What Explains the Consumption Decisions of Low-Income Households? *

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

October 15, 2024

Abstract

We empirically show that the consumption choices of low-income households are significantly distorted both downwards and upwards by frictions. We measure the impact of frictions as the wedge between actual consumption and a counterfactual "frictionless" consumption. Our approach exploits the fact that consumption, income, wealth, and expectations over income, inflation, and interest rates are sufficient statistics for these wedges in a broad class of models. To implement this approach, we use a new survey of economic beliefs, which we link to administrative consumer transactions data for disproportionately low-income US consumers. In absolute value, the median wedge is 40% of frictionless consumption. 51% of consumers have negative wedges (under-consumers) and 49% have positive wedges (over-consumers). Because financial constraints only generate negative wedges, additional or alternative frictions (such as present bias or consumption commitments) are necessary to rationalize our results.

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

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

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

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

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

1 Introduction

Financial constraints have come to play a central role in macroeconomics and household finance in explaining consumer behavior (e.g., Kaplan, Moll and Violante, 2018). A key motivation for this trend is empirical evidence that households have marginal propensities to consume (MPCs) that are much larger than predicted by frictionless models, especially among low-income and low-liquidity households (e.g., Johnson, Parker and Souleles, 2006). However, alternative frictions, such as present bias, consumption adjustment costs, and bounded rationality, have also been shown to generate large MPCs (Lee and Maxted, 2023; Maxted, Laibson and Moll, 2024; Beraja and Zorzi, 2024; Ilut and Valchev, 2023). These alternative models predict different distributional and aggregate consequences of fiscal and monetary policy as well as business cycles compared to models featuring only financial constraints.[1] To better guide theory and ultimately policy, more evidence is needed on which frictions distort consumption and by how much.

This paper empirically studies the impact of frictions on consumption for a population of predominantly low-income US households. Our approach uses a novel sufficient statistics framework and new data linking consumer expectations to administrative consumer transactions data. We measure the impact of frictions as a wedge between actual consumption and a counterfactual "frictionless" consumption. To our knowledge, we are the first to study individual-level consumption wedges. These wedges capture the effects of both economic constraints (such as liquidity constraints) and behavioral preferences (e.g., present bias) or bounded rationality that result in "as if" constrained behavior.

We find that consumption is significantly distorted for many households in our sample. While the distribution of wedges is centered near zero, the median (absolute value) distortion is 40% of frictionless consumption. The substantial size and variation in wedges indicate that frictions are important for explaining the cross-section of consumption among low-income households. Our findings highlight that studying micro-level wedges can be valuable, as aggregate wedges may fail to reveal the cross-sectional importance of frictions. Our findings also have strong implications for theories of consumer behavior. Patterns in our data reject financial constraints (the primary friction used to generate high MPCs) as being the dominant friction for nearly half of our sample. Present bias, consumption commitments, and bounded rationality are plausible additional or

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

alternative frictions that can rationalize our results.

To measure consumption wedges, we begin by characterizing frictionless consumption in a stylized model. We present a frictionless benchmark where a household chooses consumption and saving (or borrowing) via a risky asset given their realized income and wealth. The benchmark flexibly allows the household's beliefs to deviate from full-information rational expectations (FIRE). More rigidly, the benchmark assumes that households face no frictions (borrowing constraints, consumption adjustment costs, etc.) and have standard preferences.[2] As a result, the difference between a household's actual and "frictionless" consumption summarizes the impact of frictions and non-standard preferences on their consumption. That is, the consumption wedge quantifies the degree to which households over- or under-consume due frictions or behavioral preferences.

In the benchmark, frictionless consumption is determined by an Euler equation and the budget constraint. We derive a log-linear first-order approximation of frictionless consumption. This approximate measure of frictionless consumption is a function of net worth (divided by income) and beliefs about future nominal income growth, nominal interest rates, and inflation.

The simplicity of our benchmark allows our wedge measurement analysis to apply to a broad class of models. Specifically, our frictionless benchmark is a "special case" common to a rich set of models. For example, it corresponds to a model with consumption adjustment costs equal to zero, or a model with infinite borrowing constraints. For all models that nest the frictionless benchmark, the wedge between actual and frictionless consumption reflects the total impact on consumption of *all* frictions and behavioral preferences directly affecting the household.

Additionally, the consumption wedge formula applies for richer models featuring, for example, heterogeneous agents, additional optimization choices (such as portfolio choice and labor supply), and general equilibrium.[3] As long as the Euler equation and budget constraint hold when frictions and behavioral preferences are "turned off" in such models, frictionless consumption has the same characterization that we derive.[4]

To empirically measure consumption wedges, we use new data linking subjective expectations

[2]Specifically, we assume preferences are time consistent, time separable, strictly increasing, and strictly concave.

[3]For our formula to hold in the presence of general equilibrium, it must be the case that households are price takers. Otherwise, their Euler equation could feature additional terms not present in our formula.

[4]Frictions in richer models that affect the choices of agents other than households, such as firms or financial intermediaries, are not reflected in our consumption wedge formula. The consumption wedge formula takes as given wealth and income, which in such models may be influenced by frictions directly shaping other agents' behavior and prices. Therefore, our consumption wedges summarize the impact of frictions and non-standard preferences that directly affect households.

to consumer transactions. A lack of such data has likely limited the ability of prior research to conduct a similar analysis. The transactions data come from EarnIn, an American financial technology company. Through an app, EarnIn offers users earned wage access, which allows users to shorten their pay cycle and receive their wage income prior to their regularly scheduled payday. Compared to the broader US population, EarnIn users are younger, more likely to identify as female, and have lower income and net worth. EarnIn fielded a survey to its users over September 29 to October 2, 2022 that solicited one-year-ahead expectations of inflation, nominal income growth, and nominal interest rates (for both saving and borrowing). We study merged data linking 10,000 respondents' de-identified transactions data to their surveyed expectations.

Beliefs data have two important advantages for studying consumption wedges. First, subjective expectations data allow us to avoid assuming FIRE. Prior analyses of economic wedges assume FIRE because it makes it possible to estimate rational expectations by averaging realized outcomes (e.g., Chari, Kehoe and McGrattan, 2007; Berger, Bocola and Dovis, 2023). However, behaviors that appear consistent with constrained optimization can often be rationalized by some set of (possibly non-rational) beliefs. If FIRE does not hold, prior wedge analyses studying "frictions" may in fact be conflating the effects of frictions with deviations from FIRE. Our wedge analysis is the first, to our knowledge, to relax this assumption, allowing us to quantify the impact of frictions and non-standard preferences separately from any effects of deviations from FIRE.

To gauge the importance of allowing for deviations from FIRE, we test for such deviations in our sample. Under FIRE, choices made at time t should not predict errors in forecasts over time $t+1$ outcomes.[5] Motivated by this, we test whether current consumption predicts income growth forecast errors. We find that higher consumption relative to income (i.e., the average propensity to consume or "APC") is strongly associated with overestimating income growth. Specifically, a 5 pp larger overestimate (approximately 0.22 standard deviations) is associated with a 3.0 pp higher APC. We also find that underestimating income growth is associated with higher consumption as well, though to a lesser extent. Conditional on underestimating income growth, a 5 pp smaller forecast predicts a 0.21 higher APC. Our findings reject FIRE in our sample; the economically significant relationship with consumption suggests that it may be quantitatively important to use subjective expectations to isolate the influence of frictions and non-standard preferences.

[5]While forecasts will generally be correlated with household choices, forecast *errors* are not correlated with consumer decisions under FIRE. Intuitively, such a correlation reveals that households are failing to incorporate predictive information into their expectations. This could be due to either non-rational expectations (e.g., violations of the law of iterated expectations) or incomplete information about the "true model" of the world (i.e., not knowing the true predictive relationship).

The second advantage of using beliefs data is that it is possible to measure individual-level wedges, rather than aggregate wedges. Even if households have FIRE, attempting to measure individual-level wedges using realized (future) individual-level outcomes would conflate the impact of frictions or non-standard preferences with prediction error. To overcome this, prior work has assumed FIRE and averaged wedges across households with similar observable characteristics and focused on aggregate or group-level wedges (Berger et al., 2023).[6] Subjective expectations data allow us to instead directly calculate individual-level wedges.

Using our data, we estimate consumption wedges for the EarnIn sample to study their distribution and how they covary with other characteristics. Our first main result is that consumption wedges are generally large. That is, many consumers in our sample have consumption that is significantly distorted by frictions or behavioral preferences. The median person has consumption distorted 40% (in absolute value) relative to their frictionless consumption. For household finance, this implies that frictions are first-order determinants of the consumption-saving decision for low-income households. For the class of macroeconomic models where heterogeneity in consumer behavior matters, this suggests that is important to include frictions to generate realistic consumption behavior (and ultimately model predictions).

Our second main result is that there is a mix of positive and negative wedges. We find 49% of wedges are positive (over-consumption) and 51% are negative (under-consumption). A key implication of this result is that it rejects financial constraints as being sufficient to explain deviations from frictionless consumption. Financial constraints only generate negative wedges and therefore cannot account for the 49% of over-consumers. To generate positive wedges, additional or alternative frictions are necessary. One potential solution is to augment models featuring financial constraints to also include frictions, such as present bias, that can generate positive wedges. Another possible solution is to include frictions that are capable of generating both positive and negative wedges (e.g., Lee and Maxted, 2023; Maxted et al., 2024). Examples of the latter include consumption adjustment costs and bounded rationality (e.g., Beraja and Zorzi, 2024; Ilut and Valchev, 2023). In addition to these qualitative implications, our quantitative results could be used to calibrate structural models.

To evaluate the robustness of our findings, we conduct sensitivity analyses for the key assump-

[6]The prediction errors average to zero under FIRE if the group members over which the average is taken provide valid counterfactuals for each other. For example, with respect to income, one would need to assume that the group members draw idiosyncratic income shocks from the same distribution. Berger et al. (2023) discuss their averaging solution in detail.

tions necessary to quantify consumption wedges. The first set of assumptions relate to parameter choices. Consumption wedges are functions of the consumer's discount factor and coefficient of relative risk aversion. Our baseline analysis assumes standard values; we intend to relax this assumption in a future iteration by using preference parameters measured via an additional survey wave. We show that our main findings of (1) a median absolute value wedge of 40% and (2) 49% of households over-consuming remain similar under alternative parameter values. The consumption wedge also depends on the value of two of our chosen approximation points (steady state consumption-to-income and wealth-to-income). We also verify that our results are similar under a range of alternative approximation points.

The second important assumption is that there is no measurement error in consumption, income, wealth, or beliefs. To address this, we conduct three sets of analyses. First, we confirm that our results remain similar when restricting to various subsets of users with higher quality data. For example, excluding users with a larger share of spending in the form of ATM cash withdrawals or those with expectations divisible by five. Second, we obtain a similar distribution of consumption wedges using consumption, income, wealth, and beliefs data averaged within groups of similar users. The idea being that measurement error may average out in these measures within a group of similar people. We use k-prototypes clustering to group users that are similar along a large set of observable characteristics. Third, we add random noise to our data and show that our results are not significantly quantitatively affected.

Our last analyses show that consumption wedges vary intuitively with other observable characteristics, providing suggestive evidence on their origins. We document that consumption wedges are strongly correlated with proxies for financial distress. This arises for both self-reported proxies of perceived financial distress (such as anxiety about finances) as well as observable proxies such as regularly having bank account balances below $500. Specifically, we see that larger, positive consumption wedges are strongly related to financial distress.

We provide additional suggestive evidence that points to present bias and consumption commitments as plausible sources of consumption wedges. First, we show that large, positive consumption wedges are associated with higher MPCs.[7] This implies that the frictions driving overconsumption are associated with higher MPCs. Both present bias and consumption commitments could give rise to this pattern. We also find that consumption wedges are strongly, positively

[7]We calculate individual-level MPCs as the "excess" spending during March 2021 when respondents received COVID-related stimulus checks (relative to their spending in the same time period in 2023 and 2022).

related to a proxy for consumption commitments. Our proxy is the share of income devoted to two important expenses that are typically large and difficult to adjust: housing and childcare. To further gauge present bias as a plausible driver of consumption wedges, we intend to field an additional survey wave that separately solicits both discount factor and present bias preferences.

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

Our findings can help guide the design of models by providing new data points, beyond existing MPC evidence, that speak to *how much* and *what* are the dominant frictions that drive consumption among low-income households. The large frictions we document highlight the importance of including frictions in models, as they are a significant determinant of consumption for low-income households. Additionally, the heterogeneity in wedges that we document, such as the mix of positive and negative wedges and correlation of wedges with MPCs, support or reject alternative models of frictions. Quantitatively, our results may be useful calibration targets for models as well.

Additionally, the empirical approach demonstrated in our analysis may be used to study consumption wedges in other settings. In principle, such an analysis could be done with only survey data.[8] Notably, our approach does not require quasi-experimental variation, unlike the empirical literature focusing on MPCs. A useful endeavour for future research would be to measure and study these wedges to provide additional evidence to test alternative models of frictions.

Second, we add to the empirical macroeconomics literature on consumer expectations. This literature has documented the importance of beliefs, including departures from FIRE, in explaining consumer behavior. D'Acunto et al. (2023) and Weber et al. (2022) provide recent reviews of

[8]While high-quality consumption data like ours is ideal, a survey could attempt to solicit this information.

this area. We complement recent papers that have linked consumption and beliefs data using, for example, survey measures (Coibion et al., 2023; D'Acunto et al., 2022), grocery shopping data through the Nielsen panel (Weber et al., 2023), German bank data (Hackethal et al., 2023), and credit card transactions (Kanz et al., 2021). Consumer beliefs appear to deviate from FIRE. For example, inflation expectations are excessively influenced by grocery prices (D'Acunto et al., 2021) and D'Acunto et al. (2024) finds evidence of extrapolative income expectations. We innovate on this literature in several ways. We demonstrate a novel application of consumer expectations data: measuring wedges. Additionally, we document that consumption is strongly associated with departures from FIRE in our sample of low-income households.

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

Outline. We next introduce our frictionless benchmark and wedge measurement approach in Section 2. We then describe our survey and linked transactions data in Section 3. We then characterize departures from FIRE related to consumption in our sample in Section 4.1, motivating our need for beliefs data to measure wedges. Section 4 presents our main results measuring and describing wedges in our sample of low-income households. We then conduct calibration sensitivity and measurement error robustness analyses. Lastly, we correlate wedges with observables to test alternative theories about the sources and effects of distortions to consumption. Section 5 concludes.

2 Theory: Measuring Consumption Wedges

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

2.1 Frictionless Benchmark

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

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

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

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

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

We refer to the model above as our "frictionless" benchmark. There are three important features of this benchmark. First, it assumes that there are no frictions. That is, there are no borrowing

constraints, transaction or adjustment costs, etc. Second, it assumes standard preferences. As a result, real-world deviations from the benchmark can also arise from behavioral preferences that result in "as-if" constrained behavior, such as present bias and habit formation. Third, the frictionless benchmark flexibly allows for deviations from FIRE. The first two features are what will allow the consumption wedge to flexibly capture distortions due to either frictions or behavioral preferences. The third assumption is what will allow the consumption wedge to avoid conflating the influence of non-rational expectations with constraints and behavioral preferences.

Optimal consumption $C_t^\star$ in the frictionless benchmark is characterized by the budget constraint in equation (1) and the Euler equation:

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

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

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

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

We provide a derivation in Appendix C. Equation (4) relates the log APC (i.e., $\ln\left(\frac{C_t^\star}{Y_t}\right)$) to three sets of objects. First are preference parameters governing risk aversion (γ) and the discount factor (β). Second is start-of-period-wealth (divided by income): $\frac{A_tR_t}{Y_t}$. Third are expectations of gross nominal income growth $\left(G_{t+j}^Y\right)$, inflation (π_{t+j}), and nominal interest rates (R_{t+j}). The parameters, $\{\alpha_0, \alpha_1, \alpha_Y, \alpha_\pi, \alpha_R, \rho\}$, are functions of the preference parameters (γ, β) and the steady state points around which we log-linearize. The latter are steady state APC $\left(\frac{C}{Y}\right)$, initial net worth to income $\left(\frac{AR}{Y}\right)$, nominal income growth G^Y, and the nominal interest rate R. Table 1 gives expressions for the parameter values.

In general, frictionless consumption is increasing in initial wealth and expected nominal income growth. It is decreasing in the expected nominal interest rate. Holding constant expected

Table 1. Frictionless Consumption Coefficient Formulas

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

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

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

2.2 Consumption Wedges

In the presence of frictions and behavioral preferences, actual consumption can deviate from frictionless consumption. With a measure of frictionless consumption, we can calculate the distortions to consumption induced by frictions and behavioral preferences as the wedge between actual and frictionless consumption. We define this (log) "consumption wedge" below:

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

The log consumption wedge η_{it} for person i in time t is the difference between their actual and frictionless log-APC (where $\star$ denotes frictionless). Intuitively, the log consumption wedge η_{it} is a measure of how far "off" their Euler equation a consumer is. We highlight also that, because

actual income is in the denominator for both actual and frictionless APCs, these terms will cancel out. Thus η_{it} describes relative differences in terms of consumption. That is:

$$\eta_{it} = \ln\left(\frac{C_{it}}{C^{\star}_{it}}\right).$$

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

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

Example: Financial Constraints. To make the interpretation of wedges more concrete, we discuss several prominent frictions and behavioral preferences and how they are relate to consumption wedges. We start with the primary friction considered by macroeconomics and household finance: financial constraints. These are most often modeled as a constant borrowing limit (e.g., Aiyagari, 1994), an endogenous borrowing limit (e.g., Bornstein and Indarte, 2023), or "soft" constraints arising from discrepancies in borrowing and saving rates (e.g., Kaplan et al., 2018). These frictions can introduce a wedge into the Euler equation, relative to the frictionless Euler equation in Equation(3). For example, with a constant borrowing limit, this wedge would be the Lagrange multiplier on the constraint. An important feature of financial constraints is that they only generate negative consumption wedges. Therefore, a testable implication of financial constraints relates to the sign of consumption wedges. The presence of positive wedges would indicate that additional or alternative frictions are necessary 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 the agent discounts future utility by an additional factor $\beta < 1$ relative to the standard exponential discounting model. These preferences cause the Euler equation to differ from Equation (3), specifically the expectation term is scaled down by the parameter governing the degree of present bias. As a result, these preferences cause consumption to be higher relative to a "debiased" consumer and thus result in positive consumption wedges under our formula (Maxted, 2022). Similar to financial constraints, we can use the sign of empirical consumption wedges to test whether present bias is sufficient to rationalize empirical consumption choices.

Example: Inertia. Another class of frictions introduces inertia into consumption choices. One example is consumption commitments, where inertia is generated by consumption adjustment costs (e.g., Chetty and Szeidl, 2007; Beraja and Zorzi, 2024) or Calvo-style adjustment shocks (e.g., Auclert, Rognlie and Straub, 2020; Bornstein, 2021). Another is habit formation, which is a preference-based source of inertia where the utility of current consumption depends on past consumption (e.g., Fuhrer, 2000; Christiano et al., 2005; Smets and Wouters, 2007). This history dependence violates the time separability assumption of our benchmark and will therefore also be captured by our consumption wedge formula. Bounded rationality can similarly create inertia when costly cognition limits or delays consumption adjustments. For example, in Ilut and Valchev (2023), cognition costs limit households' updating of consumption decision rules, leading to inertial behavior. This class of frictions can produce positive consumption wedges when inertia limits the downward adjustment of consumption in the wake of negative shocks. Similarly, positive shocks can lead to negative wedges. Empirical findings of both positive and negative wedges could be rationalized by this class of frictions.

Empirical evidence on consumption wedges can help guide the choice and modeling of frictions. Qualitatively, the presence of both positive and negative wedges can (i.e., both over- and under-consumption) indicates that financial constraints and present bias are insufficient 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 frictions.

2.3 Model Extensions

The frictionless benchmark is a "special case" in a large class of models featuring frictions and/or behavioral preferences. The benchmark corresponds to versions of these models where the frictions and behavioral preferences are turned "off." For example, in a model with borrowing constraints, our benchmark corresponds to infinite borrowing constraints. In a model with beta-delta discounting, it corresponds to zero present bias (i.e., $\beta = 1$ in the notation of Maxted, 2022). It is the simplicity of our benchmark that allows it to simultaneously be a special case for many models. This enables consumption wedges to measure the effect of such a wide set of frictions and behavioral preferences. Our wedge formula is also robust to a variety of other model extensions. Below, we discuss these extensions and their implications for the interpretation of consumption wedges.

Additional Household Choices. The consumption wedge formula remains unchanged and its interpretation similar when including additional household choice variables. These include, for example, labor supply.[9] Adding choice variables results in additional optimality conditions. But as long as the Euler equation and budget constraint continue to hold, the consumption wedge formula is unaltered. However, if these other choices are subject to separate frictions, such as distortionary taxation on labor supply, the impact of those frictions is not captured in the consumption wedge. The wedges reflect only the frictions that alter the Euler equation. In this sense, the distortions measured are specific to the consumption decision.

Additional Assets. We can also enrich the frictionless benchmark to feature a portfolio choice problem where the household chooses a mix of assets and liabilities, including housing. To do so, one can rewrite the household problem into two stages, where the first stage has the consumption/savings decision modeled in the frictionless benchmark and the second stage has the portfolio choice. In such a model, savings A in our benchmark would correspond to the amount invested in the portfolio. Similarly to other endogenous choices like labor supply, this addition does not alter the wedge formula nor its interpretation. However, it does imply that ideally one would measure beliefs with respect to overall portfolio returns when calculating consumption wedges. Such a requirement could pose more challenges for wedge measurement in settings where households

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

hold complicated portfolios. A measurement advantage of our empirical setting is the low-income households we study likely have simpler portfolios.

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

Heterogeneity. In a model with a single representative household, the average or median wedge would be the appropriate object of interest. For models with multiple households, a distribution of consumption wedges can be calculated. If one finds empirically that there is significant heterogeneity in wedges, this would indicate that modeling household heterogeneity is important for the class of models where heterogeneous consumption behavior has aggregate implications (e.g., Kaplan et al., 2018).

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

2.4 Discussion

Consumption wedges have advantages as an object of study compared to widely-used alternatives such as MPCs and proxies for constraints. An advantage of studying consumption wedges compared to MPCs is that estimating wedges does not require quasi-random variation, only observational data. While our analysis uses administrative consumption data to minimize measurement error, in principle one could use a survey to solicit all of the necessary inputs to measure consumption wedges. Additionally, both consumption wedges and MPCs can serve as calibration targets for quantitative models. MPCs are local estimates and, as such, they are generally not sufficient statistics to fully distinguish between models. While consumption wedges are local to the context in which they are measured, they provide additional data points that can help distinguish between models in cases when MPCs cannot.

Household finance research has long relied on proxies like credit card utilization and a lack of liquid wealth to tag people as constrained. However, frictions like present bias and consumption adjustment hazard can result in high utilization and low liquid wealth without financial constraints binding. Consumption wedges instead can help differentiate between competing models while, at the same time, quantifying the consumption impact of frictions.

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

A second advantage of using subjective expectations data to measure consumption wedges is that it enables us to measure micro-level wedges. That is, individual wedges for consumers. Prior wedge analyses have focused on aggregate/macro-level wedges. One reason for this choice is the need to calculate FIRE beliefs by averaging. If one were to instead assume FIRE and use the realized outcome for an individual as an estimate of their rational expectation, the wedge would reflect the impact of both prediction error and frictions. Studying micro-level wedges in addition

to macro-level wedges is useful because even if aggregate wedges are close to zero, there may be significant heterogeneity in wedges in the cross-section. Heterogeneity in consumer behavior can be an important determinant of macro transmission (e.g., Kaplan et al., 2018).

3 Data and Survey Design

Our data come from EarnIn, a US-based financial technology company that provides earned wage access to users with regular pay schedules, a fixed work location or electronic timekeeping system, and a connected bank account.[10] Earned wage access allows users to access their earnings prior to receiving their paycheck. EarnIn maintains an administrative database which includes information about each user and their bank account transactions (categorized by Plaid, a financial services company which links users' bank accounts with EarnIn), bank account balances, earnings, and cashout activity through the smartphone application. For an overview of the structure of the EarnIn data, see Appendix B.

EarnIn users naturally skew lower-income compared to the broader US population, given their adoption of the application. With millions of users, their database enables new insights into a policy-relevant population. Data containing both subjective economic expectations and detailed, comprehensive transactions data are rare. Prior studies have linked economic expectations to data on grocery spending (D'Acunto et al., 2021) or credit card spending (Kanz et al., 2021). Our dataset is one of few that links expectations to earnings, spending, and savings data that can paint a near-comprehensive picture of a consumer's economic activity at a high frequency (link similar data for users of a German and Chinese bank, respectively Hackethal et al., 2023; D'Acunto et al., 2024).

In selecting the subset of EarnIn users to receive survey invitations, we imposed data quality requirements on the transactions data to ensure the users' linked accounts capture their economic activity. We restricted the sample to approximately 500,000 users for whom we observe earnings, regular spending, and balances in the 12 months leading up to the survey. See Appendix B for the full data-processing details including sample restrictions, earnings identification, and categorization of outflow transactions.

EarnIn sent qualifying users an invitation to complete the survey in waves spanning September 29, 2022 to October 2, 2022 using their standard email marketing channels. Users were invited to complete a five-to-ten-minute survey about their current economic well-being and their outlook

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

for the future. They were offered a $5 Amazon gift card as an incentive to complete the survey. We closed the survey after receiving approximately 10,100 responses.[11]

The survey included questions on income and basic demographics, economic expectations, household finances, financial distress, and financial literacy. We ask respondents to forecast percent changes in prices (short-run inflation) and income over the next 12 months, medium-run inflation between September 2024 to September 2025, and their estimate of inflation over the prior 12 months (September 2021 to September 2022). The phrasing of these questions was based on similar questions from the University of Michigan Survey of Consumers (MSC) and the Federal Reserve Bank of New York Survey of Consumer Expectations (SCE). We also asked respondents to report the average percentage yield they would expect to receive on any extra money saved and the average percentage rate they would pay for any additional borrowing.

Our household finances questions include household income, debt, savings, and whether households use alternative financial products such as payday loans or pawn shops. Additionally, we elicit perceived financial distress by asking whether households perceive their debt as manageable, whether they have difficulty borrowing, and whether they have anxiety relating to their finances, among other questions. For financial literacy, we ask the two of the "Big 5" financial literacy questions that ask about interest rates and inflation (Lusardi and Mitchell, 2014). EarnIn linked survey responses to each user, which enables us to study the relationship between individuals' economic beliefs and granular consumption-savings decisions.

The median completion time of our survey was seven minutes. For our analysis, we drop responses that reflect inattention or low-effort by limiting the sample to respondents who spent at least 3.5 minutes, provided internally consistent responses, and reported expectations within "reasonable" ranges.[12] We also drop users with insufficient transactions data coverage in the post-survey period (see Appendix B for details). After imposing all sample restrictions, we have 4,753 survey respondents linked with the transactions data.

We measure earnings, spending, and bank account balances in each of the twelve months preceding and following the survey, covering the period from October 2021 through September 2023. Our analysis focuses on the 12-month pre-survey period from October 2021 through September 2022. We classify inflow transactions as earnings using a combination of the observed earnings

[11] This represented a response rate of approximately four percent before we exhausted our survey incentive budget and closed the survey.

[12] We limit the survey sample to respondents who reported expectations between the 3rd and 97th percentiles (-10% and 50% for short-run inflation expectations, 0% and 25% for savings rate expectations, 1% and 53% for borrowing rate expectations, and -25% and 75% for income growth expectations).

data with the Plaid categorization, memo line, and periodicity of the transaction. We additionally identify unemployment insurance transactions as a secondary source of income. We measure income as the sum of earnings and any received unemployment insurance transactions. We define our measure of spending as outflow transactions that can be categorized as non-durable spending.[13] We focus on non-durable spending because it corresponds more closely to consumption, whereas the relationship between spending and consumption of durables depends on their rate of depreciation and the flow of services rendered. Our primary object of interest is the average propensity to consume (APC) out of income, which we define by dividing monthly non-durable spending by monthly income.

3.1 Summary Statistics, Expectations, and Realized Deviations

Table 2 presents summary statistics on the survey respondents. Respondents are 67 percent female, 46 percent non-white, 21 percent Black, 36 years old on average, and 43 percent have attained a bachelors or graduate degree. 44 percent correctly answered two financial literacy questions, 60 percent report having less than $500 in savings, and 67 percent report having an unmanageable amount of debt.

Table 2 displays summary statistics for these economic measures. Mean non-durables spending is approximately $2,500 and mean earnings are a little over $3,600, generating a mean APC of 80%. Bank account balances just $750 on average. Appendix Figure A.1 shows the earnings distribution of the EarnIn sample compared to the Current Population Survey from September and October 2022. Our sample is more likely to be low-income, financially distressed, younger, and female than the general population.

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

Table 3 presents summary statistics for the expectation measures in the survey. One-year ahead inflation expectations for our sample are around nine percent, but respondents expect inflation to come down, with three-year inflation expectations of about five percent. Respondents perceive

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

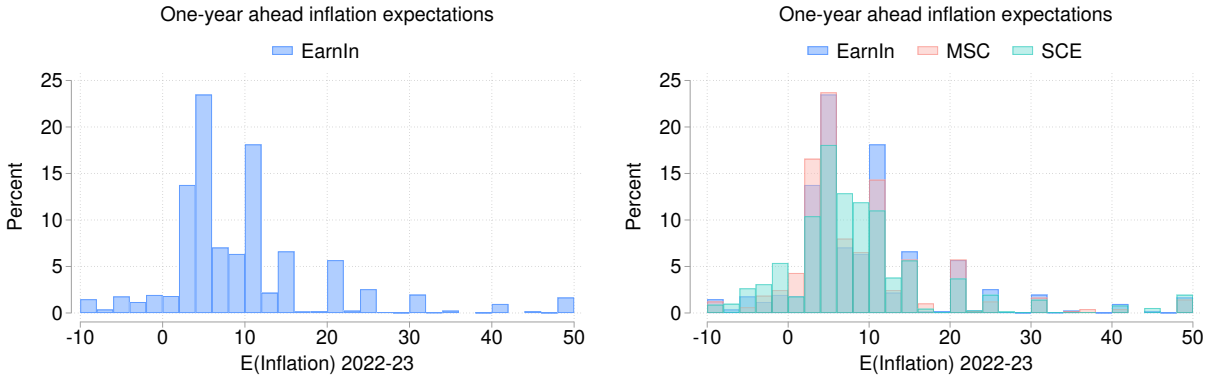
Table 2. Summary Statistics

	Mean (1)	SD (2)	P25 (3)	P50 (4)	P75 (5)	N (6)
Panel A: Demographics						
Female (%)	67.26	.	.	.	.	4,354
White (%)	53.72	.	.	.	.	4,354
Black (%)	21.66	.	.	.	.	4,354
Other race (%)	12.15	.	.	.	.	4,354
Mixed race (%)	5.32	.	.	.	.	4,354
Age	35.75	8.85	29.00	35.00	41.00	4,347
Has children (%)	50.30	.	.	.	.	4,354
Spouse or partner (%)	53.03	.	.	.	.	4,354
College (%)	42.91	.	.	.	.	4,354
Panel B: Reported household finances						
Reported income ($)	67,429	38,150	45,000	55,000	87,500	4,350
Reported savings ($)	1,971	4,517	250	250	1,750	4,352
Reported debt ($)	24,257	18,446	7,500	17,500	37,500	4,346
High financial literacy (%)	43.77	.	.	.	.	4,354
Panel C: Observed household finances						
Total spending ($)	2,893	1,740	1,678	2,475	3,651	4,354
Nondurables spending ($)	2,536	1,529	1,460	2,182	3,221	4,354
Income ($)	3,685	2,105	2,307	3,196	4,533	4,354
Nondurables C/Y (%)	77.87	45.87	47.17	66.81	94.24	4,354
Average account balance ($)	689	13,731	180	385	787	4,354
Panel D: Perceived constraints						
Difficulty borrowing (%)	45.34	.	.	.	.	4,352
Debt unmanageable (%)	67.17	.	.	.	.	4,352
Bad financial situation (%)	59.24	.	.	.	.	4,350
High financial anxiety (%)	28.62	.	.	.	.	4,351
Reported savings <$500 (%)	60.33	.	.	.	.	4,354
Used alternative financial services (%)	53.19	.	.	.	.	4,349
Panel E: Observed constraints[1]						
Balance ever <$0 (%)	61.57	.	.	.	.	4,354
Balance often <$0 (%)	5.10	.	.	.	.	4,354
Balance often <$500 (%)	77.16	.	.	.	.	4,354
Paid overdraft or late fees (%)	59.46	.	.	.	.	4,354

Notes: The table presents summary statistics of survey responses (Panels A, B, and D) and variables derived from the transactions data (Panels C and E). Columns (1) through (5) show the distribution of each variable, and column (6) shows the number of nonmissing users. Includes only users who met our survey and transactions data quality restrictions (outlined in Appendix B). All variables are at the user level except for observed household finances, which are at the user-month level and include observations from October 2021 to September 2023. Observed income and spending are are winsorized at the 1st and 99th percentiles, and the ratio of nondurables spending to earnings is winsorized at the 97.5th percentile.

[1]Users are considered to have balances "often" below $0 or $500 if their average weekly balance is below $0 or $500 for more than 52 weeks (i.e., over 50% of the sample period), respectively.

Figure 1. Distribution of Inflation Expectations



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

inflation over the prior year to be eleven percent which is higher than standard measures of inflation but could be reflective of the consumption baskets of our comparatively low-income sample.[14] Respondents forecast nominal one-year income growth of around five percent which implies forecasted real income losses of almost four percent. Respondents report quite reasonable interest rates on marginal savings and borrowing of three percent and 14 percent, respectively.

Table 3 and Figure 2 present the deviations of the realized inflation, nominal earnings growth, and interest rates from the elicited expectations. To calculate the ex-post realization of inflation, the interest rate on savings, and the interest rate on borrowing, we use annual CPI growth from BLS, the national deposit rate on savings from the FDIC, and the average commercial bank interest rate on credit cards from the Board of Governors of the Federal Reserve System.[15] Measurement of individual-level realizations for these expectations is not feasible in the transactions data, so these deviations reflect the difference between the individual's expectations and aggregate measures of inflation and interest rates. Realized inflation for 2022-2023 is 3.7 percent, 5.5 percent lower than the average respondents' forecast of 9.2 percent. Responses are in line with the relationship between perceived and realized inflation between 2021-22, where respondents reported perceived

[14]For example, the Bureau of Labor Statistics finds that lower income households have experienced higher rates of inflation from 2003 to 2021 (https://www.bls.gov/spotlight/2022/inflation-experiences-for-lower-and-higher-income-households/home.htm).

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

Table 3. Distribution of Economic Expectations

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

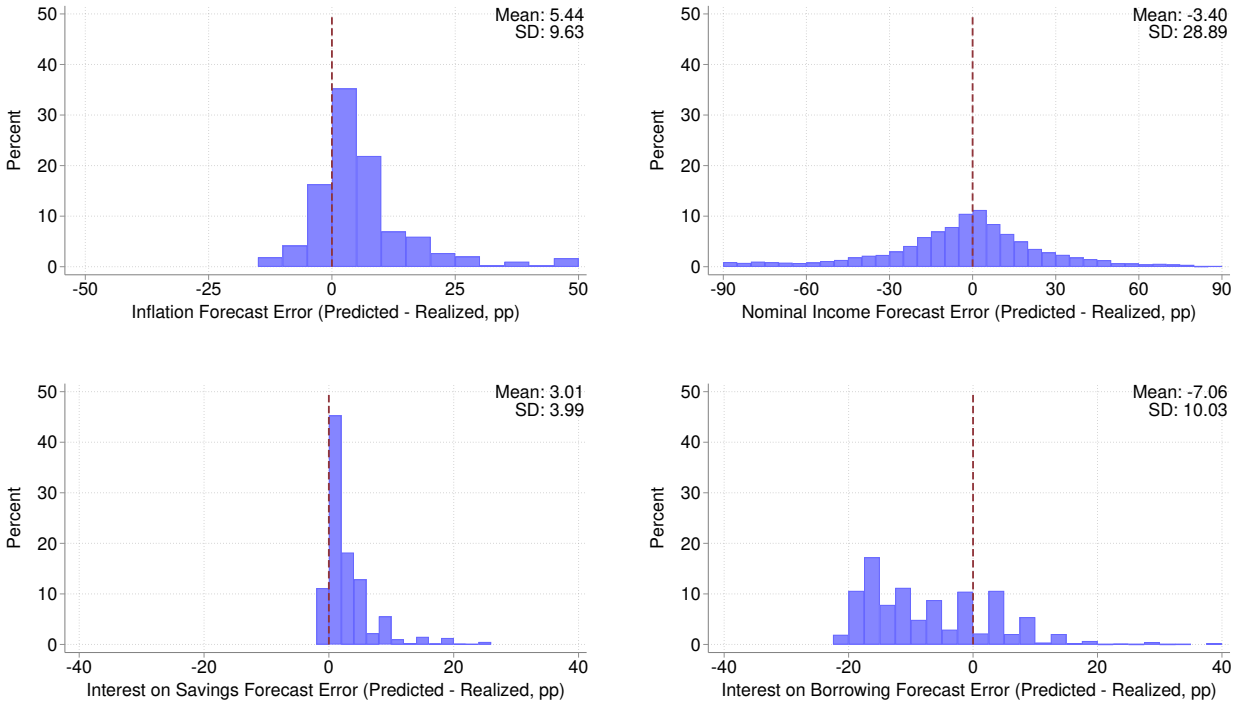
Notes: The table shows summary statistics for the economic expectations questions (Panel A) and the difference between expectations and realized values (Panel B). Includes only users who met our survey and transactions data quality restrictions (outlined in Appendix B). Column (1) reflects the realized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

inflation of 11 percent as realized inflation was 8.2 percent and could reflect higher exposure to inflation among our sample. Respondents' expected returns to savings were 3 percentage points higher than the average national deposit rate of 0.45 percent, and they underestimated the cost of borrowing as measured by the average credit card interest rate by 7 percentage points.

By leveraging the transactions data for the year following the survey, we can measure individual-level earnings realizations. This allows us to test how accurate individuals' income growth expectations are and how deviations from their expectations corresponds to their spending and savings decisions. Earnings growth expectations are remarkably accurate on average, with mean realized nominal earnings growth (measured as the percent change in earnings in the twelve months before

and after the survey) of 5.99 percent relative to expected nominal earnings growth of 5.50 percent. In addition to being approximately mean zero, the distribution of prediction errors is symmetric around zero. The 25th percentile forecasted earnings growth 15 percentage points lower than they obtained, while the 75th percentile forecasted their earnings growth by a similar magnitude of 17 percentage points higher than they obtained. We will leverage this variation in income growth forecast errors in conjunction with spending behavior to test the assumption of full-information rational expectations.

Figure 2. Deviations from Ex-Post Realizations



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

4 Results: Estimated Consumption Wedges

This section measures and analyzes consumption wedges for our survey population. To first gauge the necessity of subjective beliefs data, we test for deviations from FIRE in our sample in Section 4.1. We find that deviations from FIRE are present and strongly related to consumption.

In Section 4.2, we then calculate consumption wedges for our sample and study their distribution. We find that, while the median wedge is close to zero, there is significant heterogeneity. Notably, the median absolute value wedge is 40% of frictionless consumption. Additionally, 49% of consumers have positive wedges (over-consumption), which cannot be rationalized by financial constraints. We then show our main results are robust to calibration choices and measurement. Lastly, we correlate wedges with observables to test alternative theories about the sources and consequences of consumption distortions.

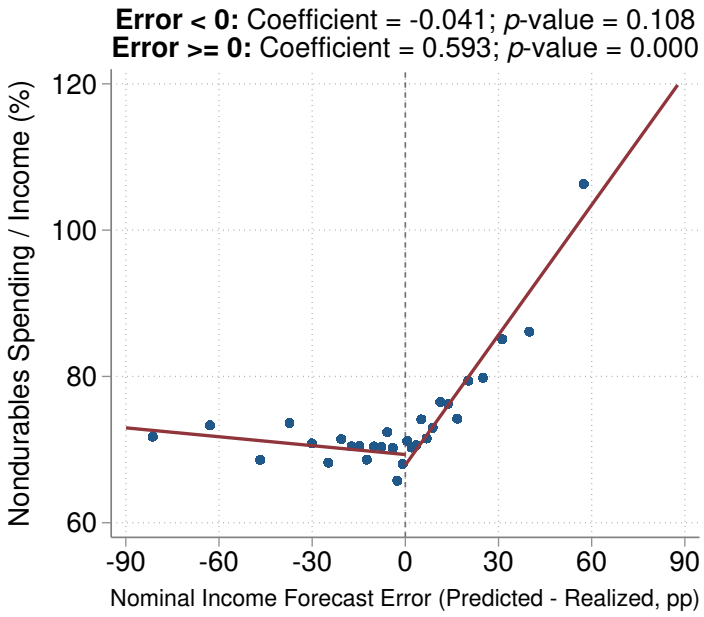
4.1 Testing for Violations of FIRE

To test for departures from FIRE, we examine whether consumption predicts income growth forecast errors. Under FIRE, time t variables should not predict errors in forecasts made at time t over time $t+1$ objects. Intuitively, if a time t variable in the consumer's information set predicts forecast errors, this indicates that information is not being fully incorporated into their expectation. Therefore, finding a correlation between current consumption and forecast errors over future income rejects FIRE. This violation could be due to a failure of rational expectations, manifesting as a violation of the law of iterated expectations. Alternatively, it could also be due to incomplete information, such as imperfect knowledge of the true stochastic processes governing the economy, resulting in the use of the "wrong" conditional expectation.

Figure 3 displays the relationship between consumption (APCs) and nominal income growth forecast errors. We measure the forecast error as the difference between the respondent's predicted nominal income growth and its realized value. In general, we see that larger forecast errors (in absolute value) are associated with higher APCs. Among those with positive forecast errors (optimists), a 5.0 pp over-estimation of income growth (0.22 standard deviations) predicts a 3.0 pp higher APC. For those with negative forecast errors (pessimists), a 5.0 pp larger under-estimation of income growth predicts a 0.21 pp higher APC. However, we note that this last relationship is marginally statistically insignificant (with a p-value of 0.108).

The correlations between nominal income growth forecast errors and consumption reject FIRE in our sample. This indicates that it is necessary to use subjective beliefs when measuring consumption wedges in order to isolate the impact of frictions and behavioral preferences from the influence of non-FIRE beliefs. If we instead used an estimate of the respondent's FIRE beliefs, the consumption wedges would reflect the difference between their actual and FIRE beliefs in addition to the influence of frictions and behavioral preferences.

Figure 3. Relationship Between Forecast Errors and Consumption



Notes: The figure presents a binned scatterplot that plots the average value of APC among quantile-based intervals of income growth forecast errors. The red lines present fitted lines for users with forecast errors above and below zero. The coefficients and p-values at the top of the figure are estimated with OLS regressions of APCs on income growth forecast errors with no fixed effects or control variables. We trim users with forecast errors outside of -90 to 90 percentage points.

The large magnitudes we estimate also have implications for theories and analysis of consumers behavior beyond our consumption wedge analysis. The strong relationship between consumption and deviations from FIRE indicates that violations of FIRE explain an economically significant amount of variation in consumption. This suggests that incorporating deviations from FIRE into models of consumer behavior is a promising direction for macroeconomics and household finance research (e.g., Auclert et al., 2020; Bhandari et al., 2022; D'Acunto et al., 2024).

4.2 Consumption Wedges

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

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

Frictionless consumption (specifically, the log-APC) is a linear function of the ratio of net worth to income $\left(\frac{A_t R_t}{Y_t}\right)$ and expectations over nominal income growth ($\widetilde{E}_t \ln G^Y_{t+k}$), inflation ($\widetilde{E}_t \ln \pi_{t+k}$), and nominal interest rates ($\widetilde{E}_t \ln \pi_{t+k}$). The α coefficients and ρ are functions of two preference parameters (discount factor β and the coefficient of relative risk aversion γ) and two approximation points (steady state APC $\frac{C}{Y}$ and net worth to income $\frac{AR}{Y}$).

4.2.1 Baseline Specification and Calibration

Throughout, we assume households face an infinite horizon ($t \to \infty$). This allows us to simplify expressions for the coefficients in the formula above. Our baseline analysis will also assume that the term structure of beliefs is flat (e.g., $\widetilde{E}_t \pi_{t+1} = \widetilde{E}_t \pi_{t+2} = \cdots$). We are currently working on a robustness analysis that will allow us to relax this assumption by instead imputing the term structure of beliefs. Under the assumption of a constant term structure and infinite horizon, we can write frictionless consumption (its log-APC) simply as:

$$\ln\left(\frac{C_t}{Y_t}\right) \approx \alpha_0 + \alpha_1 \frac{A_t R_t}{Y_t} + \overline{\alpha_Y} \widetilde{E}_t \ln G^Y_{t+1} + \overline{\alpha_\pi} \widetilde{E}_t \ln \pi_{t+1} + \overline{\alpha_R} \widetilde{E}_t \ln R_{t+1} \tag{6}$$

Table 4 details expressions for the α parameters above under these assumptions. It also reports their calibrated values, which we discuss further below.

After calculating the frictionless log-APC, we subtract it from the respondent's actual log-APC. We then exponentiate it and subtract one to obtain our preferred measure of consumption wedges. That is:

$$\text{consumption wedge} \equiv \exp(\eta_{it}) - 1 = \exp\left[\ln\left(\frac{C_{it}}{C^\star_{it}}\right)\right] - 1 = \frac{C_{it} - C^\star_{it}}{C^\star_{it}}.$$

To calculate the consumption wedge, we use both the survey and transactions data. We also make several assumptions. To measure consumption (the log-APC), we begin by calculating the monthly log-APC for non-durable goods only. That is, we sum each month's expenditures on non-durable goods and divide it by that month's total income (including any UI payments received). We calculate this measure for the 12 months prior to and including the survey month (i.e., October 2021 to September 2022).[16] To obtain a "total" APC (i.e., for all consumption), we divide each respondent's non-durable APC by the expenditure share of non-durable goods (79.37%). Under

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

Table 4. Calculated Wedge Coefficients

Coefficient	Value	Formula	Multiplicand
κ_0	0.533	$\frac{C}{Y} + \left(\frac{C}{Y} - 1\right) \frac{\rho}{1-\rho}$	Intermediate parameter
κ_1	0.581	$\kappa_0 - \frac{C}{Y} \frac{\ln \frac{C}{Y}}{1-\rho} - \left(\frac{C}{Y} - 1\right) \ln \rho \left[\frac{\rho}{(1-\rho)^2}\right]$	Intermediate parameter
α_0	0.247	$(1 - \kappa_1)\left(\frac{1-\rho}{C/Y}\right) - \frac{\ln \beta}{\gamma} \frac{\rho}{(1-\rho)}$	Intercept
α_Y	0.380	$\frac{1-\rho}{C/Y}$	Intermediate parameter
α_π	-0.161	$-\alpha_Y \frac{C}{Y}\left(1 - \frac{1}{\gamma}\right)$	Intermediate parameter
α_R	-0.219	$-\alpha_Y \left(1 - \frac{C}{Y} + \frac{C/Y}{\gamma}\right)$	Intermediate parameter
α_1	0.380	$\frac{1-\rho}{C/Y}$	$\frac{A_t R_t}{Y_t}$
$\overline{\alpha_Y}$	2.470	$\alpha_Y \frac{\rho}{(1-\rho)^2}$	$\widetilde{E}_t \ln G^Y_{t+1}$
$\overline{\alpha_\pi}$	-1.049	$\alpha_\pi \frac{\rho}{(1-\rho)^2}$	$\widetilde{E}_t \ln \pi_{t+1}$
$\overline{\alpha_R}$	-1.421	$\alpha_R \frac{\rho}{(1-\rho)^2}$	$\widetilde{E}_t \ln R_{t+1}$

Notes: This table presents formulas for the expressions used in the wedge analysis. It identifies their multiplicand and also notes their calibrated value.

the assumption that notional ("total") consumption is a Cobb-Douglas aggregate of durable and non-durable good consumption flows, this calculation yields the notional APC (for a proof, see Appendix C.2). We obtain the non-durable expenditure share from Beraja and Zorzi (2024), which calculates it using Consumer Expenditure Survey data.

For the remaining inputs, we turn to the survey data. We measure the ratio of net worth to income as reported assets minus debt divided by reported annual income. We use annual income in order to match the horizon over which expectations were solicited (one year ahead). For beliefs, we use reported one-year-ahead expectations. We assume that all inputs (APCs, the ratio of net worth to income, and expectations) are measured without error. We will later show robustness evidence suggesting that our main findings are robust to realistic violations of this "no measurement error" assumptions. Lastly, to calculate the expected interest rate, we have another degree of freedom because we solicit expectations over both the return to saving and cost of borrowing. Our baseline specification uses the cost of borrowing, as nearly all respondents have negative net worth. In our sensitivity analysis, we later show that our results are robust to using either rate as well as convex combinations.

Next, we calibrate the preference parameters and steady state values. Our baseline analysis

uses standard values for the discount factor $\beta = 0.92$ and coefficient of relative risk aversion $\gamma = 2$. For the steady state APC, we take the average of non-durable APCs across users and several years of data (when possible). We apply the same transformation to obtain the notional ("total") APC from the non-durable APC. For the steady state ratio of net worth income, we take the cross-sectional average of this ratio using the survey-reported measures. The parameter ρ is a function of the other two steady state variables. Table 5 summarizes our calibration. In a sensitivity analysis, we later examine the robustness of our main findings with respect to these calibration choices.

Table 5. Calibrated Parameter Values and their Sources

Parameter	Value	Meaning	Source
$\frac{C}{Y}$	84.92%	Steady state ratio of consumption expenditures to income	Median ratio of non-durable spending to income in EarnIn sample (67.40%) divided by non-durable share of expenditures 79.36% (calculations using the Consumer Expenditure Survey in Beraja and Zorzi, 2024)
$\frac{AR}{Y}$	-46.70%	Steady state ratio of net worth to income	Median ratio of net worth to income in EarnIn sample
$\rho = \frac{G^Y}{R}$	67.72%	Steady state ratio of income growth and return to saving	Calculation (approximating around steady state implies $\rho = \frac{1+\frac{AR}{Y}-\frac{C}{Y}}{\frac{AR}{Y}}$)
γ	2	Coefficient of relative risk aversion	Standard value
β	0.92	Annual discount factor	Standard value

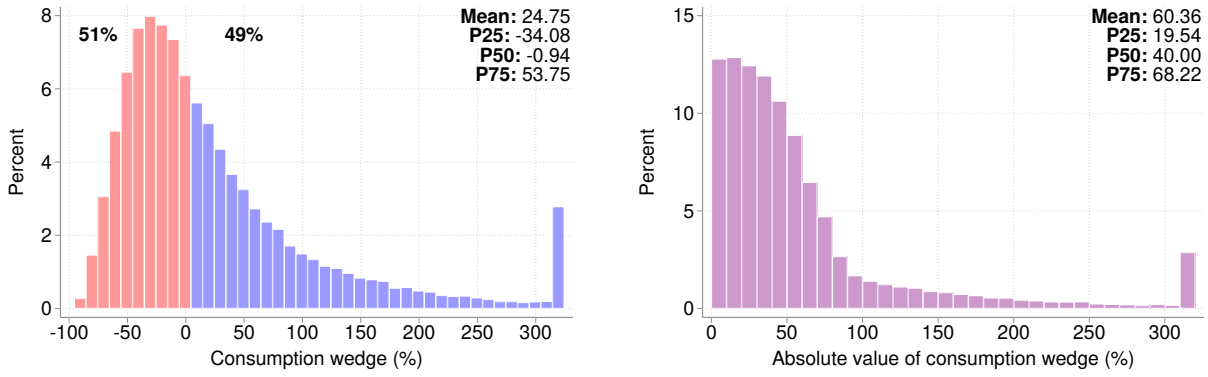
Notes: This table presents parameters used in the wedge analysis. It details the values used in our preferred specification, the economic meaning of the parameters, and the source of the chosen value.

4.2.2 Empirical Consumption Wedges

Calculating consumption wedges for each user-month, we find significant heterogeneity. Figure 4 displays histograms of measured wedges. Our first main finding is that many consumers have significantly distorted consumption. The modal wedge is approximately -25%, indicating that the modal respondent has consumption distorted downwards 25% relative to their counterfactual frictionless consumption. The mean wedge is also large, but positive: 24.7%. Taking the absolute

value of wedges, we find that the median distortion is 40%. Further examining the absolute value wedges, we see that the modal respondent has a distortion on the order of 10-20%.

Figure 4. Distribution of Consumption Friction Wedges



Notes: The figures show the distribution of consumption friction wedges (left) and their absolute values (right). Wedges are defined as the percent difference between observed APC and frictionless APC, and they are winsorized at the 97.5th percentile. Includes only users who met our survey and transactions data quality restrictions (outlined in Appendix B).

Overall, the histograms reveal that many consumers in our sample face large distortions to their consumption. These patterns indicate that frictions or behavioral preferences are a key determinant of the consumption choices of low-income households. Therefore, frictions and behavioral preferences are important to include in theories and models of consumer behavior featuring low-income households.

Moreover, our findings also highlight the value of studying distributions of wedges, as opposed to aggregate or average/median wedges. The median wedge is close to zero (0.94%). But this masks significant heterogeneity, as the median absolute value wedge is 40% of frictionless consumption. Frictions and behavioral preferences are therefore important for explaining the cross-section of consumption. In many modern macroeconomic models, heterogeneity in consumption behavior matters in the sense that it influences aggregate consumption (e.g., Kaplan and Violante, 2014; Maxted et al., 2024). For the class of models where such heterogeneity affects the aggregate economy, knowledge of the *distribution* of consumption wedges can be especially helpful for disciplining the choice and quantitative modeling of frictions and behavioral preferences. Indeed, the moments we document here may prove useful for calibrating quantitative models.

A feature of the empirical distribution of wedges that is perhaps surprising is the amount of households consuming close to their frictionless consumption. While the typical consumer has

significantly distorted consumption, around 13% of respondents are within 10 percentage points of frictionless consumption. Given the nature of the services that EarnIn provides, one might have expected respondents to be especially financially constrained. However, we note that those with near-zero wedges may be facing "offsetting" frictions and behavioral preferences. For example, a consumer may be present biased but financial constraints limit their ability to over-consume.

Our second main finding is that many households have positive wedges (over-consume). Upward and downward distortions are almost equally common in our sample. Specifically, 49% of consumers are over-consuming relative to their frictionless consumption. This finding challenges the dominant modeling paradigm in household finance and macroeconomics: that financial constraints are the primary friction driving the consumption choices of low-income and low-wealth consumers. As discussed in Section 2, financial constraints only generate negative wedges. While financial constraints could explain the 51% of consumers with negative wedges, such constraints cannot be the primary friction for the other 49%. Similarly, the presence of negative wedges rejects present bias as the sole friction/behavioral preference distorting low-income households' consumption.

What can rationalize the mix of positive and negative wedges? One possible solution is that consumers are subject to both financial constraints and present bias. Alternatively, as discussed in Section 2, frictions that generate inertia in consumption can also give rise to a mix of positive and negative wedges in equilibrium. These include consumption adjustment costs, habit formation, and bounded rationality (as in Ilut and Valchev, 2023).

A valuable endeavour for future research would be to quantitatively investigate the ability of these alternative frameworks to match the wedge distributions and other patterns that we document. We next assess the robustness of the results to calibration choices and measurement error. We then correlate wedges with other observables to provide additional insights into the sources and implications of large versus small wedges, as well as positive versus negative wedges.

4.3 Robustness

A strength of our wedge measurement approach is that can recover consumption wedges for a broad class of models using the same set of sufficient statistics. But implementing this approach requires several important assumptions that are difficult to directly test/falsify. To gauge the robustness of the main findings of (1) median absolute value wedges of 40% and (2) a 49% share of over-consumers, we conduct several sensitivity and robustness analyses. Overall, we find that

our results are strongly robust to both calibration choices and measurement error.

4.3.1 Sensitivity Analysis

Time Preference Parameter. We begin by varying our choice of annual discount factor (β) from 0.80 to 0.98. Figure D.1 displays results. Our median absolute value wedge remains nearly unchanged at 40%. The over-consumer share is slightly more sensitive. This share is increasing in the discount factor as it decreases frictionless consumption, which implies a larger wedge relative to a given amount of actual consumption. For the range of discount factors we consider, the over-consumer share ranges from approximately 40 to 50% of frictionless consumption. Thus, for even a very low discount rate, we still find many consumers are over-consuming relative to their hypothetical frictionless consumption.

Risk Preference Parameter. We next examine alternative choices for ranging from one to five for the coefficient of relative risk aversion (γ). Figure D.2 reports our results. Our median absolute value wedge is again virtually unchanged, remaining close to 40%. The share of over-consumers also remains close to half, ranging from 47% to 51%.

Preference Heterogeneity. Our baseline approach assumes homogeneous time and risk preferences. If consumers instead have heterogeneous preferences, the deviation of an individual's preference from the calibrated value would be reflected in our wedge measurements. The above analyses help alleviate concerns that the wedges we measure primarily reflect preference heterogeneity, as they show the quantitatively weak sensitivity of our main findings to these parameters.

To further alleviate this concern, we also conducted a second survey wave ("Wave 2") in July 2024 where we re-surveyed past respondents and solicited risk preferences. After applying data quality control measures, include of filtering out responses that failed a new attention check that we added, we had Wave 2 responses for 170 users. We obtained a similar distribution of wedges and found that wedges were correlated within respondents across Waves 1 and 2 (see Figures D.12 and D.13, respectively). Using survey-solicited measures of individual-level risk aversion, we recalculated consumption wedges within Wave 2 respondents. We calculate wedges using consumption data from July 2023 to June 2024. One version uses $\gamma = 2$ and the other users the survey-implied γ. Figure D.14 shows that the distribution of wedges is extremely similar. Across both versions, the Wave 2 over-consumer share is 44% and the median absolute value wedge

matches our Wave 1 result of 40%.

Steady State Values. We start by varying the steady state values of the non-durable APC from 55% to 75%. The calibrated value was 67.40%. Figure D.3 reports our results. We find that the over-consumer share remains close to 49%. The median absolute value wedge stays close to 40% over most of this range; it grows for extremely low or high values of steady state non-durable APC, peaking at nearly 55% for a steady state value of 75%. We next vary net worth to income from -75% to -33%. Its calibrated value was -46.7%. Figure D.5 displays results showing that the over-consumer share is little-changed, ranging from 44 to 57%. The median absolue value wedge ranges from 49 to 40%.

Non-Durable Expenditure Share. This share is used to convert the non-durable APC into a "total" APC. In our sensitivity analysis, we vary this share from 68% to 90%.[17] The calibrated value is 79.37%. The over-consumer share remains fairly stable, ranging from 42 to 55%. The median absolute value wedge is more sensitive. While stable for many values near our calibrated point, it eventually begins to rise quickly as the non-durable share becomes sufficiently low. This implies that our finding of a median absolute value wedge of 40% is, at worst, a conservative estimate with respect to this calibration choice.

Expected Interest Rate. Our baseline analysis uses the expected gross cost of borrowing as the expected interest rate. Our next sensitivity analysis instead uses a convex combination of the user's reported expected borrowing rate and savings rate, varying the weights from 0 to 100%. Our baseline therefore corresponds to a weighting of 100%. Figure D.6 displays results. The over-consumer share is somewhat sensitive to this choice, rising from 41% to 49% as the weight on the cost of borrowing increases. The median absolute value wedge remains stable near 40% of frictionless consumption.

Number of Pre-Survey Months. Measuring consumption wedges requires comparing the APC at time t to the frictionless APC implied by time t beliefs about $t+1$ objects. Throughout, we interpret our model as an annual model. However, to increase power, our baseline analysis uses

[17] This choice of range is motivated by other estimates of this expenditure share in the literature. Estimates in Ganong and Noel (2019) using transactions data imply a value of 68%. Laibson et al. (2022) estimate a value close to 87.5% from aggregate spending data. We obtain a value of 89.7% if we attempt to classify durable expenditure in the EarnIn transactions data.

12 months of pre-survey APC data for each respondent. Specifically, because our survey was completed at the end of September/start of October 2022, we use APCs from October 2021 to September 2022. In Figure D.7 we show how our results vary when using one, two, and up to 12 months of APC data. The over-consumer share ranges from 44 to 50%. The median absolute value wedge is nearly unchanged, ranging from 39 to 40%.

4.3.2 Measurement Error Robustness

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

Subgroups with Milder Measurement Error. We begin by studying subgroups where measurement error is plausibly milder. This helps gauge the plausibility that our results are sensitive to measurement error. Appendix D details the groups that we consider. Examples include dropping users with UI income that, despite the phrasing of our survey questions, may not have included growth in UI in their income growth forecasts. Another analysis omits users who answered one or more financial literacy questions incorrectly.[18] Such users may have a more difficult time identifying their economic expectations. Similarly, we also drop users who rounded inflation answers to a multiple of five, as such responses may exhibit rounding. We also consider several subgroups related to consumption measurement error. One example includes users whose cash withdrawals exceed 50% of non-durable spending in at least one month. Such users may have significant spending that we do not capture.

Figure D.9 reports results from restricting to these various subgroups. The share of over-consumers generally ranges from 45% to 53% across these groups. The median absolute value consumption wedge also remains similar, varying from 37% to 40% of frictionless consumption.

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

To implement this, we group respondents using k-prototype clustering, which is a combination of k-means and k-modes clustering (for more details on the algorithm, see Appendix D). We

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

group respondents based on their similarity in terms of reported age, annual income, savings, and indicators for gender, race, relationship status, presence of children, college education, and political affiliation.[19] We vary the number of clusters from 1,000 to as few as 100. With 1,000 clusters, there are approximately 47 people per cluster (and approximately 470 people per cluster with 100 clusters). The share of over-consumers remains similar with 1,000 clusters at 46%. It shrinks as low as 37% when using only 100 clusters. The median absolute value wedge is more sensitive to cluster size. With small clusters (1,000 total), the median absolute value wedge is approximately 25% of frictionless consumption. It falls as low as 12% when using only 100 clusters. However, we note that these clusters are quite large (around 470 people each). These smaller wedges may be the result of attenuation due to being more likely to aggregate wedges across over- and under-consumers as cluster size grows.

User-Level. An alternative to grouping users together is to collapse our data to the user-level. To do so, we measure a single APC for each user by taking the median APC over October 2021 to September 2022. This aggregation may smooth out both measurement error in APCs as well as seasonal fluctuations. Figure D.8 reports results. We obtain the same rate of over-consuming (49%) and a slightly lower median absolute value wedge (32%).

Adding Noise. To gain a sense of the amount of measurement error necessary to significantly alter our findings, we study the impact of adding noise to the inputs used to measure wedges. We conduct 1,000 simulations where we add random noise to each of the consumption wedge inputs (APC, ratio of net worth to income, and expectations). The noise is drawn from a distribution with mean-zero and a standard deviation of 1.5 percentage points. In each simulation, we re-calculate wedges and measure the share of over-consumers and median absolute value wedge. Figure D.11 reports histograms of our results across the simulations. Overall, the over-consumer share remains near 49—50% and the median absolute value wedge close to 40%. This suggests it would take sizable measurement error to significantly alter our findings.

4.4 Interpretation: Evidence from Wedge Correlates

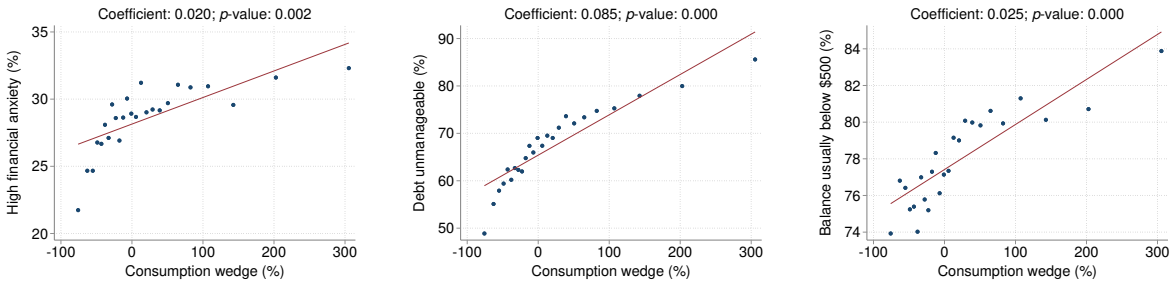
We next examine how consumption wedges vary with observable characteristics. We focus in particular on correlations that may shed light on the nature of the frictions and/or behavioral

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

preferences distorting consumption.

Financial Distress Proxies. We begin by examining how wedges co-vary with proxies for financial distress. These include subjective measures such as ratings of anxiety about finances or the manageability of one's debt. We also use objective measures such as having savings account balances below $500 most of the time. For all measures we consider, we find a strong, positive relationship between consumption wedges and financial distress. One interpretation of this result is as a validation of our consumption wedge measures. If the wedges truly captured distortions to consumption, one would expect them to be related to financial distress.

Figure 5. Relationship Between Consumption Friction Wedges and Financial Distress



Notes: The figure presents binned scatterplots that illustrate the relationship between consumption friction wedges and different measures of financial distress. Each binned scatterplot plots the average value of the financial distress indicator within quantile-based intervals of consumption wedges, with no fixed effects or control variables. Wedges are defined as the percent difference between observed APC and frictionless APC, and they are winsorized at the 97.5th percentile. Includes only users who met our survey and transactions data quality restrictions (outlined in Appendix B).

A second interpretation relates to implications for theories of consumer behavior. A perhaps surprising pattern is that the relationship is not V-shaped. That is, financial distress is not increasing in the absolute value of wedges. Rather, the most financially distressed consumers are those with positive wedges. Those with more extreme negative wedges experience less financial distress. One possible rationalization for this pattern relates to wealthy hand-to-mouth households. While our sample skews low-income, it also includes some moderate-income households and homeowners. These house-rich (but possibly cash-poor) users may experience less financial distress than non-homeowner users. At the same time, constraints on their ability to borrow against this wealth may significantly distort the consumption of homeowners downwards. Users that are both house- and cash-poor may lack the necessary collateral to be as exposed to financial constraints as homeowners.

Our findings also indicate that over-consumption in particular is associated with experiencing more severe financial distress. This suggests that the underlying frictions or behavioral preferences driving over-consumption are linked with more severe financial distress as well. To the extent that financial distress is associated with higher marginal utility of consumption, those with positive wedges would tend to have the highest marginal utility.

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

We measure individual-level MPCs based on consumers' non-durable spending responses to the March 2021 stimulus payments. These checks provided $1,400 to each eligible individual, with an additional $1,400 for each dependent.[20] Approximately 70% of the survey analysis sample received a stimulus check. Of these recipients, 94% received their stimulus checks via direct deposit. We determine each user's stimulus payment date and amount from the transactions data. For each user, we examine consumption from 28 days before to 27 days after the stimulus check was received. Days -27 through -1 are the "pre" period, and days 0 through 27 are the "post" period. We then use the same date ranges in 2022 and 2023 as comparison periods. We calculate each individual's MPC as follows:

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

where

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

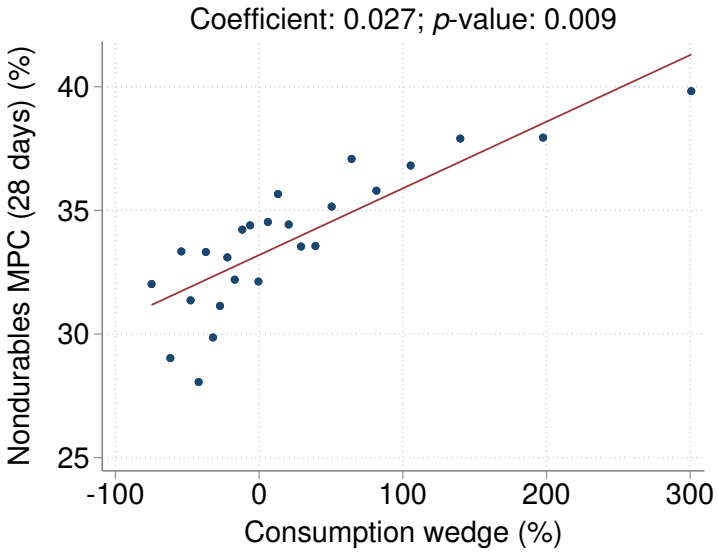
Our MPC measure captures the "excess" consumption associated with receipt of the stimulus check. We note that this measure should be interpreted as at best a proxy for an individual's MPC, as we only have three observations per person. As such, this measure is unlikely an asymptotically valid estimate of the individual's true MPC. The estimated MPCs generally range from 30-40%. 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, our next analysis excludes users with MPCs in the bottom or top 2.5 percentile.

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

Figure 6 displays a binscatter comparing individuals' MPCs against their consumption wedges. As consumption wedges increase, we observe larger MPCs on average. A 20 percentage point larger wedge is associated with a 0.54% larger MPC. A limitation of our MPC measure is that it measured for a check received over a year prior to our survey (conducted at the end of September 2022). The relationship we measure likely understates the relationship one would find if able to instead use a contemporaneous MPC measure.

These results have two implications. The first is as a validation of the consumption wedges, showing that they are strongly related to an important economic behavioral response. The second is that higher MPCs are associated with over-consumption, rather than the under-consumption. This suggests that the cause of high MPCs may be largely due to frictions or behavioral preferences that generate over-consumption.

Figure 6. Relationship Between Consumption Friction Wedges and Nondurable MPCs

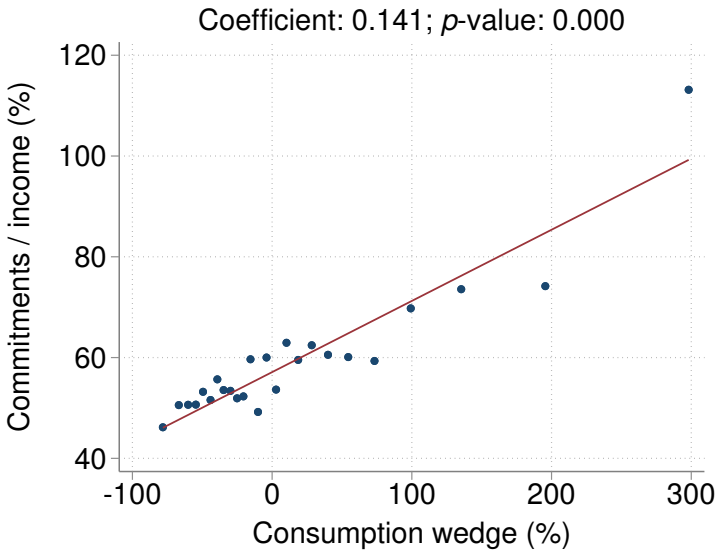


Notes: The figure illustrates the relationship between consumption friction wedges and nondurable MPCs. The binned scatterplot plots the average nondurables MPC within quantile-based intervals of consumption wedges, with no fixed effects or control variables. Wedges are defined as the percent difference between observed APC and frictionless APC, and they are winsorized at the 97.5th percentile. Each user's nondurable MPC is calculated by comparing the change in nondurable spending in the 28 days after their March 2021 stimulus check to the average change in spending during the same 28-day period in 2022 and 2023. MPCs are winsorized at the 2.5th and 97.5th percentiles. Includes wave 1 respondents who meet the sample restrictions outlined in Appendix B. Includes monthly observations from October 2021 to September 2022.

Consumption Commitments. Lastly, we examine the relationship between consumption wedges and a proxy for consumption commitments in Figure 7. This measure is available only for the

Wave 2 survey respondents. Wave 2 added questions soliciting respondents' monthly expenditures on housing and childcare. Our survey focused on these two specific expenses because they are among the largest "consumption commitments" (i.e., difficult to adjust expenditures) and can be relatively more difficult to identify in transactions data. We divide these reported monthly expenditures by median monthly income over the July 2023 to June 2024, as measured in the transactions data. We interpret a high value of this ratio as a consumer having a high degree of consumption commitments. We find a a strong, positive relationship between wedges and consumption commitments. A 10 percentage point larger wedge is associated with a 1.4% percentage point higher ratio of committed consumption to income. This pattern suggests consumption commitments are a plausible friction behind the consumption wedges we measure.

Figure 7. Relationship Between Consumption Friction Wedges and Consumption Commitments



Notes: The figure illustrates the relationship between consumption friction wedges and consumption commitments. The binned scatterplot plots the average consumption commitments (% income) within quantile-based intervals of consumption wedges, with no fixed effects or control variables. Wedges are defined as the percent difference between observed APC and frictionless APC, and they are winsorized at the 97.5th percentile. Consumption commitments (% income) is calculated as the ratio of survey-reported monthly housing plus childcare costs to median monthly income from the transactions data. Includes wave 2 respondents who meet the sample restrictions outlined in Appendix B. Includes monthly observations from July 2023 to June 2024.

5 Conclusion

This paper introduces a novel approach to measure individual-level distortions to consumption. We use a new dataset that links surveyed economic expectations to administrative transactions data for predominantly low-income households. We measure the impact of frictions as a wedge between actual consumption and a counterfactual "frictionless" benchmark. Our benchmark allows households to deviate from full-information rational expectations (FIRE), so that the wedge isolates the influence of frictions or behavioral preferences from non-FIRE beliefs. Because our benchmark is a special case in a large class of models, our formula for consumption wedges would allow users to measure distortions due to a wide variety of frictions or behavioral preferences.

We document two main results. First, the average consumer exhibits large distortions to consumption. The median wedge is 40% in absolute value. Second, we observe a mix of positive and negative wedges: 49% of wedges are positive (over-consumption) and 51% are negative (under-consumption). Since financial constraints can only generate negative distortions to consumption, this finding implies that additional or alternative frictions are necessary to explain the consumption choices of low-income households. Our results are robust to varying the wedge parameters and to reducing measurement error through additional sample restrictions or data transformations.

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

consumption choices of low-income households.

References

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

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

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

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

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

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

Bhandari, Anmol, Jaroslav Borovička, and Paul Ho, "Survey data and subjective beliefs in business cycle models," *Available at SSRN 2763942*, 2022.

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

Appendix

Contents

A Appendix Figures **44**

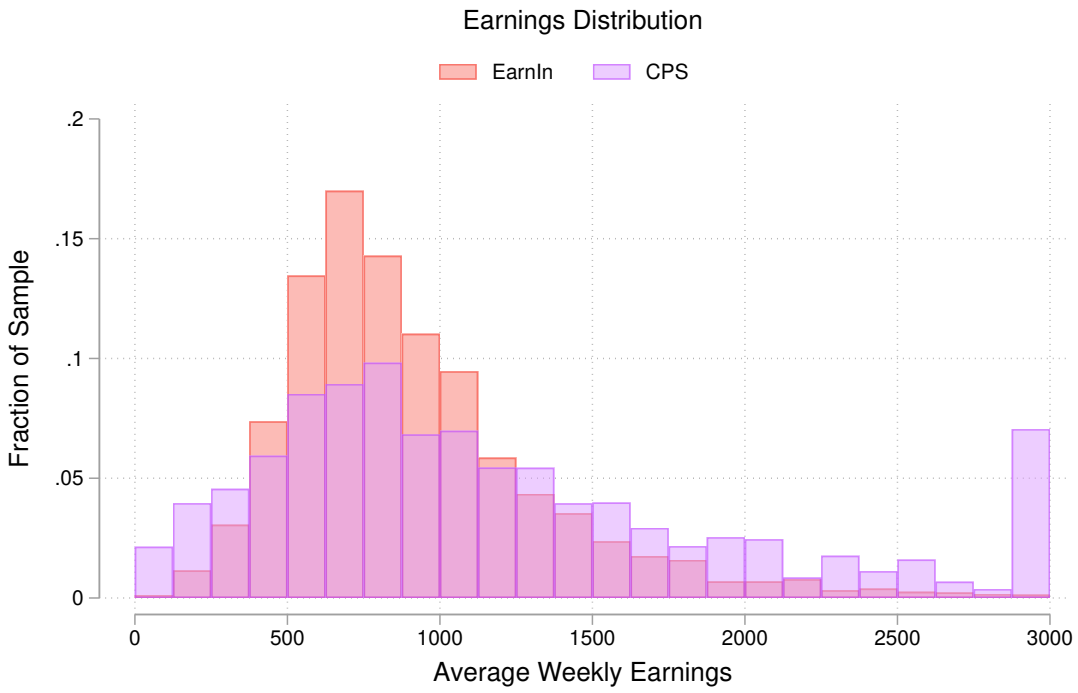
B Data Construction **44**
B.1 Structure of Transactions Data . . . 44
B.2 Identifying Earnings and UI Payments . . . 45
B.3 Categorizing Transaction Outflows . . . 45
B.4 Defining the Sampling Frame . . . 47
B.5 Creating a Linked Survey-Transactions Data Panel . . . 47
B.6 Defining the Analysis Sample . . . 47
B.7 Data Transformations . . . 48

C Theory Derivations and Extensions **49**
C.1 Deriving Log-linearized Frictionless Consumption . . . 49
C.1.1 Log-linearized Euler Equation . . . 49
C.1.2 Log-Linearized Budget Constraint . . . 52
C.1.3 Combining the Euler Equation and Budget Constraint . . . 54
C.2 Extension: Durable and Non-Durable Goods . . . 59
C.3 Calibration of Wedge Parameters . . . 60

D Sensitivity Analysis **63**

A Appendix Figures

Figure A.1. Distribution of Average Weekly Earnings: Earnin vs. CPS



Notes: The figure shows the distribution of average weekly earnings for the EarnIn sample versus the September and October Current Population Survey.

B Data Construction

B.1 Structure of Transactions Data

We receive anonymized transactions data from EarnIn that covers bank transactions, daily checking and savings account balances, transactions classified as earnings, and user information in the form of "tags." We use data from October 1, 2021 through September 30, 2023, which covers the 12 months before and after the survey. None of the data we receive contains personally identifiable information, and all data is stored and processed on secure servers.

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

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

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

The earnings data is a direct subset of the transactions data, covering the earnings inflows from the jobs each user reports to EarnIn. These data include the date of payment, posted date of the transaction, the amount of earnings, and whether those earnings are from unemployment benefits.

B.2 Identifying Earnings and UI Payments

We leverage the transactions and observed earnings datasets to 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. If a user has only one source of earnings within a week, the first earnings variable reflects this amount of earnings from this source, and the remaining two variables are 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. This excludes X individuals.

B.3 Categorizing Transaction Outflows

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

The transaction outflows data comprises over 500 categories from Plaid. We start by mapping these Plaid categories to 36 broader categories that can be mapped to spending, payments, or

transfers. These categories are as follows:

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

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

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

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

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

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

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

B.4 Defining the Sampling Frame

We send survey invitations to a restricted list of EarnIn users with adequate transactions data coverage in the 12 months leading up to the survey. We apply the sample restrictions listed below. 500,804 of 3,527,031 EarnIn users met these sample restrictions. We sent invitations to this full set of users and closed the survey once 10,103 respondents fully completed the survey.

- Non-missing balances data in each bi-weekly period from September 1, 2021 through August 30, 2022
- Non-missing earnings data at least once between September 1, 2021 through August 30, 2022
- First recorded transaction before September 1, 2021 and latest recorded transaction after August 15, 2022
- At least 5 outflows per month between September 2021 and August 2022
- Non-missing bank connection date

B.5 Creating a Linked Survey-Transactions Data Panel

We merge the survey responses to the cleaned transactions data, which includes user tags, balances, and categorized inflows and outflows for each user at the weekly level. One hundred survey respondents are dropped during this merge because they deleted their EarnIn accounts and no longer appear in the transactions data.

For our analysis, we collapse this data to the monthly level, ranging from October 2021 through September 2023. Because weeks do not perfectly align to months, we adjust inflows and outflows by the number of end-of-week dates (Fridays) in each month using the equation below.

$$Flow_t^{adj} = \sum_{w=1}^{eow_t} (Flow_w \times \frac{52}{12} \times \frac{1}{eow_t}) \tag{9}$$

In this equation, t indexes months, w indexes weeks, $Flow$ represents either inflows or outflows, and eow_t represents the number of Fridays in month t. The expression within the summation annualizes $Flow_t$ by multiplying by 52, converts to a monthly average by dividing by 12, then converts to an adjusted weekly amount by dividing by the number of Fridays. For months with 5 Fridays, weekly inflows and outflows are adjusted downwards. Adjusted monthly inflows and outflows are calculated by summing the adjusted weekly amounts for each Friday within the month.

The weekly transactions data includes the average bank account balance each week and indicators for whether balances fall below $0 or $500 any day during the week. For the monthly panel, we save the average, first-week value, and last-week value of average weekly balances. We also calculate the number of weeks each month where balances fall below $0 and $500.

B.6 Defining the Analysis Sample

We apply several sample restrictions to the monthly panel to arrive at our analysis sample. We first impose survey-based restrictions that drop users who completed the survey in an unreasonably short time frame or provided contradictory or unrealistic responses. Then, we apply restrictions based on the transactions data that incorporate the 12-month periods pre- and post-survey. These

latter 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.

The survey quality restrictions are listed below. 7,924 of 10,003 respondents (79%) meet these restrictions.

- Survey duration at least 3.5 minutes (approximately the 5^{th} percentile)
- Reported debt amounts are consistent (i.e., users who report zero debt must report N/A for debt manageability, and vice versa)
- Expectations for one-year inflation, income growth, interest on savings, and interest on debt are within the 3^{rd} and 97^{th} percentiles:
  - E(Inflation, 12M): -10% to 50%
  - E(Income Growth): -25% to 70%
  - E(Interest on Savings): 0% to 25%
  - E(Interest on Borrowing): 1% to 60%

The transactions data quality restrictions are listed below. 5,864 of 10,003 respondents (59%) meet these restrictions.

- Sufficient transaction activity: 20+ outflows per month for all 24 months
- Sufficient balances data: Non-missing balances each week for 18+ months
- Sufficient categorizable spending: $\frac{\text{Consumption}}{\text{Outflows - Internal Transfers}} \geq 20\%$ for 18+ months
- Reasonable balance of inflows and outflows: $\frac{\text{Outflows}}{\text{Inflows}} \in [50\%, 150\%]$ for 18+ months
- Informative memos: $< 1\%$ of memos are "CREDIT", "DEBIT", or missing across months

After applying these restrictions, we drop all user-month observations where earnings and UI payments are both zero. The resulting analysis sample contains 102,368 observations across 4,354 users.

B.7 Data Transformations

We make a number of adjustments to the transactions data to standardize variables and reduce noise. As highlighted above, we aggregate data to the monthly level, which involves adjustments for the number of weeks per month. We deflate inflows, outflows, and balances using the CPI, with September 2022 as the base month. To account for outliers, we winsorize the variables used for our analysis. Before calculating APC, forecast errors, and consumption wedges, we winsorize earnings and spending at the 1^{st} and 99^{th} percentiles. After calculating these variables, there are still significant outliers driven by observations with small denominators. Thus, we winsorize APCs and wedges at the 97.5^{th} percentile.

C Theory Derivations and Extensions

C.1 Deriving Log-linearized Frictionless Consumption

C.1.1 Log-linearized Euler Equation

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

Lemma 1: Log-linearized Original Euler Equation

A first order log-linear approximation of the Euler equation

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

yields:

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

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

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

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

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

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

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

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

Note that this implies

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

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

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

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

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

Taking a first order approximation of the RHS yields:

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

where

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

Plugging in and rearranging gives

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

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

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

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

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

This allows us to obtain:

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

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

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

Rearranging, we finally obtain:

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

□

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

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

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

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

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

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

Remark 2: Log-Linearized Multi-Period Euler Equation

Iterating Equation (11) forward to j periods gives:

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

C.1.2 Log-Linearized Budget Constraint

We next turn our attention to the budget constraint.

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

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

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

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

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

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

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

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

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

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

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

Rearranging, we isolate initial wealth:

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

Forward-iterating the above expression gives

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

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

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

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

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

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

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

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

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

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

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

With this, begin approximating the RHS:

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

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

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

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

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

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

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

□

C.1.3 Combining the Euler Equation and Budget Constraint

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

Proposition 1: Characterization of Frictionless Consumption

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

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

where

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

With these definitions, we write our approximation as:

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

□

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

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

Lemma 3: Frictionless Consumption Under an Infinite Horizon

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

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

where

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

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

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

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

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

We next turn to κ_1:

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

We now simplify the remaining parameters with infinite sums:

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

and

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

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

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

where the α parameters are defined above. □

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

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

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

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

where

$$\overline{\alpha_Y} = \alpha_Y \frac{\rho}{(1-\rho)^2} \tag{20}$$

$$\overline{\alpha_R} = \alpha_R \frac{\rho}{(1-\rho)^2} \tag{21}$$

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

C.2 Extension: Durable and Non-Durable Goods

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

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

We make two assumptions.

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

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

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

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

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

where

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

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

Assumption 2: Cobb Douglas Aggregation.

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

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

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

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

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

Lemma 4: APC Calculation Including Consumption of Durables.

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

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

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

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

C.3 Calibration of Wedge Parameters

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

Remark 4: Simplification of κ_0 and ρ

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

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

This implies $\kappa_0 = \frac{C}{Y}$ and

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

Table C.1. Calibrated Parameter Values and their Sources

Parameter	Value	Meaning	Source
$\frac{C}{Y} \equiv \tilde{c}$	84.92%	Steady state ratio of consumption expenditures to income	Median ratio of non-durable spending to income in EarnIn sample (67.40%) divided by non-durable share of expenditures 79.36% (calculations using the Consumer Expenditure Survey in Beraja and Zorzi, 2024)
$\frac{AR}{Y}$	-46.70%	Steady state ratio of net worth to income	Median ratio of net worth to income in EarnIn sample
$\rho = \frac{g^Y}{R}$	67.72%	Steady state ratio of income growth and return to saving	Calculation (approximating around steady state implies $\rho = \frac{1+\frac{AR}{Y}-\frac{C}{Y}}{\frac{AR}{Y}}$)
γ	2	Coefficient of relative risk aversion	Standard value
β	0.92	Annual discount factor	Standard value

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

Table C.2. Calculated Wedge Coefficients

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

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

D Sensitivity Analysis

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

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

Second, we vary the assumed steady state parameters. These parameters include the steady state nondurable APC, the nondurable share of spending, the steady state AR/Y, and whether the Euler equation incorporates interest rates on borrowing versus saving. In our baseline specification, we assume a 67.4% nondurable APC, a 79.4% nondurable share of spending. and a -46.7% ratio of net worth to income. Appendix Figures D.3 and D.4 vary the nondurable APC from 0.55 to 0.75 and the nondurable share of spending from 0.68 to 0.90. Appendix Figure D.5 varies the AR/Y assumption from -75% to -33%. Our baseline specification also assumes that the interest rate in the Euler equation is the interest rate on borrowing. Appendix Figure D.6 tests various linear combinations between the interest rates on borrowing and saving.

Third, we vary the amount of data we incorporate as well as the observation level. Our baseline specification uses 12 months of pre-survey data, from October 2021 to September 2022, and presents wedges at the user-month level. Appendix Figure D.7 shows the sensitivity of our results to varying the number of pre-survey months from one to 12. Additionally, Appendix Figure D.8 presents user-level wedges, calculated by taking the median wedge across the 12 pre-survey months for each user.

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

Fifth, we test the robustness of our main wedge results to clustering users before calculating the wedge, taking the within-cluster median of each input. These results are shown in Appendix Figure D.10. We cluster users on reported age, income, and savings, and indicators for gender (male, female, or other), race (white, non-white, or not specified), spouse or partner, having children, college education, and political affiliation (Democrat, Republican, or other). To accommodate both numeric and categorical variables, we employ with the k-Prototype clustering algorithm, which combines the k-means and k-modes clustering algorithms.[21] We vary the number of clusters from

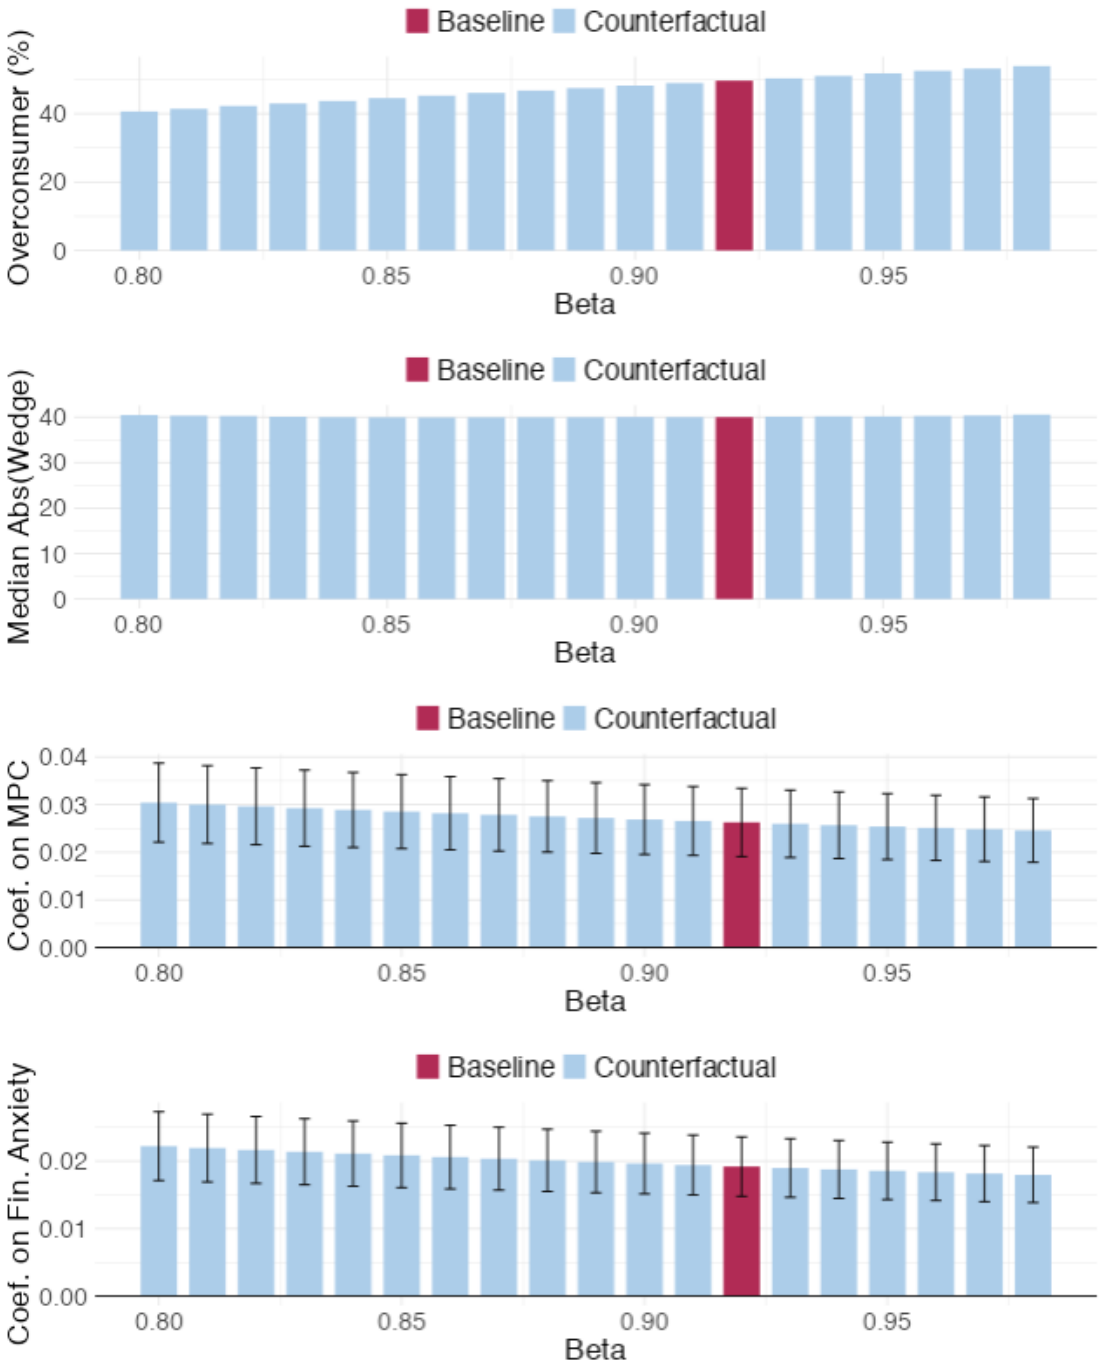
[21]The k-Prototype algorithm combines the k-means and k-modes clustering algorithms to accommodate both numeric and categorical variables. For numeric variables, distance is calculated using standard Euclidean distance, and the prototype (or center) of each cluster is the mean of all points within the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

100 to 1,000.

Sixth, we test the robustness of our wedge results to adding random noise to each input of the wedge calculation (nondurables APC, reported wealth-to-income, and expectations). We conduct a Monte Carlo simulation with 1,000 iterations. In each iteration, we generate six normal random variables (corresponding to each wedge parameter) with a mean of 0 percentage points and a standard deviation of 1.5 percentage points, and then add these to each corresponding input. In line with the observation levels of the data, noise is generated at the user-month level for APCs and at the user level for initial wealth and expectations. After adding the noise, we calculate wedges and reproduce our key results.

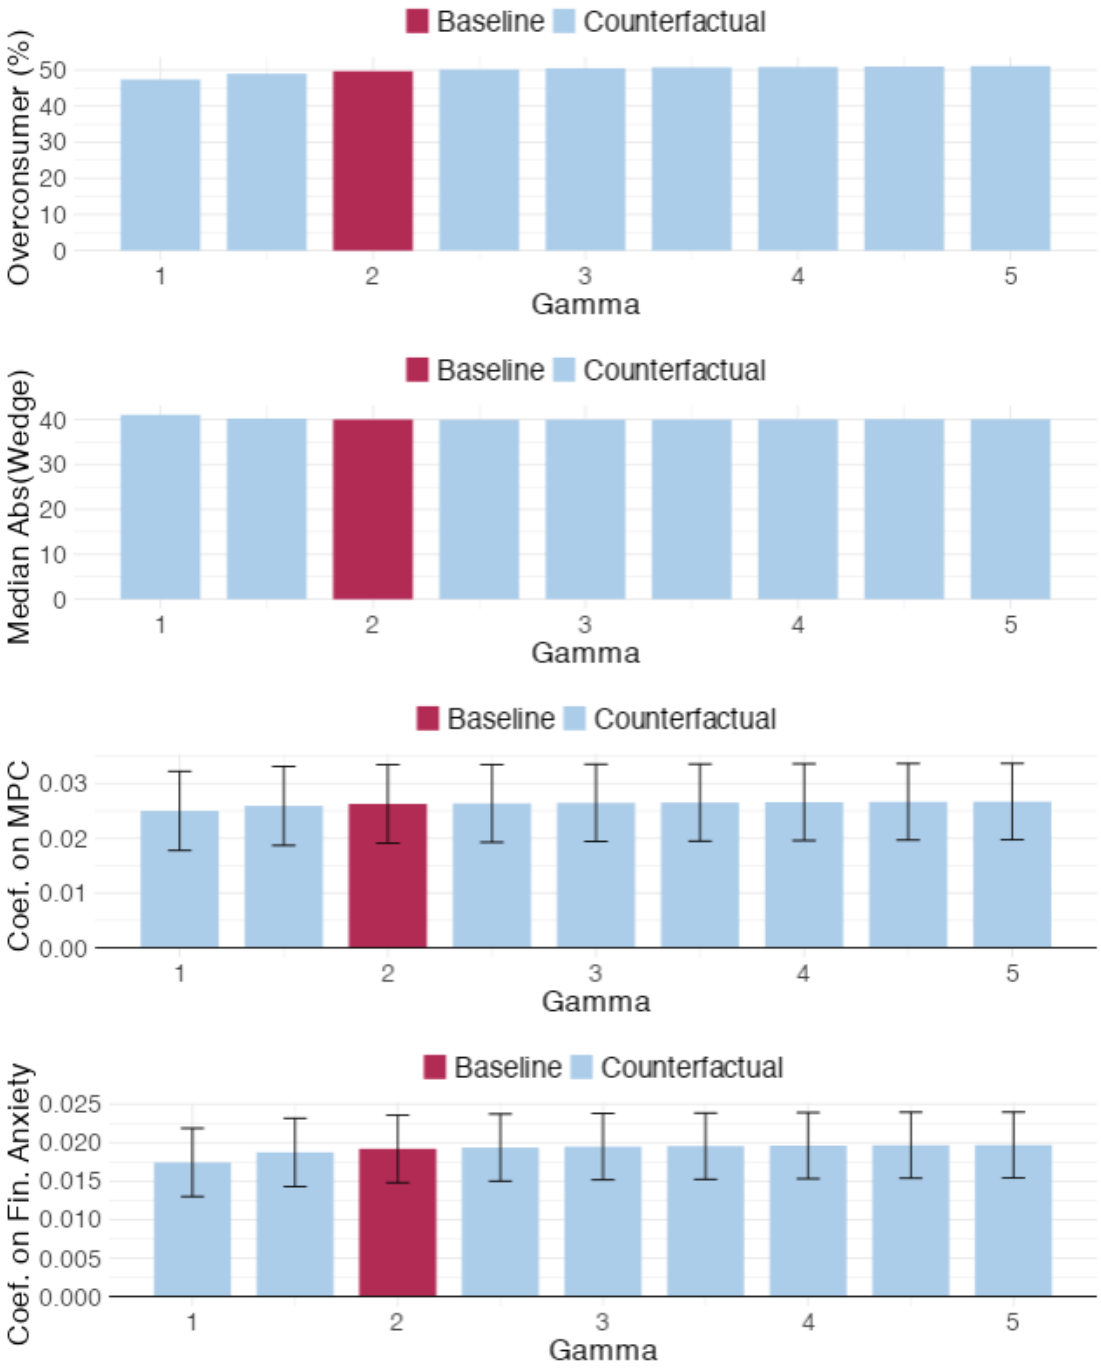
Finally, we present consumption wedges for the second-wave survey. Appendix Figure D.12 presents that baseline consumption wedges for second-wave respondents. A key element of the second-wave survey is that we measure risk aversion by eliciting respondents' certainty equivalent to a lottery with a 50-50 chance of receiving \$0 or \$450. We use this certainty equivalent measure to calculate user-specific gamma values. This allows us to plug risk preferences directly into the wedge calculation, rather than assuming a standard gamma value. Appendix Figure D.14 presents these results.

Figure D.1. Sensitivity of Wedges to Beta



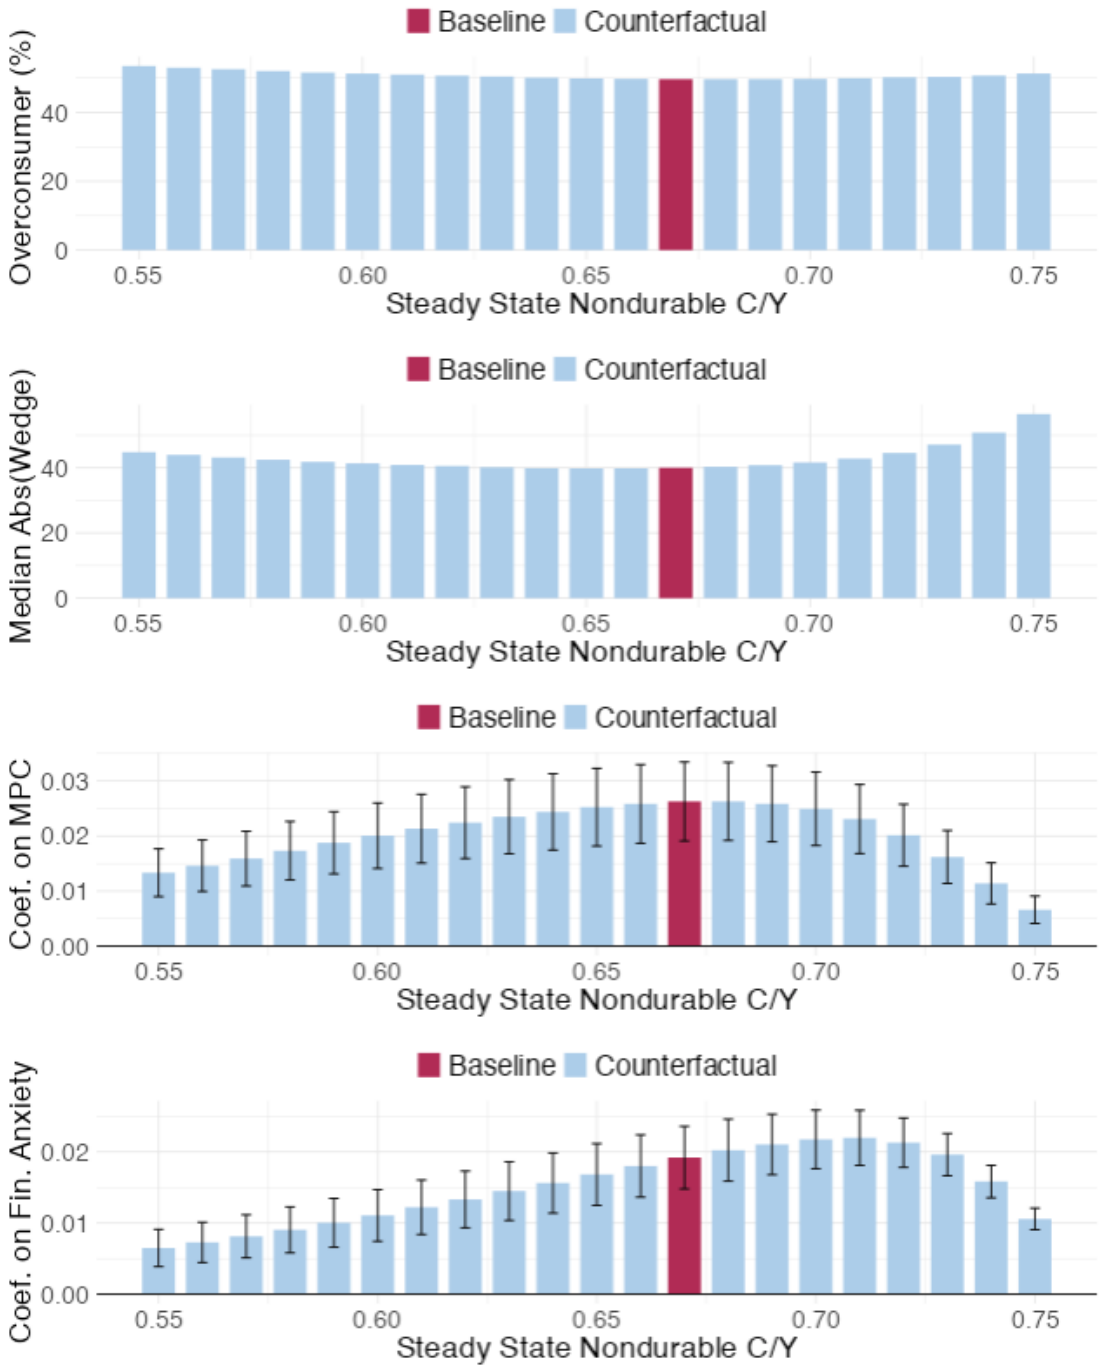
Notes: The figure presents the sensitivity of four estimated results to our assumed value of beta in the wedge calculation, with the assumed value ranging from 0.80 to 0.98 in increments of 0.01 (our baseline assumption is 0.92). These results include the percent of users with a positive wedge (Panel A), the median absolute value wedge (Panel B), and the estimated coefficients from separate regressions of nondurables MPC and an indicator for financial anxiety on the consumption wedge (Panels C and D, respectively). In the regressions, we cluster standard errors at the user level. We hold all other parameters constant at our baseline values (for the steady state parameters, these are rounded to the nearest percentage point). Includes wave 1 respondents who meet the sample restrictions outlined in Appendix B. Includes monthly observations from October 2021 to September 2022.

Figure D.2. Sensitivity of Wedges to Gamma



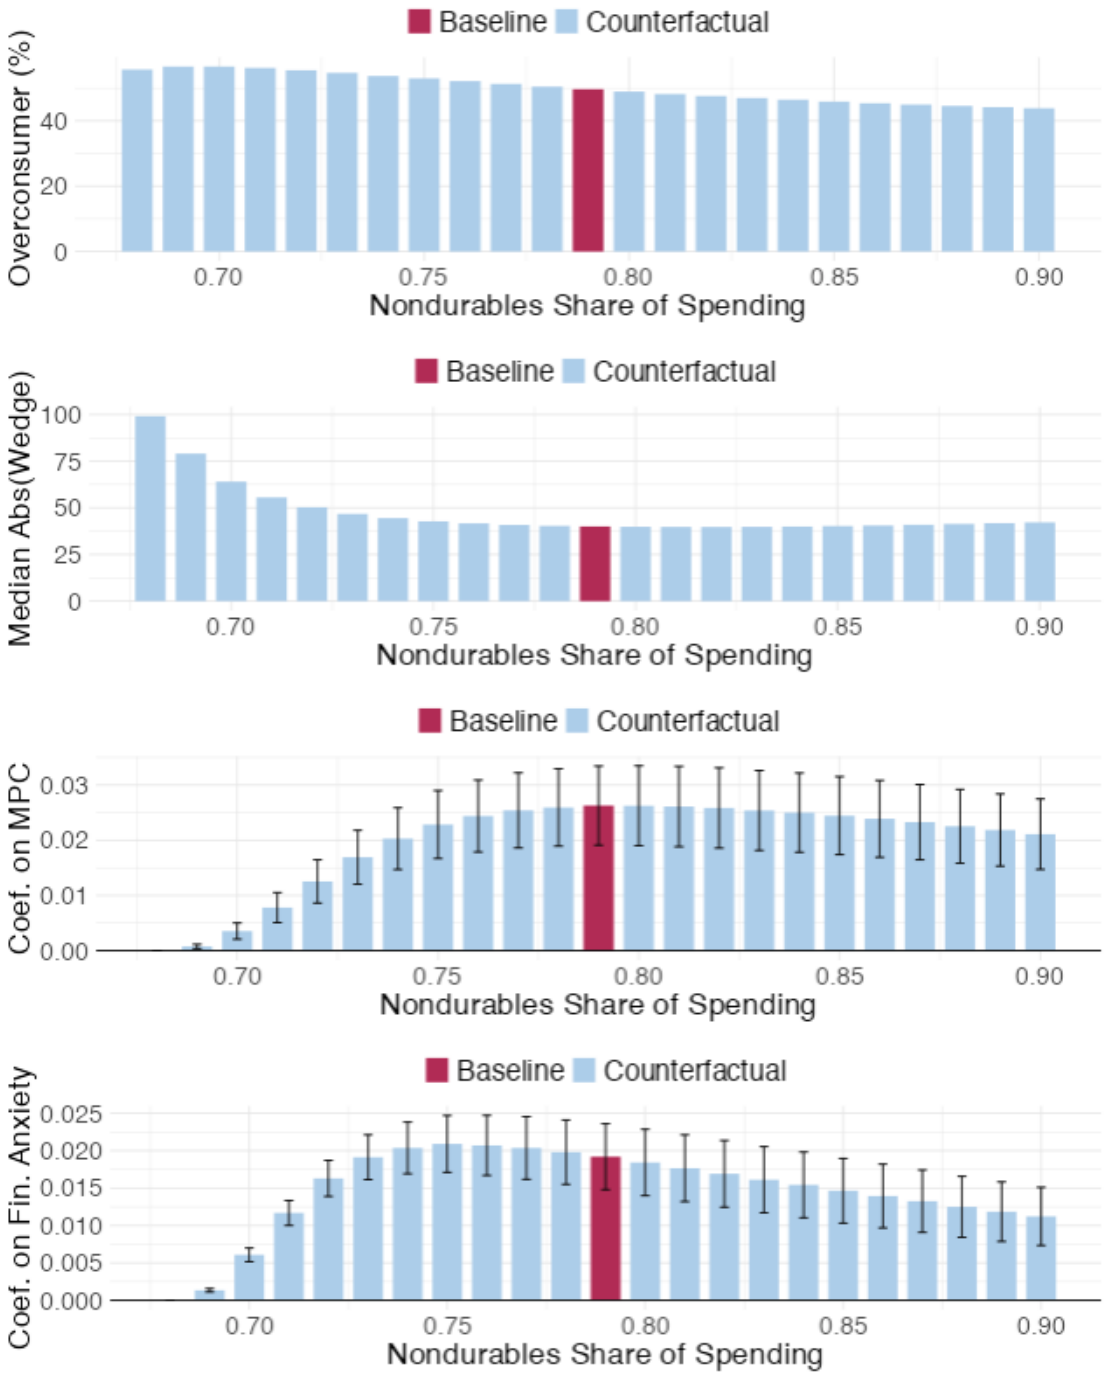
Notes: The figure presents the sensitivity of four estimated results to our assumed value of gamma in the wedge calculation, with the assumed value ranging from 1.0 to 5.0 in increments of 0.5 (our baseline assumption is 2.0). These results include the percent of users with a positive wedge (Panel A), the median absolute value wedge (Panel B), and the estimated coefficients from separate regressions of nondurables MPC and an indicator for financial anxiety on the consumption wedge (Panels C and D, respectively). In the regressions, we cluster standard errors at the user level. We hold all other parameters constant at our baseline values (for the steady state parameters, these are rounded to the nearest percentage point). Includes wave 1 respondents who meet the sample restrictions outlined in Appendix B. Includes monthly observations from October 2021 to September 2022.

Figure D.3. Sensitivity of Wedges to Steady State Nondurable C/Y



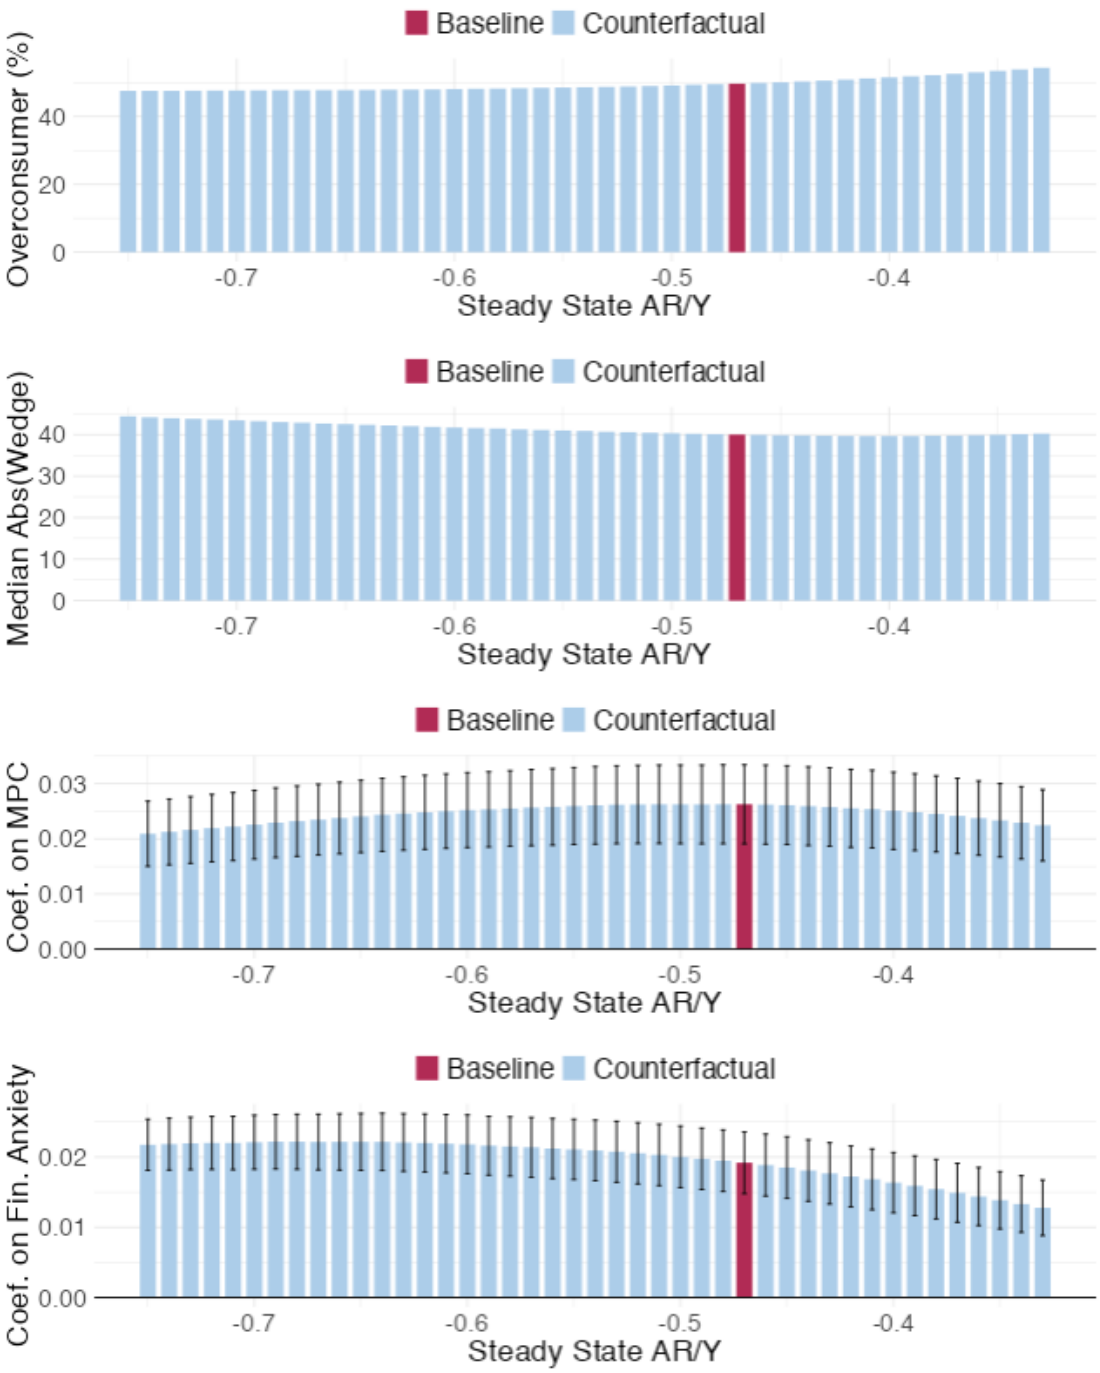
Notes: The figure presents the sensitivity of four estimated results to our assumed value of steady state nondurable C/Y in the wedge calculation, with the assumed value ranging from 55% to 75% in increments of 1 percentage point (our baseline assumption is 67%). These results include the percent of users with a positive wedge (Panel A), the median absolute value wedge (Panel B), and the estimated coefficients from separate regressions of nondurables MPC and an indicator for financial anxiety on the consumption wedge (Panels C and D, respectively). In the regressions, we cluster standard errors at the user level. We hold all other parameters constant at our baseline values (for the steady state parameters, these are rounded to the nearest percentage point). Includes wave 1 respondents who meet the sample restrictions outlined in Appendix B. Includes monthly observations from October 2021 to September 2022.

Figure D.4. Sensitivity of Wedges to Nondurables Share of Spending



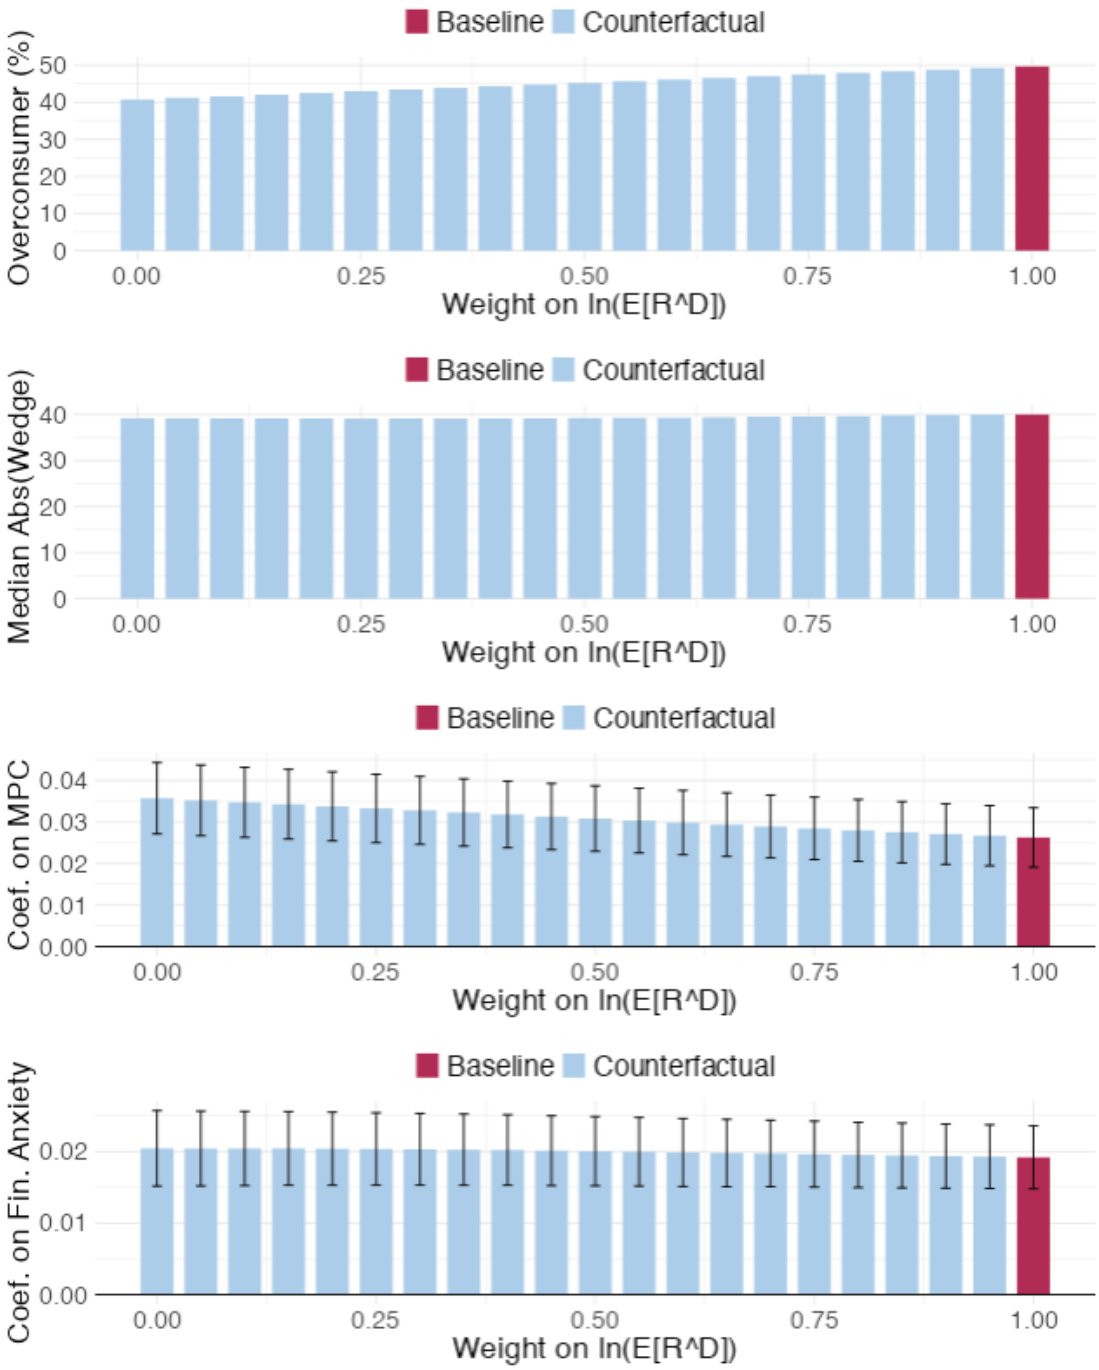
Notes: The figure presents the sensitivity of four estimated results to our assumed nondurables share of spending in the wedge calculation, with the assumed value ranging from 68% to 90% in increments of 1 percentage point (our baseline assumption is 79%). These results include the percent of users with a positive wedge (Panel A), the median absolute value wedge (Panel B), and the estimated coefficients from separate regressions of nondurables MPC and an indicator for financial anxiety on the consumption wedge (Panels C and D, respectively). In the regressions, we cluster standard errors at the user level. We hold all other parameters constant at our baseline values (for the steady state parameters, these are rounded to the nearest percentage point). Includes wave 1 respondents who meet the sample restrictions outlined in Appendix B. Includes monthly observations from October 2021 to September 2022.

Figure D.5. Sensitivity of Wedges to Steady State AR/Y



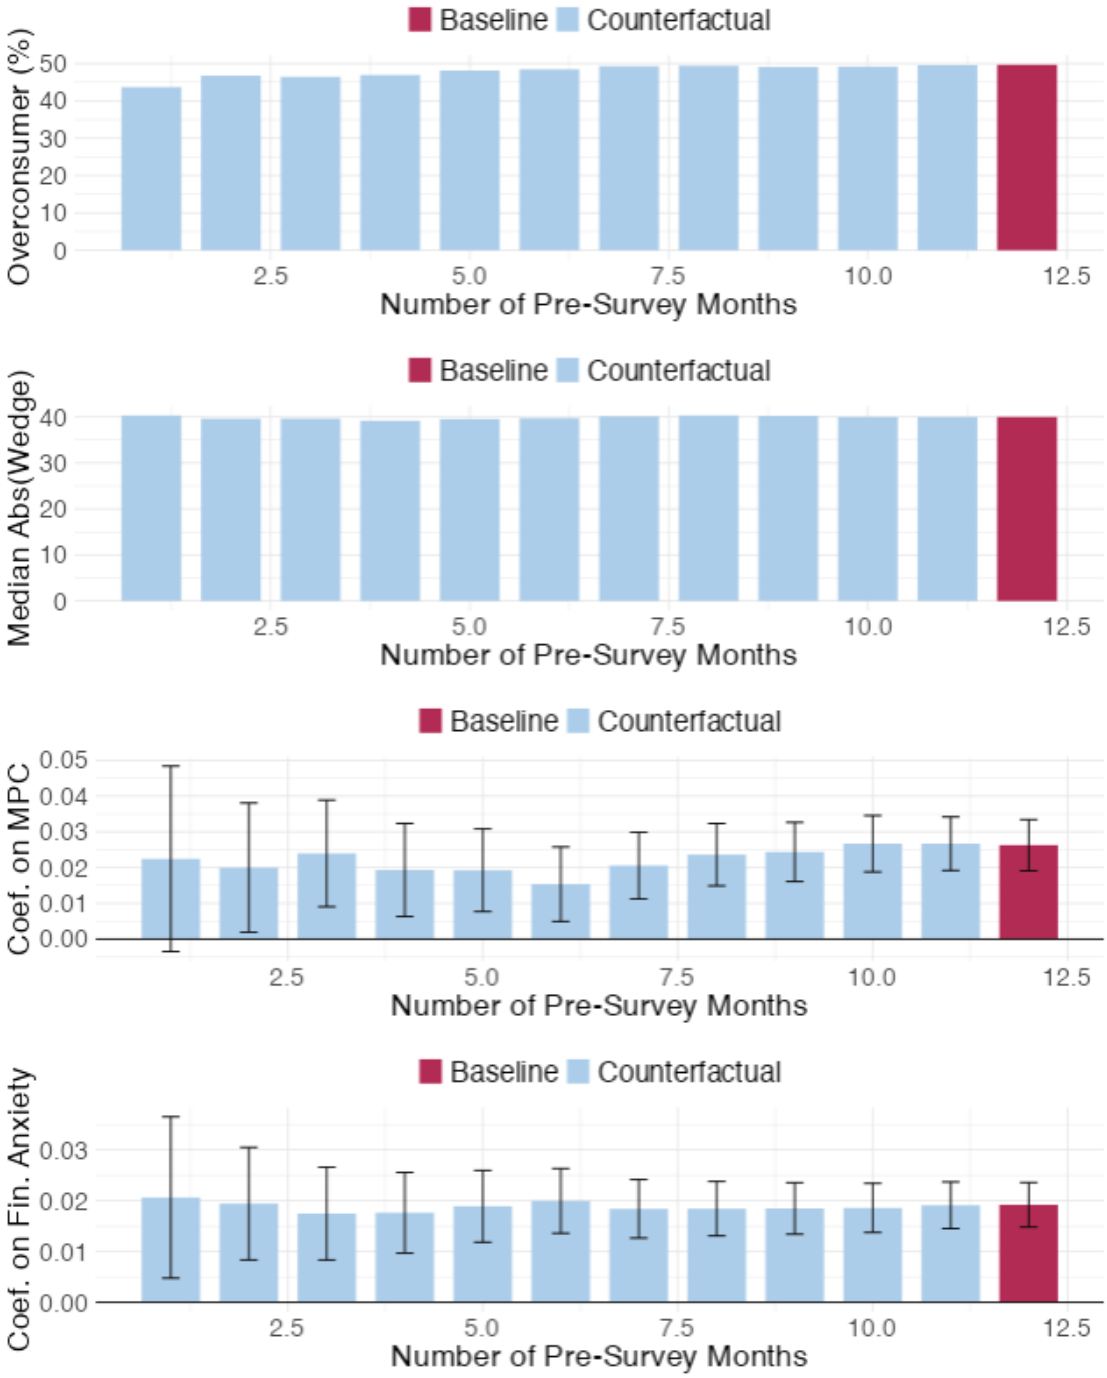
Notes: The figure presents the sensitivity of four estimated results to our assumed value of steady state AR/Y in the wedge calculation, with the assumed value ranging from -75% to -33% in increments of 1 percentage point (our baseline assumption is -47%). These results include the percent of users with a positive wedge (Panel A), the median absolute value wedge (Panel B), and the estimated coefficients from separate regressions of nondurables MPC and an indicator for financial anxiety on the consumption wedge (Panels C and D, respectively). In the regressions, we cluster standard errors at the user level. We hold all other parameters constant at our baseline values (for the steady state parameters, these are rounded to the nearest percentage point). Includes wave 1 respondents who meet the sample restrictions outlined in Appendix B. Includes monthly observations from October 2021 to September 2022.

Figure D.6. Sensitivity of Wedges to Weight on $\ln E(R^D)$



Notes: The figure presents the sensitivity of four estimated results to our assumed weight toward $\ln E[R^D]$ (vs. $\ln E[R^S]$) in the wedge calculation, with the assumed value ranging from 0% to 100% in increments of 5 percentage points (our baseline assumption is 100%). These results include the percent of users with a positive wedge (Panel A), the median absolute value wedge (Panel B), and the estimated coefficients from separate regressions of nondurables MPC and an indicator for financial anxiety on the consumption wedge (Panels C and D, respectively). In the regressions, we cluster standard errors at the user level. We hold all other parameters constant at our baseline values (for the steady state parameters, these are rounded to the nearest percentage point). Includes wave 1 respondents who meet the sample restrictions outlined in Appendix B. Includes monthly observations from October 2021 to September 2022.

Figure D.7. Sensitivity of Wedges to Number of Pre-Survey Months



Notes: The figure presents the sensitivity of four estimated results to the number of pre-survey months of data we include in the wedge calculation, with the number ranging from 1 to 12 in increments of 1 month (our baseline number of months is 12, spanning October 2021 to September 2022). These results include the percent of users with a positive wedge (Panel A), the median absolute value wedge (Panel B), and the estimated coefficients from separate regressions of nondurables MPC and an indicator for financial anxiety on the consumption wedge (Panels C and D, respectively). In the regressions, we cluster standard errors at the user level. We hold all other parameters constant at our baseline values (for the steady state parameters, these are rounded to the nearest percentage point). Includes wave 1 respondents who meet the sample restrictions outlined in Appendix B.

Figure D.8. Sensitivity of Consumption Wedge to Within-User Median

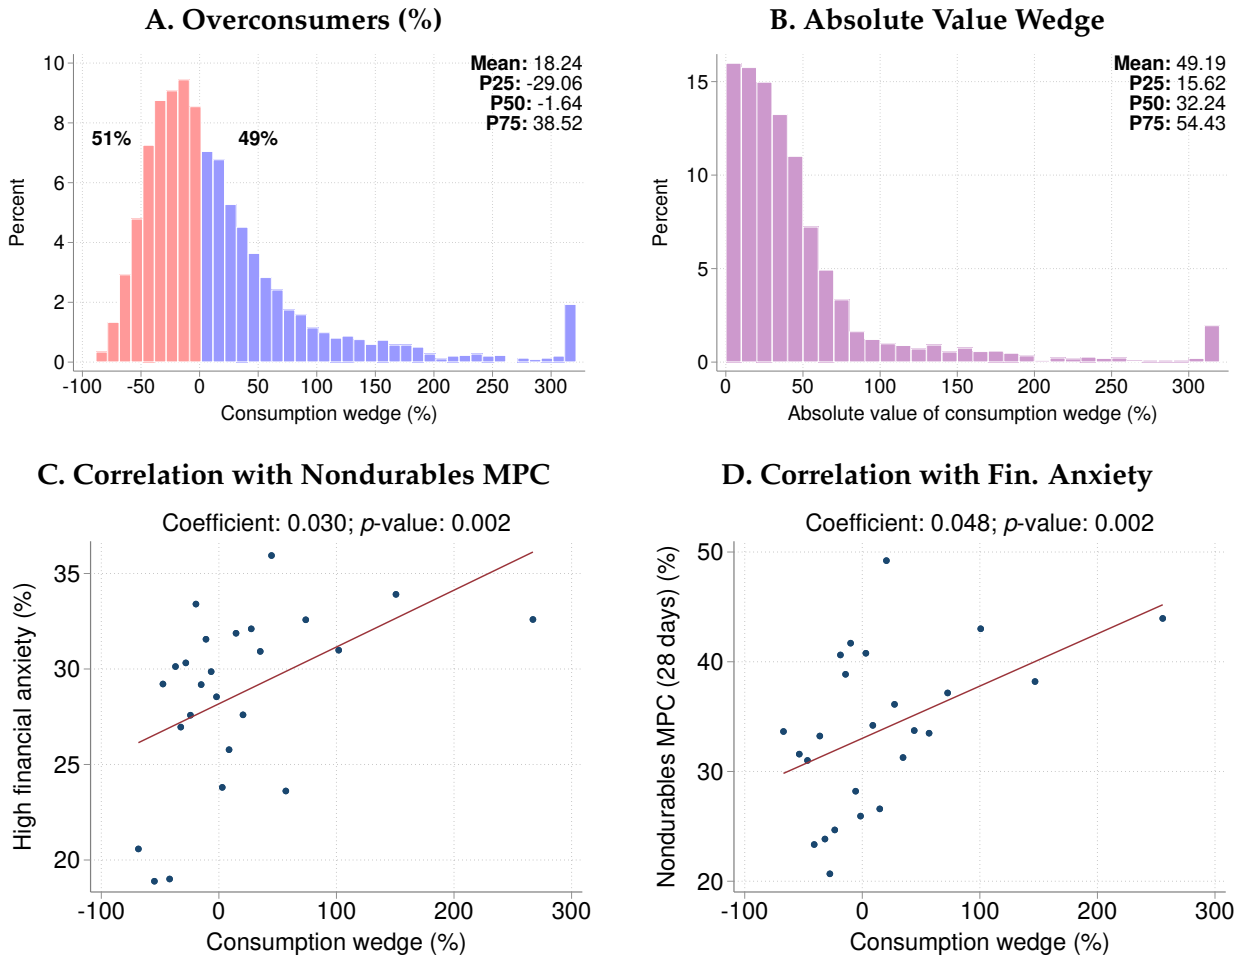


Figure presents the sensitivity of our main wedge results to taking the median wedge within-user across the 12 pre-survey months (October 2021 to September 2022). These results include the percent of users with a positive wedge (Panel A), the median absolute value wedge (Panel B), and the estimated coefficients from separate regressions of nondurables MPC and an indicator for financial anxiety on the consumption wedge (Panels C and D, respectively). Includes wave 1 respondents who meet the sample restrictions outlined in Appendix B.

Figure D.9. Sensitivity of Consumption Wedge to Dropping Users

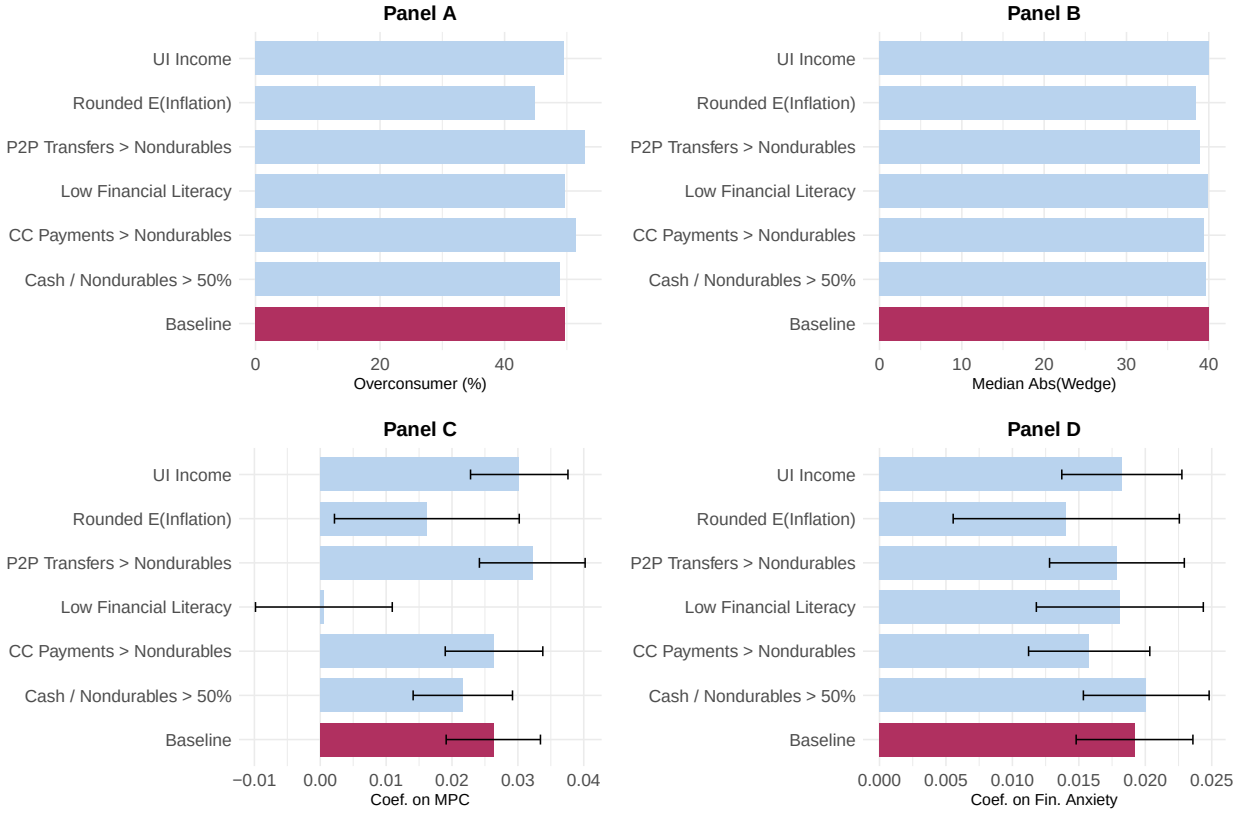


Figure presents the sensitivity of our main wedge results to dropping users with high measurement error. These results include the percent of users with a positive wedge (Panel A), the median absolute value wedge (Panel B), and the estimated coefficients from separate regressions of nondurables MPC and an indicator for financial anxiety on the consumption wedge (Panels C and D, respectively). In the regressions, we cluster standard errors at the user level. For reference, our baseline results are shown in red. Baseline sample includes wave 1 respondents who meet the sample restrictions outlined in Appendix B. Includes monthly observations from October 2021 to September 2022.

Figure D.10. Robustness of Consumption Wedges to Clustering Users

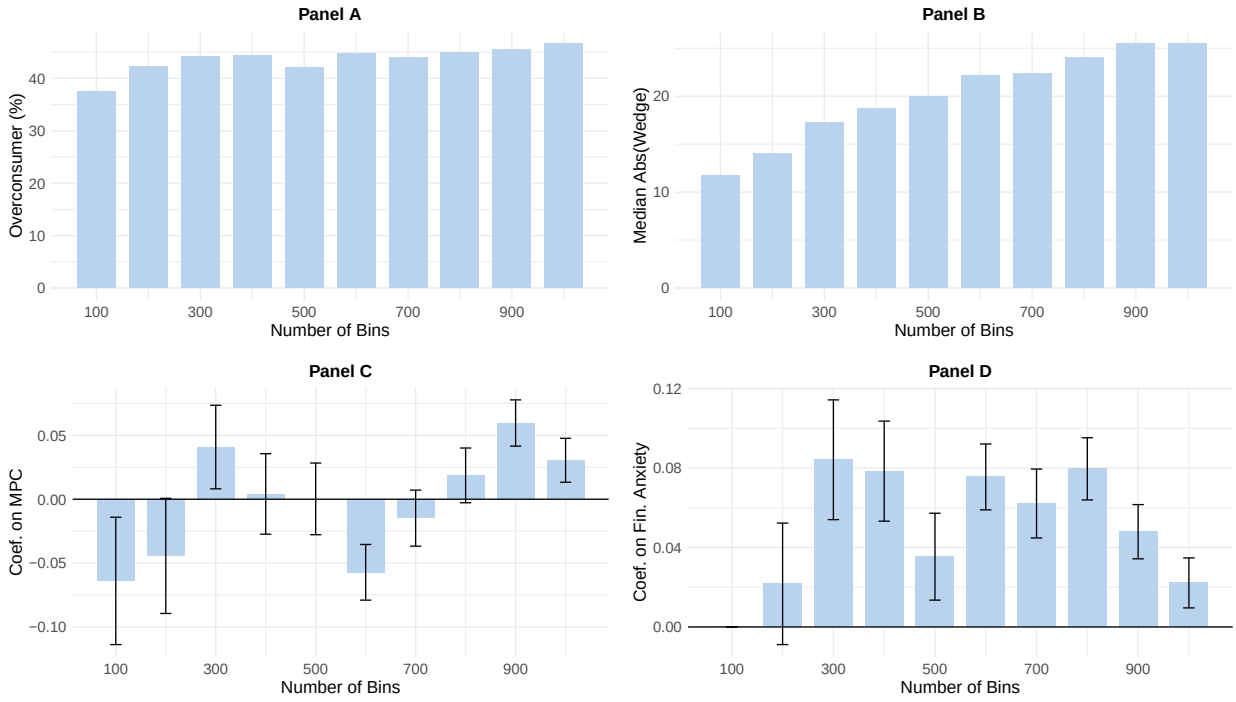


Figure presents the sensitivity of our main wedge results to clustering users and taking the within-cluster median of each wedge input. These results include the percent of users with a positive wedge (Panel A), the median absolute value wedge (Panel B), and the estimated coefficients from separate regressions of nondurables MPC and an indicator for financial anxiety on the consumption wedge (Panels C and D, respectively). In the regressions, we cluster standard errors at the user level. Includes wave 1 respondents who meet the sample restrictions outlined in Appendix B. Includes monthly observations from October 2021 to September 2022.

Figure D.11. Sensitivity of Consumption Wedge to Adding Random Noise

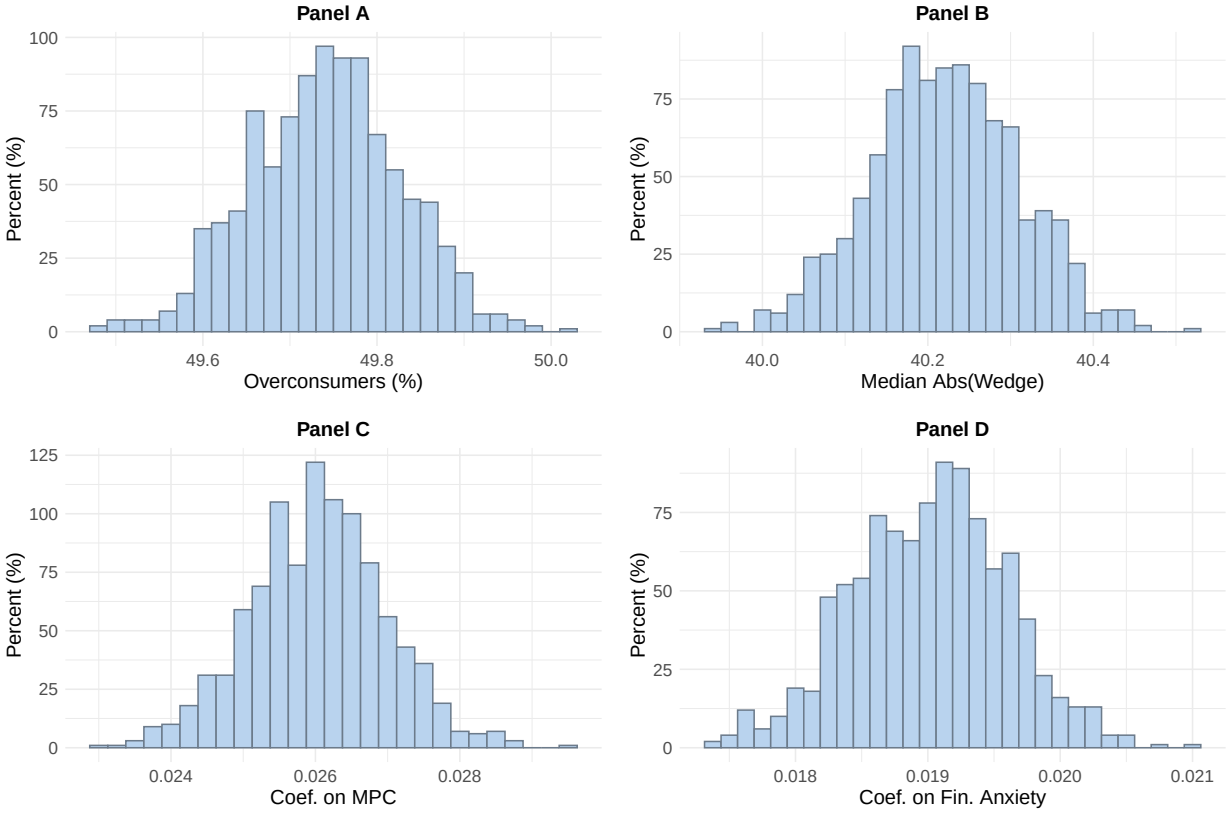


Figure presents the sensitivity of our main wedge results to adding random noise with a mean of 0pp and standard deviation of 1.5pp to the wedge inputs. These results include the percent of users with a positive wedge (Panel A), the median absolute value wedge (Panel B), and the estimated coefficients from separate regressions of nondurables MPC and an indicator for financial anxiety on the consumption wedge (Panels C and D, respectively). Each panel represents a histogram of the estimated parameter from a Monte Carlo simulation with 1,000 iterations. Includes wave 1 respondents who meet the sample restrictions outlined in Appendix B. Includes monthly observations from October 2021 to September 2022.

Figure D.12. Consumption Wedges from the Second Wave Survey

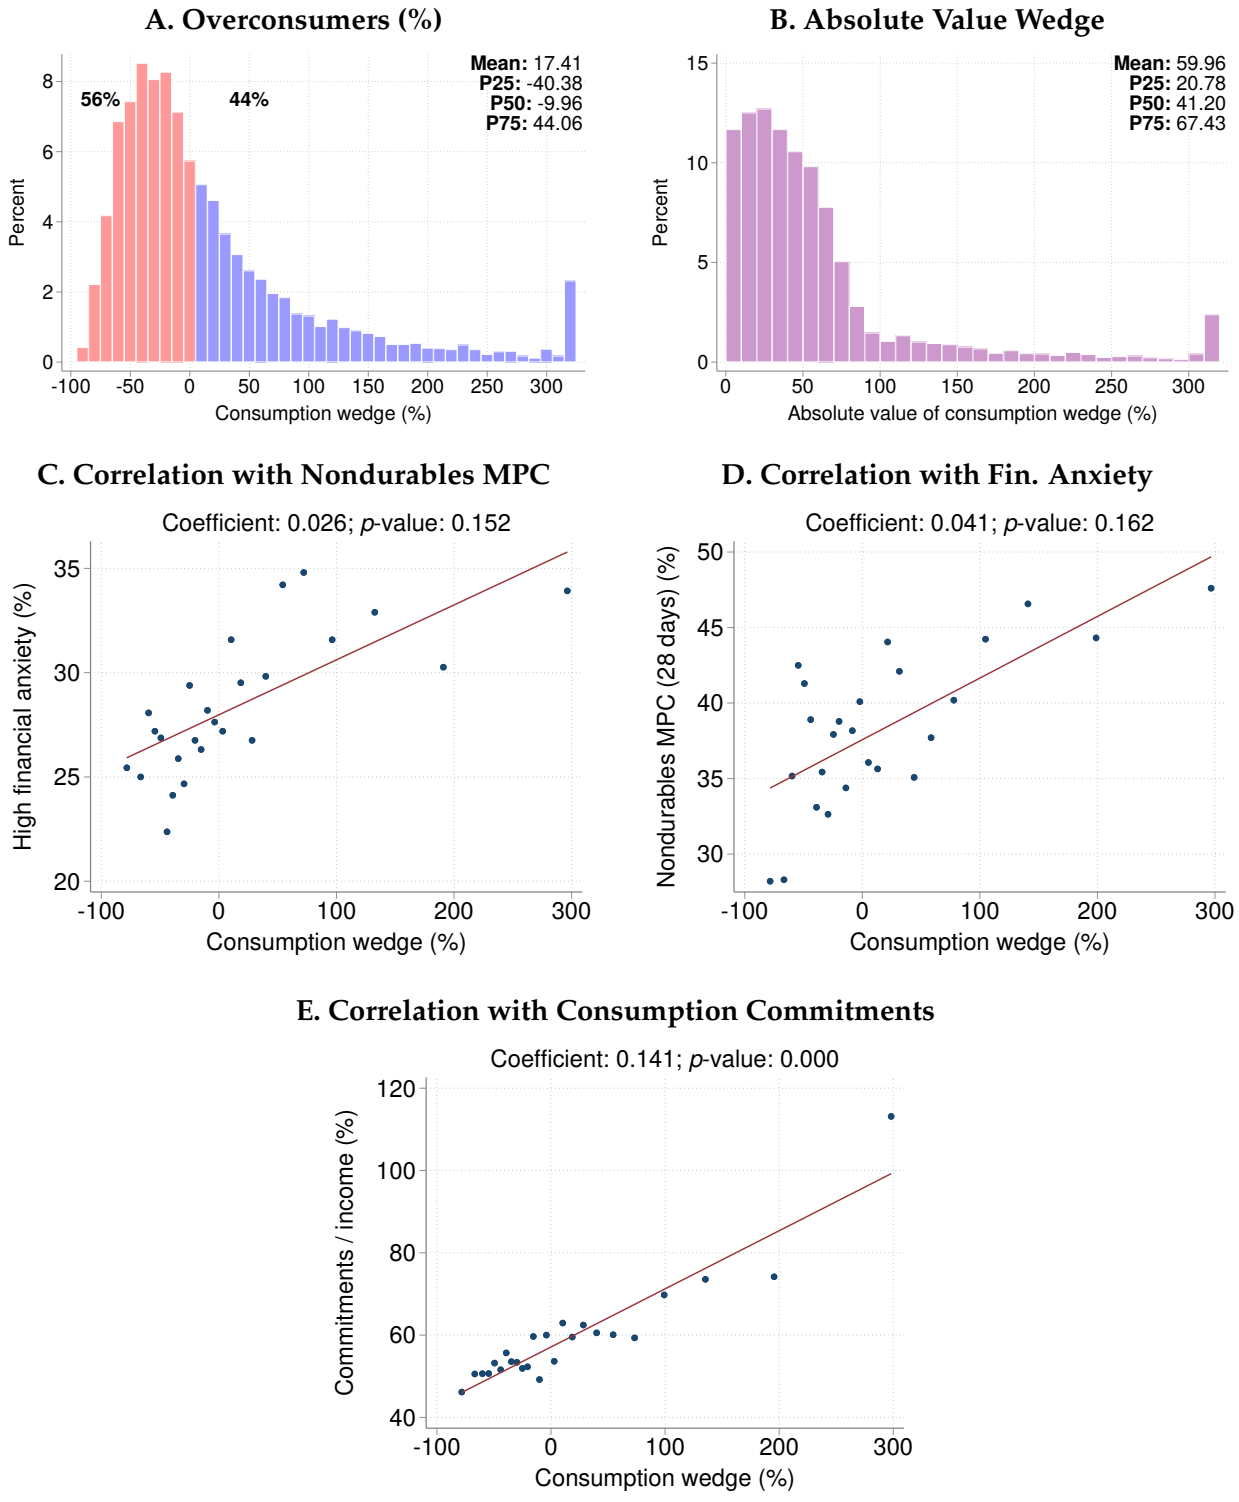


Figure presents the sensitivity of our main wedge results to using data from the second wave survey. These results include the percent of users with a positive wedge (Panel A), the median absolute value wedge (Panel B), and the estimated coefficients from separate regressions of nondurables MPC and an indicator for financial anxiety on the consumption wedge (Panels C and D, respectively). In the regressions, we cluster standard errors at the user level. Includes wave 2 respondents who meet the sample restrictions outlined in Appendix B. Includes monthly observations from July 2023 to June 2024.

Figure D.13. Comparison of First- and Second-Wave User-Level Consumption Wedges

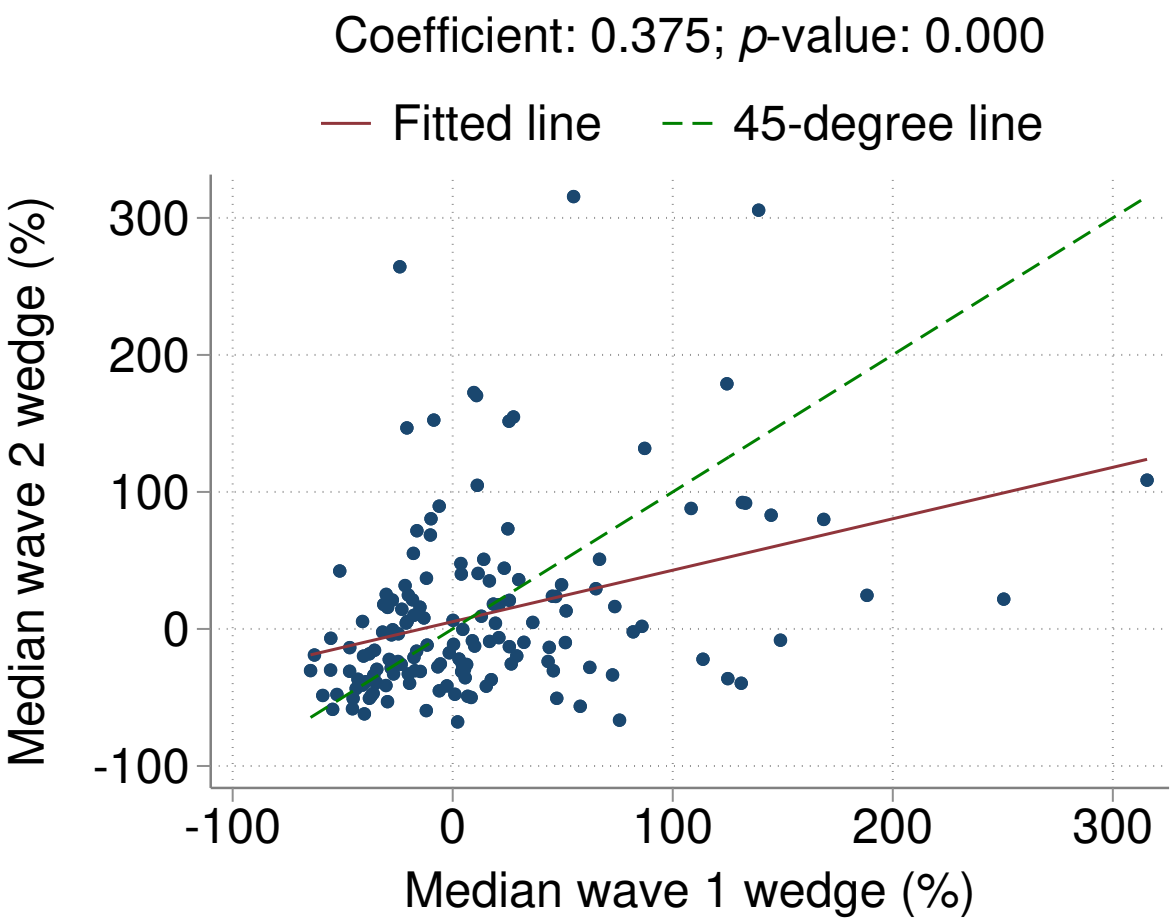


Figure presents the correlation between user-level wedges across survey waves. Includes wave 1 and wave 2 respondents who meet the sample restrictions outlined in Appendix B. Wave 1 wedges include data from October 2021 to September 2022, and wave 2 wedges include data from July 2023 to June 2024. Left column presents a binscatter with 25 quantile bins. Right column presents a scatterplot with all observations, overlaid with a red fitted line and a green 45-degree line.

Figure D.14. Consumption Wedge with Observed Gamma

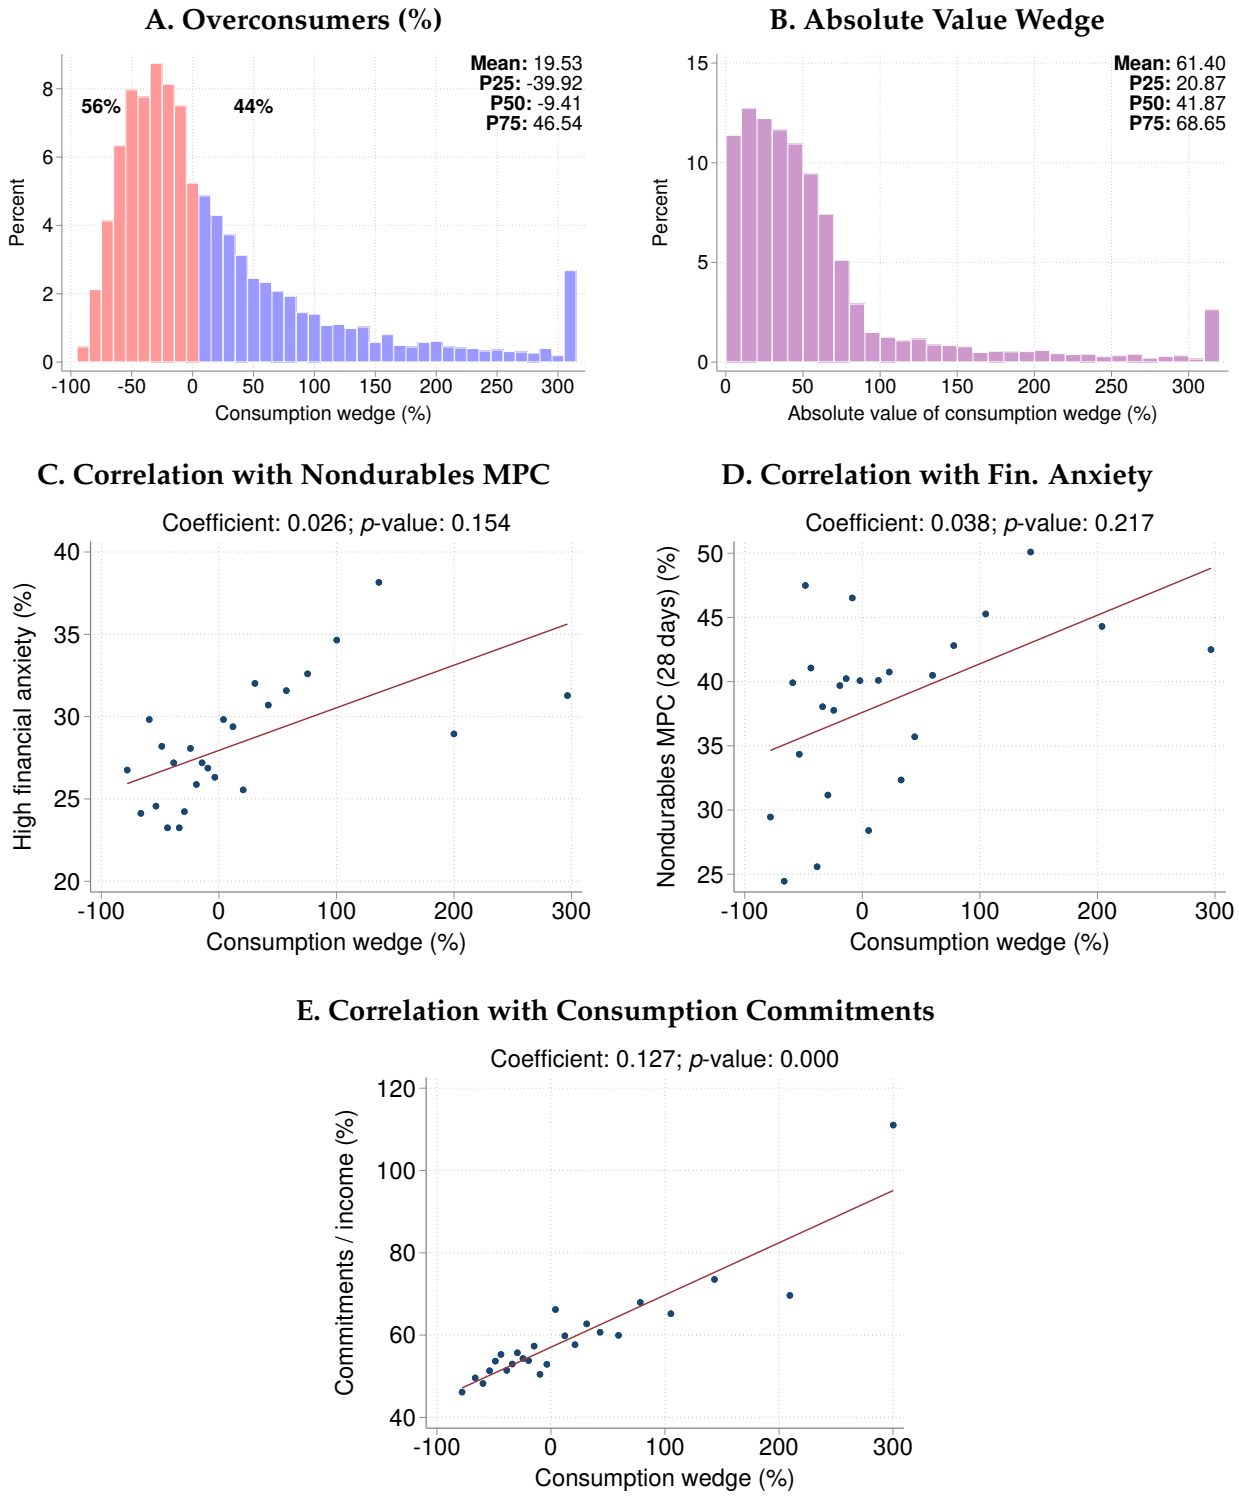


Figure presents the sensitivity of our main wedge results to incorporating observed user-specific gamma values from the second wave survey. These results include the percent of users with a positive wedge (Panel A), the median absolute value wedge (Panel B), and the estimated coefficients from separate regressions of nondurables MPC and an indicator for financial anxiety on the consumption wedge (Panels C and D, respectively). In the regressions, we cluster standard errors at the user level. Includes wave 2 respondents who meet the sample restrictions outlined in Appendix B. Includes monthly observations from July 2023 to June 2024.