Consumption Wedges: Measuring and Diagnosing Distortions *

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

January 4, 2026

Abstract

Deviations from canonical consumption-savings models are attributed to a variety of distortions, including financial constraints and behavioral preferences. We develop a sufficient-statistics approach to measure the impact of distortions as a wedge between observed and counterfactual "frictionless" consumption. Different distortions imply different wedge properties, hence wedges provide a diagnostic to distinguish between models. We measure wedges using data linking individuals' economic beliefs to transactions data for a population of US consumers with predominantly low liquid wealth. By using subjective beliefs data, we can isolate the influence of frictions and behavioral preferences from deviations from full-information rational expectations (FIRE). The median consumer's wedge is 40% of frictionless consumption in absolute value, with 49% of consumers under-consuming (negative wedges) and 51% over-consuming (positive wedges). Because borrowing constraints only generate negative wedges, they cannot rationalize our data. Models combining present bias with borrowing constraints, or featuring consumption adjustment costs, can 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, 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 the Wharton School's Jacobs Levy Center.

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

‡Entrepreneurial Management Unit, Harvard Business School and NBER, rkluender@hbs.edu

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

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

1 Introduction

Borrowing constraints play a central role in theories of consumption-savings behavior in macroeconomics and household finance (e.g., Kaplan and Violante, 2014). A key motivation for the focus on borrowing constraints is extensive empirical evidence of deviations from the permanent income hypothesis, such as high marginal propensities to consume (MPCs) out of transitory income changes, especially among consumers with low liquid wealth (Johnson, Parker and Souleles, 2006).[1] However, alternative distortions—such as present bias (Attanasio, Kovacs and Moran, 2024; Maxted, Laibson and Moll, 2024), consumption adjustment costs (Berger and Vavra, 2015), and bounded rationality (Ilut and Valchev, 2023)—can also account for these empirical patterns. Importantly, these alternative models imply different distributional and aggregate consequences of fiscal policy, monetary policy, and business cycle fluctuations.[2] 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.

We develop a new sufficient statistics approach to measure consumer-level "wedges" between observed consumption and the hypothetical "frictionless" consumption implied by an Euler equation and budget constraint in the absence of distortions. These consumption wedges quantify the total net impact of distortions on consumption, including both frictions (such as borrowing constraints) and behavioral preferences that result in "as if" constrained behavior (e.g., present bias or bounded rationality). Different distortions have different predictions for the properties of consumption wedges, such as their sign, correlates, and responses to shocks. As a result, consumption wedges can be used as a diagnostic to discriminate between models of consumer behavior.

We make two innovations relative to existing research that 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.[3] This proves to be important: the median wedge is 0.9% of frictionless consumption, but the median *absolute value* wedge is 40%, with similar shares over-consuming (positive wedges) and under-consuming (negative wedges). Second, we use individual-level subjective beliefs data,

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

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

[3] This innovation is similar in spirit to Hsieh and Klenow (2009), which estimates firm-level productivity wedges (misallocation).

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

We estimate consumption wedges using a sample of 5,028 consumers with de-identified administrative banking data linked to their surveyed expectations. The consumers in this sample are customers of EarnIn, an American financial technology company that offers their users early access to their wages prior to their regularly scheduled payday. EarnIn fielded three surveys to its customers between 2022–2024 that elicited subjective expectations over inflation, nominal income growth, and nominal interest rates for both saving and borrowing.

The EarnIn sample is not representative of the full US population, but this subpopulation of wage workers with low liquid wealth is of particular economic and policy relevance. Our sample covers a significant swathe of the income distribution, with over-representation between the 30th and 80th percentiles and under-representation only in the tails. As expected given the revealed demand for early access to their paychecks, our sample has low liquid wealth: 64% of our respondents report liquid wealth below $500 compared to 16% of respondents to the 2022 Survey of Consumer Finances. 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). Consistent with this, many of the largest tax-expenditure and social-insurance programs in the United States—including the Earned Income Tax Credit, the Child Tax Credit, health insurance premium subsidies, and unemployment insurance—are explicitly targeted at low- and moderate-income wage workers with limited liquid assets. In addition to these public programs, households in this group increasingly rely 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 consump-

4 For example, inflation expectations are excessively influenced by grocery prices (D'Acunto et al., 2021) and D'Acunto et al. (2024) finds evidence of extrapolative income expectations.

tion and the hypothetical frictionless level they would choose absent distortions (given their beliefs and resources). We formulate a frictionless benchmark model in which a consumer chooses consumption and saving (or borrowing) via a risky asset given their realized income, wealth and expectations. In the benchmark, consumers face no distortions: there are no frictions (e.g., constraints) and consumers have standard preferences.[5] In the benchmark, frictionless consumption is characterized by an Euler equation and a budget constraint.

The purpose of the frictionless benchmark is not to explain observed consumption. 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.

Despite our use of a highly stylized benchmark, the same set of sufficient statistics can be used to calculate wedges in 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).

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. In terms of dollars the median absolute value wedge is \$13,301 per year, which is 34% of respondents' median annual post-tax income of \$39,615. The large scale of the wedges implies that distortions are significant determinants of the consumption decisions of low-liquidity consumers. In contrast, the mean and median wedges are significantly 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. The discrepancy between aggregate and individual-level wedges suggests that modeling distortions in households' consumption savings decisions is important when heterogeneity in consumer behavior matters for the model's predictions (e.g., Kaplan et al., 2018) but less so for questions where a

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

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 existence of a large share of under-consumers and over-consumers implies that the consumption choices of our sample cannot be rationalized by a model that generates distortions with a single sign. Because borrowing constraints can only generate under-consumption, our results indicate that borrowing constraints cannot be the dominant distortion for a majority of our sample.

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 generate positive wedges, such as present bias (e.g., Maxted, Laibson and Moll, 2024; Attanasio, Kovacs and Moran, 2024). The second is to include frictions that generate both positive and negative wedges, such as consumption adjustment costs and bounded rationality (e.g., Fuster et al., 2021; Berger and Vavra, 2015; Beraja and Zorzi, 2024; Ilut and Valchev, 2023). Both are sources of consumer inertia, which can create both positive and negative wedges by limiting the consumption response to shocks.

To assess the robustness of our results, we provide evidence on their sensitivity to two important sets of assumptions: the modeling of preferences and the magnitude of measurement error. Starting with preferences, we generally find quantitatively insignificant sensitivity of our results to the choice of the inverse IES (γ). The discount factor β is a more important choice, however we still find a large (20%) share of over-consumers even for a very low annual discount factor (0.80). Our next exercises are motivated by the fact that heterogeneity in preferences, specifically deviations from our calibrated values, would also manifest as a wedge. We modify our approach to allow for a plausible degree of preference heterogeneity, setting preferences to match one of the three types of Aguiar et al. (2024) such that they minimize wedges. We obtain similar results: 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 not statistically nor economically significant predictors of wedges. These (and additional) exercises suggest that distortions, rather than preference heterogeneity, are likely the driver of our main findings.

Across four sets of robustness exercises, we find that measurement error—in either consumption, income, or beliefs data—is unlikely to significantly distort our findings. In the first exercise, we drop respondents who are likely to exhibit more measurement error—such as users with low financial literacy, who spent less than 6 minutes on the survey, or who rounded inflation expectations to a multiple of five percentage points—to test whether the inclusion of these users skews our results. Second, we address the possibility that respondents may receive income from sources not observable in the transactions data, which would bias our measured wedges upwards, by recalculating wedges using an alternative survey-based income measure instead of transactions income. Third, we use *k*-prototypes clustering on economic and demographic characteristics to group similar users. We then 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 a similar magnitude median absolute value wedge.

To evaluate the importance of subjective beliefs data for these properties of the wedges, and as a determinant of consumption, we calculate an alternative set of wedges instead assuming that consumers form full-information rational expectations (FIRE). We impute FIRE over income, inflation, and returns for each person using their realizations. Observed consumption can then be decomposed into three terms: frictionless consumption under FIRE, a wedge induced by subjective beliefs deviating from FIRE, and the remaining wedge induced by distortions. A covariance decomposition reveals that frictionless consumption computed with FIRE beliefs and realized income and wealth explains 57% of the variation in observed consumption, while deviations in beliefs from FIRE explain 25% and distortions explain the remaining 18%. This indicates that distortions and subjective beliefs are similarly important in explaining the cross-section of consumption. Without data on subjective beliefs, one could not separate the wedge due to non-FIRE beliefs versus the wedge due to distortions.

Our last analyses provide additional evidence on the source of wedges and, more broadly, demonstrate how to apply wedges as a diagnostic tool. We illustrate this in the context of both reduced-form and structural applications. Our reduced-form analyses examines correlations between wedges and other variables of economic interest, providing evidence on the nature of distortions and helping to validate the interpretation of wedges as measuring distortions . We begin by correlating the wedges with survey responses to a hypothetical scenario that is independent of the wedge measurement. We find that when asked how they would adjust saving if they expected higher inflation, over-consumers are more likely to report saving less. This reveals an internal consistency between consumers' own anticipated behavioral tendencies and their over-/under-consumer type, as indicated by their wedge.

We next correlate wedges with MPCs (measured out of the 2021 stimulus checks) and proxies for consumption commitments, financial distress, and homeownership. We generally find strong correlations, which provides additional validation as one would expect such variables to be related to wedges if they indeed capture the severity of the consumption response to distortions. The signs of the correlations also provide further diagnostic evidence. For example, more binding borrowing constraints generate more negative wedges and higher MPCs. Therefore, we would expect a negative correlation between wedges and MPCs if borrowing constraints were the dominant friction. In contrast, present bias can generate a positive correlation between wedges and high MPCs.

The correlations we find are inconsistent with borrowing constraints being the dominant distortion and instead point to either present bias and inertia. MPCs are positively correlated with wedges—the opposite of what we would expect if borrowing constraints were the main reason MPCs are high. We proxy for consumption commitments with the ratio of spending on housing and childcare to income, and find a strong positive correlation with wedges. This pattern is consistent with consumption commitments as a dominant distortion. Over-consumption is also strongly related to financial distress (perceptions of financial anxiety, unmanageable debt, and difficulty borrowing). These patterns fit present-biased behavior and inertia but are at odds with borrowing constraints. We also 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 that generate over-consumption tend to dominate for those without such assets. Our structural application studies 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 50.6% over-consumer share and 40.1% median absolute value wedge found in the empirical wedge distribution. Quantitatively, we find that present bias combined with a borrowing constraint is best able to match these empirical moments. This model can account for 40% of the observed over-consumer share and 73% of the median absolute value wedge. We discuss how non-rational expectations can be incorporated into such models such that they can further improve their ability to account for the empirical distribution of wedges. The implication that non-rational beliefs may be necessary to fully account for consumer behavior is consistent with our covariance decomposition, which indicates that deviations from FIRE explain a substantial share of the cross-sectional variation in consumption.

Taken together with the substantial mass of positive wedges in the cross-section, these facts 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. First, we innovate on the macroeconomics literature on quantifying frictions and distortions by measuring wedges at the individual level using high-quality transactions data linked to subjective beliefs. This unique dataset enables us to isolate frictions and behavioral preferences from deviations from FIRE. 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 on wedges has focused on quantifying the importance of misallocation across firms and risk-sharing across households for growth and business cycles (e.g. Hsieh and Klenow, 2009; Baqaee and Farhi, 2020; Berger, Bocola and Dovis, 2023). We take this approach in a different direction by combining survey data on subjective beliefs with administrative transactions data to calculate consumer-level wedges, which allow us to disentangle the influence of biased beliefs from frictions and behavioral preferences. Imputing consumer-level FIRE beliefs is difficult, but doing so further enables us to decompose the variance in consumption explained by the frictionless benchmark under FIRE (57%), biased beliefs (23%), and distortions (20%).

Second, we contribute to the empirical macro literature studying the determinants of consumption (Attanasio and Weber, 2010; Krueger et al., 2016; Kaplan and Violante, 2022). Early work in this area focused on testing the permanent income hypothesis (PIH) through Euler equation estimation (Hall, 1978; Flavin, 1981; Hansen and Singleton, 1982; Campbell and Mankiw, 1989). This literature found systematic departures from the predictions of frictionless PIH models, including excess sensitivity to predictable income changes and violations of the orthogonality conditions implied by FIRE. The advent of high-quality spending data gave rise to a related body of research documenting large MPCs, especially among consumers with low liquidity (Johnson, Parker and Souleles, 2006; Baker, 2018; Ganong and Noel, 2019; Gross, Notowidigdo and Wang, 2020; Fagereng, Holm and Natvik, 2021). These cross-sectional patterns have served as important motivation for the inclusion of wealth heterogeneity and financial frictions in macro models (e.g., Kaplan and Violante, 2014; Koşar, Melcangi, Pilossoph and Wiczer, 2023). However, recent work has also found high MPCs among high-earning and high-liquidity consumers. These findings have motivated a growing literature proposing behavioral explanations, such as bounded rationality and present bias (Ilut and Valchev, 2023; Boutros, 2022; Lian, 2023; Attanasio, Kovacs and Moran, 2024; Maxted, Laibson and Moll, 2024; Ganong, Greig, Noel, Sullivan and Vavra, 2024). Our analysis contributes to this ongoing debate by documenting the need for explanations that generate over- as well as under-consumption. The prevalence of over-consumption, in particular, indicates that it is important to incorporate distortions like present bias or inertia.

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

Third, we build on work in empirical macroeconomics showing that consumer beliefs, including departures from FIRE, are instrumental to understanding consumer behavior (see recent reviews by Weber et al., 2022; D'Acunto et al., 2023). Recent papers have linked beliefs to consumption decisions using surveys (Coibion et al., 2023; D'Acunto et al., 2022), grocery purchases (Weber et al., 2023), German bank data (Hackethal et al., 2023), and credit cards (Kanz et al., 2021). Consumer beliefs systematically deviate from FIRE by, for example, overweighting grocery prices in inflation expectations (D'Acunto et al., 2021) and extrapolative income forecasts (D'Acunto et al., 2024). These findings motivate our use of subjective beliefs data. Additionally, we provide new evidence on the relative importance of beliefs and 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 applies the wedges as a diagnostic tool. Section 7 concludes.

2 Theory: Measuring Consumption Wedges

We derive a formula for frictionless consumption by solving a stylized consumption-savings model. The formula generalizes to much richer environments. We show that data on income, wealth, and beliefs over future inflation, income growth, and returns are sufficient statistics for identifying frictionless consumption in a large class of models.

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. 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.

The consumer chooses consumption $C_{i,t}$ and savings $A_{i,t+1}$ to maximize her expected utility subject to a budget constraint, solving:

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

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

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, it has no economic frictions (borrowing constraints, adjustment costs, etc.). Second, it assumes standard preferences and thus precludes behavioral preferences like present bias or habit formation, which can result in "as if" constrained behavior. The benchmark *intentionally* omits both of these types of distortions because doing so makes it possible to measure their impact as a wedge between frictionless and observed consumption. The purpose of the benchmark is to facilitate this measurement; the benchmark is not to intended to provide a realistic model of observed consumer behavior. Third, we depart from the prior wedge measurement literature (e.g., Chari et al., 2007; Berger et al., 2023) by letting beliefs deviate from FIRE. This 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. We obtain the following expression (see

Appendix A for the derivation):

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

Frictionless consumption $C_{i,t}^{\star}$ 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 $G_{i,t,t+j}^{Y}$ (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 intertemporal elasticity of substitution (IES) γ and the discount factor β. The resulting equation reflects the familiar logic of the textbook permanent income hypothesis, where the numerator corresponds to (approximate) expected lifetime wealth and the denominator dictates the share of that lifetime wealth to be consumed in period t.

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

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

Frictionless consumption represents a counterfactual level of consumption. Specifically, it holds constant the consumer's wealth and the prices they expect to face, but it removes any distortions. Therefore, the consumption wedge $\eta_{i,t}$ quantifies the total impact on observed consumption of all distortions omitted from the frictionless benchmark (holding wealth and prices constant).[6] 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 here by describing implications for the sign of wedges for some commonly used distortions. First, consider borrowing constraints, which can only reduce consumption. Therefore, a testable implication of borrowing constraints is that they generate negative consumption wedges. A second example is present bias, which produces positive wedges by increasing the preference for con-

[6] In general equilibrium, "turning off" distortions will generally imply a different wealth distribution and prices. Hence the counterfactual we study is partial equilibrium in the sense that it holds wealth and prices constant. The wedge captures a "direct" effect of distortions as opposed to a "total" effect (i.e, both direct and indirect effects). In other words, the wedge describes how distortions shift the consumer's consumption policy function.

suming more in the present. Finally, distortions producing inertia can generate either positive or negative wedges. Examples of such distortions include consumption commitments (adjustment costs), habit formation (reference-dependent preferences), and 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. See Appendix A.2 for further details on these examples.

Robustness to Model Extensions. The frictionless benchmark abstracts away from additional choice variables, such as labor supply, multiple assets, and durable goods. This is without loss of generality. Adding choice variables adds further necessary conditions for optimality to characterize the overall optimum (i.e., *all* optimal choices). But because 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 follows even 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). Appendix A.3 discusses extensions featuring additional choices, assets, and durable goods.

3 Data and Survey Design

3.1 EarnIn Administrative Data

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

We use the transactions data to construct two variables central to our analysis: non-durable spending and income. We define outflow transactions as non-durable spending based on a mapping from more than 500 transaction categories associated with the transactions to spending categories building on Ganong and Noel (2019). We measure non-durable spending using data from

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

the 12 months preceding the survey (i.e., October 2021 to September 2022 for Wave 1).[8] We focus on non-durable spending because it corresponds more closely to consumption, whereas the relationship between spending and consumption of durables depends on their rate of depreciation and the flow rate of consumption. We obtain total or ”notional” consumption (i.e., the argument of the utility function) by dividing each respondent’s non-durable consumption by the typical expenditure share of non-durable goods (79.37%). Under the assumption that notional consumption is a Cobb-Douglas aggregate of durable and non-durable good consumption flows, this calculation yields notional consumption (for more, see Appendix A.3.1). We obtain the non-durable 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. We classify inflow transactions as either earnings (observed post-tax), UI benefits, or “other” using the observed earnings data together with the transaction category, memo line, and periodicity of the transaction. Our baseline analysis covers the 12-month period before each survey wave (e.g., annual non-durable spending and income between October 2021 to September 2022 for wave 1). For more details on the EarnIn data, sample construction, and variable definitions see Appendix B.

3.2 Survey Data

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

EarnIn fielded three survey waves between 2022 and 2024. Wave 2 was a follow-up survey sent to respondents from wave 1. In each wave, qualifying users were invited via EarnIn’s standard email marketing channels to complete a short survey about their current economic well-being and future outlook. Survey windows spanned from September 29 to October 2, 2022 (wave 1); July 12 to July 19, 2024 (wave 2); and November 22 to December 4, 2024 (wave 3). Surveys took five

[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 analysis of wedges obtains similar results when aggregating over fewer than 12 pre-survey months (see Appendix Figure D.3).

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

to ten minutes (wave 1) or ten to fifteen minutes (waves 2 and 3) to complete. As an incentive, respondents received a $5 (waves 1 and 3) or $10 (wave 2) Amazon gift card upon completion.

We construct the sampling frame for each wave by imposing data quality requirements on the transactions data to ensure that users' linked accounts plausibly capture the bulk of their economic activity. For waves 1 and 3, we restricted to users for whom we observe earnings, regular spending, and balances in the 12 months leading up to the survey. For the wave 2 follow-up survey, we invited only wave 1 respondents who met additional data quality restrictions. Across all three surveys, EarnIn sent approximately 475,000 survey invitations and received approximately 15,866 responses, yielding an aggregate response rate of 3.4%.[10] The median survey completion time among respondents was 7.5 minutes. Survey respondents are similar to all users eligible to take the survey on account balances, inflows, and outflows, but women were more likely to respond (see Appendix Table B.2). See Appendix B for additional details on sample construction and response rates by survey wave.

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 percentage yield they would expect on additional savings over the next 12 months, and the percentage rate they would expect to pay on new borrowing over the same period. 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 (Lusardi and Mitchell, 2014), demographics, and total household income, savings, and debt (across categories).[11]

In waves 2 and 3, we added several 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 added measures of committed consumption by asking respondents whether they own or rent their home, monthly rent/mortgage payments, and childcare spending. In wave 3, we elicited respondents' hypothetical marginal propensity to spend, repay debt, and save out of a $1,000 one-time income shock using questions based on (Fuster et al., 2021; Colarieti et al., 2024).

[10] Note that this response rate is artificially low because we hit our budget constraint and ended the survey for waves 1 and 3.

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

Sample Restrictions for Analysis. We impose a series of sample restrictions for our analysis, detailed in Appendix B.3. We exclude responses flagged for inattention or low effort by requiring a minimum survey duration of 3.5 minutes, internally consistent responses, and trim economic expectations at reasonable bounds.[12] We also drop users with insufficient transactions data coverage in the 12 months surrounding the survey, and we trim outliers for key financial variables (nondurables spending, income, average propensity to consume, wealth-to-income, and expected levered return). After all restrictions, our analysis sample includes 5,263 responses from 5,028 unique users.[13]

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 (the inter-quartile range of age is 31-43), female (69% of respondents), 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

[12] We retain respondents with inflation and income growth expectations between the 3$^{\text{rd}}$ and 97$^{\text{th}}$ percentiles and interest rate expectations between the 1$^{\text{st}}$ and 97$^{\text{th}}$ percentiles.

[13] There are 238 repeat users who responded in both Waves 1 and 2, giving us a total of 5,263 linked survey responses.

Figure 1. Sample Benchmarks

Panel A. Distribution of Income Percentiles

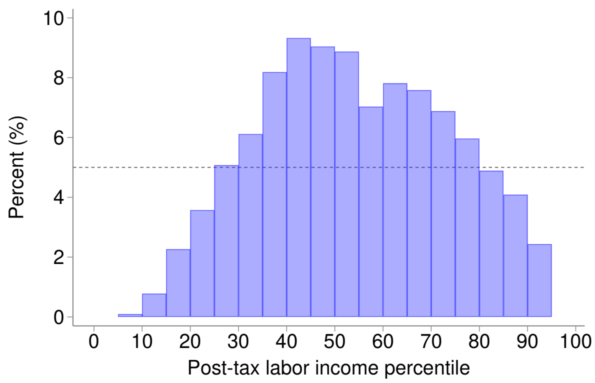

Panel B. Distribution of Liquid Wealth

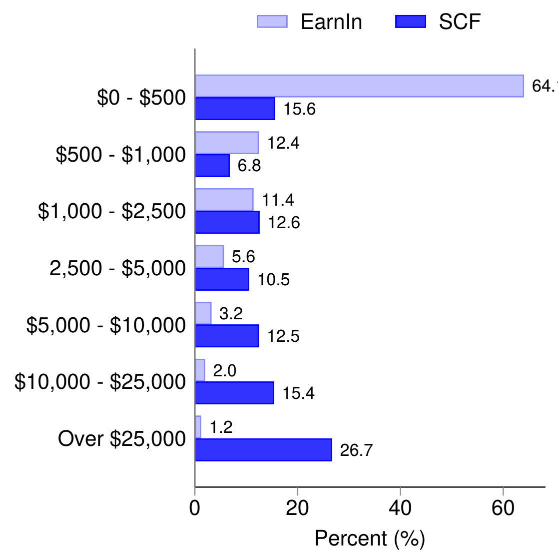

Panel C. Comparison with National Surveys

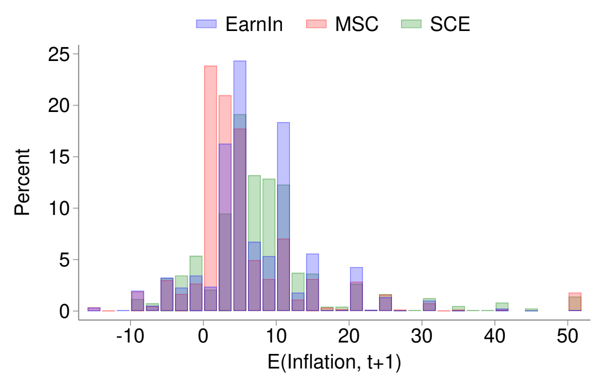

Notes: Panel A presents the distribution of post-tax labor income percentiles among our EarnIn analysis sample, estimated using CPS ASEC data (United States Census Bureau, 2025) and the CPS ASEC Tax Model (Lin, 2022). PS ASEC sample restricts to adults age 18 or older who received wage or salary earnings during the reference year. Black dashed line represents the uniform distribution. Panel B presents a bar graph showing the distribution of liquid assets among EarnIn survey respondents compared to respondents in the 2022 Survey of Consumer Finances. Following Kaplan and Violante (2014), we define liquid assets as the sum of assets held in transactions accounts. Includes users from all three survey waves that meet the restrictions outlined in Appendix B. Panel C shows the relative to the distributions in the Michigan Survey of Consumers (MSC) and NY Federal Reserve Survey of Consumer Expectations (SCE) (right). To match the timing of the EarnIn surveys, we use MSC and SCE data from September 2022, July 2024, and November 2024 (the SCE data is unavailable in November 2024) and reweight the data based on the share of the EarnIn sample in each survey wave.

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 after trimming wedges at the 1st and 95th percentiles, which consists of 5,263 EarnIn users. Columns (1) through (5) show the distribution of each variable, and column (6) shows the number of survey responses with non-missing data. See Appendix B.4 for variable definitions.

income decline of -2%. Reported interest rate expectations on marginal savings and borrowing are broadly reasonable, with medians of 2% and 15%, respectively.

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 ($\gamma = 2$).[14] In Section 5, we perform sensitivity analyses to assess 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/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 D.10).

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). Our approach is motivated by the observation that the ratio of one- to three-year-ahead beliefs is similar to that of the swaps-implied beliefs (measured in the same month of each survey wave).[15] 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 important empirical properties: (1) large but temporary variation in expected growth rates of income in the short-run variation due to anticipated shocks (e.g., job changes) and (2) lifecycle dynamics. Using the SCE, we find significant reversion in income expectations over a year and use this to discipline the relationship between one-year- versus two-year-ahead expected income. We use the MSC to estimate a lifecycle profile of income growth expectations, which we apply to our sample for three-year-ahead (and later) expectations. We detail all of our

[14] 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).

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

imputation procedures in Appendix C and discuss the robustness of the wedge estimates to these assumptions in Section 5.

4 Results: Consumption Wedges

This section reports the empirical distribution of consumption wedges and discusses how their properties relate to those implied by different 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 the frictionless benchmark level of consumption and Panel B plots the distribution of their absolute value. The mean wedge is 15%, implying observed consumption is 15% above frictionless consumption on average, while the median absolute value wedge is 40%. Appendix Figure D.16 reports wedges in terms of dollars. The median wedge is \$259 while the median absolute value wedge is \$13,301, which indicates that distortions to consumption are significant relative to the median post-tax income of \$39,614 in our sample.

Figure 2. Distribution of Consumption Wedges

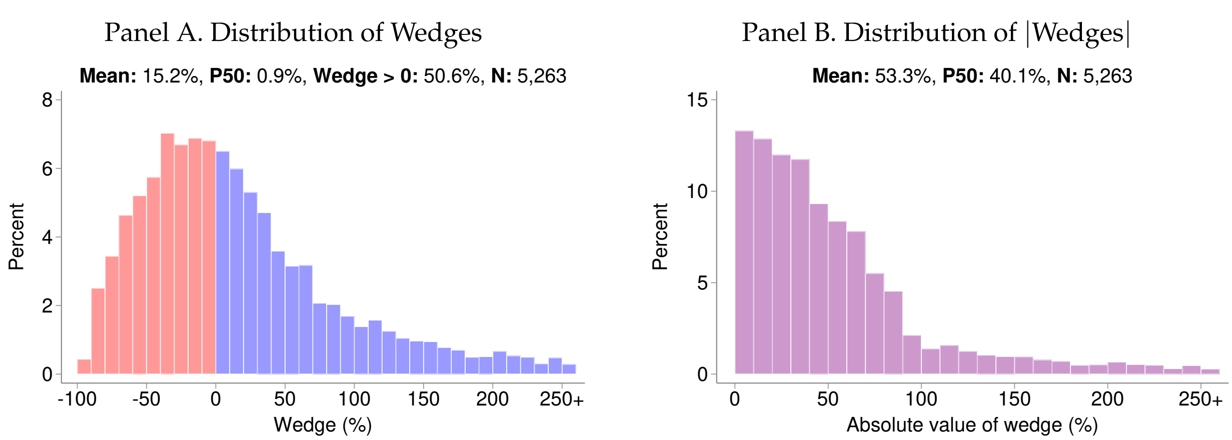

Notes: The figures show the distribution of dynamic consumption wedges (left) and their absolute values (right). Wedges are reported here as the percent deviation of observed consumption from frictionless consumption, and they are trimmed at the 1st and 95th percentiles.

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 (e.g., Kaplan et al., 2018), these patterns imply that 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 dominant modeling paradigm in household finance and macroeconomics: that borrowing constraints are the key friction shaping the consumption choices of low-liquidity consumers. As discussed in Section 2, borrowing constraints only generate negative wedges. While borrowing constraints could explain the 49% of consumers with negative wedges, such constraints cannot be the dominant friction for the other 51%. Borrowing constraints are therefore unable to account for the behavior of a large fraction of consumers in our sample. We identify two promising alternative models. 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 also allows us to assess how valuable subjective beliefs data are for both measuring wedges and, through a novel exercise, as a determinant of consumption. We begin to explore this by calculating each respondent's "frictionless FIRE consumption" via the same benchmark formula (Equation 4) but instead using imputed FIRE beliefs rather than their subjective beliefs. We impute FIRE beliefs by setting them equal to their realized value, applying the same term structure modeling approach described in Section 3.4 for more distant values. We restrict our analysis to wave 1 because we have realized income only for this wave. We measure realizations for income using the transactions data, the CPI for inflation, and the typical interest rates on savings (FDIC) and credit cards (Federal Reserve Board's *Consumer Credit*, G.19 statistical release).[16]

[16] 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 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

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 ($C_{i,t}^{\star,\text{FIRE}}$). Therefore, the SE wedge measures the impact of deviations in subjective expectations from FIRE on consumption (relative to the frictionless FIRE benchmark). The distortion wedge ($\eta_{i,t}^{\text{distortions}}$) is the same wedge from Equation (5) that has so far been our focus, here we add the "distortions" superscript to distinguish it from the SE wedge.

Appendix Figure F.3 reports the sum of the SE and distortion wedges as a percentage of frictionless FIRE consumption. This corresponds to the distribution we would have measured if we did not have subjective beliefs data and had instead assumed that consumers form FIRE. Without incorporating subjective expectations into the benchmark, we find a higher share of overconsumers (68% versus 51%) and smaller absolute value distortions (28% versus 40%).

Next, we calculate a covariance decomposition of Equation (7) to quantify how much of the variation in observed consumption is explained by its three components. This enables us to quantify the relative importance of deviations from FIRE versus frictions/behavioral preferences (i.e., the distortions that are the focus of the rest of the paper) in explaining the cross-section of consumption. Starting with its first component, we find that frictionless FIRE consumption ($C_{i,t}^{\star,\text{FIRE}}$) accounts for 57% of the cross-sectional variation in consumption. Given the significant variation in consumption and the limited wealth of our sample, the fact that a textbook model explains the majority of the cross section of consumption with a parsimonious set of inputs (wealth, income, and expectations imputed from realizations) is striking and provides support that it is a reasonable modeling foundation. Subjective beliefs also emerge as an important determinant of consumption, explaining more than half of the residual variation in consumption (25%). Distortions account for the remaining 18% of the variation in consumption. Hence, subjective expectations and distortions are similarly important in accounting for variation in the cross-section of consumption. While the frictionless benchmark with subjective expectations can explain 82% of the cross-section of consumption, we reiterate that the typical distortion wedge is large (40.1% of frictionless consumption), and in this sense distortions exert a large influence on consumption.[17]

balance and weighted by the institution's share of domestic deposits.

[17] The covariance decomposition exercise and the distribution of wedges speak to the importance of distortions in

5 Robustness

In this section, we examine the robustness of our two main findings: (1) a median absolute value wedge of 40% and (2) a 51% share of over-consumers. 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 D.1.

Table 2. Sensitivity Analysis: Impact of Alternative Calibration Choices

	Parameter Range			Overconsumer (%)		P50 Abs(Wedge) (%)	
	Calibration (1)	Min (2)	Max (3)	Min (4)	Max (5)	Min (6)	Max (7)
β	0.92	0.80	0.98	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 two results to our parameter calibration: (1) the percent of users who over-consume and (2) the median absolute value wedge. Under the "Parameter Range" heading, the "Calibration" column shows our baseline calibrated value and the "Min" and "Max" columns show the range of values that we test. Under the "Over-consumer (%)" and "P50 Abs(Wedge) (%)" headings, the "Min" and "Max" columns show the minimum and maximum value of each result that we get across each parameter range. When we vary one parameter, we hold all other parameters at their baseline calibrated values.

Discount Factor (β). Naturally, the wedges are sensitive to the discount factor (β): if a consumer discounts future consumption more strongly, their frictionless consumption is higher, leading to lower 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 it combines the

distinct ways and distortions can be large without explaining the cross-section of consumption. For example, if every consumer had large but identical distortion wedges, distortions would be an important determinant of consumption but they would not explain any of the cross-section of consumption.

wedges on one side of the distribution becoming larger while attenuating those on other side. These results illustrate that there remain both a meaningful share of over-consumers and large wedges even when everyone is assumed to have extremely low discount rates.

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 decreases from 60% to 37% going from a γ of 1 to 2, thereafter stabilizing near 40% for larger values of γ (see Appendix Figure D.1 Panel B).

Non-Durable Expenditure Share. In contrast to β and γ, the non-durable expenditure share does not affect the frictionless benchmark but it does affect the measurement of total observed consumption. Recall that we divide observed non-durable spending by the non-durable expenditure share to infer total consumption; therefore, assuming a lower non-durable expenditure share translates to higher consumption and more over-consumers. Varying the non-durable expenditure share from 72% to 90% monotonically decreases the share of over-consumers from 56.7% to 42.0% .[18] 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 (relatively extreme) alternative assumption of constant beliefs in Appendix Figure D.4. 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 D.3 reports the share of over-consumers and the median absolute value wedge re-estimated using 1, 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.

[18] This choice of range is motivated by other estimates of this expenditure share in the literature. Our baseline value is 79.37% (obtained 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. We obtain a value of 88.8% when we classify non-durable and durable expenditure in the EarnIn transactions data.

5.2 Preference Heterogeneity

Our baseline parameterization assumes homogeneous preferences. If consumers' preferences differ from the homogeneous parameters of $\beta = 0.92$ and $\gamma = 2$ in our baseline, this would manifest as wedge. 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. A large "Euler equation estimation" literature has sought to infer preferences from consumption data while assuming that consumers face no distortions (e.g., Hansen and Singleton, 1982; Attanasio and Weber, 1995). In contrast, we instead impose restrictions on preference heterogeneity in order to infer distortions. We adopt this approach because our goal is to inform the design of a large and important class of models that rely primarily on distortions, rather than flexible preference heterogeneity, to study macroeconomic phenomena. This class has proven useful in terms of its ability to explain empirical patterns like the wealth distribution and high MPCs, and remains a widely-used tool for studying policy, business cycles, and inequality (e.g., Kaplan et al., 2018; Maxted et al., 2024; 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.[19] 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 Three Preference Types. We modify our baseline results to allow three heterogeneous types, using the three preference types estimated in Aguiar et al. (2024). 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. The median absolute value wedge declines only modestly from 40% to 39%, and the share of over-consumers falls from 51% to 39% (see Appendix Tables D.1 and Appendix Figure D.2).

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

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.[20] 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 D.14 shows that the survey measures have very weak relationships with the zero-wedge preferences (Panel A). Following similar logic, we also examine within-individual, cross-wave consistency of the zero-wedge preferences for our subset of 238 repeat-responders in Panel B. 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 first test the sensitivity of the wedge estimates to measurement error in the inputs (beliefs, income, wealth, and consumption) by dropping respondents who are likely to exhibit larger measurement error. To focus on a subgroups more likely to effectively communicate their economic expectations, we sequentially drop users who finish the survey in less than six minutes, 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. We focus most of our attention on consumption measurement error and 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 non-durable spending. Appendix Figure D.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 range from 40–55%. The insensitivity of our wedge estimates to the exclusion of these higher measurement error groups suggests that our findings cannot be explained by differential measurement error across these subgroups.

[20] 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. Appendix Figure D.13 plots the distributions of the "zero-wedge" β and γ.

Repeated Measurement. The second survey wave re-surveyed the same users from wave one, which allows us to estimate wedges for the same consumers at two different time periods. In addition to eliciting new expectations, we also measure income, wealth, and consumption independently almost two years after the initial survey. While the distortions consumers face are likely to evolve somewhat over time, we expect that consumers borrowing constrained or present biased in 2022 may also be more likely to be subject to those 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 D.15 shows that the wedges are highly correlated across waves ($p < 0.001$) and demonstrate strong test-retest reliability. This suggests that the distortions a consumer faces tend to be persistent over time and our wedge estimates are capturing their impacts.

***k*-Prototype Clustering.** We next attempt to reduce measurement error by aggregating over wedge inputs for clusters of similar respondents. We group similar respondents using *k*-prototype clustering (for details, see Appendix D.1) based on their similarity across demographics, income, assets, expectations, and other financial characteristics.[21] Under the assumptions that (1) respondents within a cluster have the same data generating process for their wedge inputs (wealth, expectations, etc.) and (2) measurement in inputs is zero-mean noise, 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 D.8 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) 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 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 D.3. Appendix

[21] 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.

Figure D.9 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 56% and the median absolute value wedge 43%. Intuitively, the robustness of these two moments (the over-consumer share and median absolute value wedge) reflects the fact that they are central moments, which are less sensitive to mean-zero perturbations. Centralization limits the first-order influence of noise, so even moderate measurement error tends not to meaningfully distort these distributional features.

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 two measures are highly correlated (see Appendix Figure D.11) and the distribution of wedge estimates are remarkably similar (see Appendix Figure D.12).

Static Wedges. As a final robustness, we measure a distinct but related consumption wedge that we refer to as a "static" wedge. These wedges are measured using only the one-period-ahead Euler equation. An appealing feature of static wedges is that they can be calculated using only: consumption and one-year-ahead expectations over consumption growth, returns, and inflation. They 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 distortions are currently affecting consumption and thus provide an additional way to validate our conclusions from studying dynamic wedges, while requiring much less data and fewer assumptions.

Appendix D.2 details our measurement and analysis 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 directly comparable to those of dynamic wedges. The median absolute value static wedge is 6.1%; when compared to the typical dynamic wedge of 40%, this suggests that beliefs about future distortions are important determinants of consumption.

5.4 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.[22] Approximation bias increases the frictionless consumption benchmark, which 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 median bias ranges from -10.1 pp to -30.5 pp (e.g., we find a wedge of 50 pp instead of 61.1 pp). 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, generally ranging from 3-8pp. We 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).

This has two important implications. First, approximation bias induces a level shift in the wedges (rather than noise) that leads us to understate the level of over-consumption. Second, the limited cross-sectional variation means that approximation bias is unlikely to explain the dispersion in the distribution of wedges.

5.5 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

[22] In general equilibrium models (i.e, with endogenous income, interest rates, etc.), the covariance of expected consumption growth and the real rate also affects approximation bias. 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).

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 F.1. 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 D.5. 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 they would "save less" when future 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 the 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 non-durable spending response to the March 2021 stimulus payments.[23] 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.

Financial Distress Proxies. Another prediction of a model with only a simple borrowing limit is that financial distress should be most frequent among under-consumers, who cannot access the resources to consume at their frictionless benchmark. We examine how wedges vary with financial distress using self-reports from survey questions and observed low savings balances. Figure 3 Panel C shows the relationship with financial anxiety and Appendix Figure F.2 shows six additional measures.

For all measures, financial distress is more common among over-consumers and less common among under-consumers—the reverse of what a model with borrowing constraints alone would predict. This points to distortions that generate over-consumption, such as present bias and inertia, as better candidates to explain the joint patterns we document between wedges and both MPCs and financial distress. Policymakers concerned with financial distress may therefore want to direct attention to alleviating forces driving over-consumption, including present bias and high costs for committed consumption categories like housing and childcare.

[23] 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

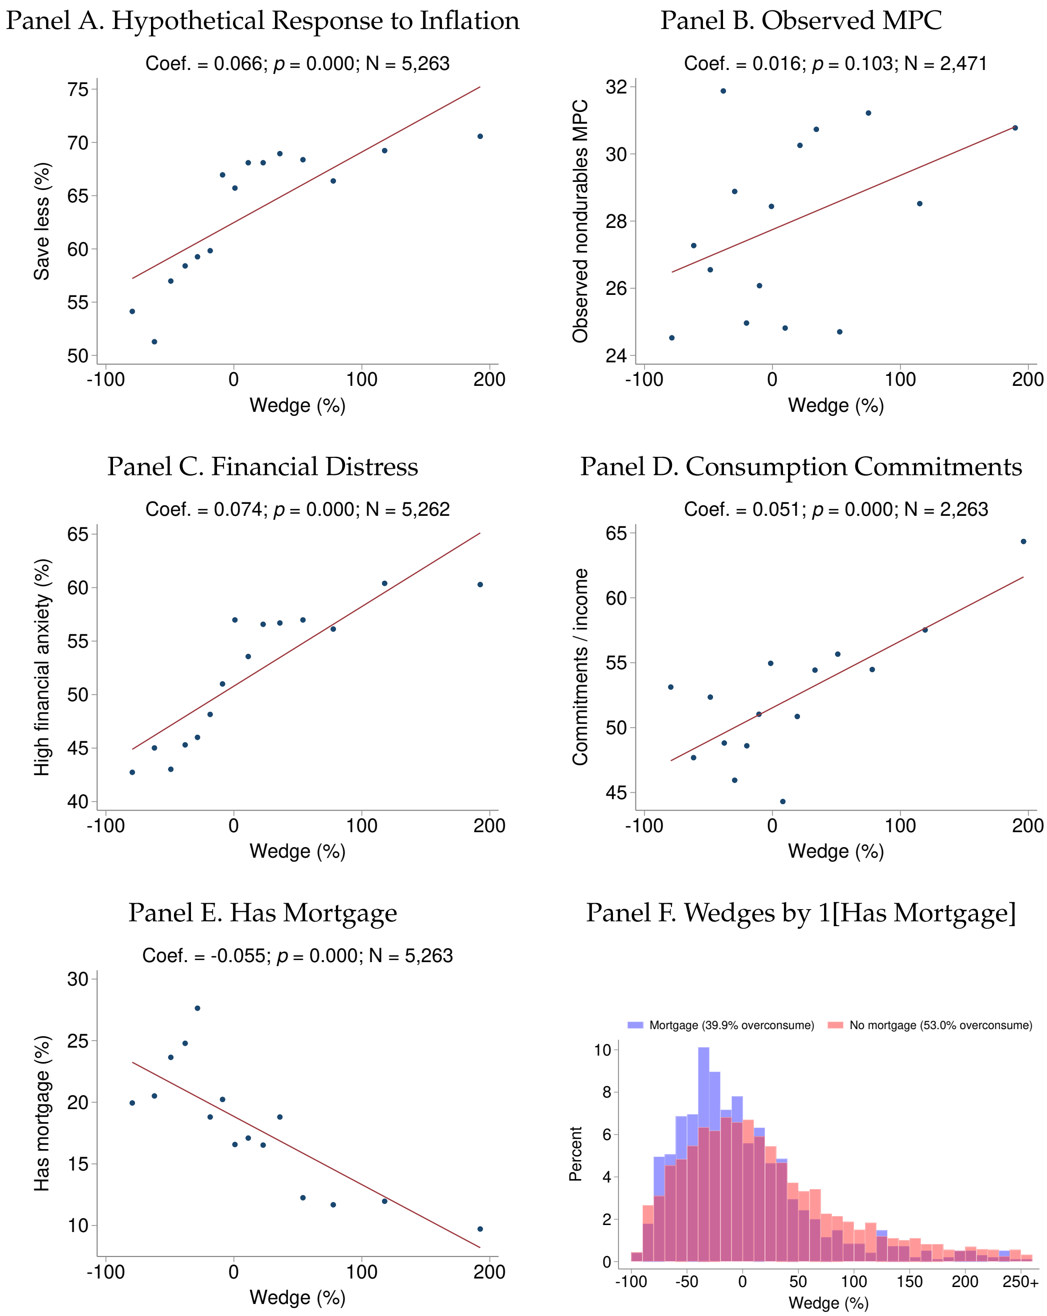

Notes: The figure illustrates the relationship between dynamic consumption wedges and observed nondurable MPCs, consumption commitments (% income), an indicator for whether the user would hypothetically "save less" in response to higher inflation, and an indicator for whether the user reports "high" or "very high" financial anxiety. The binned scatterplot plots the average value within quantile-based intervals of consumption wedges. Variable definitions are outlined in Appendix B.4. MPCs are observed for 59% of users in our analysis sample. Wedges are trimmed at the 1st and 95th percentiles and MPCs are trimmed at the 10th and 90th percentiles.

Consumption Commitments. We next ask 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 the dominant distortions behind 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)), despite being less financially distressed.

We test whether borrowing constraints driven by illiquid assets can reconcile the lower levels of financial distress among under-consumers. 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 next 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). We consider two kinds of a borrowing constraints: a simple lower-bound on how much agents can borrow (i.e., $A_{t+1} \geq \bar{A}$) and a "soft" borrowing constraint (SBC) as in Lee and Maxted (2023) where the interest rate becomes discontinuously higher when agents borrow as opposed to save. We also consider two forms of present bias: naive beta-delta hyperbolic discounting (as in Lee and Maxted, 2023) and temptation preferences (as in Gul and Pesendorfer, 2001; Attanasio et al., 2024). The latter is 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. Finally, We model consumption adjustment costs as a non-pecuniary cost, as in Fuster et al. (2021).

We parameterize each model to have a similar economic environment so that differences across the models can be attributed to differences in their distortions. This includes using the same income process, utility function, and interest rate(s) in each model. We calibrate distortions to best match the share of over-consumers and median absolute value wedge in the EarnIn sample. For full details on each model and the calibration of the economic environment and distortions, see Appendix E.

Table 3 reports statistics on the distribution of wedges across various models in Panel A. We report the same statistics from our empirical sample for comparison in Panel B. As expected, the constant limit and ”soft” borrowing constraint (SBC) models generate no positive wedges. The SBC generates larger and more disperse wedges overall, better resembling the empirical data. For example, the median absolute value wedge is 15.7% versus 10.8% in the SBC model compared to the constant limit model. The remaining models are able to generate a mix of positive and negative wedges.

The model best able generate both a large share of over-consumers and large median absolute value wedge features present bias (specifically, beta-delta hyperbolic discounting) and an SBC. This model delivers 20.2% of households over-consuming and a median absolute value wedge of 29.3%, versus 50.6% and 40.1% for our EarnIn sample (respectively). As such, this model accounts for 40% of the observed over-consumer share and 73% of the median absolute value wedge. When combined with a ”soft” borrowing constraint, both forms of present bias are similarly successful in generating a large share of over-consumers: around 20%. The beta-delta model succeeds in producing larger typical wedges primarily by generating more extreme negative wedges (rather than larger positive wedges).

Consumption adjustment cost (CAC) models, with or without an SBC, are less able to generate a large share of over-consumers. Without an SBC, the CAC model features 13.9% of agents over-consuming, with a median absolute value wedge of 13.1%. Adding an SBC results in fewer over-consumers (9.6%) but a larger median absolute value wedge (23.9%). A key reason that CAC models struggle to generate large positive wedges is that, when consumers have rational expectations, larger adjustment cost increases consumers' precautionary savings motive. This results

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	10.8	-13.9	10.5	-27.9	-17.9	-10.8	-6.7	-4.4
Soft BC (SBC)	0.0	15.7	-20.4	16.0	-45.7	-32.7	-15.7	-7.1	-3.5
Cons. adj. cost (CAC)	13.9	13.1	-13.4	13.4	-29.8	-21.0	-12.8	-5.4	2.5
Beta-Delta Pref. + SBC	20.2	29.3	-25.8	30.1	-60.9	-45.8	-27.9	-8.2	12.4
Temptation + SBC	22.2	21.8	-19.6	24.6	-49.3	-36.5	-19.8	-4.7	13.2
CAC + SBC	9.6	23.9	-23.3	17.3	-44.1	-34.7	-23.7	-12.9	-0.6
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. 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. Panel B reports the same statistics calculated using wedges from our empirical analysis.

in lower consumption and more negative wedges relative to the frictionless benchmark. We observed in our calibration that increasing the cost generally increased both the over-consumer share and typical wedge size, but to a decreasing extent (plateauing near the calibrated cost). A larger adjustment cost makes consumers more willing to tolerate their consumption drifting above or below the level to which they would like to adjust, increasing positive wedges. However, a countervailing force arises as a larger cost also decreases the desired level of consumption because there is a risk that this level will become infeasible and force a costly adjustment.[24]

It is worth noting that a similar properties arise in models with expense shocks, which we do explicitly incorporate into our analysis. Expense shocks, such as the medical expenditure shocks of Bornstein and Indarte (2023), can create positive wedges. But a similar countervailing force arises where the anticipation of such shocks increases the precautionary savings motive.

While some models generate more realistic wedges than others, all fall meaningfully short of the 50.6% over-consumer share and 40.1% median absolute value wedge. 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.

The first feature we propose incorporating is non-rational expectations. A fundamental tension in FIRE models is that frequent and large over-consumption will eventually 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.

Persistently biased beliefs can relax this tension. Consider, a consumer who is persistently

[24] This increase in the precautionary savings motive is even stronger in a CAC model with a pecuniary cost of adjustment, as the ability to use this wealth to cover the cost increases incentives to save.

overly-pessimistic about her permanent income, meaning that her beliefs are such that they 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 end up with a large positive wedge in equilibrium. And because is she persistently wrong about her permanent income, she is slower to run down her wealth or face a binding borrowing constraint.

Further evidence supporting the potential necessity of non-rational expectations comes from Section 4.2. The covariance decomposition indicated that distortions explain just under half of the cross-sectional variation in consumption unaccounted for by frictionless *FIRE* consumption, with the majority attributable to deviations in beliefs from FIRE. The models we study here use only distortions under FIRE and do not fully account for the observed over-consumer share and median absolute value wedge. While we lack a formal result relating these moments to the covariance decomposition, the large explanatory power of deviations from FIRE suggests that incorporating this feature into models with CAC or present bias and an SBC could generate wedge distributions more similar to our empirical findings.[25]

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 of the world.

Taking Stock. Qualitatively, consumption adjustment costs (CACs) or present bias combined with a borrowing constraint can account for the mix of positive and negative wedges we find. In our quantitative illustration, present bias and borrowing constraints is best able to generate similarly large shares of over-consumers and typical wedges. Incorporating persistently biased beliefs is likely necessary to produce a model that generates larger and more positive wedges.

[25] 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.

7 Conclusion

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

Our main findings indicate that distortions play an important role in driving consumption and call into question the dominant role played by borrowing constraints in explaining the consumption of low-liquidity consumers. Most individuals in our sample have large distortions; the median (absolute value) distortion stands at 40% of frictionless consumption. The average distortion is 15%, but this belies significant heterogeneity in the cross section. In particular, there is a mix of positive and negative wedges. 51% of wedges are positive (over-consumption) while the remaining 49% are negative. Because borrowing constraints can only generate negative wedges, the 51% of over-consumers cannot be explained by borrowing constraints. Additional or alternative mechanisms are necessary to explain the consumption choices of low-liquidity consumers. We identify two promising alternatives. A combination of borrowing constraints and present bias could potentially generate a similar distribution of wedges, as present bias creates positive wedges. Additionally, consumer inertia (e.g., consumption commitments) can give rise to both positive and negative wedges. Correlating wedges with MPCs, financial distress proxies, homeownership, and proxies for consumption commitments—further suggests the two alternative models hold promise to better explain the consumption choices of low-liquidity consumers. Studying wedges in a set of quantitative models, present bias combined with borrowing constraints yields a distribution most similar to our sample's.

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

References

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

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

Albanesi, Stefania and Jaromir Nosal, "Insolvency after the 2005 bankruptcy reform," Technical Report, National Bureau of Economic Research 2018.

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

Attanasio, Orazio, 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.

Attanasio, Orazio P 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," Technical Report, National Bureau of Economic Research 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. https://www.trade.gov/ecommerce-sales-size-forecast.

Bornstein, Gideon, "Entry and Profits in an Aging Economy: The Role of Consumer Inertia," Technical Report, National Bureau of Economic Research 2025.

__ **and Sasha Indarte**, "The Impact of Social Insurance on Household Debt," *Available at SSRN 4205719*, 2023.

Boutros, Michael, "Windfall Income Shocks with Finite Planning Horizons," Technical Report, Bank of Canada 2022.

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

Campbell, John Y and N Gregory Mankiw, "Consumption, Income, and Interest Rates: Reinterpreting the Time Series Evidence," *NBER Macroeconomics Annual*, 1989, *4*, 185–216.

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

Chetty, Raj and Adam Szeidl, "Consumption Commitments and Risk Preferences," *The Quarterly Journal of Economics*, 2007, *122* (2), 831–877.

Christiano, Lawrence J, Martin Eichenbaum, and Charles L Evans, "Nominal Rigidities and the Dynamic Effects of a Shock to Monetary Policy," *Journal of political Economy*, 2005, *113* (1), 1–45.

Coibion, Olivier, Dimitris Georgarakos, Yuriy Gorodnichenko, and Maarten Van Rooij, "How Does Consumption Respond to News about Inflation? Field Evidence From a Randomized Control Trial," *American Economic Journal: Macroeconomics*, July 2023, *15* (3), 109–52.

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

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

__ **, Ulrike Malmendier, and Michael Weber**, "What Do the Data Tell Us About Inflation Expectations?," in "Handbook of Economic Expectations," Elsevier, 2023, pp. 133–161.

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

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

__ **, Ulrike Malmendier, Juan Ospina, and Michael Weber**, "Exposure to Grocery Prices and Inflation Expectations," *Journal of Political Economy*, 2021, *129* (5), 1615–1639.

Fagereng, Andreas, Martin B Holm, and Gisle J Natvik, "MPC Heterogeneity and Household Balance Sheets," *American Economic Journal: Macroeconomics*, 2021, *13* (4), 1–54.

Falk, Armin, Anke Becker, Thomas Dohmen, Benjamin Enke, David Huffman, and Uwe Sunde, "Global Evidence on 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.

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*, 103834.

Gross, Tal, Matthew J Notowidigdo, and Jialan Wang, "The Marginal Propensity to Consume Over the Business Cycle," *American Economic Journal: Macroeconomics*, 2020, *12* (2), 351–384.

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

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: Journal of the Econometric Society*, 1982, pp. 1269–1286.

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

Ilut, Cosmin and Rosen Valchev, "Economic Agents as Imperfect Problem Solvers," *The Quarterly Journal of Economics*, 2023, *138* (1), 313–362.

Indarte, Sasha, "The costs and benefits of household debt relief," *INET Private Debt Initiative Technical Report*, 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," 2021.

Kaplan, Greg and Giovanni L Violante, "A Model of the Consumption Response to Fiscal Stimulus Payments," *Econometrica*, 2014, *82* (4), 1199–1239.

_ and _ , "The Marginal Propensity to Consume in Heterogeneous Agent Models," *Annual Review of Economics*, 2022, *14* (1), 747–775.

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

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," Technical Report 2023.

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

Laibson, David, Peter Maxted, and Benjamin Moll, "A Simple Mapping From MPCs to MPXs," Technical Report, National Bureau of Economic Research 2022.

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

Lian, Chen, "Mistakes in Future Consumption, High MPCs Now," *American Economic Review: Insights*, 2023, *5* (4), 563–581.

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

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

_ and Olivia S Mitchell, "The Economic Importance of Financial Literacy: Theory and Evidence," *American Economic Journal: Journal of Economic Literature*, 2014, *52* (1), 5–44.

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," *Available here: https://static1. squarespace. com/static/6186b3b155561c2ab5fe4957*, 2022, *62*, 1659150753483.

_ , David Laibson, and Benjamin Moll, "Present Bias Amplifies the Household Balance-Sheet Channels of Macroeconomic Policy," *The Quarterly Journal of Economics*, 2024, p. qjae026.

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

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

Weber, Michael, Francesco D'Acunto, Yuriy Gorodnichenko, and Olivier Coibion, "The Subjective Inflation Expectations of Households and Firms: Measurement, Determinants, and Implications," *Journal of Economic Perspectives*, 2022, *36* (3), 157–184.

_ , Yuriy Gorodnichenko, and Olivier Coibion, "The Expected, Perceived, and Realized Inflation of US Households Before and During the COVID19 Pandemic," *IMF Economic Review*, 2023, *71* (1), 326–368.

Zeldes, Stephen P, "Consumption and Liquidity Constraints: An Empirical Investigation," *Journal of political economy*, 1989, *97* (2), 305–346.

Online Appendix

Contents

A Theory Derivations and Extensions

A.1 Frictionless Consumption Derivation

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

Lemma 1: Euler Equation Approximation

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

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

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

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

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

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

yields

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

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

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

Under non-CRRA preferences, we obtain the same expression using the a log-linear first-order approximation of the marginal utility function below:

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

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

The derivation for the multi-period Euler equation:

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

is similar, yielding:

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

□

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

Lemma 2: Budget Constraint Approximation

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

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

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

Proof. We begin with the period t budget constraint:

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

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

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

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

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

And then we take expectations:

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

and then a first-order approximation:

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

□

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

Proposition 1: Frictionless Consumption

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

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

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

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

Using the above expression, we substitute in for $\widetilde{E}_{i,t}C^{\star}_{i,t+j}$ in the approximate budget constraint equation from Lemma 2:

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

□

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

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 a constant borrowing limit (e.g., Aiyagari, 1994), an endogenous borrowing limit (e.g., Bornstein and Indarte, 2023), or "soft" constraints arising from discrepancies in borrowing and saving rates (e.g., Kaplan et al., 2018). These frictions can introduce a wedge into the Euler equation, relative to the frictionless Euler Equation (3). For example, consider a constant borrowing limit such as $A_{i,t+1} \geq \overline{A}$. The Euler equation would be:

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

where $\mu_{i,t} \geq 0$ is the Lagrange multiplier. The Lagrange multiplier $\mu_{i,t}$ is positive if and only if the constraint is binding. All else equal, a binding constraint reduces consumption $C_{i,t}$. An important feature of 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 additional factor $\overline{\beta} < 1$ relative to the standard exponential discounting model (where δ is the exponential discount factor, corresponding to β in our notation above). Under these preferences, the expectation term in the Euler equation is scaled down by the degree of present bias ($\overline{\beta}$):

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

As a result, these preferences cause consumption to be higher relative to a "debiased" ($\overline{\beta} = 1$) consumer (Maxted, 2022), all else equal. Hence, present bias creates positive wedges. Similar to 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:

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

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

Empirical evidence on consumption wedges can help guide the choice and modeling of frictions. Qualitatively, the presence of both positive and negative wedges would indicate that neither 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

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.

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.3.1 Extension: Durable and Non-Durable Goods

We next show how to extend our wedge measurement results to accommodate durable goods. Durable goods present several complications: it's difficult to measure their consumption and depreciation directly, holdings of durable goods constitute a source of wealth, and some are financed with debt. While these feature do not complicate or alter the derivation of the wedge, the present measurement challenges that complicate applying our wedge formula. To overcome these challenges, we make assumptions that imply that the expenditure share of non-durable goods is a constant, known fraction. The key assumption is that notional consumption is a Cobb Douglas aggregate of both types of consumption goods.

Notation. Let $n_{i,t}$ and $d_{i,t}$ denote i's real period t consumption flows of non-durable and durable goods (respectively). We continue to denote the total nominal value of net worth by $A_{i,t}$. Total wealth includes net positions in durables (e.g., the value of vehicles net of the loans used to finance their purchase). The consumer has preferences over notional consumption flows $c_{i,t}$, which are an aggregate of non-durable and durable consumption flows (i.e., utility $u(c_{i,t})$ is the per-period utility flow). We make two assumptions.

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

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

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

We let non-durables, $n_{i,t}$, be the numeraire good. Under Assumption 1, we can write the consumer's budget constraint simply as

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

where

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

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

Assumption 2: Cobb Douglas Aggregation.

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

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

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

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

The intratemporal optimality conditions indicate that expenditure on each good is a constant share of total expenditures on consumption goods. As a result, we can infer nominal notional consumption $C_{i,t}$ from nominal non-durable consumption $n_{i,t}$ and the expenditure share α. This is formalized in the lemma below.

Lemma 3: Consumption Calculation Including Consumption of Durables

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

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

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

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

A.4 Static Wedges

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

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

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

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

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

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:

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

where

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

Proof. We begin by establishing notation. Let

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

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

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

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

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

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

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

□

B Data Construction

B.1 Transactions Data

B.1.1 Data Structure

We receive anonymized data from EarnIn that covers user-level information, bank transactions, daily checking and savings account balances, and transactions classified as earnings. All data are de-identified and stored on secured servers. The dataset spans January 2021 through November 2024, covering at least 12 months before each survey wave. Users remain in the dataset until they either delete their accounts or disconnect their linked bank account. Among the users in our analysis sample, 97% had at least one EWA cashout in the 12 months preceding the survey, with a median annual cashout amount of approximately $4,600.

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

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

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

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

B.1.2 Categorizing Transaction Inflows

We leverage the transactions and observed earnings datasets to construct a measure of income, which we define as the sum of post-tax labor earnings and unemployment insurance (UI).

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

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

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

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

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

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

B.1.3 Categorizing Transaction Outflows

Our analysis focuses on nondurables spending. To construct this measure, we apply an outflows categorization algorithm that separates durables and nondurables spending from other types of outflows, including payments (e.g., interest and principal on loans, bank fees), internal transfers (i.e., transfers across checking, savings, or other accounts), and external transfers (i.e., transfers to other individuals or entities through Zelle, Venmo, or other platforms). This algorithm builds on the approaches of Ganong and Noel (2019) and Lusardi (1996), with modifications tailored to the structure of our data and the goals of our analysis.

The Plaid transaction taxonomy included in the EarnIn database comprises over 500 granular categories. We first map these categories to 33 broader categories that can be grouped under three overarching types: spending, payments, and transfers.

- **Spending:** Auto parts & repair, cash, department stores, discount stores, drug stores, digital entertainment, other entertainment, food services, gas stations, grocery stores, healthcare, home improvement, insurance, personal care services, professional services, taxis, transportation, travel, utilities, wholesale stores, other durables, other nondurables, other retail

- **Payments:** Auto loans, non-auto loans, buy now pay later, EarnIn earned wage access, other earned wage access, housing, overdraft & late fees, other payments

- **Transfers:** Checks, transfers across bank accounts, transfers to investment accounts, credit card payments, peer-to-peer transfers, other transfers

This mapping has three key limitations. First, Plaid's categorization is based on merchant types rather than the specific products and services purchased, making it difficult to distinguish between durables and nondurables in certain cases. For example, a transaction at a department store may include both a mattress (durable) and makeup (nondurable). Second, some Plaid categories are too broad or ambiguous to be mapped reliably, such as "Purchase," "Shopping," or "Transfer." Finally, we observe frequent misclassifications in Plaid's categorization.

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

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

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

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

- **Durables:** Auto parts & repair, home improvement, insurance, other durables

- **Nondurables:** Cash, digital entertainment, other entertainment, food services, gas stations, groceries, healthcare, personal care services, professional services, taxis, transportation, travel, utilities, other nondurables

Table 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 ecommerce revenue in 2020 Bledsoe (2024).

- **Payments:** Auto loans, non-auto loans, buy now pay later, EarnIn earned wage access, other earned wage access, housing, overdraft & late fees, other payments
- **Internal transfers:** Transfers across bank accounts, transfers to investment accounts, credit card payments, other internal transfers
- **External transfers:** Checks, peer-to-peer transfers, other external transfers

B.2 Survey Outreach and Response

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

- Wave 1 (September 2022)
  - Non-missing earnings data at least once between September 2021 and August 2022
  - Non-missing balances data in each bi-weekly period from September 2021 through August 2022
  - First recorded transaction before September 1, 2021 and latest recorded transaction after August 15, 2022
  - At least 5 outflows per month between September 2021 and August 2022
  - Non-missing bank connection date
- Wave 2 (July 2024; resampled Wave 1 users)
  - Completed the wave 1 survey
  - Still in the EarnIn database as of June 2024
  - Took at least 3.5 minutes to complete the wave 1 survey
  - Reported consistent debt amounts in the wave 1 survey (i.e., users who report zero debt must report N/A for debt manageability, and vice versa)
  - At least 20 outflows per month each month between June 2023 and May 2024
  - Non-missing balances data each week for at least 9 months between June 2023 and May 2024
  - Sufficient categorizable spending ($\frac{\text{Consumption}}{\text{Outflows - Internal Transfers}} \geq 20\%$) for at least 9 months between June 2023 and May 2024
  - Reasonable balance of inflows and outflows ($\frac{\text{Outflows}}{\text{Inflows}} \in [50\%, 150\%]$) for at least 9 months between June 2023 and May 2024

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

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

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

For wave 1, EarnIn sent invitations in waves and closed the survey after 240,000 were sent, at which point 10,103 respondents had completed the survey and our incentive budget was fully spent. For waves 2 and 3, EarnIn sent invitations to the full sampling frames and closed each survey after two reminder emails, resulting in 875 responses for wave 2 and 4,888 responses for wave 3. As a result, we received and aggregate response rate of 3.4% – 4% for wave 1, 19% for wave 2, and 2% for wave 3.

We compare summary statistics for survey respondents against the eligible sample in Appendix Table B.2. Users invited to the survey look broadly similar to users who responded on all the observable financial outcomes (monthly inflows, outflows, and balances across checking and savings accounts), but respondents are more likely to be female in survey waves 1 and 3.

B.3 Sample Restrictions

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

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	6,677	6,609	6,380	6,285
Monthly Outflows	6,245	6,039	6,865	6,903	6,534	6,529	6,296	6,212
Balance (Checking)	517	526	618	717	550	659	522	573
Balance (Savings)	296	353	277	386	274	427	293	378
Female (%)	67.2	52.4	73.3	69.0	70.6	52.5	67.8	52.4
N (Users)	10,090	578,123	875	4,652	2,160	315,982	12,250	894,105

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

- First, when merging survey respondents to the EarnIn database we drop respondents who deleted their EarnIn accounts or delinked their bank accounts. This reduces the number of users from 14,991 to 14,817.
- Second, we apply the survey data quality restrictions listed below. This reduces the number of users from 14,817 to 14,386.
  - Survey duration at least 3.5 minutes (approximately the 5$^{\text{th}}$ percentile)
  - Reported debt amounts are consistent (i.e., users who report zero debt must report N/A for debt manageability, and vice versa)
- Third, we apply the transactions data quality restrictions listed below. These restrictions are designed to drop users who do not primarily consume through the bank accounts connected to EarnIn, which limits the extent to which we observe their consumption. We apply restrictions using the 12 months prior to each survey and up to 12 months after each survey (we observe 12 post-survey months for wave 1, 5 for wave 2, and 0 for wave 3). This reduces the number of users from 14,386 to 8,944.
  - Sufficient transaction activity: 20+ outflows per month for all months
  - Sufficient balances data: Non-missing balances each week for at least 75% of months
  - Sufficient categorizable spending: $\frac{\text{Consumption}}{\text{Outflows - Internal Transfers}} \geq 20\%$ for at least 75% of months
  - Reasonable balance of inflows and outflows: $\frac{\text{Outflows}}{\text{Inflows}} \in [50\%, 150\%]$ for at least 75% of months
  - Informative memos: $< 1\%$ of memos are "CREDIT", "DEBIT", or missing across months

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

Table B.3. Successive Survey Sample Restrictions

	Wave 1	Wave 2	Wave 3	Total
Users who recived an email	240,000	4,652	218,615	463,267
Total responses	14,434	926	5,342	20,702
Finished responses	10,103	875	4,888	15,866
Can merge w/ 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

Notes: Table shows the resulting number of unique individuals after successive sample restrictions from top to bottom in each column. Since Wave 2 is comprised of a subset of respondents to Wave 1, the ”Total” column represents the number of user $\times$ survey wave observations.

Table B.4. Successive Transactions Sample Restrictions

	Dynamic Wedge	Static Wedge
N after survey restrictions	12,525	4,593
At least 20 outflows per month	9,592	4,368
Consumption ¿= 20%, at least 75% of months	8,740	3,798
Outflows between 50–150% inflows, at least 75% of months	8,369	3,583
Less than 1% of outflows with uninformative memos	8,205	3,546
Non-missing bank balances, at least 75% of months	7,891	3,541
Trimmed non-durables spending (P1–P99)	7,737	3,473
Trimmed income (P1–P99)	7,611	3,416
Trimmed non-durables APC (P2.5-P97.5)	7,302	3,274
Trimmed net worth to income ratio (P5–P95)	6,631	2,973
Trimmed expected levered rate (P3–P97)	6,238	2,791
Wedge computation restrictions	5,595	2,782
Wedge trimming restrictions (Dyn. P1–P95, Stat. P1–P99)	5,261	2,727

Notes: The table shows the resulting number of unique survey date and hash ID combinations after successive sample restrictions from top to bottom. Columns are listed separately for static and dynamic wedges. N after survey restrictions corresponds to the sum of Waves 1, 2, and 3 for the dynamic wedge and Waves 2 and 3 for the static wedge. These sample sizes correspond with the sample sizes listed at the bottom of Table B.3.

B.4 Variable Measurement

B.4.1 Consumption, Income, and APCs

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

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

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

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

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

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.

B.4.4 Marginal Propensity to Consume

Observed MPC We measure individual-level MPCs based on consumers' non-durable spending responses to the March 2021 stimulus payments. These checks provided \$1,400 to each eligible individual, with an additional $1,400 for each dependent.[27] 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 30%. There are a few extreme outliers (e.g., below -500% or above 500%), likely due to large, one-time purchases. Given this feature of the data, we trim MPCs at the 10th and 90th percentiles.

C Imputations

C.1 Wealth

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

Table C.1 presents the median absolute error of our model, weighted by the liquid wealth bins of the EarnIn sample. This error captures the discrepancy between our predictions and the actual values reported in the SCF data. The magnitude of error in this model stems from the limited information we have on illiquid assets in the survey.

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

Table C.1. Median Absolute Error, Weighted by EarnIn Sample

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

C.2 Term Structure of Beliefs

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

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

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

C.2.1 Inflation Expectations

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

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

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

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

We impute missing expectations using the following approach:

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

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

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

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

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

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

C.2.2 Interest Rate Expectations

We assume that levered interest rate expectations are constant over the horizon.

C.2.3 Income Growth Expectations

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

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

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

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

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

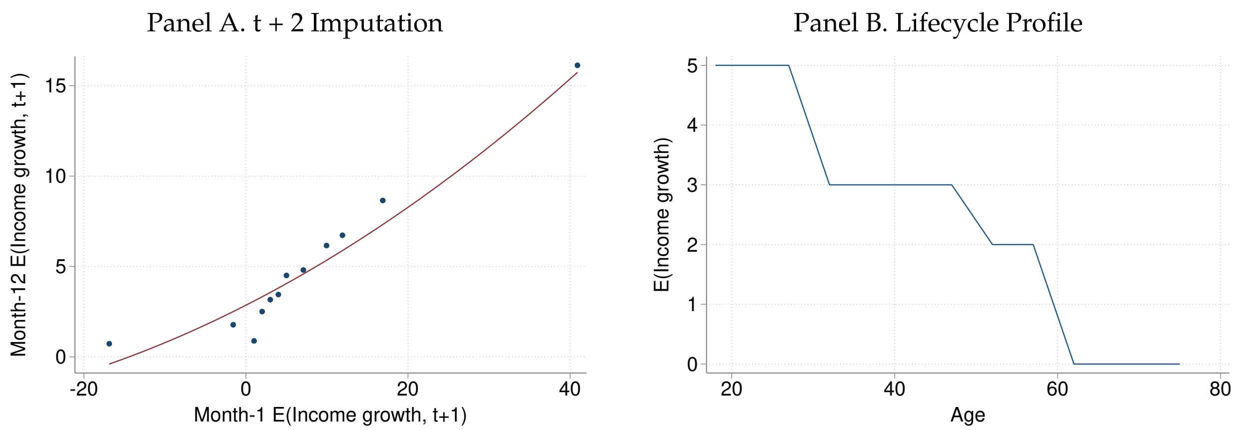

Notes: The figures describe the term structure of imputed income growth expectations. Panel A shows data from the Survey of Consumer Expectations (SCE), which measures one-year-ahead expectations in a 12-month rotating panel. From this, we use the t + 1 expectation in month 12 to impute the t + 2 expectation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

D Robustness

D.1 Sensitivity Analysis

Table D.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%	34.9%
II	$\beta = 0.94, \gamma = 1.05$	33.7%	27.6%
III	$\beta = 0.72, \gamma = 0.35$	21.6%	21.8%

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

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

Panel A. Sensitivity of Dynamic Wedges to Beta

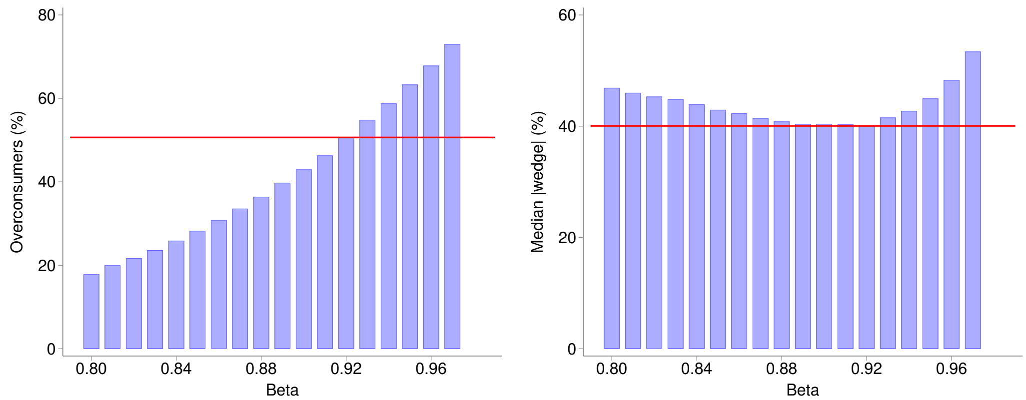

Panel B. Sensitivity of Dynamic Wedges to Gamma

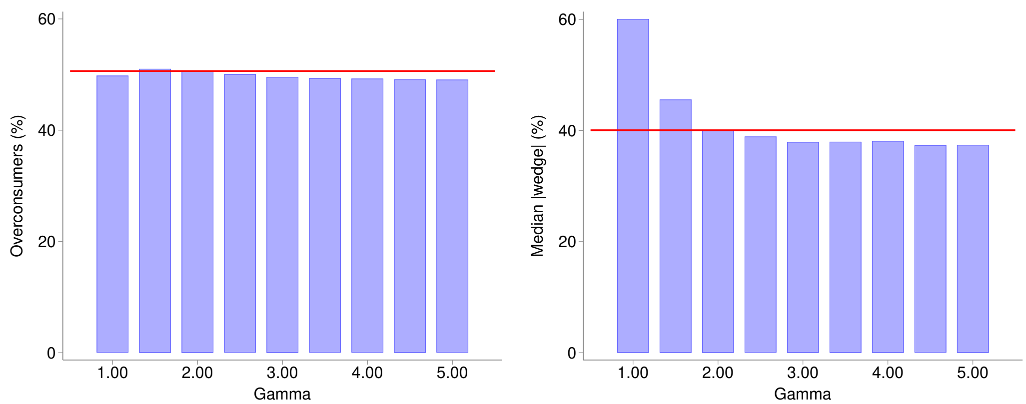

Panel C. Sensitivity of Dynamic Wedges to Nondurables Share of Spending

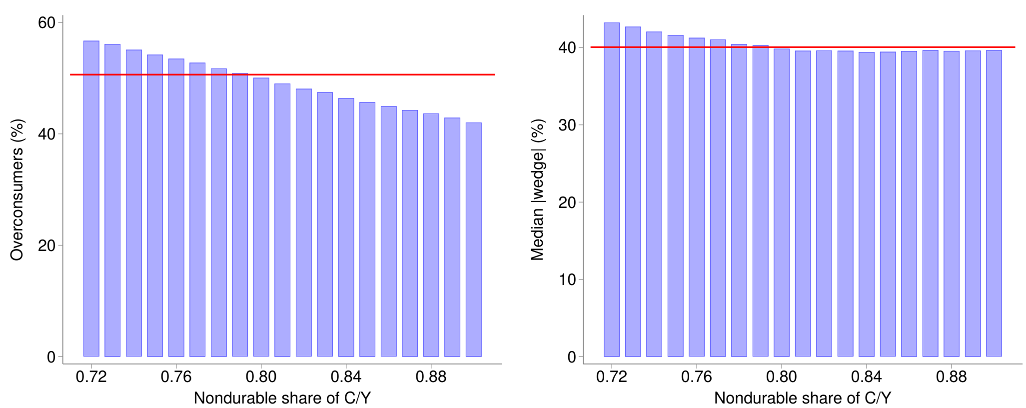

Notes: The figure presents the sensitivity of two main wedge moments to our assumed value of beta, gamma, and non-durable share of spending in the dynamic wedge calculation. We vary the assumed value of beta ranging from 0.80 to 0.98 in increments of 0.01 (our baseline calibration is 0.92) and the assumed value of gamma ranging from 1.0 to 5.0 in increments of 0.5 (our baseline calibration is 2.0). We also vary the assumed value of non-durable share ranging from 72% to 90% in increments of 1 percentage point (our baseline calibration is 79.37%). These results include the percent of users with a positive wedge and the median absolute value wedge. We hold all other parameters constant at our baseline values. Wedges are trimmed at the 1st and 95th percentiles.

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

Panel A. Raw wedges | Panel B. Absolute value of wedges

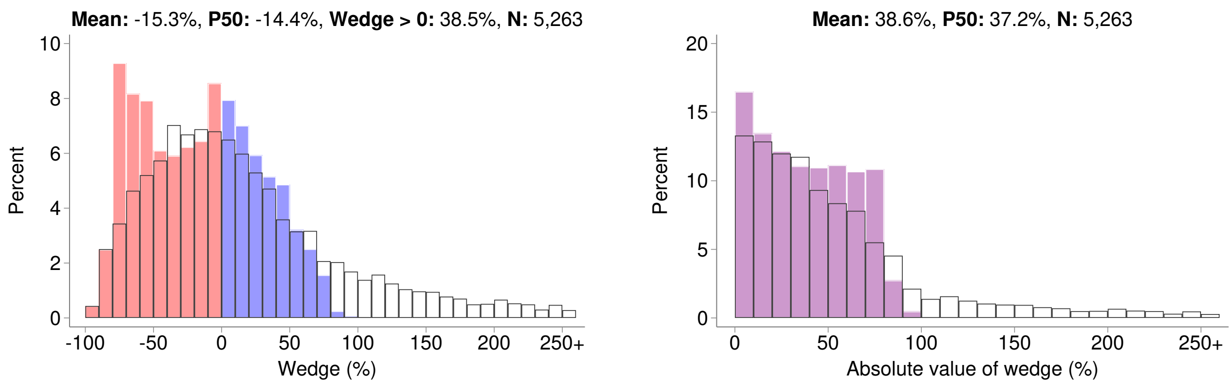

Notes: The figures show the distribution of dynamic consumption wedges (left) and their absolute values (right), assuming users have one of the three preference types outlined in Table 11 of Aguiar et al. (2024). We assign users to preference types based on which type minimizes the absolute value of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

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

Panel A. Share of Over-consumers | Panel B. Median |Wedge|

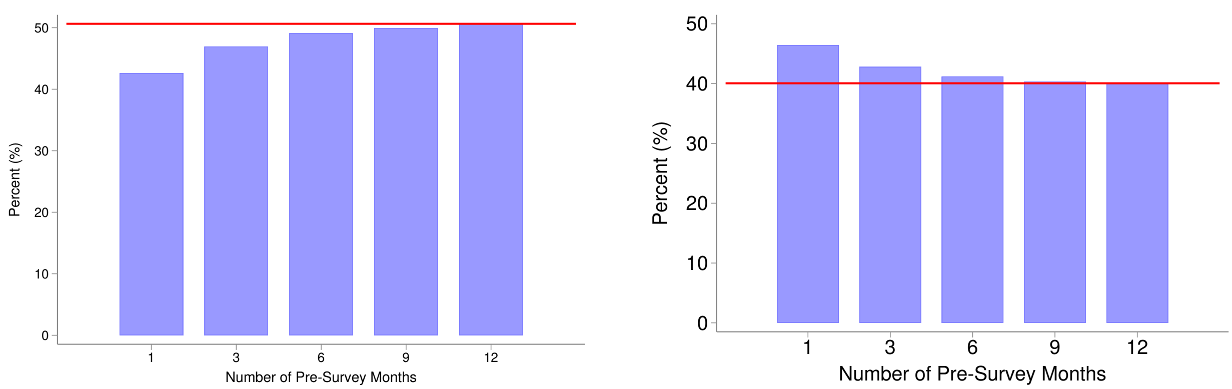

Notes: Figure presents the sensitivity of two estimated results to varying the number of months of pre-survey data. These results include the percent of users with a positive wedge (Panel A) and the median absolute value wedge (Panel B). Wedges are trimmed at the 1st and 95th percentiles. Our baseline specification uses 12 months of pre-survey data. When using less than 12 months of data, we annualize income when calculating the wealth-to-income ratio.

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

Panel A. Raw wedges | Panel B. Absolute value of wedges

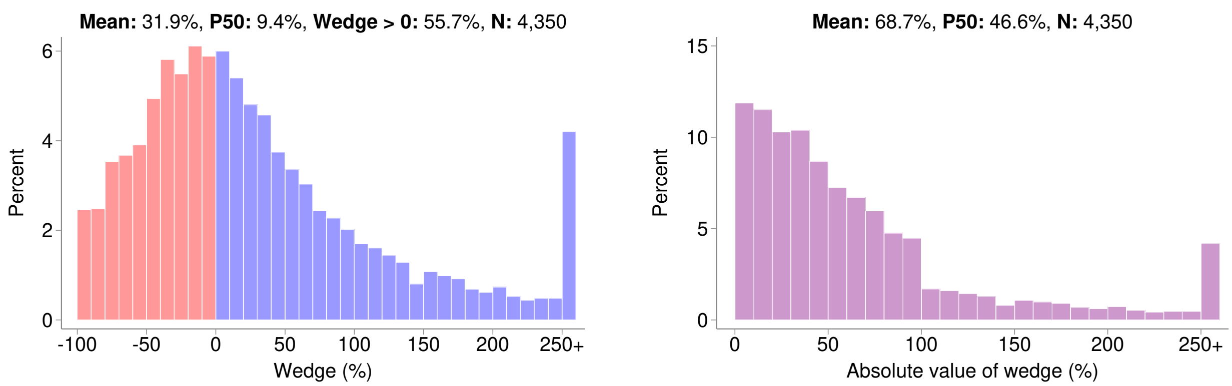

Notes: The figures show the distribution of dynamic consumption wedges (left) and their absolute values (right), assuming users have a constant term structure of beliefs. Wedges are reported here as the percent deviation of observed consumption from frictionless consumption, and they are trimmed at the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

Figure D.5. Sensitivity of Consumption Wedge to Lifecycle Differences

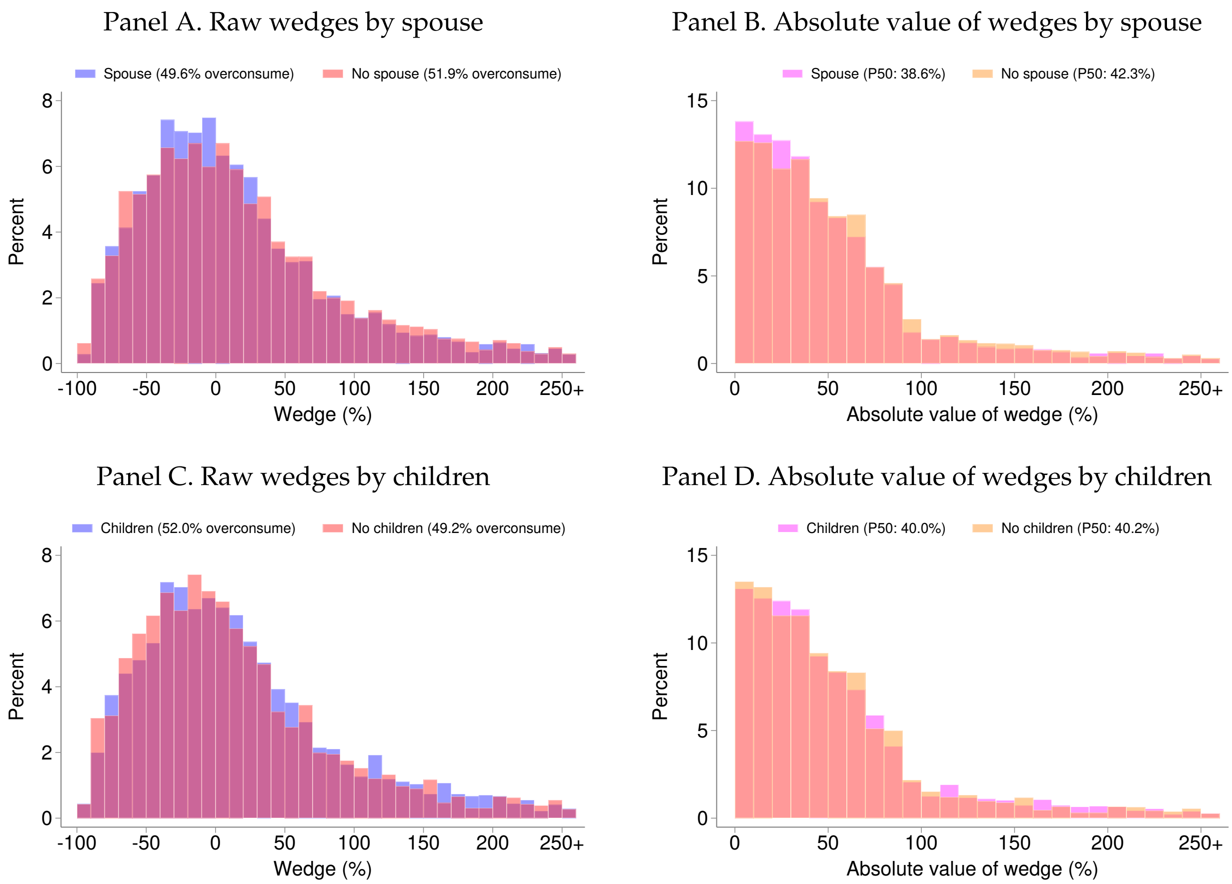

Notes: The figures show the distribution of dynamic consumption wedges (left) and their absolute values (right), broken down by various lifecycle differences. Wedges are reported here as the percent deviation of observed consumption from frictionless consumption, and they are trimmed at the 1st and 95th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

D.2 Static versus Dynamic Consumption Wedges

The wedges we have focused on thus far are "dynamic" in the sense that they reflect the impact of both current and expected future distortions on consumption. Here, we define and measure a related "static" wedge, which is calculated using expected consumption growth and does not require knowledge of the full term structure of beliefs. In contrast to the dynamic wedge, the static wedge only reflects the impact on consumption of distortions experienced at the time of consumption; it omits the effect of expected future distortions. Dynamic wedges are a geometric sum of current and expected future static wedges (see Appendix A.4).

Examining static wedges is valuable for two reasons. First, it serves as a robustness check regarding the importance of over-consumption by providing an alternative measure of distortions that only requires one-year-ahead expectations. Second, comparing an individual's dynamic and static wedges enables us to quantify the relative importance of current versus expected future distortions on consumption.

The static wedge is calculated using only the approximate one-period-ahead Euler equation. The key difference between the static and dynamic wedge is that the static wedge is measured using expected *actual* consumption ($\widetilde{E}_{i,t}C_{i,t+1}$), instead of expected *frictionless* consumption ($\widetilde{E}_{i,t}C^{\star}_{i,t+1}$). Expected frictionless consumption is determined via the frictionless benchmark's budget constraint (and expectations of its components). We define static frictionless consumption to be

$$C^{\star,\text{static}}_{i,t} \approx \frac{\widetilde{E}_{i,t}C_{i,t+1}}{\widetilde{E}_{i,t}\pi_{t+1}} \left(\beta \frac{\widetilde{E}_{i,t}R_{i,t+1}}{\widetilde{E}_{i,t}\pi_{t+1}}\right)^{-1/\gamma}. \tag{11}$$

The static consumption wedge is

$$\nu_{i,t} = C_{i,t} - C^{\star,\text{static}}_{i,t} = C_{i,t}\left[1 - \frac{\widetilde{E}_{i,t}G^{C}_{i,t+1}}{\widetilde{E}_{i,t}\pi_{t+1}} \left(\beta \frac{\widetilde{E}_{i,t}R_{i,t+1}}{\widetilde{E}_{i,t}\pi_{t+1}}\right)^{-1/\gamma}\right] \tag{12}$$

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

Figure D.6 displays the distribution of static consumption wedges. We find a slightly larger share of over-consumers as measured by their static wedge: 56% (versus 51%). Panel B of Appendix Figure D.15 shows that individuals' static and dynamic wedges are highly correlated. The difference in typical absolute value magnitudes—6.1% for static wedges versus 40% for dynamic wedges—implies that expectations about future distortions are an important determinant of consumption. Taking the ratio of individuals' static and dynamic wedges, for the median consumer, the static wedge accounts for 2.4% of their dynamic wedge. The discrepancy between the magnitude of the static and dynamic wedges suggests that beliefs about future distortions play an even more important role in determining consumption than current distortions.

Figure D.6. Distribution of Static Consumption Wedges

Panel A. Distribution of Static Wedges

Panel B. Distribution of |Static Wedges|

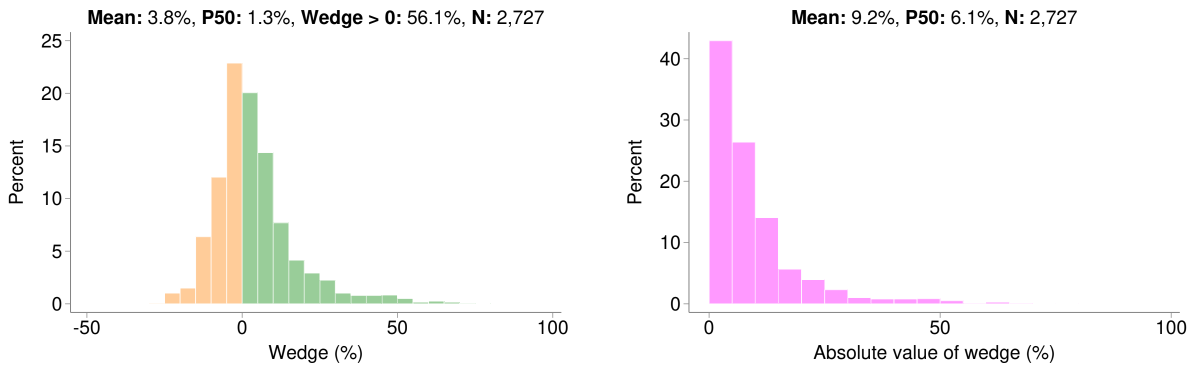

Notes: The figures show the distribution of static consumption wedges (left) and their absolute values (right). Wedges are reported as the percent deviation of observed consumption and frictionless consumption, and they are trimmed at the 1st and 99th percentiles. Includes users from waves 2 and 3 only,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

D.3 Measurement Error

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

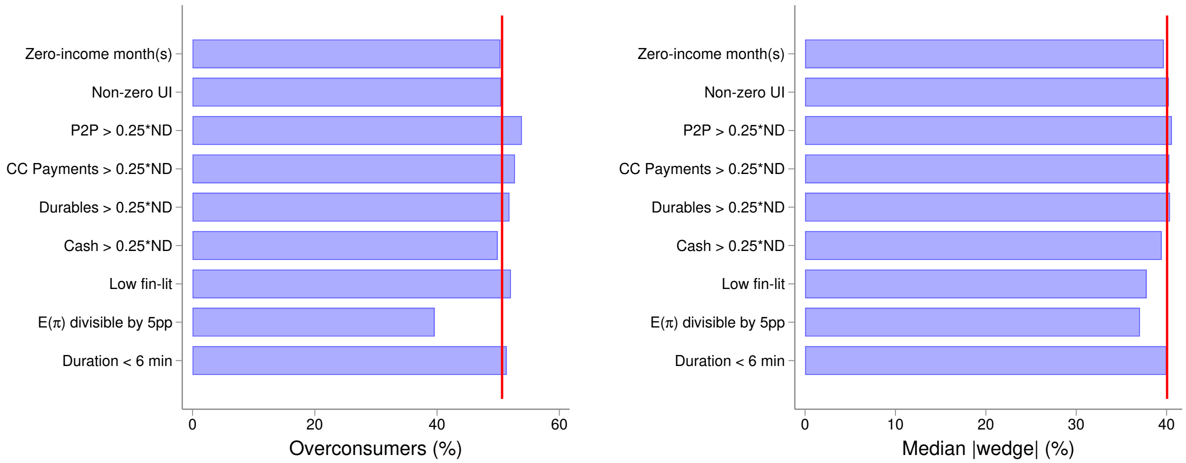

Notes: Figure presents the sensitivity of our two main results to dropping users with plausibly high measurement error. These results include the percent of users with a positive wedge (Panel A) and the median absolute value wedge (Panel B). For reference, our baseline results are shown in red. Sample sizes associated with dropping users in each of the categories are as follows: any zero-income month(s) ($N = 5,486$); non-zero UI ($N = 6,036$); P2P $> 0.25 * ND$ ($N = 3,794$); credit card payments $> 0.25 * ND$ (has)$N = 5,485$); durables $> 0.25 * ND$ ($N = 5,606$); cash $> 0.25 * ND$ ($N = 5,716$); low financial literacy ($N = 2,665$); expectations divisible by 5 ($N = 2,444$); and survey durations $< 6minutes$ ($N = 4,546$). Wedges are trimmed at the 1st and 95th percentiles.

Figure D.8. Robustness of Consumption Wedges to Clustering Users

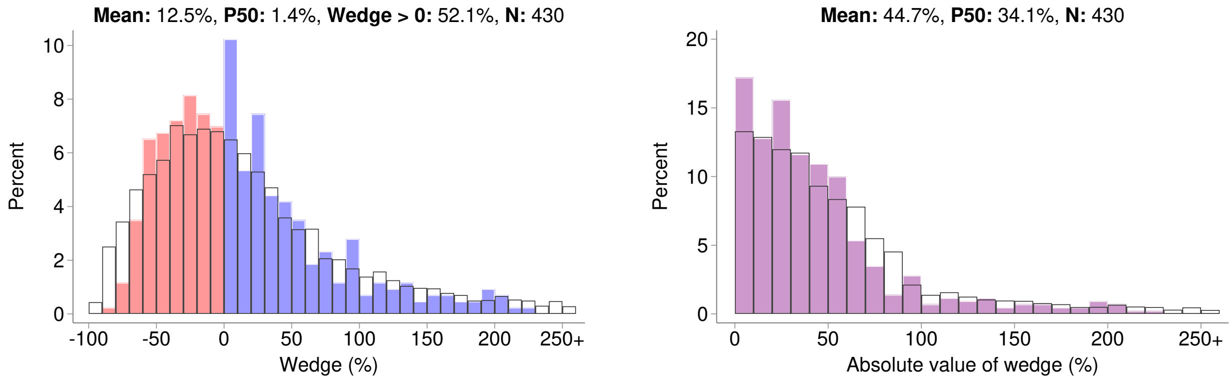

Notes: Figure presents the sensitivity of the wedge distribution to clustering users and taking the within-cluster median of each wedge input before calculating wedges. We use the k-prototypes clustering algorithm with 500 bins. Wedges are trimmed at the 1st and 95th percentiles. The baseline distribution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

Figure D.9. Sensitivity of Consumption Wedge to Noisy Expectations Inputs

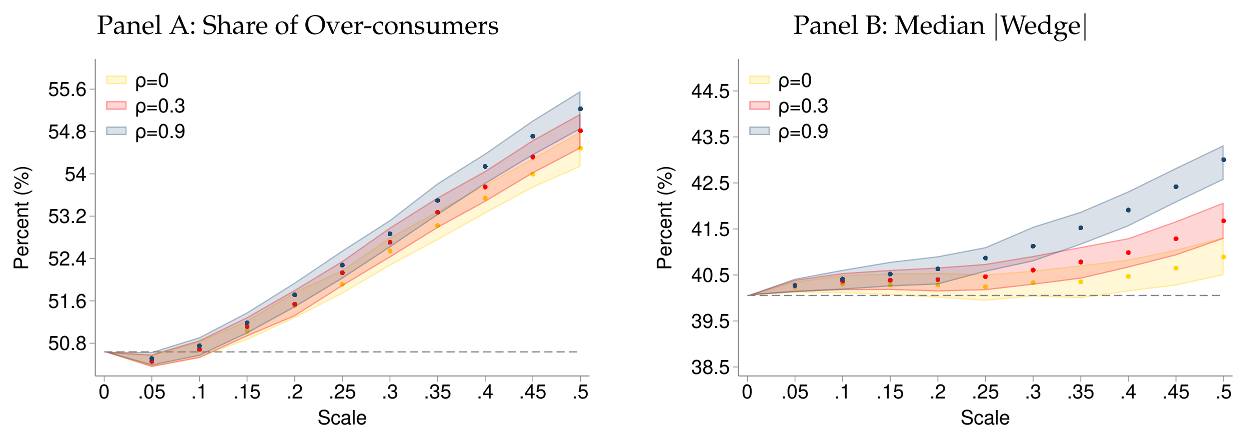

Notes: Figure presents the sensitivity of the two main wedge moments to gradually increasing noise on all expectations inputs over $n = 500$ simulations. Values of scale correspond with the scaling factor of the standard deviation of each expectations input. The solid dot represents the median at each grid point, while the shading represents the inter-quartile range of estimates. Estimates in yellow represent a correlation of $\rho = 0$, estimates in red represent a correlation of $\rho = 0.3$, and estimates in blue represent a correlation fo $\rho = 0.9$.

Figure D.9 presents the results of of $n = 500$ simulations across correlation coefficients $\rho = \{0, 0.3, 0.9\}$ and a grid of noise levels defined by $\sigma_x = \{0, 0.05, 0.10, ..., 0.50\}$. For each scalar and correlation level, we calculate $\tilde{\epsilon}_{ix} = \phi_i\sqrt{\rho} + e_{ix}\sqrt{1-\rho}$ such that $\phi \sim N(0,1)$ and $e_{ix} \sim N(0,1)$. Then, we scale the noise so that the standard deviation of $\tilde{\epsilon}_{ix}$ matches $s \times SD(x)$. This yields $\epsilon_{ix} = \tilde{\epsilon}_{ix} \times s \times SD(x)$. Then, we compute $x_{ix}^{noisy} = x_{ix} + \epsilon_{ix}$. After this, we recompute the term structure and levered interest rate. Finally, we are able to compute our wedge and plot the median and IQR of our two moments of interest across $n = 500$ simulations.

The results of Figure D.9 indicate that even under relatively extreme magnitude and correlation of noise on survey-based expectations inputs, our two main wedge moments remain stable.

Figure D.10. Robustness of Consumption Wedges to Reflect Beliefs about Default

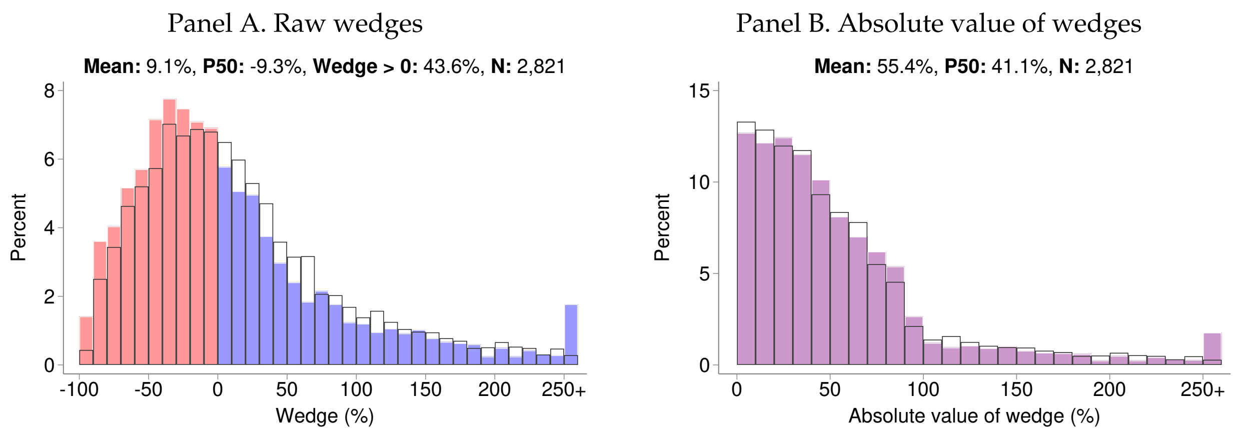

Notes: The figures show the distribution of dynamic consumption wedges (left) and their absolute values (right). Colored bars represent the distribution of consumption wedges where we recalculate R^D_{it}, the gross rate on debt, such that $R^{D*}_{it} = [(1 + R^D_{it}) \times \theta] - 1$ where $\theta = 0.9$. 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 allows us to estimate the distribution of wedges conditional on repaying. Although this is a relatively extreme assumption, it is useful to examine the extent to which are main results are robust to this belief. Meanwhile, the transparent bars represent the baseline wedge distribution. Wedges are reported here as the percent deviation of observed consumption from frictionless consumption, and they are trimmed at the 1$^{\text{st}}$ and 95$^{\text{th}}$ percentiles.

Figure D.10 shows the distribution of dynamic consumption wedges along with their absolute values such that our expected borrowing rate, R^D is recalculated downwards, to reflect actual expected borrowing costs, rather than expected borrowing costs *conditional on repaying*, as we do in our main wedge estimation.

Specifically, we set $R^{D*}_{it} = [(1 + R^D_{it}) \times \theta] - 1$ such that $\theta = 0.9$ to reflect that consumers in our sample expect to default on 10% of all debt payments in every period for the rest of their lives. This is a relatively extreme assumption, but it is useful to examine the extent to which our main results would change under this belief. After calculating R^{D*}_{it}, we recalculate our consumption wedges to generate the above distributions.

As we can see, even under relatively extreme beliefs about borrowing rate conditional on debt repayment, our main wedge moments remain stable.

D.4 Survey-Based Income Wedge Recalculation

To evaluate the sensitivity of our results to our observed income measure, we re-calculate our wedge distribution by using an approximation of post-tax labor income using reported pre-tax labor income from our survey data.

To map the reported pre-tax labor income in our survey to post-tax labor income, we use CPS data and compute $Post\_Tax\_Labor\_Income = PreTax\_Labor\_Income - FICA - Labor\_Pct(FedTax + StateTax)$. Then, we calculate post-tax labor income percentiles using survey weights and assign users to the 11 total income bins which we use in our survey. Then, this is collapsed to the median of each total income bin so that we are able to merge with our main analysis sample based on the each total income bin. We present the distribution of this new imputed reported post-tax labor income against our standard observed post-tax income measure in Figure D.11.

Figure D.11. Distribution of Observed and Reported Income

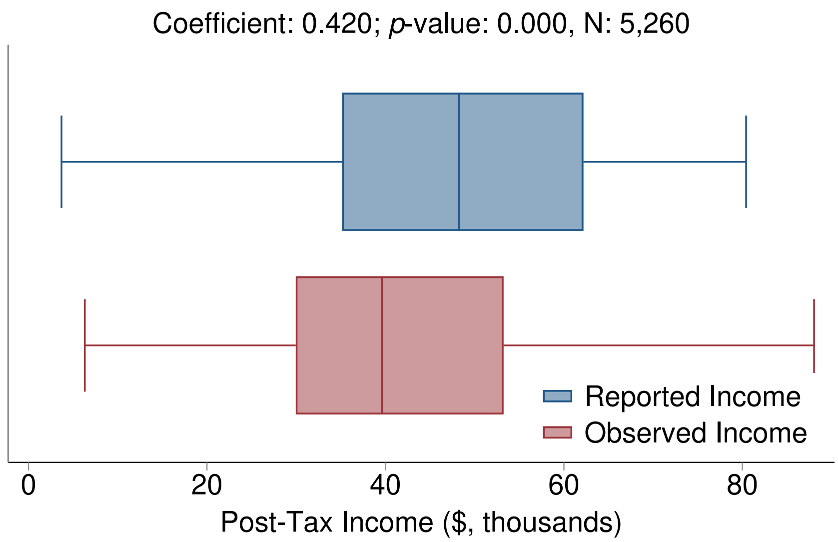

Notes: The figure a box and whisker plot of reported post-tax labor income and observed post-tax labor income. Values below the 25th percentile minus 1.5 × IQR and above the 75th percentile plus 1.5 × IQR are suppressed. I include the regression coefficient resulting from a regression of observed income on reported income.

Figure D.12 presents the results of our re-calculated perfect foresight dynamic wedge under our imputed measure of reported post-tax labor income.

Figure D.12. Consumption Wedge under Reported Income

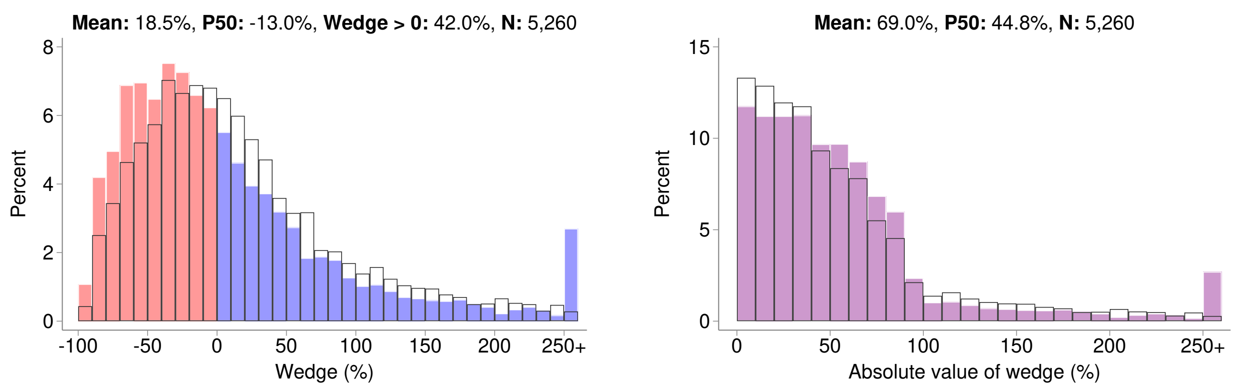

Notes: The figures show the distribution of dynamic consumption wedges (left) and their absolute values (right). Colored bins represent the distribution under reported income while the clear bins represent the standard distribution using observed income. Wedges are reported here as the percent deviation of observed consumption from frictionless consumption. Observed income wedges are trimmed at the 1st and 95th percentiles.

Table D.2. Crosswave Transition Matrix

	W2 Underconsumer	W2 Overconsumer
W1 Underconsumer	.6875	.3125
W1 Overconsumer	.4634146	.5365854

D.5 Preference Heterogeneity

Figure D.13. Distribution of Zero Wedge Beta/Gamma

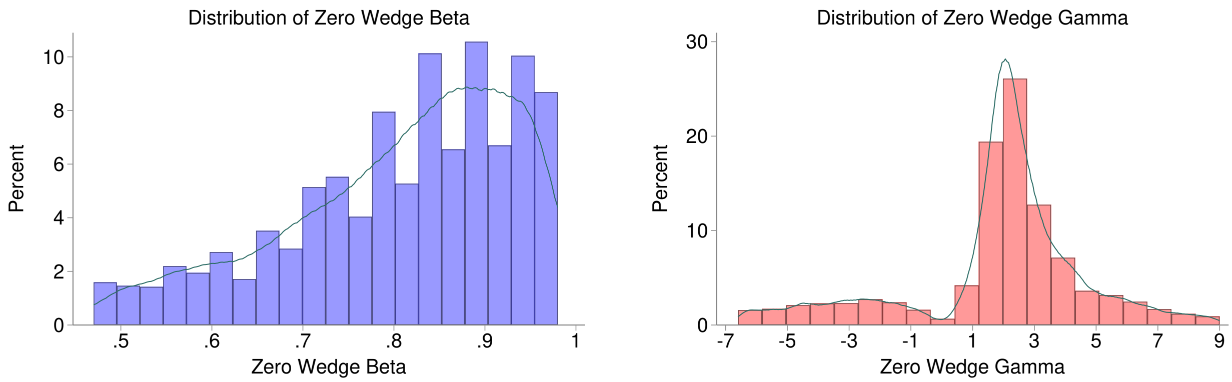

D.6 Wedges In Dollars

D.7 Crosswave Transition

Table D.2 is the overconsumption transition matrix for individuals in both waves 1 and 2. It is comprised of $N = 235$ individuals. Specifically, it highlights the probabilities of maintaining/transitioning into over/underconsumption across waves. Since Waves 1 and 2 took place approximately two years apart, we expect to see some adjustment over time.

Figure D.14. Correlations with Survey-Elicited Preferences

Panel A. Zero Wedge Beta/Gamma vs. Survey Elicited Beta/Gamma

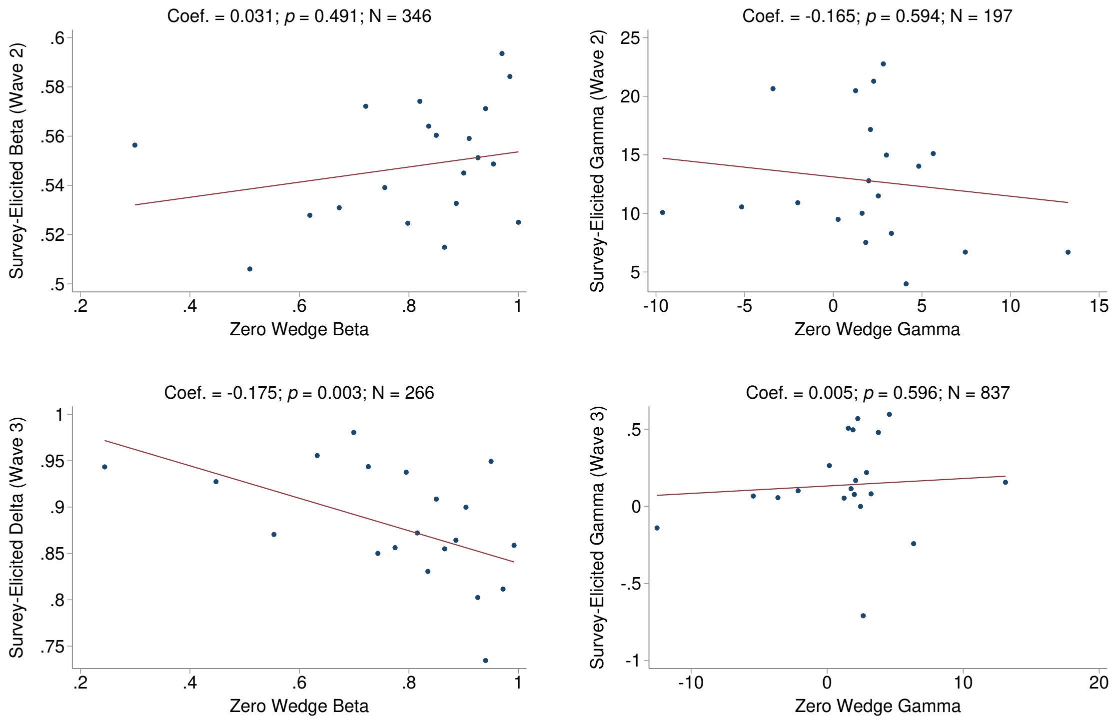

Panel B. Correlation Between Crosswave Preferences

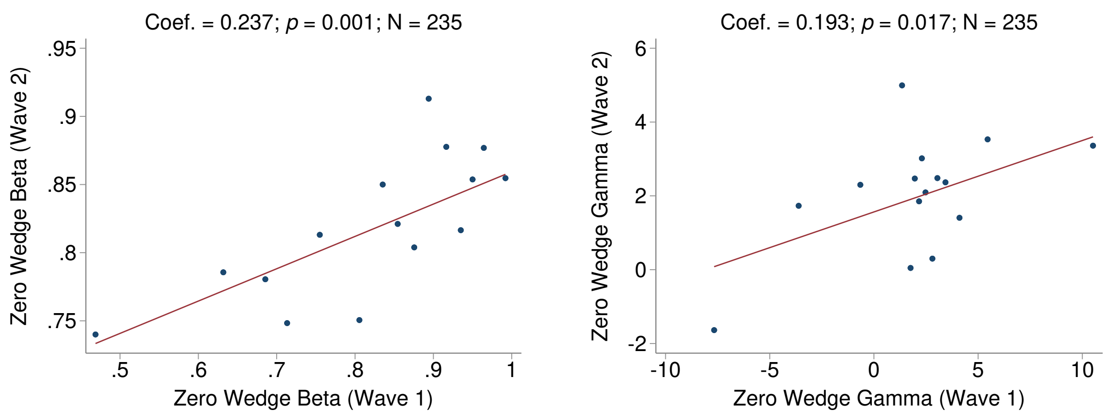

Notes: Plots in Panel A compares joint grid search zero wedge preferences and survey-elicited preferences. Plots in Panel B compare survey-elicited preferences and the size of our main dynamic consumption wedge. Survey-elicited delta in Wave 3 refers to the exponential discount factor, which has been trimmed at 0 and 1. Wedges are trimmed at the 1st and 95th percentiles.

Figure D.15. Within-Respondent Wedge Correlations

Panel A. Cross-Wave (Dynamic Wedges) Panel B. Dynamic and Static Wedges

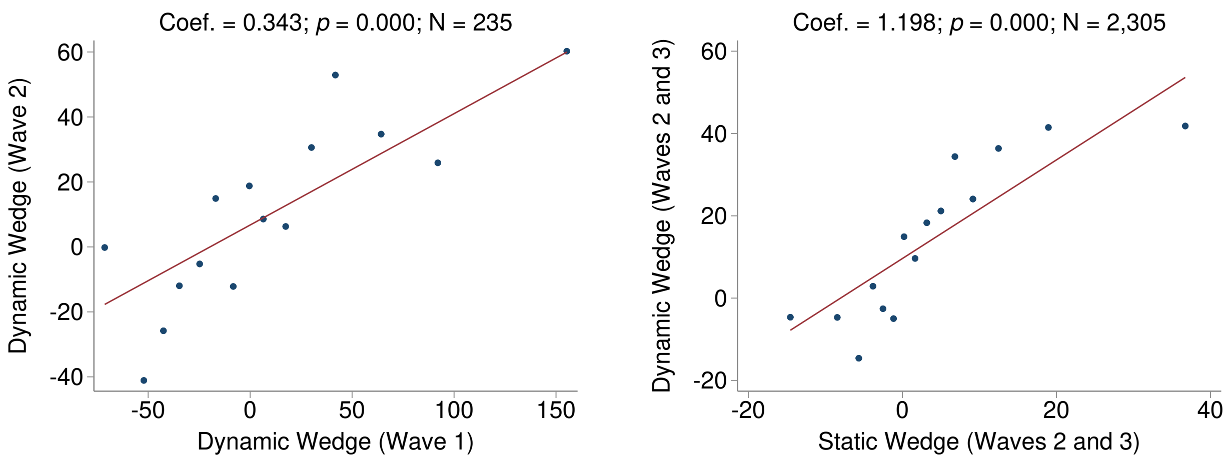

Notes: Panel A shows the correlation between dynamic wedges in waves 1 and 2 for the subset of repeat responders in our main analysis sample. Panel B shows the correlation between the dynamic and static wedges across Waves 2 and 3.

Figure D.16. Consumption Wedges in Dollars

Panel A. Raw wedges Panel B. Absolute value of wedges

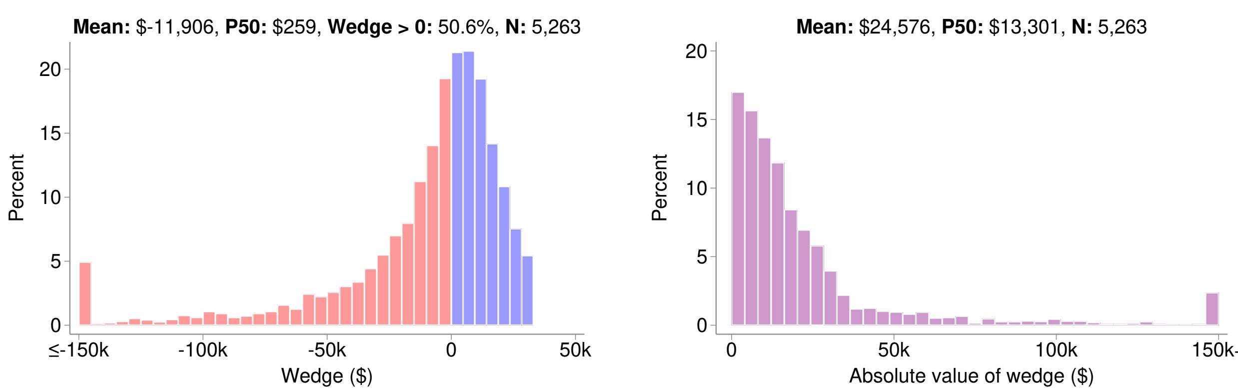

Notes: The figures show the distribution of dynamic consumption wedges (left) and their absolute values (right). Colored bins represent the distribution of consumption wedges in dollar terms. Wedges are trimmed at the 1^{st} and 95^{th} percentiles.

E Quantitative Model Appendix

E.1 Quantitative Model Specifications

We solve a series of simple Bewley (1980) style quantitative models. 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 consider two forms of borrowing constraints. The first is a simple lower bound on wealth:

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

with $\bar{A}$ denoting the borrowing constraint. The second is a "soft" borrowing constraint (SBC) where interest rates jump discontinuously when agents switch from borrowing to saving. That is:

$$R(a_{i,t}) = \begin{cases} R^S & (a_{i,t}) \geq 0 \\ R^D & (a_{i,t}) < 0 \end{cases}$$

where $R^S << R^D$.

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

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

E.2 Quantitative Model Calibrations

Table E.1 summarizes our calibration. Panel A presents parameters governing the fundamental economic environment that are common through all model variants we study. For this, we draw on a standard parameterization. Panel B gives the interest rate parameters, which vary depending on whether the model features an SBC. Panel C 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%.

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). For calibrating the income process, we follow Guerrieri and Lorenzoni (2017) and Maxted et al. (2024), using the estimated process from Floden and Lindé (2001).

For interest rates, we set the gross annual interest rate to match the average real expected levered return in the EarnIn sample (1.0692). In models not featuring an SBC, this is the single prevailing interest rate. In models with an SBC, the net return to saving is set to 2.5%, which is the value targeted in Guerrieri and Lorenzoni (2017). We set the spread $R^D - R^S = 0.1032$, which is the average expected spread in the EarnIn sample.

Table E.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)
Panel B. Interest Rates			
R	1.0692	Gross ann. interest rate (non-SBC)	EarnIn Sample mean
R^S	1.025	Gross ann. return to savings (SBC)	Guerrieri and Lorenzoni (2017)
$R^D - R^S$	0.1032	SBC borrowing/saving spread	EarnIn Sample mean
Panel C. Non-SBC Distortions			
$\bar{A}$	-0.25	Borrowing limit	Kaplan et al. (2018)
b	0.397	Hyperbolic discounting factor	Calibrated
λ	0.985	Temptation cost scalar	Calibrated
ψ	0.1 (0.5)	Cons. adj. cost (with SBC)	Calibrated

This table summarizes the calibration choices for the quantitative models.

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 non-SBC distortions such that they minimize the distance between the model-implied over-consumer share and median absolute value wedge and their empirical counterparts. We calibrate high levels of present bias in both the beta-delta and temptation preferences formulations; these values are higher than those in Maxted et al. (2024) and Attanasio et al. (2024), respectively. 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 0.1 when not including borrowing constraints and 0.5 when including an SBC. These are also relatively large values; they arise because both the over-consumer share and median absolute value wedge are generally increasing in the size of the cost (but at a decreasing rate). They plateau near these calibrated values and we truncate the candidate parameter value range at these points as higher costs imply annual rates of adjustment below 10%.

Ergodic Wedge Distributions. For comparison, below we provide the distribution of exact and approximate wedges under each model's original ergodic distribution. This differs from Table 3, which reweights the distribution to resemble the EarnIn sample, as detailed in Section 4.1.

Table E.2. Wedge Distribution under Various Models (with original ergodic distribution)

Model	% Pos.	|50th|	Mean	SD	Min	25th	50th	75th	Max
Panel A. Exact Wedges									
Borrowing constraint (BC)	0.0	45.7	-45.9	15.5	-85.9	-57.5	-45.7	-34.4	-10.3
Present bias (PB)	100.0	9.14	9.14	0.004	9.14	9.14	9.14	9.14	9.17
BC + PB	0.0	44.6	-45.7	17.4	-85.9	-58.2	-44.6	-32.7	-2.6
Cons. adjustment (CA)	96.4	30.3	60.5	62.2	-64.2	30.2	30.3	73.2	375.7
BC + CA	25.8	49.2	-15.9	58.5	-94.9	-55.6	-36.7	2.2	212.7
Hand-to-mouth (HTM)	5.4	46.6	-43.1	22.1	-89.8	-60.2	-46.6	-29.1	91.5
Panel B. Approx. Wedges									
Borrowing constraint (BC)	0.0	55.4	-55.5	15.0	-91.2	-67.1	-55.4	-44.1	-17.4
Present bias (PB)	0.0	25.1	-25.6	7.3	-49.4	-30.7	-25.1	-20.3	-3.5
BC + PB	0.0	55.1	-55.4	16.5	-91.2	-67.2	-55.1	-43.8	-10.5
Cons. adjustment (CA)	35.0	15.0	8.4	39.4	-68.3	-15.0	-6.6	18.2	132.3
BC + CA	20.4	53.7	-32.0	46.2	-95.2	-64.7	-47.0	-13.2	103.3
Hand-to-mouth (HTM)	1.7	58.7	-55.4	19.9	-93.0	-70.8	-58.7	-42.9	57.1
Panel C. Empirical Wedges									
Data	50.6	40.1	15.2	70.6	-91.8	-35.7	0.9	48.3	259.5

Notes: This table reports summary statistics for wedges that arise in different models. The first two statistics are the share of consumers with positive wedges and the median absolute value wedge, respectively. The calibration and income process is held constant across all models here. We also hold constant the ergodic distribution by reweighting each model's distribution to resemble the EarnIn sample (in terms of its marginal distributions of assets and income). Wedge statistics for models are calculated using the respective ergodic distribution of each. Panel C reports the same statistics calculated using wedges from our empirical analysis.

F Additional Figures

Figure F.1. Relationship Between Dynamic Consumption Wedges and EarnIn EWA Usage

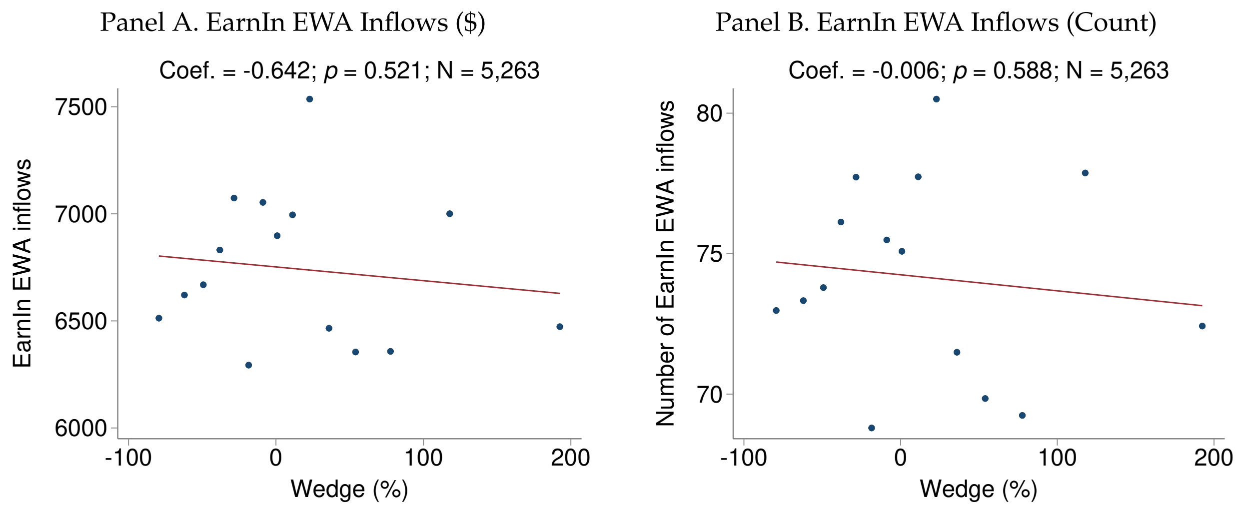

Notes: The figures illustrate the relationship between dynamic consumption wedges and two measures of EWA usage over the 12 months preceding the survey: (1) the dollar amount of EarnIn EWA inflows and (2) the number of EarnIn EWA inflows over the sample period. Each binned scatterplot plots the average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

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

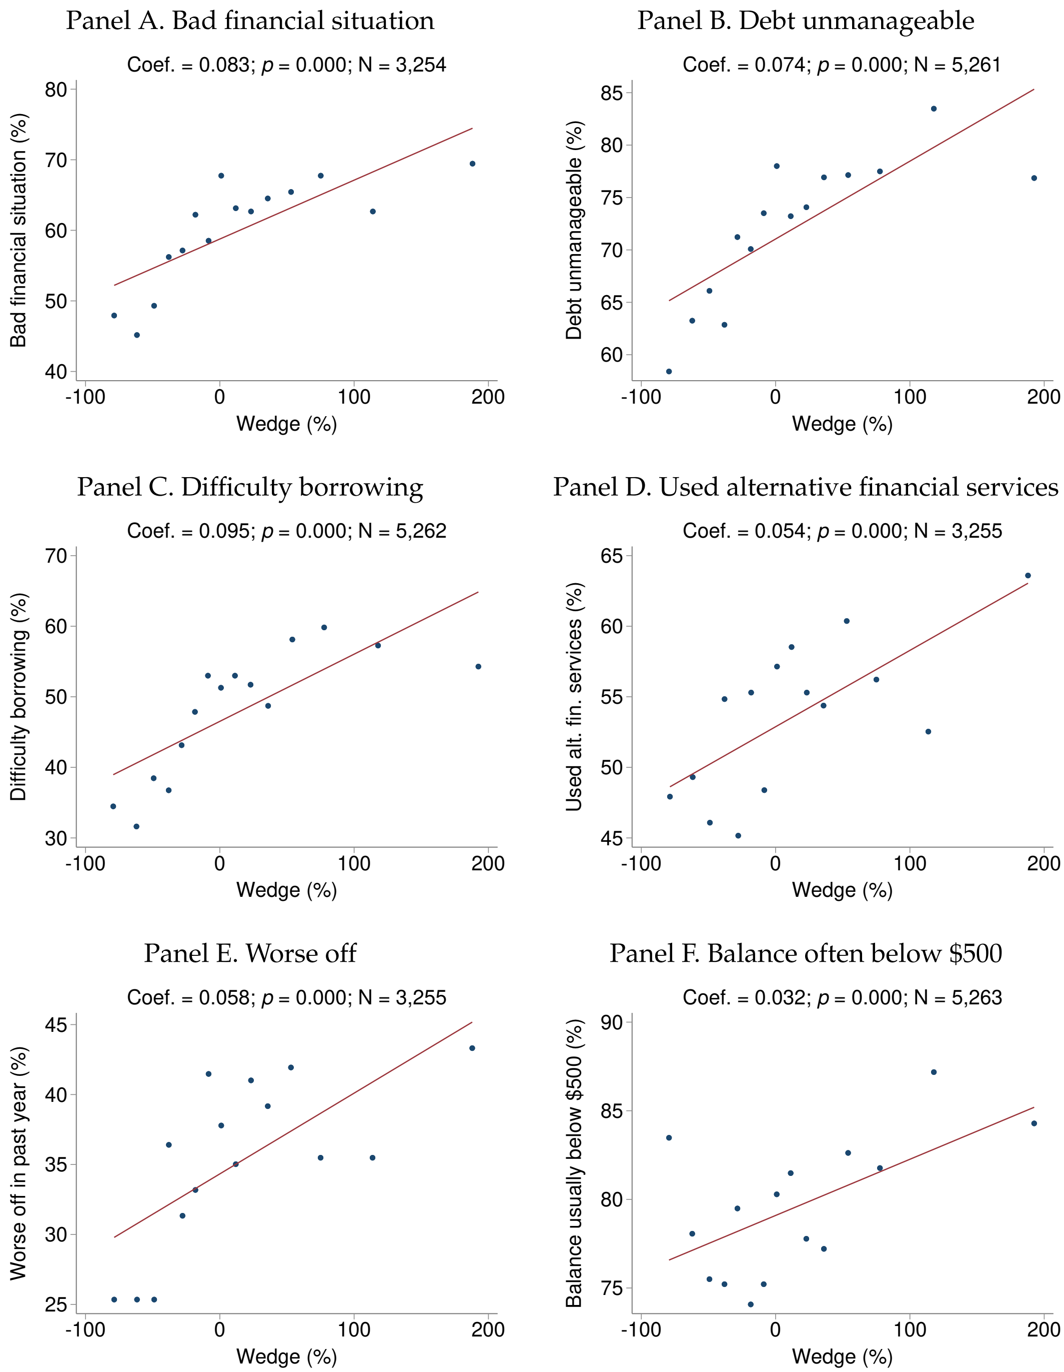

Notes: The figures illustrate the relationship between dynamic consumption wedges and six indicators of financial distress: (1) whether the user reports "just getting by" or "finding it difficult to get by;" (2) whether the user reports having "a bit more" or "far more" debt than is manageable; (3) whether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

Figure F.3. Distribution of Wedges with FIRE Beliefs

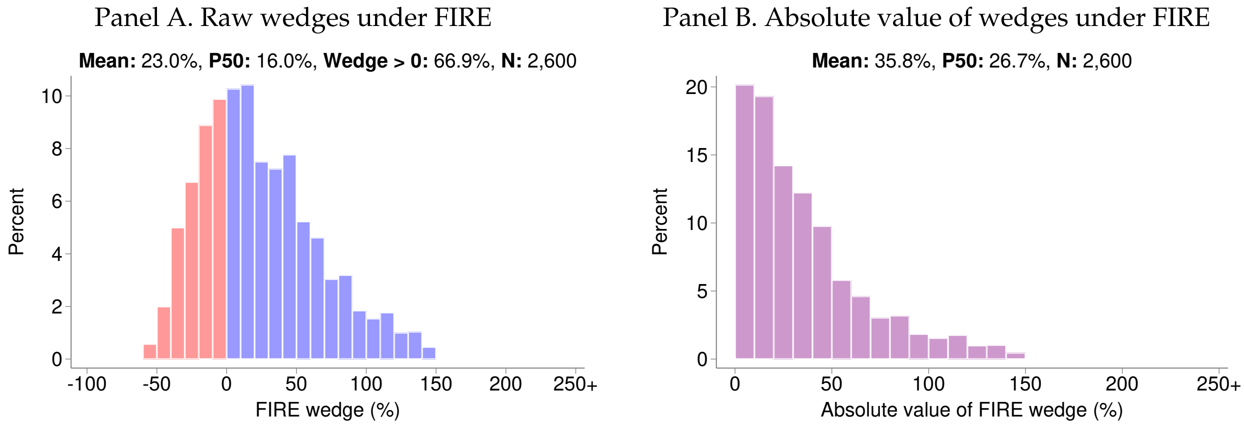

Notes: The figure plots the distribution of wedges calculated assuming FIRE. Because we do not have realized income data for our two 2024 waves, we restrict to wave 1 users. All wedges in the graphs above are expressed as a percent deviation from frictionless FIRE consumption. All wedges are trimmed at the 1^{st} and 95^{th} percentiles.