specifically, we cluster on survey-reported age, pre-tax annual income, savings, observed consumption, inflation expectations (1 and 3 years), income growth expectations, savings rate expectations, borrowing rate expectations, liquid assets, and indicators for gender, race, relationship status, presence of children, college education, and political affiliation. We also cluster on debt as a share of liquid net worth, post-tax income, and implied total debt. We z-score continuous variables so that they exert equal influence in cluster assignment.

simulations that differ in two dimensions: (1) the degree of correlation in noise across inputs and (2) the variance of each inputs' noise. Each simulation entails 500 draws. We detail our procedure in Appendix E.3. Appendix Figure E.10 summarizes our findings. We simulate noise in increments of 0.05 standard deviations (SDs) of each wedge input from 0 to 0.50 and show results for noise that is correlated across inputs with correlation coefficients $\rho = 0, 0.3, 0.9$. With uncorrelated noise, as the SD increases from 0 to 0.5 SDs for each wedge input, the share of over-consumers rises modestly from 51% to 54% and the median absolute value wedge from 40% to 41%. The impact of correlated noise is also modest. Even in the 0.5 SD and 0.9 correlation case, the over-consumer share reaches at most 55% and the median absolute value wedge 43%.

Alternative Income Measure. We measure income in the transactions data as our baseline wedge input, which requires identifying paychecks and UI payments and separating them from other inflows (see Appendix B.1.2). We re-estimate consumption wedges using an independent, complementary measure of post-tax income based on the user's survey response. Reassuringly, the distribution of wedge estimates are remarkably similar (see Appendix Figure E.11).

5.4 Static Wedges

This next analysis speaks to both measurement error and parameter sensitivity concerns, including our term structure assumptions. We measure a distinct but related consumption wedge, which we refer to as a "static" wedge. This wedge is measured using only the one-period-ahead Euler equation and is calculated using only: consumption and one-year-ahead expectations over consumption growth, returns, and inflation. Static wedges do not require measuring income, wealth, nor the term structure of expectations. An important limitation of static wedges is that they measure only the impact of current distortions but not the impact of expected future distortions, whereas the "dynamic" wedge we have so far studied captures both effects. Nonetheless, the sign of the static wedge can still speak to which *current* distortions are affecting consumption and thus provide an additional way to validate our conclusions from studying dynamic wedges, while requiring less data and fewer assumptions. Indeed, if one is only interested in measuring wedges in percentage terms, even consumption data is not necessary to measure static wedges.

Appendix E.4 details our measurement of static wedges. We find a similar share of over-consumers (56%) as indicated by their static wedge, reinforcing our conclusions regarding the prevalence of over-consumption. Because static wedges capture only a portion of the impact of distortions on consumption, their magnitudes are not expected to have the same magnitude as the dynamic wedges. The median absolute value static wedge is 6.1%; when compared to the

typical dynamic wedge of 40%. The much larger dynamic wedge implies that beliefs about future distortions are important determinants of consumption.

5.5 Approximation Bias

Our formula for frictionless consumption relies on a first-order approximation of the Euler equation, resulting in approximation bias. In Bewley (1980) style models, concave utility implies that a first-order approximation biases upwards the measurement of frictionless consumption by omitting the higher-order terms related to the precautionary savings motive.[18] Approximation bias overstates frictionless consumption and leads to a *downward* bias in the wedge estimates. Hence, approximation bias makes our conclusions regarding the size and prevalence of over-consumption conservative.

We examine the magnitude and variability of approximation bias in a set of simple Bewley (1980) style models featuring a variety of distortions. Across the models, the average bias ranges from -5 pp to -24 pp (e.g., we find a wedge of 50 pp instead of 55 pp to 74 pp).[19] While the magnitude of approximation bias is non-negligible, the direction of its effect is well-understood and the bias exhibits limited cross-sectional variation: its standard deviation ranges from 2-4 pp across the models. Importantly, this lack of cross-sectional variation means that approximation bias is unlikely to explain the dispersion in the distribution of wedges. In line with this, we also find that a higher-order moment that is a component of approximation bias—respondents' uncertainty about their inflation forecasts—is uncorrelated with their wedges (Appendix Figure E.13).

We also demonstrate, in a quantitative illustration in Section 6.2 how, even in the presence of approximation bias, researchers can use wedges as a diagnostic by using the same approximate formula when calculating wedges within a model. Intuitively, this is similar to how one can still use a biased estimand, like a local average treatment effect (LATE) to calibrate a model so long as the appropriate model moment is targeted in estimation (i.e., LATE versus ATE, as in Braxton et al., 2025).

[18] For general equilibrium models (i.e, with endogenous income and interest rates), there is also approximation bias stemming from a first-order approximation of the budget constraint. Lemma 2 of Appendix A.1 shows that this bias depends on the covariance of expected consumption growth and the real interest rate. Empirically this covariance is positive, but tends to be small in US data (Campbell and Mankiw, 1989) and in HANK models, where the intertemporal substitution channel is weak (Kaplan et al., 2018) and thus is not likely a significant source of bias.

[19] Differences in the distribution of approximation bias across these models are due purely to differences in each model's ergodic wealth distribution. Approximation bias is the same function of wealth and income in each model.

5.6 External Validity

As we interpret the wedges, it is useful to consider that our sample is middle-income but low-liquidity, and tends to be younger and more female than the overall population. Their demand for EWA may indicate an above-average demand for liquidity; which could reflect a pattern of over-consumption as in Garber et al. (2024) or under-consumption driven by borrowing constraints as in Kluender (2024). In either case, selecting into our sample could indicate that they experience larger distortions than the average consumer. We test whether more frequent and higher dollar EWA usage are associated with positive or negative wedges in Appendix Figure E.14. The relationship between EWA usage and wedges is economically small and statistically insignificant, which suggests it is unlikely that demand for EWA itself indicates selection on positive or negative wedges.

The concentrated age range (inter-quartile range of 31 to 43) suggests that lifecycle differences in the incidence of distortions are also unlikely to explain the cross-sectional variation in wedges we find. To better gauge the degree to which different lifecycle circumstances could drive the distribution of wedges, we estimate wedges separately for the subgroups (1) with and without children and (2) with and without a spouse/partner in Appendix Figure E.15. In both cases, the wedge distributions are similar across the sample splits which suggests that these lifecycle considerations are not responsible for the dispersion of wedges that we find.

Applying our wedge measurement approach to different populations is an exciting direction for future research. While our sample is not representative of the broader US population, they are of particular interest for macro transmission (given their high MPCs), public policy and tax expenditure programs (given their middle incomes and financial precarity), and regulation of financial technology products (given their demand for EWA).

6 Diagnostic Applications of Consumption Wedges

In this final section, we illustrate two diagnostic applications of consumption wedges. First, we examine how wedges correlate with other variables of economic interest. This serves to test whether the wedges reflect individual behavioral tendencies and to identify which distortions best explain the data. Second, we illustrate how to measure wedges in quantitative models and compare the ability of different models to generate distributions of wedges similar to our results.

6.1 Evidence from Wedge Correlates

Hypothetical Spending/Saving Behavior. Our surveys ask respondents how they would adjust their saving behavior if they expected higher inflation. We use this hypothetical to test whether consumers' self-reported behavioral tendencies align with the economic behavior revealed by their consumption wedges. The majority of respondents report that they would "save less" when expected inflation is higher, and Figure 3 Panel A shows that this response is positively correlated with their consumption wedge. That is, over-consumers are more likely to report they would save less. This internal consistency suggests the wedges capture real behavioral tendencies. The survey also asked respondents to rationalize their response to the hypothetical. Among those selecting "save less," a large majority (86%) attributed this dis-saving to an *inability* to reduce spending—a rationale that points most naturally to inertia as a mechanism.

MPCs. If borrowing constraints are the dominant distortion to consumption and mechanism behind high MPCs, we would expect to see the highest MPCs for the most negative consumption wedges. We use the transactions data to estimate consumer-level MPCs based on their nondurable spending response to the March 2021 stimulus payments.[20] These checks provided $1,400 to each eligible individual, with an additional $1,400 for each dependent. We estimate the MPC as the "excess consumption" in the 28 days following the receipt of the stimulus check relative to the 28 days preceding, compared to the same calendar dates in 2022, 2023, and 2024 via a difference-in-difference-style estimator. The measure is as a noisy estimate of each individual's MPC given the limited number of observations and the gap in time between their measurement and our surveys (which occur 1.5 to 3.5 years later).

Figure 3 Panel B shows that consumption wedges are *positively* correlated with MPCs: a 25 pp larger wedge is associated with a 0.4 pp larger MPC (with a p-value of 0.103). This provides some additional validation that wedges are correlated with observed consumption behavior and points away from borrowing constraints as the dominant distortion (and dominant mechanism behind high MPCs).

Financial Distress Proxies. We next examine how wedges vary with financial distress using self-reports from survey questions and observed low savings balances. We focus on self-reported financial anxiety in Figure 3 Panel D and show additional outcomes in Appendix Figure D.3. For all measures we consider, we find a strong, positive relationship between consumption wedges

[20] We observe the relevant time period and can identify the stimulus check for approximately 66% of the analysis sample. The stimulus payment dates range from March 12, 2021 to May 28, 2021.

Figure 3. Relationship Between Consumption Wedges and Nondurable MPCs



Notes: This figure illustrates the relationship between consumption wedges and: (Panel A) an indicator for whether the user reports they would "save less" in response to hypothetically higher inflation; (Panel B) the observed nondurable MPC out of the March 2021 stimulus payments; (Panel C) an indicator for whether the user reports "high" or "very high" financial anxiety; (Panel D) the ratio of consumption commitments (monthly housing and childcare costs) to monthly income; (Panel E) an indicator for whether the user has a mortgage; and (Panel F) the distribution of wedges separately for users with and without a mortgage. Variable definitions are outlined in Appendix B.4. MPCs are observed for 59% of users in our analysis sample. MPCs are trimmed at the 10th and 90th percentiles.

and financial distress. In a model with only a simple borrowing limit (i.e., requiring wealth remain above some number), we would instead expect the most financially distressed to be those with the largest negative wedges. Indeed, we even see that reports of difficulty accessing credit are *positively* correlated with the wedges (Panel C of Appendix Figure D.3). In a model with only a simple borrowing limit, one would expect financial distress would be more severe and pervasive among under-consumers with larger (in magnitude) negative wedges. Thus the lack of such a patterns aligns with our other results, pointing to distortions that generate over-consumption as important determinants of consumption for low-liquidity consumers. Researchers and policymakers concerned with financial distress among low-liquidity consumers may therefore want to direct attention to forces generating over-consumption (equivalently, *under-saving*), including high costs for committed consumption categories (such as housing and childcare) and present bias.

Consumption Commitments. We next examine whether consumption adjustments costs (specifically, consumption commitments) are indeed a good candidate to explain the wedges. In survey waves 2 and 3, we asked consumers to report their monthly housing and childcare costs, 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 to measure the share of income allocated to "committed consumption." Figure 3 Panel D shows a strong positive relationship between wedges and committed consumption as a share of income. A 25 pp larger wedge is associated with a 1.3 pp higher ratio of committed consumption to income. This pattern of "high commitments" consumers exhibiting larger, positive wedges is consistent with consumption commitments being a key force behinds their wedges.

Homeownership. The preceding results point consistently away from borrowing constraints as the dominant distortion. But homeowners may be different. They often hold substantial illiquid wealth while remaining financially constrained (e.g., the wealthy hand-to-mouth of Kaplan and Violante, 2014). Motivated by this scope for potential differences, we compare wedges across homeowners and non-homeowners, proxying for homeownership with whether the respondent has a mortgage (18% of our sample). Panels E and F of Figure 3 show the results. A majority of homeowners have negative wedges and over-consumption is 33% more common among non-homeowners. Further, there is a strong negative correlation between having a mortgage and the consumption wedge. This is consistent with borrowing constraints being the dominant distortion for consumers with substantial illiquid wealth. However, over-consumption is still relatively important even for this population. For non-owners, distortions like present bias and inertia that

generate over-consumption dominate.

6.2 Quantitative Model Illustration

We have documented large consumption wedges, substantial over-consumption, an important role for subjective expectations, and correlational evidence points to present bias and inertia over borrowing constraints. In this final exercise, we assess the ability of simple models featuring these distortions to *quantitatively* match the distribution of wedges.

We solve a set of standard heterogeneous-agent incomplete-markets models where the consumer's problem features at least one distortion. The distortions we include are borrowing constraints, present bias, and consumption adjustment costs. For each model, we calculate the distribution of wedges, measured using the same approximate formula for frictionless consumption (Equation 4). For borrowing constraints, we model a simple lower-bound on agents' net wealth (i.e., $A_{t+1} \geq \bar{A}$). For present bias, we model both naive beta-delta hyperbolic discounting (as in Lee and Maxted, 2023) and temptation preferences (as in Gul and Pesendorfer, 2001; Attanasio et al., 2024), a time-consistent form of present bias where agents are discouraged from saving by anticipating disutility from fighting temptation to consume all of their wealth in the future. We model consumption adjustment costs as a non-pecuniary cost, as in Fuster et al. (2021). We parameterize each model using a similar economic environment—the same income process, utility function, and interest rate—so that differences across the models can be attributed to differences in their distortions. We calibrate distortions to best match the share of over-consumers and median absolute value wedge in the EarnIn sample. Full details on each model and their calibration are in Appendix F.

Table 3 reports the distribution of wedges across models in Panel A, with the empirical distribution in Panel B for comparison. The model best able to generate both a large share of over-consumers and a large median absolute value wedge features present bias (specifically, beta-delta hyperbolic discounting) and a borrowing constraint. This model delivers 39.6% of households over-consuming and a median absolute value wedge of 26.5%, versus 50.6% and 40.1% for our EarnIn sample. It thus can account for 78% of the observed over-consumer share and 66% of the median absolute value wedge. Both forms of present bias (combined with a borrowing constraint) are similarly successful in generating a large share of over-consumers: around 40%. Each produces a similar distribution of negative wedges, but the beta-delta models succeeds in producing larger positive wedges in particular.

As expected, the borrowing constraint model has no positive wedges. The wedges it generates are also quite small in magnitude. The median wedge is -4.2% and the 10th percentile is -12.4%.

Table 3. Wedge Distribution under Various Models

Model	% Pos.	|50th|	Mean	SD	10th	25th	50th	75th	90th
Panel A. Wedges									
Borr. constraint (BC)	0.0	4.2	-6.2	6.1	-12.4	-7.7	-4.2	-2.5	-1.8
Beta-Delta Pref. + BC	39.6	26.5	-11.2	25.5	-39.8	-27.6	-13.8	-1.5	28.4
Temptation + BC	39.8	18.4	-11.4	24.9	-39.8	-27.6	-13.8	-1.5	18.4
Cons. adj. cost (CAC)	16.7	5.1	-5.4	5.5	-12.4	-9.0	-4.9	-1.8	1.4
CAC + SBC	21.2	6.2	-5.7	7.0	-14.8	-10.4	-5.4	-0.8	-3.6
Panel B. Empirical Wedges									
Data	50.6	40.1	15.2	70.8	-62.0	-35.7	0.9	48.3	115.9

Notes: This table reports summary statistics for wedges that arise in different models (Panel A) and for our empirical wedge distribution (Panel B). The first two statistics are the share of consumers with positive wedges and the median absolute value wedge. These statistics are calculated using each model's respective ergodic distribution.

Consumption adjustment cost (CAC) models, with or without a borrowing constraint, do not generate a large share of over-consumers. Without a borrowing constraint, the CAC model features 16.7% of agents over-consuming, with a median absolute value wedge of 5.1%. Adding a borrowing constraint results in a slightly larger share of over-consumers (21.2%) and median absolute value wedge (6.2%). A key reason CAC models struggle to generate large positive wedges is due to how they alter consumers' precautionary savings motive. As the size of adjustment costs increase, two forces compete. A larger adjustment cost makes consumers more willing to tolerate their consumption drifting from their target consumption, generating larger wedges (positive and negative). But rational consumers also lower their target consumption in anticipation of these costs, since a higher target carries greater risk of becoming infeasible and forcing costly adjustment. This precautionary response shifts consumption downward for all consumers, limiting the share of over-consumption arising in a CAC model with rational expectations.[21]

While some models generate more realistic wedges than others, all fall short of the 50.6% over-consumer share and 40.1% median absolute value wedge. A fundamental tension in FIRE models is that frequent and large over-consumption will tend to impoverish a rational agent, reducing their wealth and ability to over-consume. This makes is difficult for such models to generate many agents with significant over-consumption. We identify several promising directions for future research to build on our analysis by incorporating additional features that may improve the ability of models to generate large positive wedges.

We identify three features that may help models close this gap. The first feature is persistently non-rational expectations. Consider a consumer who is persistently overly-pessimistic about her

[21] Similar properties also arise in models with expense shocks, such as the medical expenditure shocks of Bornstein and Indarte (2023). These shocks can create positive wedges, but a similar countervailing force arises where the anticipation of such shocks increases the precautionary savings motive.

permanent income, so that her beliefs imply a low level of frictionless consumption. For a given level of observed consumption, greater pessimism implies a larger wedge. If consumption adjustment costs limit her ability (or if present bias limits her willingness) to select this low level of consumption, she may have a large positive wedge in equilibrium. But because she is persistently wrong about her permanent income, she is slower to run down her wealth or face a binding borrowing constraint. As such, if she remains pessimistic she can have a similarly large positive wedge in subsequent periods. Additional evidence that non-rational expectations may be crucial to match the empirical wedge distribution comes from the covariance decomposition results in Section 4.2, which found a large role for non-rational expectations in explaining the cross-section of consumption.[22]

A second feature that may improve the ability of a CAC model to generate larger wedges is naivety with respect to adjustment frictions. If consumers fail to anticipate adjustment being costly in the future, this eliminates the strengthening of the precautionary savings motive that would otherwise curb over-consumption in a rational CAC model.

Finally, a third feature that may improve the ability of the various models to generate large wedges is to incorporate default. Bankruptcy, which discharges unsecured debts, typically improves credit access soon after filing, despite filers incurring a bankruptcy flag that is visible to creditors on their credit report (Albanesi and Nosal, 2018; Indarte, 2022). The ability to discharge debt and experience improved credit access can reduce how persistently consumers face a binding borrowing constraint. Nonpayment of debt (delinquency) may also give rise to more positive wedges by allowing consumers to increase their consumption in high-debt states.

7 Conclusion

We introduce consumption wedges as a novel approach to measure individual-level distortions to consumption. We leverage a new dataset that links surveyed economic expectations to administrative transactions data. The expectations data allow us to isolate the impact of frictions and behavioral preferences separately from deviations from FIRE, and the transactions data provide a granular measure of consumption for a sample of middle-income, low-liquidity US consumers. There are valuable opportunities for future work measuring consumption wedges in other samples (e.g., at different lifecycle stages or in a broader, more representative population), relating them to other consumer behaviors (such as the four household types of Colarieti et al., 2024), us-

[22] We do not study quantitative models with deviations from FIRE because our beliefs data is almost entirely cross-sectional. This limits our ability to discipline the modeling of the dynamic behavior of beliefs. Panel belief data would enable future research to explore this hypothesis more fully.

ing surveys alone (either by focusing on static wedges or soliciting consumption via a survey), and studying the wedges produced by quantitative structural models.

Our central finding is that consumption wedges are large and heterogeneous: the median consumption wedge is close to zero (0.9%) but the median absolute value wedge is 40.1%. 51% of wedges are positive (over-consumption) while the remaining 49% are negative (under-consumption). Because borrowing constraints can only generate negative wedges, the prevalence of over-consumption challenges the paradigm that treats borrowing constraints as the primary friction shaping low-liquidity consumers' behavior. Our diagnostic analyses—both correlational evidence and quantitative model illustrations—point toward present bias and inertia (e.g., consumption commitments) as promising alternative frictions. The model closest to replicating the empirical wedge distribution quantitatively combines present-biased preferences with a borrowing constraint, yet even this specification falls short. Incorporating non-rational expectations may help close the gap: a covariance decomposition shows that deviations from FIRE explain 25% of the cross-sectional variation of consumption, versus 18% for distortions. If consumers hold persistently pessimistic beliefs, this implies a low level of frictionless consumption, generating large positive wedges for a given level of observed consumption. This pattern is consistent with the recent "vibecession" phenomenon, in which consumer sentiment diverged sharply from traditional economic indicators (Scanlon, 2022).

The prevalence of over-consumption, and its association with financial distress, suggest that policymakers concerned with consumer financial well-being should direct attention beyond credit access to forces that generate excessive spending. Over-consumption indicates that consumers are struggling to save, rather than struggling to borrow. The strong correlation between consumption commitments and over-consumption connects our findings to the broader "affordability crisis" (Glaeser and Gyourko, 2025; U.S. Department of the Treasury, 2021): when housing and childcare consume large shares of income, consumers may be locked into unsustainable spending patterns that drive financial distress. More broadly, policies designed and models specified assuming borrowing constraints are the primary friction may be incomplete; present bias and inertial forces warrant greater consideration.

References

Aguiar, Mark, Mark Bils, and Corina Boar, "Who are the Hand-to-Mouth?," *The Review of Economic Studies*, 2025, pp. 1293–1340.

Aiyagari, S Rao, "Uninsured Idiosyncratic Risk and Aggregate Saving," *The Quarterly Journal of Economics*, 1994, *109* (3), 659–684.

Albanesi, Stefania and Jaromir Nosal, "Insolvency After the 2005 Bankruptcy Reform," Working Paper w24934, National Bureau of Economic Research Aug 2018.

Andreoni, James and Charles Sprenger, "Estimating Time Preferences from Convex Budgets," *American Economic Review*, 2012, *102* (7), 3333–3356.

Attanasio, Orazio P, Agnes Kovacs, and Patrick Moran, "Temptation and Commitment: A Model of Hand-to-Mouth Behavior," *Journal of the European Economic Association*, 2024, *22* (4), 2025–2073.

_ **and Guglielmo Weber**, "Consumption Growth, the Interest Rate and Aggregation," *The Review of Economic Studies*, 1993, *60* (3), 631–649.

_ **and** _ , "Is Consumption Growth Consistent with Intertemporal Optimization? Evidence from the Consumer Expenditure Survey," *Journal of Political Economy*, 1995, *103* (6), 1121–1157.

_ **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, "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," Working Paper w32080, National Bureau of Economic Research Jan 2024.

Berger, David and Joseph Vavra, "Consumption Dynamics During Recessions," *Econometrica*, 2015, *83* (1), 101–154.

_ **, 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.

Bornstein, Gideon, "Entry and Profits in an Aging Economy: The Role of Consumer Inertia," Working Paper w33820, National Bureau of Economic Research May 2025.

_ **and Sasha Indarte**, "The Impact of Social Insurance on Household Debt," Working Paper 2023.

Boutros, Michael, "Windfall Income Shocks with Finite Planning Horizons," *Journal of Financial Economics*, 2026, *176*.

Braxton, Carter, Nisha Chikhale, Kyle Herkenhoff, and Gordon Phillips, "Intergenerational mobility and credit," Technical Report, National Bureau of Economic Research 2025.

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.

Carroll, Christopher D, "How Does Future Income Affect Current Consumption?," *The Quarterly Journal of Economics*, 1994, *109* (1), 111–147.

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.

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*, Jul 2023, *15* (3), 109–152.

Colarieti, Roberto, Pierfrancesco Mei, and Stefanie Stantcheva, "The How and Why of Household Reactions to Income Shocks," Working Paper w32191, National Bureau of Economic Research Mar 2024.

D'Acunto, Francesco, Daniel Hoang, Maritta Paloviita, and Michael Weber, "IQ, Expectations, and Choice," *The Review of Economic Studies*, 2023, *90* (5), 2292–2325.

_ **, Michael Weber, and Xiao Yin**, "Subjective Income Expectations and Household Debt Cycles," Working Paper w32715, National Bureau of Economic Research Jul 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.

_ **,** _ **, 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 Economic 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.

Friedman, Milton, "The Permanent Income Hypothesis," in "A Theory of the Consumption Function," Princeton University Press, 1957, pp. 20–37.

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.

_ , Fiona Greig, Pascal Noel, Daniel M Sullivan, and Joseph Vavra, "Spending and Job-Finding Impacts of Expanded Unemployment Benefits: Evidence From Administrative Micro Data," *American Economic Review*, 2024, *114* (9), 2898–2939.

Garber, Gabriel, Atif Mian, Jacopo Ponticelli, and Amir Sufi, "Consumption Smoothing or Consumption Binging? The Effects of Government-led Consumer Credit Expansion in Brazil," *Journal of Financial Economics*, 2024, *156*.

Glaeser, Edward and Joseph Gyourko, "America's Housing Supply Problem: The Closing of the Suburban Frontier?," *Brookings Papers on Economic Activity*, 2025, *Spring*, 375–425.

Guerrieri, Veronica and Guido Lorenzoni, "Credit Crises, Precautionary Savings, and the Liquidity Trap," *The Quarterly Journal of Economics*, 2017, *132* (3), 1427–1467.

Gul, Faruk and Wolfgang Pesendorfer, "Temptation and Self-Control," *Econometrica*, 2001, *69* (6), 1403–1435.

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*, 1982, *50* (5), 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.

Indarte, Sasha, "The Costs and Benefits of Household Debt Relief," Working Paper 2022.

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," Working Paper w28281, National Bureau of Economic Research Dec 2020.

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.

Kluender, Raymond, "Pay-As-You-Go Insurance: Experimental Evidence on Consumer Demand and Behavior," *The Review of Financial Studies*, 2024, *37* (4), 1118–1148.

Koşar, Gizem, Davide Melcangi, Laura Pilossoph, and David G Wiczer, "Stimulus Through Insurance: The Marginal Propensity to Repay Debt," Working Paper 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," Working Paper w29664, National Bureau of Economic Research Jan 2022.

Lee, Sean Chanwook and Peter Maxted, "Credit Card Borrowing in Heterogeneous-Agent Models: Reconciling Theory and Data," Working Paper 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," Working Paper 2022.

Lusardi, Annamaria, "Permanent Income, Current Income, and Consumption: Evidence from Two Panel Data Sets," *Journal of Business & Economic Statistics*, 1996, *14* (1), 81–90.

Marek, Lynne, Shaun Lucas, and Julia Himmel, "EWA Chases Regulatory Clarity," *Payments Dive*, Oct 2025.

Maxted, Peter, "Present Bias Unconstrained: Consumption, Welfare, and the Present-Bias Dilemma," *The Quarterly Journal of Economics*, 2025, *140* (4), 2963–3013.

_ , David Laibson, and Benjamin Moll, "Present Bias Amplifies the Household Balance-Sheet Channels of Macroeconomic Policy*," *The Quarterly Journal of Economics*, 2025, *140* (1), 691–743.

Patterson, Christina, "The Matching Multiplier and the Amplification of Recessions," *American Economic Review*, 2023, *113* (4), 982–1012.

Scanlon, Kyla, "The Vibecession: The Self-Fulfilling Prophecy," Kyla's Newsletter (Substack) Jun 2022. Accessed January 6, 2026.

Smets, Frank and Rafael Wouters, "Shocks and Frictions in US Business Cycles: A Bayesian DSGE Approach," *American Economic Review*, 2007, *97* (3), 586–606.

Stantcheva, Stefanie, "How to Run Surveys: A Guide to Creating Your Own Identifying Variation and Revealing the Invisible," *Annual Review of Economics*, 2023, *15*, 205–234.

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

U.S. Department of the Treasury, "The Economics of Child Care Supply in the United States," Technical Report, U.S. Department of the Treasury Sep 2021.

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.

Online Appendix

Contents

A Theory Derivations and Extensions **43**
- A.1 Frictionless Consumption Derivation . . . 43
- A.2 Example Distortions and Implied Wedge Propeties . . . 46
- A.3 Model Extensions . . . 48
  - A.3.1 Additional Choices . . . 48
  - A.3.2 Additional Assets . . . 48
  - A.3.3 Durable and Nondurable Goods . . . 48
- A.4 Static Wedges . . . 50

B Data Construction **53**
- B.1 Transactions Data . . . 53
  - B.1.1 Data Structure . . . 53
  - B.1.2 Categorizing Transaction Inflows . . . 53
  - B.1.3 Categorizing Transaction Outflows . . . 54
- B.2 Survey Outreach and Response . . . 57
- B.3 Sample Restrictions . . . 58
- B.4 Variable Measurement . . . 60
  - B.4.1 Consumption, Income, and APCs . . . 60
  - B.4.2 Wealth-to-Income Ratio . . . 61
  - B.4.3 Beliefs . . . 62
  - B.4.4 Marginal Propensity to Consume . . . 63

C Imputations **63**
- C.1 Wealth . . . 63
- C.2 Term Structure of Beliefs . . . 64
  - C.2.1 Inflation Expectations . . . 64
  - C.2.2 Interest Rate Expectations . . . 65
  - C.2.3 Income Growth Expectations . . . 65

D Supplementary Results **67**

E Robustness **69**
- E.1 Sensitivity Analysis . . . 69
- E.2 Preference Heterogeneity . . . 72
- E.3 Measurement Error . . . 74
  - E.3.1 Simulating Additional Measurement Error . . . 76
  - E.3.2 Survey-Based Income Wedge Recalculation . . . 77

E.4 Static versus Dynamic Consumption Wedges 78
E.5 External Validity 79

F Quantitative Model Appendix **81**
F.1 Quantitative Model Specifications 81
F.2 Quantitative Model Calibrations 82

A Theory Derivations and Extensions

A.1 Frictionless Consumption Derivation

We approximate frictionless consumption using a 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 yields $\widetilde{E}_t[f(x)] \approx f[\widetilde{E}_t(x)]$. 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 $\gamma = -\frac{u''(c^\star_{i,t})}{u'(c^\star_{i,t})} c^\star_{i,t}$, is consumer i's inverse IES (coefficient of relative risk aversion) evaluated at frictionless consumption at time t, $c^\star_t$.

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 $\gamma = -\frac{u''(\bar{c})}{u'(\bar{c})}\bar{c}$.

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:*

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

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*.*

Note that when the interest rate term structure is flat (i.e., $\widetilde{E}_{i,t}R_{i,t+k} = R_{i,t}$*), the formula is exact.*

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:

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

□

Remark 1: Preference Heterogeneity

Our notation above does not explicitly allow preference heterogeneity, that is: $\beta = \beta_{i,t}$ and $\gamma = \gamma_{i,t} = \frac{u''(c^{\star}_{i,t})}{u'(c^{\star}_{i,t})}c^{\star}_{i,t}$. However, it is straightforward to rewrite the above equations with preferences that can vary both by individual i and over time t.

A.2 Example Distortions and Implied Wedge Propeties

Example: Borrowing 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: borrowing constraints. These are most often modeled as constant borrowing limits (e.g., Aiyagari, 1994), endogenous borrowing limits (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 borrowing constraints is that they only generate negative consumption wedges. Therefore, a testable implication of borrowing constraints is the sign of the consumption wedges. The presence of positive wedges would indicate that borrowing 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 addi-

tional 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, 2025), all else equal. Hence, present bias creates positive wedges. Similar to borrowing 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, either pecuniary (Chetty and Szeidl, 2007) or non-pecuniary (Bornstein, 2025). 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:

$$\begin{aligned}
&u'\left(\frac{C_{i,t}}{P_t}\right) - P_t\phi'(C_{i,t} - C_{i,t-1}) \\
&= \beta\widetilde{E}_{i,t}\left\{\frac{R_{i,t+1}}{\pi_{t+1}}\left[u'\left(\frac{C_{i,t+1}}{P_{t+1}}\right) - \left(1 - \frac{1}{R_{i,t+1}}\right)P_{t+1}\phi'(C_{i,t+1} - C_{i,t}) + P_{t+1}\beta\phi'(C_{i,t+2} - C_{i,t+1})\right]\right\}
\end{aligned}$$

Under these 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 borrowing 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.

A.3 Model Extensions

A.3.1 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 require 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.

A.3.2 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.

A.3.3 Durable and Nondurable 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 de-

preciation 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, they present measurement challenges that complicate applying our wedge formula. To overcome these challenges, we make assumptions that imply that the expenditure share of nondurable 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 nondurable 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 nondurable 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 nondurables, $n_{i,t}$, 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 nondurable and durable consumption.

Assumption 2: Cobb Douglas Aggregation.

The consumer's notional consumption good is a Cobb Douglas aggregate of nondurable 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 nondurable 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 nondurable share of expenditures.

In our baseline analysis, we recover notional consumption by dividing observed nondurable consumption by the nondurable expenditure share α.

A.4 Static Wedges

Below, we reproduce the approximate Euler equation but instead write it in terms of *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* notion 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 $C_{i,t}^{\star,\text{static}}$ would include the effect of this distortion. But because $C_{i,t}^{\star,\text{static}}$ 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

$$\nu_{i,t} = C_{i,t} - C_{i,t}^{\star,\text{static}} = C_{i,t}\left[1 - \frac{\widetilde{E}_{i,t}G_{i,t+1}^{C}}{\widetilde{E}_{i,t}\pi_{t+1}}\left(\beta\frac{\widetilde{E}_{i,t}R_{i,t+1}}{\widetilde{E}_{i,t}\pi_{t+1}}\right)^{-1/\gamma}\right] \tag{8}$$

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:

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

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 expression and applying the law of iterated expectations, we obtain:

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

□

B Data Construction

B.1 Transactions Data

B.1.1 Data Structure

We receive anonymized data from EarnIn that covers user-level information, detailed bank account and debt 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 unless they terminate their account with EarnIn or disconnect their bank account. Among the users in our analysis sample, 98% had received a cashout from EarnIn's Earned Wage Access product at least once in the 12 months preceding the survey.

User-Level Data. We receive de-identified 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 the earnings are from paid work or unemployment benefits. Earnings are observed net of taxes and payroll deductions when deposited into users' linked bank accounts; we do not receive information on gross pay or withheld amounts.

B.1.2 Categorizing Transaction Inflows

We define our measure of income as the sum of post-tax labor earnings and unemployment insurance (UI). To use EarnIn's Earned Wage Access product, users submit their employment information to EarnIn and EarnIn identifies and validates the payroll inflows associated with that user-submitted job. The dataset of EarnIn-validated earnings is the starting point for our income measure, but it is not comprehensive of all income. Users with multiple jobs can choose to submit employment information for one job and use Earned Wage Access for only one of their sources

of income. Users may also choose not to update their employment information during time periods when they are not actively using the Earned Wage Access product. We therefore supplement EarnIn-validated earnings by identifying additional sources of income using transactions memos and categories. The procedure for constructing this comprehensive measure of income from transaction data is described below.

We start by cleaning transaction memos to remove any non-alphabetic characters. This helps aggregate transactions from the same source, even where memos include dates of payment or a reference number.

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 associated with 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. Each source of earnings (up to three) is stored as a unique weekly earnings variable. For example, if a user has two sources of earnings within a week, the first two earnings variables reflect 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. To avoid the false positives that we observe with these tags, 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.

B.1.3 Categorizing Transaction Outflows

We apply an outflows categorization algorithm that separates durables and nondurables spending from other types of outflows, including debt payments and fees (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 to overcome three challenges inherent to transactions data. 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, such as "Purchase," "Shopping," or "Transfer." Finally, we observe some misclassifications in Plaid's categorization that we must reallocate based on the transaction memos.

To address the first challenge, we proportionally allocate 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 e-commerce revenue in 2020 (Bledsoe, 2024). Appendix Table B.1 summarizes these allocations.

To address the second (broad/ambiguous categories) and third (potential misclassified categories) challenges, 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 allocate transactions these catch-all categories to more specific ones and to re-categorize misclassified transactions in other categories.

In addition to these category allocation decisions, our data also face other standard limitations inherent to bank account transactions datasets. Transactions are only observable and categorizable to the extent that they are made through linked bank accounts. Cash withdrawals and external transfers are observed in the data, but we cannot directly observe where they are spent. We exclude external transfers and treat cash as spending on nondurables. Mortgage and rent payments are unobservable or difficult to identify for many users due to being paid by check, peer-to-peer transfers, or other transactions with uninformative memos.

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

Table B.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 e-commerce revenue in 2020 (Bledsoe, 2024).

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

B.2 Survey Outreach and Response

EarnIn users were eligible to be invited to complete the survey if they met minimum data quality thresholds based on their transactions data over the 12 months preceding each survey. We list the sampling criteria for each survey wave below.

- Wave 1 (September 2022)
  - At least one earnings observation between September 2021 and August 2022
  - Non-missing biweekly balances data between September 2021 through August 2022
  - First recorded transaction before September 1, 2021
  - Recently active data (i.e., latest recorded transaction after August 15, 2022)
  - At least 5 outflows per month between September 2021 and August 2022
- Wave 2 (July 2024; Wave 1 Follow-Up Survey)
  - Completed the wave 1 survey and still in the EarnIn database as of June 2024
  - Met minimum wave 1 survey attention requirements (i.e., at least 3.5 minutes to complete the survey, consistent responses about their debt across questions)
  - Non-missing weekly balances data for at least nine months between June 2023 and May 2024
  - At least 20 outflows per month each month between June 2023 and May 2024
  - Appear to use account for consumption (i.e., $\frac{\text{Consumption}}{\text{Outflows - Internal Transfers}} \geq 20\%$ for at least nine months between June 2023 and May 2024)
  - Reasonable balance of inflows and outflows ($\frac{\text{Outflows}}{\text{Inflows}} \in [50\%, 150\%]$) for at least nine 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 (these eligible users were surveyed just 5 months prior in wave 2)
  - Non-missing earnings data at least once between October 2023 and September 2024
  - Non-missing weekly balances data for at least nine months between October 2023 and September 2024
  - First recorded transaction before September 1, 2021
  - Appear to use account for consumption (i.e., $\frac{\text{Consumption}}{\text{Outflows - Internal Transfers}} \geq 20\%$ for at least nine months between October 2023 and September 2024)
  - Reasonable balance of inflows and outflows ($\frac{\text{Outflows}}{\text{Inflows}} \in [50\%, 150\%]$) for at least nine months between October 2023 and September 2024
  - At least 20 outflows per month each month 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)

Appendix Table B.2 compares key summary statistics between users eligible for the survey and respondents. EarnIn sent invitations to 240,000 users in wave 1, closing the survey after receiving 10,090 responses and exhausting the incentive budget. EarnIn sent invitations to a randomly drawn subset of the 578,123 eligible users until reaching the target number of respondents. EarnIn sent invitations to 4,652 users in wave 2 and 218,615 in wave 3, receiving 875 and 4,888 initial responses, respectively. For wave 3, EarnIn sent the survey invitations to the subset of users who were eligible for email marketing outreach. We further process the survey responses to weed out low-effort responses and incomplete transactions data, which we describe in Appendix B.3. Users eligible for the survey look very similar on all observable financial outcomes and there appears to be minimal selection into the survey on these variables. Respondents are, however, more likely to be female than the full eligible sample. Overall, the survey achieved an aggregate response rate of 3.4%—4.2% for wave 1, 18.8% for wave 2, and 2.2% for wave 3.

B.3 Sample Restrictions

Measuring consumption wedges requires high-fidelity data across all of the inputs (income, assets, consumption, and beliefs). We impose sample restrictions that isolate the survey respondents for whom we have both high-quality survey and transactions data. We describe each of the five sample restrictions we make to arrive at our analysis sample from the original 15,853 total combined survey responses.

- First, we lose a very small number of survey respondents who deleted their EarnIn accounts or de-linked their bank accounts when we merge survey respondents to the EarnIn database. This reduces the number of users from 15,853 to 15,690.

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

Row	W1 Resp.	W1 Samp.	W2 Resp.	W2 Samp.	W3 Resp.	W3 Samp.	All Resp.	All Samp.
Monthly Inflows	6,316	6,107	7,030	7,075	6,733	6,610	6,452	6,286
Monthly Outflows	6,245	6,039	6,865	6,903	6,598	6,530	6,360	6,213
Balance (Checking)	517	526	618	717	544	658	525	573
Balance (Savings)	296	353	277	386	290	425	294	377
Female (%)	67.2	52.4	73.3	69.0	72.3	52.7	68.9	52.5
N (Users)	10,090	578,123	875	4,652	4,888	318,710	14,978	896,833

Notes: This table reports statistics for the survey respondents and the full sampling frame from which they were 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). For each, we report average monthly outflows, inflows, checking and savings balances, and gender. Prior to computing the means of the monthly variables, observations were winsorized at the 1st and 99th percentiles.

- Second, we drop users who completed the survey implausibly quickly or who offered inconsistent responses across questions as described below. This reduces the number of users from 15,960 to 15,243.
  - Survey duration at least 3.5 minutes (approximately the 5th 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 trim users with expectations outside of the percentiles listed below. Outlier expectations could indicate weaker understanding or attention to the survey questions and can produce implausible implied consumption-savings behavior. We drop users based on percentile cutoffs calculated separately for each survey wave. This reduces the number of users from 15,243 to 12,525.
  - $E_t G^Y_{t+1}$ (P3-P97)
  - $E_t \pi_{t+1}$ (P3-P97)
  - $E_t \pi_{t+3}$ (P3-P97)
  - $E_t R^S_{t+1}$ (P1-P97)
  - $E_t R^D_{t+1}$ (P1-P97)
- Fourth, we apply the transactions data quality restrictions designed to identify and drop users who may not primarily consume through the bank accounts connected to EarnIn. We make these sample restrictions to ensure we observe a complete picture of users' consumption to the extent possible. 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 is our most consequential set of sample restrictions and reduces the sample from 12,525 to 7,891.

- Sufficient transaction activity: 20+ outflows per month for all months
- Non-missing balances: Non-missing balances each week for at least 75% of months
- Categorizable spending: $\frac{\text{Consumption}}{\text{Outflows - Internal Transfers}} \geq 20\%$ for at least 75% of months
- Balanced 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

- Fifth, we trim the remaining variables listed below to minimize the likelihood that we mismeasure any inputs into wedge measurement. To calculate each threshold, we restrict to the 12-month pre-survey period, deflate to September 2022 prices, and collapse to the annual level separately for each survey wave. This set of sample restrictions reduces the number of users from 7,891 to 6,238.
  - 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)
- Sixth, we restrict the sample to users for whom the wedge can be successfully computed. We drop users with missing inputs required for the calculation, users whose inputs violate necessary mathematical constraints (e.g., negative consumption or division by zero in valuation formulas), and users for whom their beliefs imply divergent sums in the wedge formula. This final set of sample restrictions reduces the number of users from 6,238 to 5,263.
  - Requiring valid inputs: Non-missing values for APC, wealth-to-income, and expectations (income growth, inflation, and returns)
  - Requiring mathematical bounds: $APC > 0$, nonzero denominators, and convergent geometric series
  - Trimming outlier wedges (P1-P95; the asymmetric thresholds are applied because the wedges are bounded below by -100% but are not bounded not above)

B.4 Variable Measurement

B.4.1 Consumption, Income, and APCs

Consumption: The transactions data reflect spending, which does not generally correspond to consumption for *durable* goods. We measure notional (i.e., total) consumption by scaling observed nondurables spending from the transactions data by our assumed nondurables share of spending $\alpha = 79.37\%$, which we obtain from an estimate in Beraja and Zorzi (2024) that uses Consumer

Table B.3. Successive Survey and Transactions Sample Restrictions

	Wave 1	Wave 2	Wave 3	Total
Finished responses	10,090	875	4,888	15,853
Can merge with transactions	9,998	873	4,819	15,690
Debt amount contradiction	9,860	863	4,777	15,500
Survey length too short	9,668	857	4,718	15,243
E(Inflation, 1YR) (P3-P97)	9,131	793	4,413	14,337
E(Inflation, 3YR) (P3-P97)	8,831	778	4,311	13,920
E(Income Growth) (P3-P97)	8,379	742	4,103	13,224
E(Savings Rate) (P1-P97)	8,134	732	3,991	12,857
E(Borrowing Rate) (P1-P97)	7,932	709	3,884	12,525
At least 20 outflows per month	5,224	592	3,776	9,592
Consumption $\geq$ 20%, at least 75% of months	4,942	545	3,253	8,740
Outflows between 50–150% inflows, at least 75% of months	4,786	511	3,072	8,369
Less than 1% of outflows with uninformative memos	4,659	496	3,050	8,205
Non-missing bank balances, at least 75% of months	4,350	494	3,047	7,891
Trimmed non-durables spending (P1–P99)	4,264	486	2,987	7,737
Trimmed income (P1–P99)	4,195	479	2,937	7,611
Trimmed non-durables APC (P2.5–P97.5)	4,028	456	2,818	7,302
Trimmed net worth to income ratio (P5–P95)	3,658	413	2,560	6,631
Trimmed expected levered rate (P3–P97)	3,447	389	2,402	6,238
Wedge computation restrictions	3,091	359	2,145	5,595
Wedge trimming restrictions	2,911	346	2,006	5,263

Notes: This table shows the resulting number of unique individuals after successive sample restrictions. Because Wave 2 is comprised of a subset of respondents to Wave 1, the "Total" column represents the number of user $\times$ survey wave observations. Wedge trimming restrictions and wedge computation restrictions correspond to that of the dynamic wedge. Of the 346 Wave 2 respondents, 235 also met the Wave 1 sample restrictions.

Expenditure Survey data. Under the assumption that notional consumption is a Cobb Douglas aggregate of durables and nondurables, this procedure identifies notional consumption, as discussed in Appendix A.3.3. 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.

B.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.[23] Liquid assets are reported in the following bins: \$0 to \$499, \$500 to \$999, \$1,000

[23] 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 C.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.

B.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 C.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 C.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 C.2.2.

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

B.4.4 Marginal Propensity to Consume

Observed MPC We measure individual-level MPCs based on consumers' nondurable 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 66% 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:

$$MPC_i = \frac{1}{StimulusAmount_i} \times (\Delta Spend_i^{2021} - \frac{\Delta Spend_i^{2022} + \Delta Spend_i^{2023}}{2}) \tag{9}$$

where

$$\Delta Spend_i^t = Spend_i^{Post,t} - Spend_i^{Pre,t} \tag{10}$$

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.

C Imputations

C.1 Wealth

In the survey, we ask respondents to report their liquid assets and debt by select one of multiple binned ranges. Because these variables are discrete, censored at the top, and do not capture illiquid assets, we impute total assets and total debt. To do so, we train an XGBoost model to predict total debt and total assets (separately), where our model is trained using Survey of Consumer Finances (SCF) data. 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 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 observed or calculated in the SCF (income range, has mortgage, has auto loan, has credit card debt, has student debt, age, gender, marital status, number of children, race, and education). We randomly partition the SCF sample into a training set (80%) and a test set (20%). Hyperparameters are selected using 5-fold

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

cross-validation within the training set. After hyperparameter selection, we obtain the final model by re-estimating the model on the full training set. Table C.1 reports out-of-sample performance in terms of the median absolute error of our model within the test set.

Table C.1. Median Absolute Error of Wealth Imputations

Variable	Median Absolute Error
Debt	\$868
Assets	\$9,586
Net Worth	\$9,503

Notes: This table reports the median absolute error of the wealth component prediction model described in this section. These statistics are calculated for a 20% out-of-sample subsample of the SCF (i.e., a subsample not used for estimation nor hyperparameter tuning). The statistics are calculated using weights corresponding to the distribution of liquid wealth bins of the EarnIn sample.

C.2 Term Structure of Beliefs

To calculate dynamic wedges, we need subjective beliefs over an infinite horizon. We measure the following subjective beliefs in all three survey waves:

- 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)

Frictionless consumption in the dynamic wedge formula depends on the entire term structure of beliefs over all future periods. Because we do not observe the full term structure, we impute the remaining term structure. We detail our imputation approaches here.

C.2.1 Inflation Expectations

For inflation expectations, we start with two points on the term structure: one-year-ahead and three-year-ahead expected inflation rates. We begin by imputing the two-year-ahead ($t + 2$) expectation as the average of the $t + 1$ and $t + 3$ expectation for each observation i, that is:

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

For the remaining term structure ($t + k > t + 3$), our imputation is disciplined by two factors. The first is the term structure of expected inflation for $k = 1, ..., 30$ (i.e., 30-year-ahead beliefs), which is produced by the Federal Reserve Bank of Cleveland and derived from inflation swaps data. These can be thought of as capturing the inflation expectations of (sophisticated) investors

participating in these markets. This dataset is updated monthly; for each survey wave, our imputation uses the term structure at the time of the survey.

The second factor is an observation-specific bias that helps capture, for example, a persistent tendency to over-estimate inflation. For each observation, we calculate an average bias (relative to the swaps-implied beliefs) based on the one- and three-year-ahead inflation expectations:

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

For the $t+4$ to $t+30$ expectations, we impute them as:

$$\widetilde{E}_{i,t}\pi_{t+k} = Bias_{i,t} \times \widetilde{E}_{s,t}\pi_{t+k} \quad \text{for } k = 4, 5, ..., 30$$

where $\widetilde{E}_{s,t}$ is the swap-implied point on the term structure. Intuitively, our procedure rescales the swaps-implied term structure to reflect the typical over- or under-estimation of inflation that respondent i exhibits.

For horizons beyond $t+30$, we assume that the term structure becomes flat. That is:

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

C.2.2 Interest Rate Expectations

We assume that levered interest rate expectations have a flat term structure.

C.2.3 Income Growth Expectations

Our imputation of income expectations 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 $t=1$ to $t=2$ as of time $t=0$ (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 the same individual 12 months apart in time. Our imputation procedure assumes $\widetilde{E}_{i,t}G^{Y}_{i,t+1,t+2} = \widetilde{E}_{i,t+1}G^{Y}_{i,t+1,t+2}$. For example, if we observe responses in January 2025 and December 2025 for the same individual, we treat the former as their one-year-ahead expectation and the latter as their two-year-ahead expectation as of January 2026. This allows us to observe one-year- and two-year-ahead income growth expectations within the same individual, that is, a $\widetilde{E}_{i,t}G^{Y}_{i,t,t+1}$ and $\widetilde{E}_{i,t+1}G^{Y}_{i,t+1,t+2}$.

We find 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 C.1 below). We estimate this quadratic relation and then, for the EarnIn sample, calculate expected income growth

through period $t+2$ as:

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

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 C.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 C.1. Term Structure of Income Growth Expectations

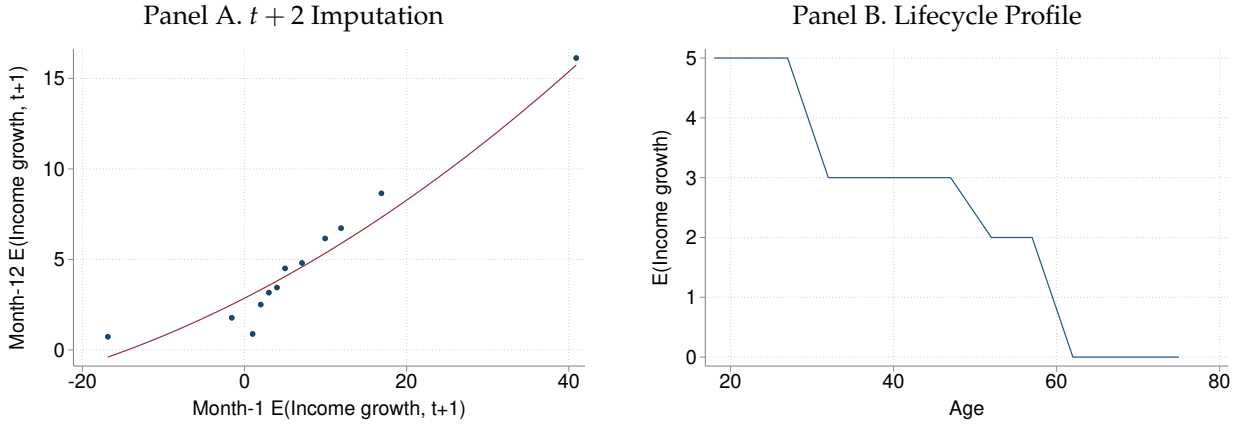

Notes: These graphs visualize the patterns used to discipline the imputation of the term structure of income growth. Panel A reports data from the Survey of Consumer Expectations (SCE), which measures one-year- and two-year-ahead expectations in a 12-month rotating panel. We approximate this relationship with a quadratic function. Panel B shows the lifecycle profile of income growth expectations, calculated using the Michigan Survey of Consumers. We take the median expectation within 5-year age bins and then interpolate across bins to smooth expectations.

D Supplementary Results

Figure D.1. Consumption Wedges in Dollars

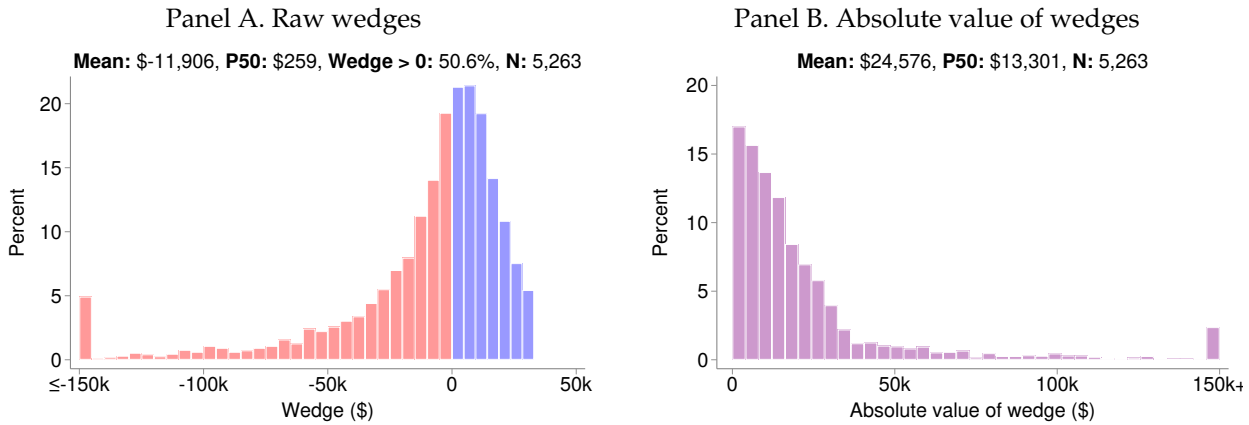

Notes: The graphs show the distribution of consumption wedges (left) and their absolute values (right) in terms of dollars. Colored bins represent the distribution of consumption wedges in dollar terms.

Figure D.2. Distribution of Wedges with FIRE Beliefs

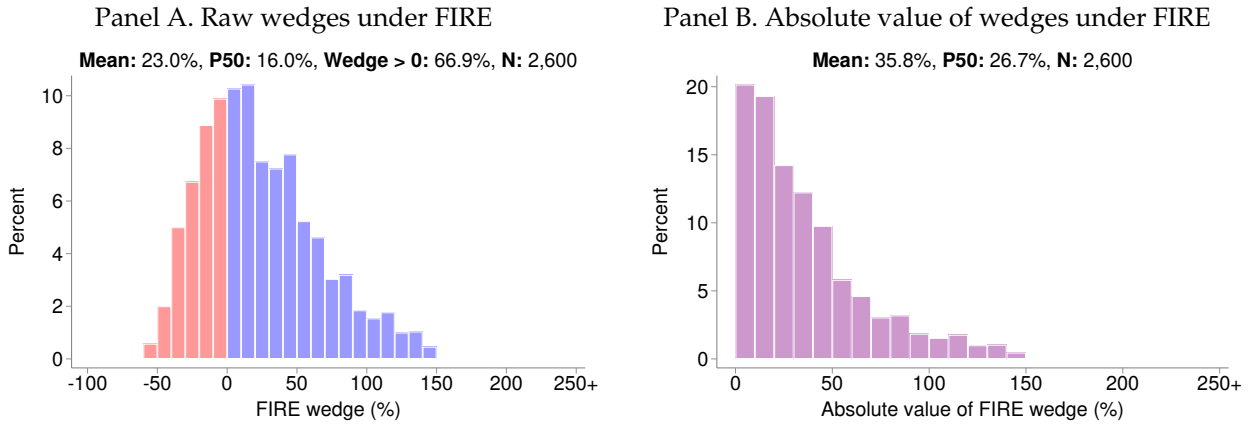

Notes: The graphs show the distribution of wedges calculated assuming consumers form FIRE. Because we do not have realized income data for our two 2024 waves, we restrict to wave 1 users. All wedges in the graphs above are expressed as a percent deviation from frictionless FIRE consumption.

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

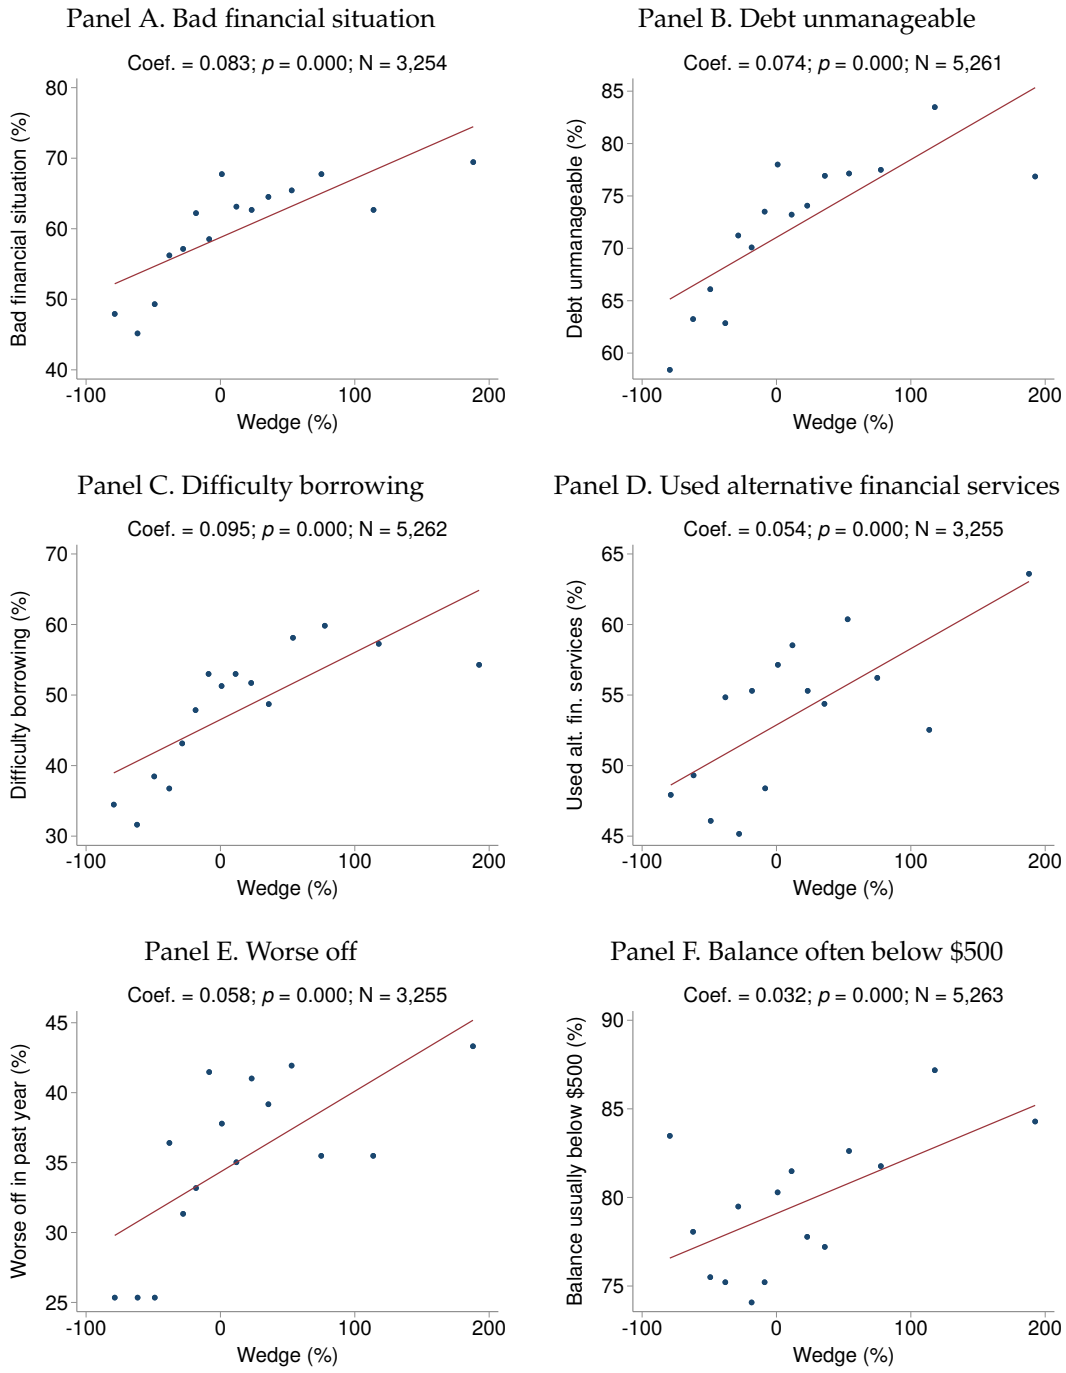

Notes: The graphs illustrate the relationship between consumption wedges and six indicators of financial distress: (1) whether the user reports "just getting by" or "finding it difficult to get by;" (2) whether the user reports having "a bit more" or "far more" debt than is manageable; (3) whether the user reports difficulty borrowing "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.

E Robustness

E.1 Sensitivity Analysis

Figure E.1. Robustness of Consumption Wedges to Reflect Beliefs about Default

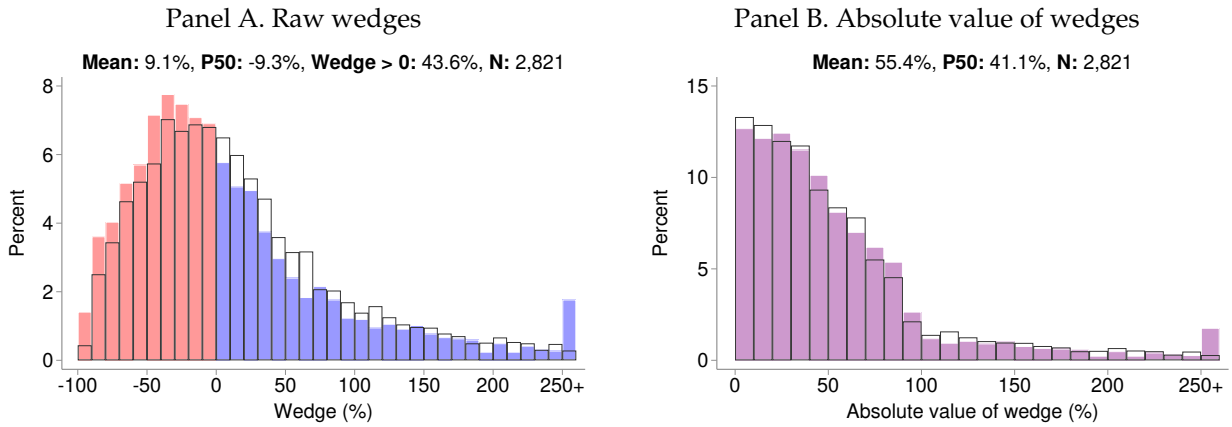

Notes: The graphs show the distribution of consumption wedges (left) and their absolute values (right). Colored bars represent the distribution of consumption wedges where we recalculate $R^D_{i,t}$, the gross rate on debt, such that $R^{D*}_{i,t} = [(1 + R^D_{i,t}) \times \theta] - 1$ where $\theta = 0.9$. This corresponds to a 10% reduction in the gross rate, to reflect a default and/or nonpayment on 10% of all future debt payments in every period. This is a relatively extreme assumption and therefore provides a plausible bound on the extent to which our results are robust to excluding default expectations in the calculation of the wedge. The transparent bars represent the original wedge distribution. Wedges are reported here as the percent deviation of observed consumption from frictionless consumption.

Figure E.2. Sensitivity of Dynamic Wedges to Preference Parameters

Panel A. Sensitivity of Dynamic Wedges to Beta

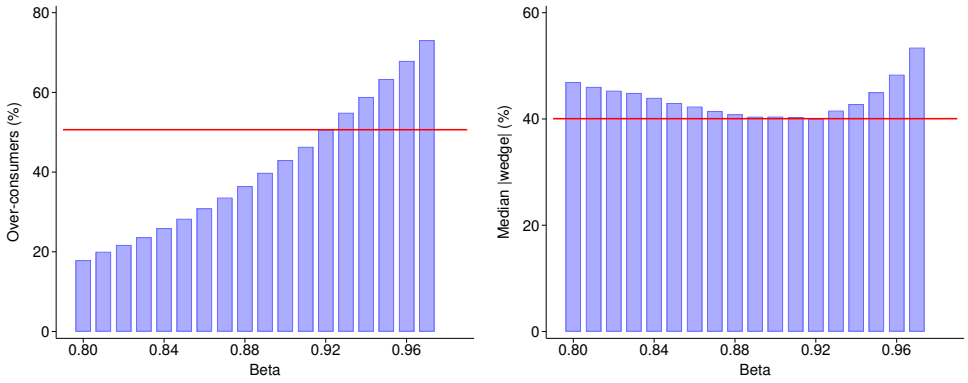

Panel B. Sensitivity of Dynamic Wedges to Gamma

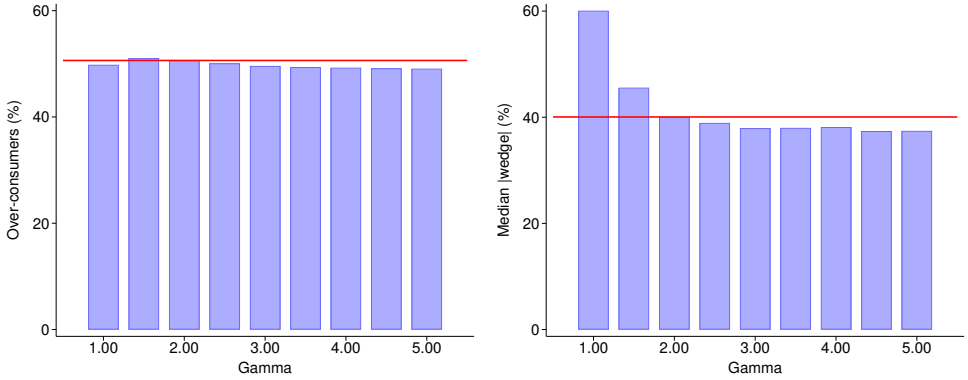

Panel C. Sensitivity of Dynamic Wedges to Nondurables Share of Spending

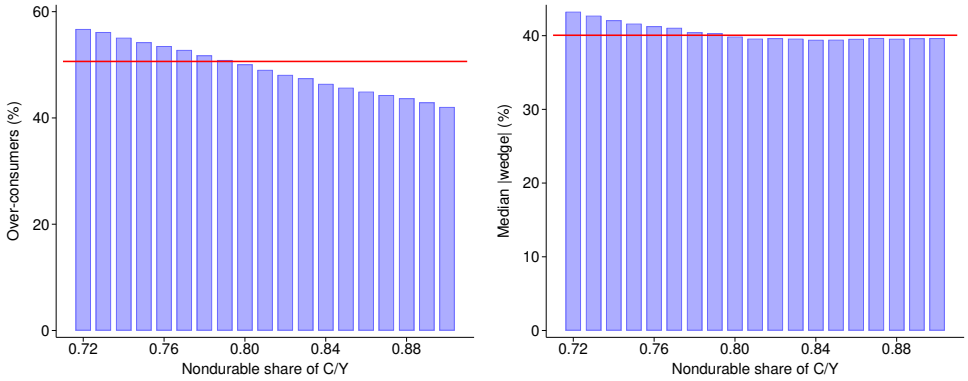

Notes: This figure reports the sensitivity analysis varying three wedge parameters: β (Panel A), γ (Panel B), and the nondurable share of spending (Panel C). Recall that the latter is used to scale up nondurable consumption to recover notional consumption. The left graphs report the sensitivity of the over-consumer share to these parameters. The right graphs report the sensitivity of median absolute value wedge. We vary one parameter at a time in each analysis.

Figure E.3. Sensitivity of Consumption Wedge to a Constant Term Structure of Beliefs

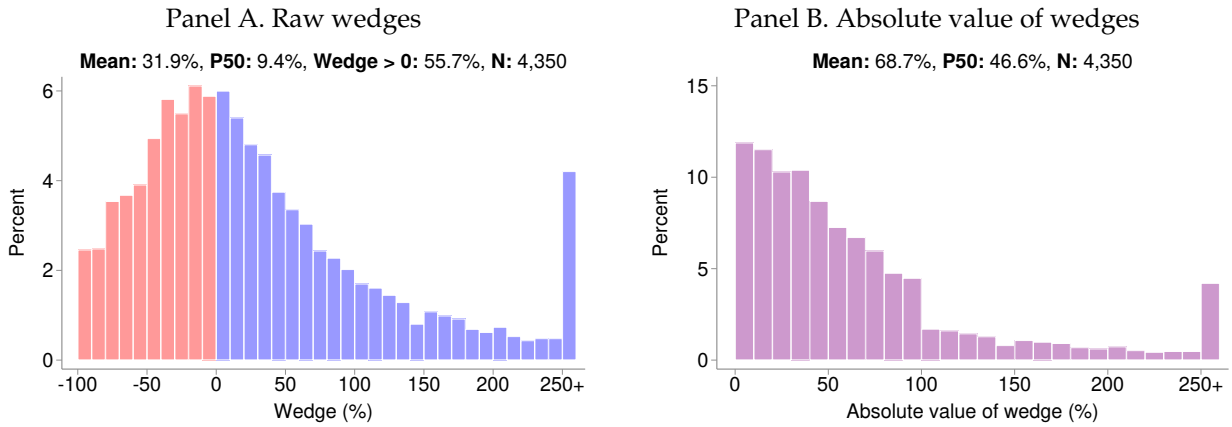

Notes: The graphs show the distribution of dynamic consumption wedges (left) and their absolute values (right), assuming users have a constant term structure of beliefs. Wedges are reported here as the percent deviation of observed consumption from frictionless consumption.

Figure E.4. Sensitivity of Consumption Wedge to the Number of Months of Data

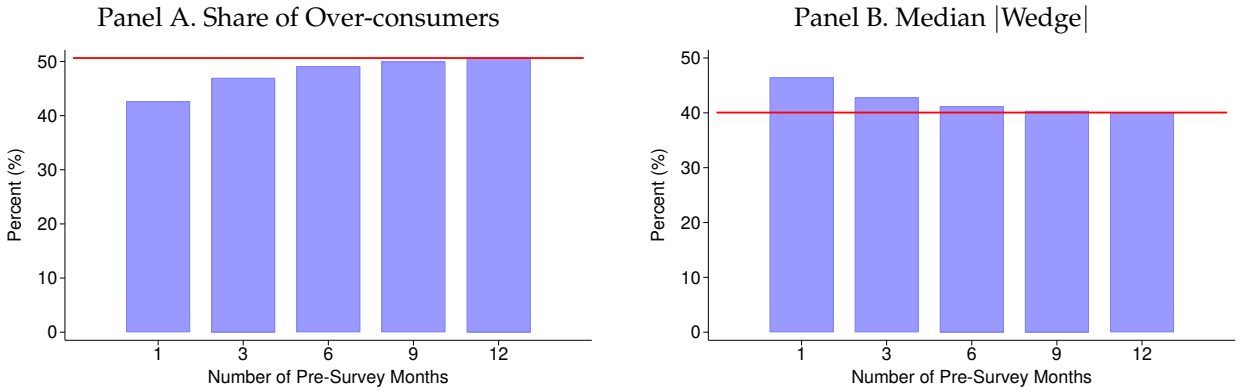

Notes: This figure reports the sensitivity of the over-consumer share (Panel A) and median absolute value wedge (Panel B) to the number of pre-survey months used to measure consumption when calculating the wedge. Our baseline uses 12 months. For each horizon, we annualize consumption to keep the units comparable to the same (annual) measure of frictionless consumption.

E.2 Preference Heterogeneity

Table E.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%	41.3%
II	$\beta = 0.94, \gamma = 1.05$	33.7%	32.8%
III	$\beta = 0.72, \gamma = 0.35$	21.6%	25.9%

Notes: The table lists the three preference types outlined in Table 11 of Aguiar et al. (2025) in columns 1-2. Note that γ in our notation corresponds to $1/\sigma$ in their notation. Column (3) presents the distribution of preference types within their sample. Column (4) reports 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 E.5. Distribution of Dynamic Consumption Wedges Using Wedge-Minimizing Preference Types

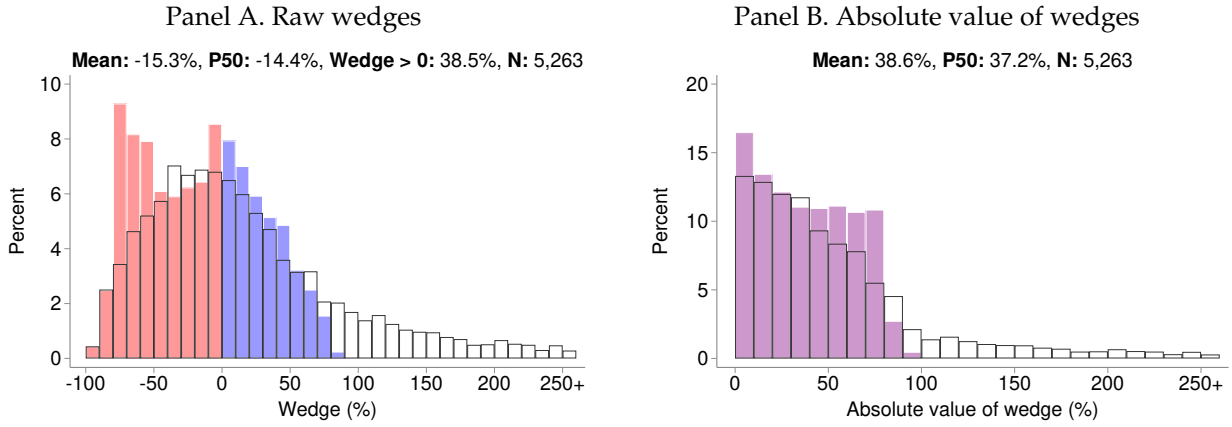

Notes: The graphs show the distribution of consumption wedges (left) and their absolute values (right), assuming users have one of the three preference types outlined in Table 11 of Aguiar et al. (2025). We assign users to preference types based on which type minimizes the absolute value of their wedge (see Table E.1). Wedges are reported here as the percent deviation of observed consumption from frictionless consumption. The distribution of wedges under our baseline's homogeneous preferences in transparent bars.

Figure E.6. Correlations with Survey-Elicited Preferences

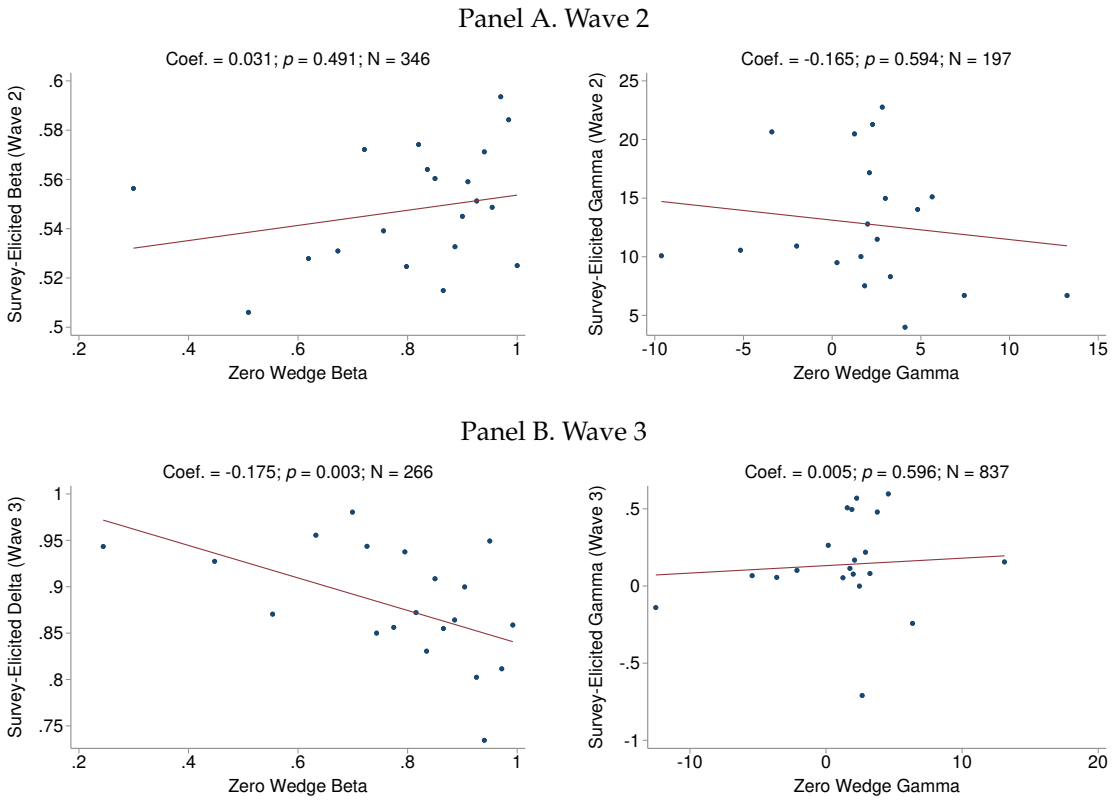

Notes: This figure reports correlations between preferences solicited via survey questions and zero-wedge-implied preferences for beta (left) and gamma (right). Results are reported separately for waves 2 and 3 in Panels A and B (respectively). Wave 3's exponential discount factor is denoted here by "delta" following the traditional notation of beta-delta discounting; thus it corresponds to β in our main text's notation.

E.3 Measurement Error

Figure E.7. Sensitivity of Consumption Wedge to Dropping Users

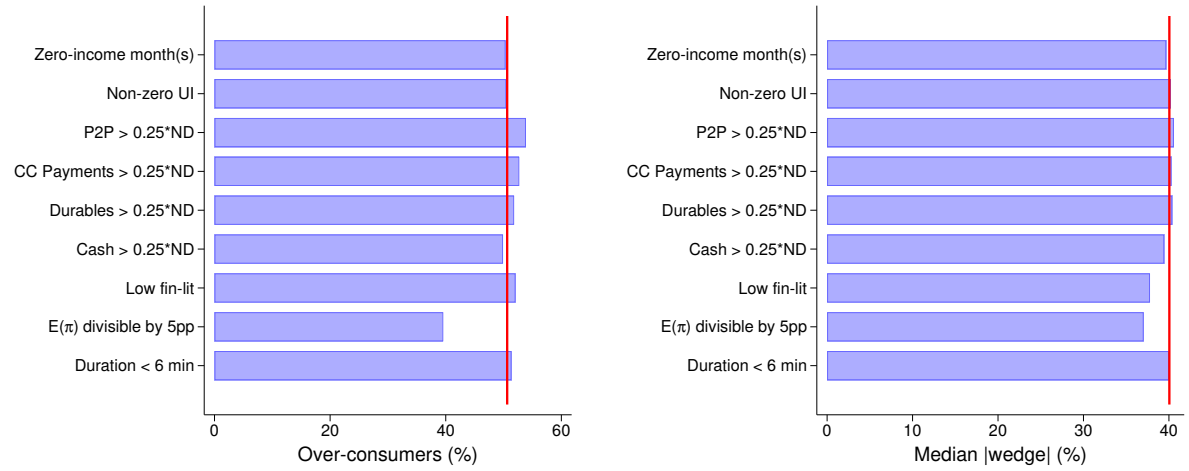

Notes: Figure presents the sensitivity of the over-consumer share (Panel A) and median absolute value wedge (Panel B) to dropping users with plausibly high measurement error. For reference, our baseline results are shown in red. Sample sizes associated with dropping users in each of the categories are as follows: any zero-income month(s) ($N = 4,615$); non-zero UI ($N = 5,077$); P2P $> 0.25 * ND$ ($N = 4,710$); credit card payments $> 0.25 * ND$ (has)$N = 2,260$); durables $> 0.25 * ND$ ($N = 2,333$); cash $> 0.25 * ND$ ($N = 3,824$); low financial literacy ($N = 4,808$); expectations divisible by 5 ($N = 4,584$); and survey durations < 6 minutes ($N = 3,211$).

Figure E.8. Within-Respondent Wedge Correlations

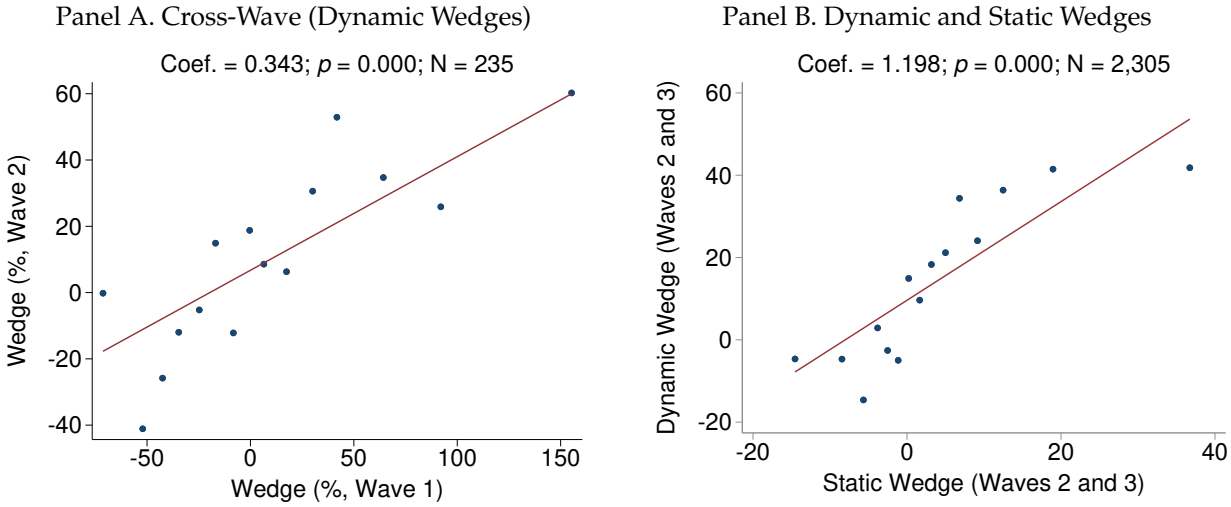

Notes: The figure binscatters wedges measured in wave 2 against the same user's wave 1 wedge in Panel A, while Panel B binscatters static wedges in waves 2 and 3 against the same user's wave 2 or 3 dynamic wedge. Note that the sample for Panel A is comprised of all repeat respondents whose wave 1 and 2 wedges met all trimming and convergence restrictions while Panel B is comprised of all individuals whose wave 2 and 3 wedges met all trimming and convergence restrictions. Dynamic wedges are trimmed at the 1st and 95th percentiles, while static wedges are trimmed at the 1st and 99th percentiles.

Table E.2. Crosswave Transition Matrix

	Wave 2 Under-consumer	Wave 2 Over-consumer
Wave 1 Under-consumer	68.8%	31.2%
Wave 1 Over-consumer	46.3%	53.7%

Notes: This table details the transition matrix between under- and over-consumption for the subset of the $N = 235$ repeat responders in both waves 1 and 2. It reports the share of respondents whose under- or over-consumer status remains the same or changes between the two waves, which take place approximately two years apart.

Figure E.9. Robustness of Consumption Wedges to Clustering Users

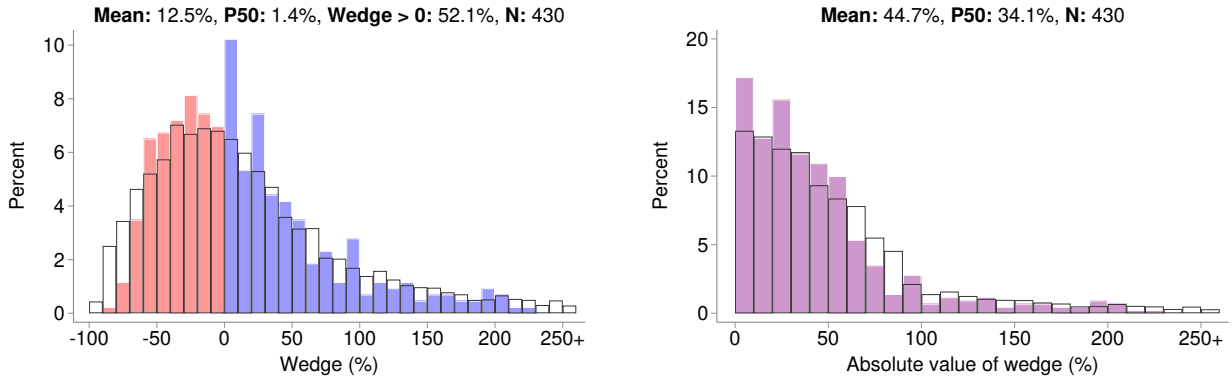

Notes: This figure reports the wedge distribution obtained when measuring wedges after clustering users and taking the within-cluster median of each wedge input before calculating the wedges. We use the *k*-prototypes clustering algorithm with 500 bins described in the main text. The baseline distribution of user-level dynamic wedges (from Figure 2) is overlaid with hollow black bars.

E.3.1 Simulating Additional Measurement Error

The effect of zero-mean measurement error in wedge inputs on the over-consumer share and median absolute value wedge is not a priori obvious. This is because the wedge is a nonlinear object. To gauge the potential impact of such measurement error, we conduct a series of simulations.

Figure E.10 summarizes the results of simulation analysis. Each dot on the graphs below corresponds to a set of 500 simulations. What we vary across the simulations is the standard deviation of the measurement error (along the x-axis of each plot) and the degree of correlation in the measurement error (the different colored points).

We simulate noise as follows. For each wedge input x, we specify a $\sigma_x \in \{0, 0.05, ..., 0.50\} \times SD(x)$. This corresponds to the standard deviation of noise; which is set to be some percentage of the underlying input's standard deviation. In each of the 500 draws, we simulate noise for each observation i's wedge inputs $\varepsilon_{i,x} \sim N(0, \sigma^2 x)$. We create their "noisy" input $x_i^{\text{noisy}} = x_i + \varepsilon_i$ for each input x. We then recalculate their wedges using the noisy x_i's. We then recompute the two wedge moments of interest.

When we allow noise to be correlated across wedge inputs within respondents, we specify a correlation level $\rho \in \{0, 0.3, 0.9\}$. In order to control both the correlation and standard deviation of noise, we proceed as follows. In each simulation, we draw a respondent-specific common factor $\phi_i \sim N(0,1)$. For each wedge input, we calculate $\widetilde{\varepsilon}_{i,x} = \phi_i \sqrt{\rho} + e_{i,x}\sqrt{1-\rho}$, where $e_{i,x} \sim N(0,1)$. We then rescale this noise to maintain the desired standard deviation: $\varepsilon_{i,x} = \widetilde{\varepsilon}_{i,x}\sigma_x$. We do this for each input x, and once again recalculate wedges and the two moments of their distribution.

The results of Figure E.10 indicate that even when noise is extremely large and highly correlated across inputs our two main wedge moments remain stable.

Figure E.10. Sensitivity of Consumption Wedge to Noisy Expectations Inputs

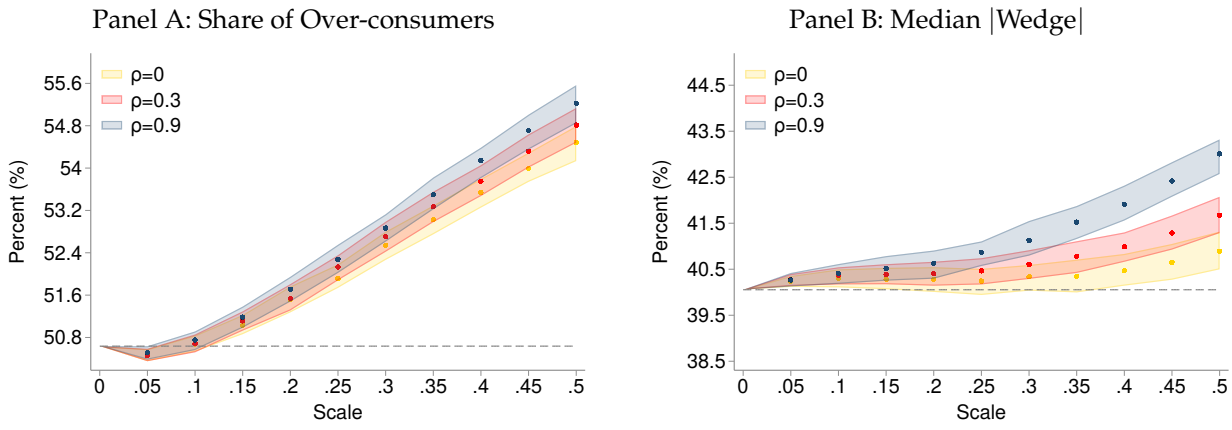

Notes: This figure reports the simulation results from adding noise to wedge inputs. Each dot corresponds to a set of 500 simulations. The x-axis reports the scaling factor controlling the variance of the noise to all inputs. The colors indicate different correlation coefficients used (ρ). Shaded bars indicate the interquartile range of simulation results. The dashed line indicates our baseline analysis's values.

E.3.2 Survey-Based Income Wedge Recalculation

Figure E.11. Consumption Wedge under Reported Income

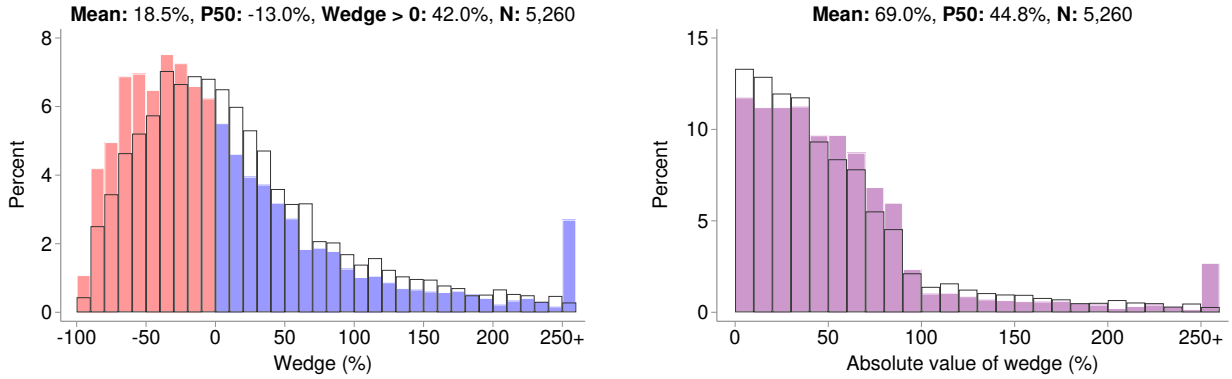

Notes: The graphs show the distribution of consumption wedges (left) and their absolute values (right). Colored bins represent the distribution under survey-reported income, while the clear bins represent the original distribution using observed income from the administrative data. To construct the a post-tax labor income measure from the survey data, we first calculate post-tax labor income for CPS respondents as: $PostTax_Labor_Income = PreTax_Labor_Income - FICA - Labor_Pct(FedTax + StateTax)$, where $PreTax_Labor_Income$ is wage and salary income as reported in the CPS, $FICA$ is Social Security and Medicare payroll taxes, $FedTax$ and $StateTax$ are federal and state income tax liabilities, and $Labor_Pct$ is the share of taxable income from wages and salary. Post-tax labor income is then winsorized at the 1st and 99th percentiles. Next, we assign each CPS respondent to one of 11 total pre-tax income bins and calculate the median post-tax labor income within each bin. Then, we assign this CPS-based median post-tax labor income to survey respondents in the corresponding total pre-tax income bin and recalculate wedges. Wedges are reported here as the percent deviation of observed consumption from frictionless consumption.

E.4 Static versus Dynamic Consumption Wedges

Figure E.12. Distribution of Static Consumption Wedges

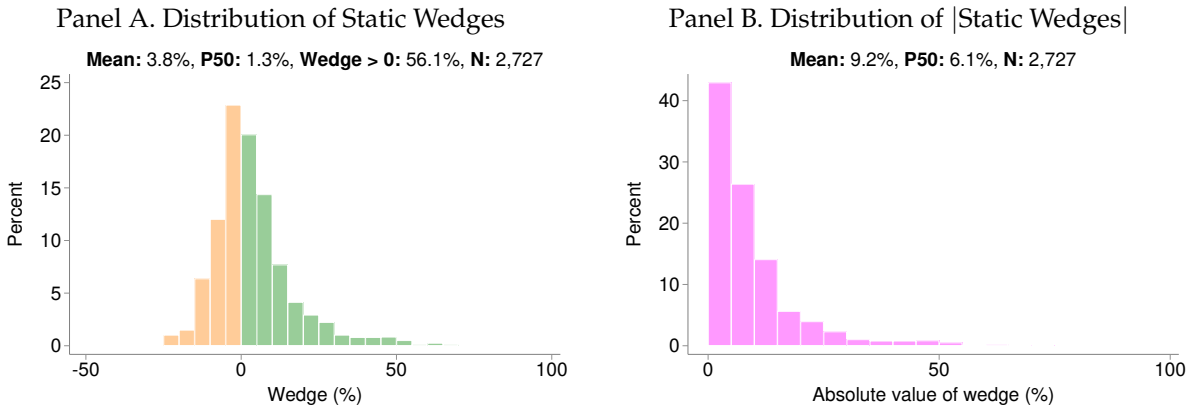

Notes: The graphs show the distribution of static consumption wedges (left) and their absolute values (right). Wedges are reported as the percent deviation of observed consumption and frictionless consumption. Includes users from waves 2 and 3 only, as we did not solicit spending growth expectations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

E.5 External Validity

Figure E.13. Relationship Between Dynamic Consumption Wedges and Inflation Certainty

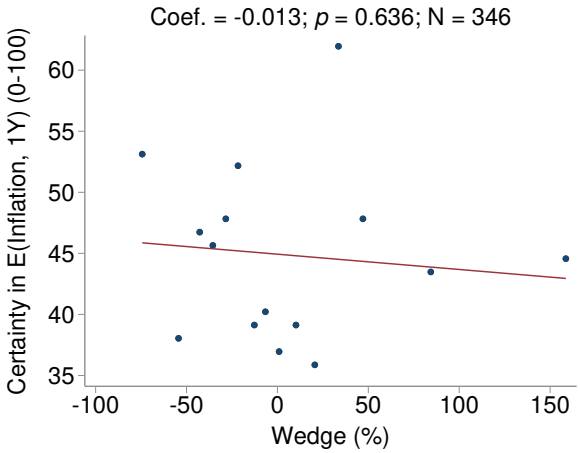

Notes: The graph illustrates the relationship between consumption wedges and a measure of certainty in inflation expectations elicited in survey wave 2. Respondents in wave 2 were asked to rate their uncertainty on a scale of 0-100. Each binned scatterplot plots the average value of the inflation certainty scale within quantile-based intervals of consumption wedges. Note that the sample is comprised of all individuals whose wave 2 wedges met our sample criteria (this sample is larger than the population for whom both their wave 1 and wave 2 wedges met the sample criteria.

Figure E.14. Relationship Between Dynamic Consumption Wedges and EarnIn EWA Usage

Panel A. EarnIn EWA Inflows ($)

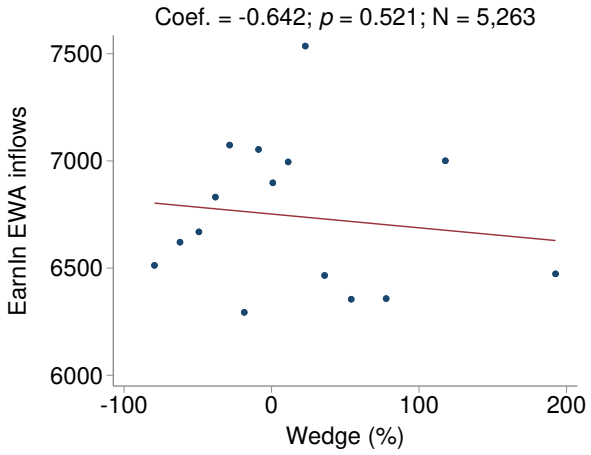

Panel B. EarnIn EWA Inflows (Count)

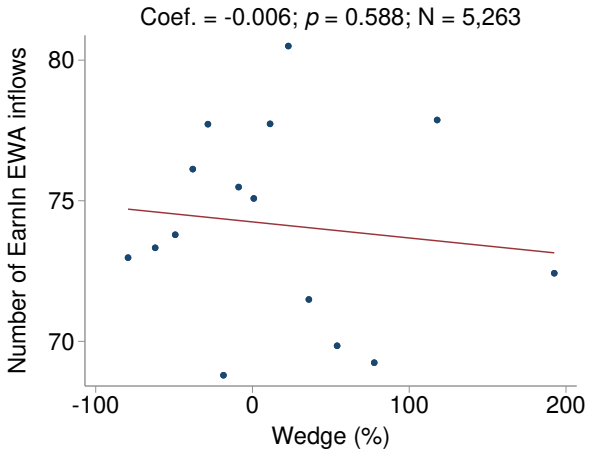

Notes: The graphs illustrate the relationship between consumption wedges and two measures of EWA usage over the 12 months preceding the survey: (1) the dollar amount of EarnIn EWA inflows and (2) the number of EarnIn EWA inflows over the sample period. Each binned scatterplot plots the average value of the EWA usage variable within quantile-based intervals of consumption wedges.

Figure E.15. Sensitivity of Consumption Wedge to Lifecycle Differences

Panel A. Raw wedges by spouse

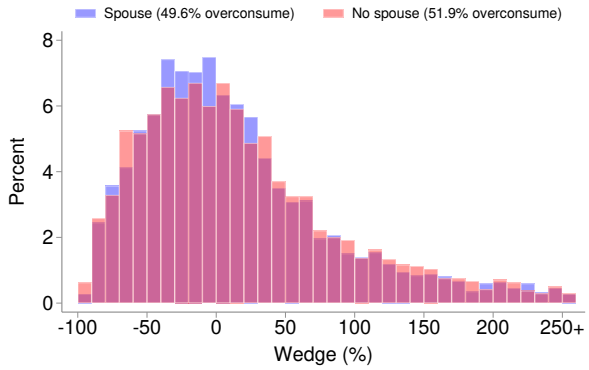

Panel B. Absolute value of wedges by spouse

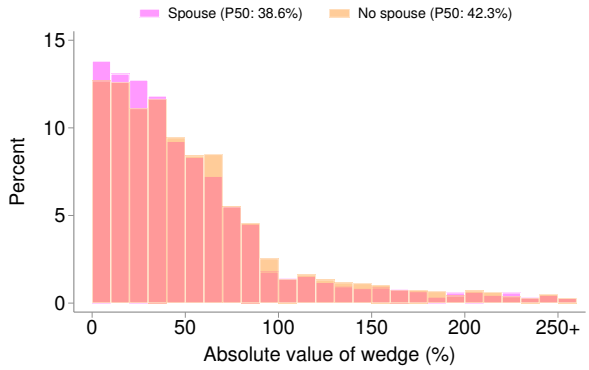

Panel C. Raw wedges by children

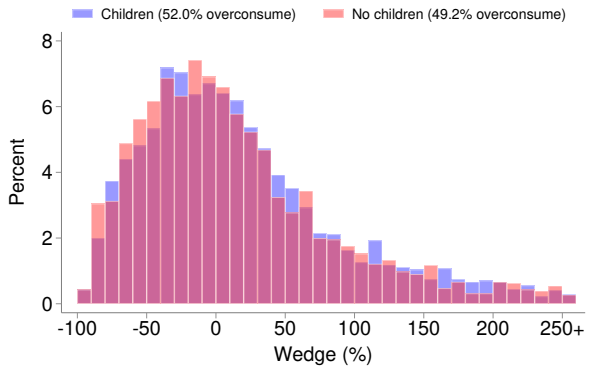

Panel D. Absolute value of wedges by children

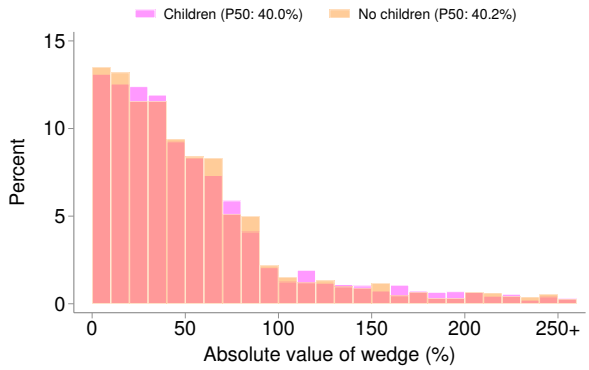

Notes: The graphs show the distribution of consumption wedges (left) and their absolute values (right), broken down by various lifecycle differences. Wedges are reported here as the percent deviation of observed consumption from frictionless consumption.

F Quantitative Model Appendix

F.1 Quantitative Model Specifications

We solve a series of simple quantitative models in the style of Bewley (1980). That is, partial equilibrium, heterogeneous agent, incomplete markets models. Each model features a single asset, a borrowing constraint, idiosyncratic (but persistent) income shocks, and infinitely-lived consumers that are heterogeneous in terms of income and wealth. The components of the consumer's problem common across all models are:

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

The three distortions we add, either individually or in combinations, are a borrowing constraint, present bias, and a consumption adjustment friction.

Borrowing constraints. We model borrowing constraints as a simple lower bound on net worth:

$$A_{i,t+1} \geq \bar{A}$$

with $\bar{A}$ denoting the borrowing constraint.

Present Bias. We study two forms of present bias. The first is naive beta-delta hyperbolic discounting (as in Lee and Maxted, 2023). To retain our notation, we continue to use β to denote the usual exponential discount factor; $b \leq 1$ denotes the present bias discount factor. That is, in period t consumers discount expected utility in period $t+1$ by a factor of $b\beta$ as opposed to β. The second is Gul and Pesendorfer (2001) temptation preferences. Our formulation most closely follows Attanasio et al. (2024). These preferences modify the value function to include an additional term, shown below:

$$V(y_{i,t}, a_{i,t}) = \max_{\{a_{i,t+1}, c_{i,t}\}} u(c_{i,t}) + \beta E\left[V(y_{i,t+1}, a_{i,t+1}) | y_{i,t}, a_{i,t+1}\right] + \lambda\left[u(c_{i,t}) - u(\tilde{c}_{i,t})\right].$$

The parameter $\lambda > 0$ governs the degree of temptation. Here, $\tilde{c}_{i,t}$ is the most tempting feasible consumption choice. Specifically,

$$\tilde{c}_{i,t} = y_{i,t} + a_{i,t} R - \bar{a}_{i,t+1}$$

where $\bar{a}_{i,t+1}$ is the maximum feasible borrowing (i.e., the borrowing limit if there is one or the natural borrowing limit). The term $\lambda\left[u(c_{i,t}) - u(\tilde{c}_{i,t})\right]$ embodies the cost of temptation; it is the disutility from exerting self control and choosing ci, t instead of consuming as much as is feasible

($\tilde{c}_{i,t}$). These preferences create present bias by making saving less rewarding as it leads to future disutility from temptation.

Consumption Adjustment Costs. We model consumption adjustment costs as a non-pecuniary cost ψ (we draw on the formulation of Fuster et al., 2021). Consumers choose whether or not to adjust their consumption relative to the previous period, incurring a cost $\psi > 0$ when doing so. In this variant, their value function is:

$$V(a_{i,t}, y_{i,t}, c_{i,t-1}) = \max \left\{ V^A(a_{i,t}, y_{i,t}) - \psi, V^N(a_{i,t}, y_{i,t}, c_{i,t-1}) \right\}$$

where

$$\begin{aligned} V^A(a_{i,t}, y_{i,t}) &= \max_{\{a_{i,t+1}, c_{i,t}\}} u(c_{i,t}) + \beta E\left[V(a_{i,t}, y_{i,t}) | y_{i,t}, a_{i,t+1}\right] \\ V^N(a_{i,t}, y_{i,t}, c_{i,t-1}) &= u(c_{i,t-1}) + \beta E\left[V(a_{i,t}, y_{i,t}) | y_{i,t}, a_{i,t+1}\right]. \end{aligned}$$

F.2 Quantitative Model Calibrations

Table F.1 summarizes our calibration. Panel A presents parameters governing the fundamental economic environment that are common through all model variants we study. Panel B reports the distortion parameters we obtain by calibrating the models to target the median absolute value wedge of 40.1% and over-consumer share of 50.6%.

Table F.1. Quantitative Model Parameters

Parameter	Value	Meaning	Source
Panel A. Fundamental Environment			
β	0.92	Annual discount factor	Auclert, Rognlie and Straub (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	1.0793	Real gross annual interest rate	EarnIn Sample median levered real return
Panel B. Distortions			
$\bar{A}$	-0.25	Borrowing limit	Kaplan et al. (2018)
b	0.375	Hyperbolic discounting factor	Calibrated
λ	0.4375	Temptation cost scalar	Calibrated
ψ	0.0062 (0.0225)	Cons. adj. cost (with BC)	Calibrated

This table presents the calibration for the quantitative models. Panel A reports parameters that are common across all models. Panel B reports parameters related to distortions.

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 0.87-0.95 range used in Auclert et al. (2024). $\gamma = 2$ is a widely used parameterization of the inverse IES. For calibrating the income process,

we follow Guerrieri and Lorenzoni (2017) and Maxted (2025), using the estimated process from Floden and Lindé (2001). The interest rate we use is the median real levered return reported in the EarnIn sample.

For the simple lower-bound borrowing limit, we follow Lee and Maxted (2023) and use the value of -0.25 from Kaplan et al. (2018). We opt not to calibrate this parameter as it is not possible to generate positive wedges with borrowing constraints alone. We calibrate the remaining distortions such that they minimize the distance between the model-implied over-consumer share and median absolute value wedge and their empirical counterparts. Relative to Maxted (2025) and Attanasio et al. (2024), we calibrate high levels of present bias in both the beta-delta and temptation preferences formulations. This is likely because the EarnIn population has a higher-than-average demand for liquidity (as revealed by their participation in EarnIn) and therefore may experience larger distortions than the average consumer. For the consumption adjustment cost parameter ψ, we calibrate a value of 0.0062 when not including borrowing constraints and 0.0225 when including a borrowing constraint. These values are relatively large in that they imply fewer than 20% of agents adjust consumption in each period.