

Mental Accounting Through Two-stage Budgeting Under Bounded Rationality

Nick Pretnar^{12*} Christopher Y. Olivola^{3†} Alan Montgomery^{4‡}

¹*Laboratory for Aggregate Economics and Finance, UCSB*

²*University of Austin Texas (UATX)*

³*Kelley School of Business, Indiana University*

⁴*Tepper School of Business, Carnegie Mellon University*

June 7, 2026[§]

Abstract

We propose a model that generalizes two-stage budgeting to incorporate frictions which induce budgetary stickiness and mental accounting behavior. Inattention limits the fungibility of money across mental accounts, leading to boundedly rational behavior. Financial attentiveness is cognitively costly, so consumers might re-assess only a subset of spending budgets every period. Over- or under-spending affects future budgeting and expenditure decisions. We discipline the model with agent-level weekly expenditure data from low-income consumers and explore its quantitative implications for the degree to which budgeting behavior reflects bounded rationality, as well as the potential consequences of interventions that encourage budgeting discipline. Attentiveness constraints can have a disciplinary effect on finances: under the model's assumptions, relaxing constraints leads some consumers to over-spend and become financially worse off.

Keywords: budgeting, mental accounting, bounded rationality, expenditure, savings

JEL Classification: D11, D90, Z19

*UCSB, LAEF; 2112 North Hall, Santa Barbara, CA 93106-9215; npretnar@ucsb.edu.

†IU, Kelley; 801 W. Michigan Street, Indianapolis, IN 46202; colivola@iu.edu.

‡CMU, Tepper; 4765 Forbes Ave., Pittsburgh, PA 15213; alanmontgomery@cmu.edu.

[§]Nick acknowledges support from the National Science Foundation Graduate Research Fellowship under Grant No. DGE1252522. We would like to thank the PNC Center for Financial Services for funding. We would also like to thank Laurence Ales, Manel Baucells, William Bednar, Siddhartha Chib, Tony Cookson, Erik Eyster, Finn Kydland, Xiao Liu, Ryan Oprea, Fernando Payro, Ali Shourideh, Stephen Spear, and Kannan Srinivasan for valuable insights, along with anonymous referees. We also thank attendees of the 2017 INFORMS Marketing Science Conference in Los Angeles, the 2019 Association for Consumer Research Conference in Atlanta, the Certified Financial Planner Board of Standards' 2020 Academic Research Colloquium in Washington D.C., the 2020 NBER-NSF SBIES Virtual Conference, the 2020 INFER AC Virtual Conference, the 2022 FUR Conference in Ghent, Belgium, and seminars at the University of California, Santa Barbara, Carnegie Mellon University, Tepper School of Business, and Columbia Business School.

1 Introduction

Budgeting is often touted as a way for consumers to discipline their expenditure decisions. However, there is still much we do not know about how consumers use or adhere to budgets in their everyday lives. Mental accounting theory, in part, posits that consumers use budgets to discipline their spending by tagging money for specific purposes (Thaler, 1999). Meanwhile, in classical expenditure models money is assumed perfectly fungible and freely transferable (Deaton and Muellbauer, 1980a). Our goal is to build a general model of spending and budgeting behavior that allows for individual agents to exhibit varying degrees of bounded rationality. We then take this micro-founded model to agent-level expenditure data to explore the degree to which observed expenditure patterns are consistent with mental accounting behavior under the model’s assumptions.

The model implies rich heterogeneity with respect to the degree to which consumers’ budgeting appears boundedly rational. We define a pure mental accounter as one who does not rebalance their budgets but abides by the dictates of their initial allocation. The model’s quantitative implications suggest that consumer decision-making, within its framework, is neither fully rational nor fully irrational, but rather falls along a rationality continuum as previously documented in Stango and Zinman (2023).¹ In our set-up pure mental accounting and perfect fungibility are opposing, extreme cases of a fungibility continuum. Our approach nests multiple features from the behavioral economics literature such as: 1) fungibility frictions that result from timing constraints (i.e., ex-ante budgets are only equal to ex-post expenditure on a measure-zero set of expenditure/preference shock realizations); 2) narrow choice bracketing as in Gabaix (2014) and Lian (2021) (i.e., consumers may update some budgets for certain expenditure categories in one period while leaving other budgets unchanged); 3) numeracy constraints with respect to setting optimal budgets (i.e., consumers may only update a budget if the optimal candidate budget is significantly different from the underlying budget with which they enter the period); 4) mental accounting (i.e., future budgets may be informed by past behavior with respect to over- and under-expenditure relative to past budgets).

To explore the quantitative relevance of each behavioral feature, we discipline the model with weekly, agent-level expenditure data for a set of low-income consumers. Our quantitative strategy is designed to answer several questions through the model’s lens. First, are observed expenditure patterns consistent with consumers using budgets to discipline and constrain their expenditure decisions? If so, how do the ways in which people use budgets vary? Do consumers regularly revise all of their budgets or only a subset? Finally, what are the implications for both welfare and savings rates when such rigidities are either enhanced or relaxed?

Our approach is best understood as quantitative theory: we build a model rich enough to capture multiple behavioral mechanisms, discipline it with observed expenditure data, and then interrogate the model’s quantitative implications for latent budgeting decisions. We show that we

¹Our findings jibe with recent work in Alos-Ferrer and Garagnani (2023) who show that some consumers are rational belief updaters and some are not.

can invert the model’s first-order conditions in order to tractably recover latent budgeting decision variables conditional upon observed expenditure data and the theoretical structure of the model. One major technical contribution we provide is the tractable (but probabilistic) inference of latent behavioral objects underlying observed expenditure. The results are model-dependent, but applying structured quantitative models to high-frequency field data represents a natural complement to the theory and lab experiments that have thus far dominated the study of consumer budgeting behavior.

Under our preferred parameterization, the model implies that consumers re-evaluate budgets for approximately half of their consumption categories in any given period. However, there is substantial heterogeneity across consumers with respect to how they use budgets and how they respond to expenditure/preference shocks. Under this parameterization, a plurality of consumers use mental accounting and sticky budgets to discipline their total spending and thus savings behavior. A substantial subset of these consumers exhibit a fascinating, albeit financially perplexing, pattern of behavior: they maintain spending discipline by updating their budgets *infrequently* due to the cognitive costs of doing so. The model predicts that for these consumers, lowering the cognitive costs of updating their budgets may push them to raise their budgets to avoid the disutility of over-spending, yielding a repetitive budgeting/spending cycle that ultimately leads to debt accumulation and, in extreme cases, bankruptcy.

Our methodological approach treats the model as an experiment: we construct a theoretical framework specifically designed to be taken to data, so that once its parameters are disciplined by observed expenditure, the parameterized model can serve as a laboratory for quantitative thought experiments. This design imperative shapes the model’s construction — it must be rich enough to nest multiple behavioral mechanisms while remaining tractable enough for estimation on agent-level data. A critical requirement is that the framework accommodate individual heterogeneity, given extensive evidence that consumers exhibit varying degrees of rationality and behavioral tendencies (Choi et al., 2014; Stango and Zinman, 2023). As with all structural models, the results are conditional on the maintained assumptions. We view the model not as a literal description of consumer cognition, but as a disciplined framework for asking: *if* consumers budget in a manner consistent with our theory, *what would the data-consistent parameterization imply* about the degree and heterogeneity of bounded rationality, and about the consequences of policy interventions?

The paper proceeds as follows. In §2 we position our work in the broader literature of two-stage budgeting, bounded rationality, and mental accounting. §3 builds a model of two-stage budgeting that incorporates features from behavioral economics. §4 describes our unique dataset of consumer-level weekly spending and income. This dataset motivates our underlying model construction from §3 since we want to allow for endogenous variation and heterogeneity in observed weekly spending patterns at the consumer level. §5 estimates the model and discusses the results of our estimation and their implications for inference with regards to mental accounting theory, budgeting, and cognitive/financial attentiveness, including as pertaining to numeracy constraints. We conclude with a discussion in §6 of the model’s implications and the limitations

of this research.

2 Motivating Literature

2.1 Two-stage Budgeting and Classical Demand

The theory of two-stage budgeting, which posits that consumers first allocate expenditure shares to broad commodity categories prior to making individual spending decisions, provides micro-foundational justification for estimating demand systems derived from separable utility models. The idea behind two-stage budgeting can be traced back to theories of aggregation described independently in [Hicks \(1936\)](#) and [Leontief \(1936\)](#). Each author shows that if the marginal rate of substitution for two goods in commodity group j is independent of the marginal rate of substitution for two goods in group k , then preferences are at least weakly separable over those broad commodity groups.² Building on the two-stage budgeting framework, [Deaton and Muellbauer \(1980a\)](#) argue that if prices for commodities within a group move in parallel then weak separability can be assumed. The consumption decision process can then be framed around these broad aggregates, where consumers are assumed to allocate resources for consumption expenditure not by optimally responding to posted prices of individual commodities, but only by considering price levels for broad groups of them. Indeed, tractable demand systems, like the linear expenditure systems of [Geary \(1950\)](#), [Stone \(1954\)](#), and [Houthakker \(1960\)](#), the Rotterdam models of [Barten \(1964\)](#), [Theil \(1965\)](#), [Barten \(1967\)](#), [Theil \(1976\)](#), and [Barten \(1977\)](#), or the almost ideal demand system of [Deaton and Muellbauer \(1980b\)](#), are generated from a utility function with some degree of separability, exploiting these aggregation results.

Such models implicitly rely on two strong assumptions about consumer behavior. First, consumers are assumed to exhibit perfect fungibility; that is, they can freely transfer money that has been allocated to different commodity-group budgets. Second, first-stage budgets and second-stage spending are assumed to always equate, so that there are no timing frictions whereby consumers engage in a level of spending that violates their ex-ante plans. Our contribution is to consider both theoretical and quantitative results from a model environment in which these assumptions do not necessarily hold. Hence, our framework relaxes the traditional two-stage budgeting assumptions, where these relaxations can be interpreted as allowing for mental accounting behavior.

2.2 Bounded Rationality and Rational Inattention

Among our additions to the two-stage budgeting framework, we allow for heterogeneous sparse maximization, similar to the setup in [Gabaix \(2014\)](#). This places our research into conversation with contemporary economics research, popularized by [Sims \(2003\)](#) and thoroughly reviewed

²[Gorman \(1959\)](#) (later extended by [Blackorby and Russell \(1997\)](#)) recognized the importance of this result in order to generate preference structures that aggregate.

in [Maćkowiak et al. \(2023\)](#), where consumers make decisions under limited information or with limited cognitive resources to optimize. This feature of our framework follows from the well-established finding in cognitive psychology that people can only attend to, and process, a limited amount of information at any given time ([Miller, 1956](#); [Simon, 1957](#); [Cowan, 2000](#)). In our sparse-max framework, consumers might only re-evaluate a subset of their first-stage budgets in any given period, thereby engaging in narrow choice bracketing ([Kahneman and Lovallo, 1993](#); [Read et al., 1999](#); [Rabin and Weizsäcker, 2009](#); [Felső and Soetevent, 2014](#); [Koch and Nafziger, 2016, 2019, 2020, 2021](#)). By allowing for heterogeneous rational inattention and heterogeneous non-fungibility in this way, our approach reflects a broad literature which suggests there is no reason to assume that all consumers will regard all attributes of a decision problem as uniformly salient to the same degree ([Chetty et al., 2009](#); [Sussman and Olivola, 2011](#); [Kőszegi and Szeidl, 2013](#); [Bordalo et al., 2014](#); [Schwartzstein, 2014](#); [Caplin and Dean, 2015](#); [Abeler and Marklein, 2017](#); [Read et al., 2017](#); [Bordalo et al., 2020](#); [Kőszegi and Matějka, 2020](#); [Bordalo et al., 2022](#)). Under such a structure, we will use our model to show that, when disciplined by data, the framework implies that different consumers appear rationally constrained to varying degrees over different time intervals.³

It is important to note that we are not the first to consider how choice bracketing, budgeting, and substitutability are linked. [Kőszegi and Matějka \(2020\)](#) consider a similar, theoretical environment to ours, where high attention costs lead to essentially fixed budgets. A key result in their model predicts that broad budgeting behavior depends on the underlying substitutability of the commodities, so that highly-substitutable commodities are more likely to be budgeted together. This is an inherent result of a consumer’s inattentiveness: they lack the intrinsic ability to think about how to optimally distribute resources amongst commodities each of which is substitutable with another to varying degrees. Thus, the model in [Kőszegi and Matějka \(2020\)](#), using behavioral insights, yields the classical aggregation result in [Hicks \(1936\)](#), [Leontief \(1936\)](#), and [Deaton and Muellbauer \(1980a\)](#) — that separability is supported if the elemental commodities that comprise broad commodity-groupings are highly substitutable.

2.3 Mental Accounting and Budgeting Commitment

Our approach to modeling mental accounting builds on, but also departs from, prior work. On one hand, we motivate our model by including budgeting behavior consistent with some aspects of the planner/doer models of [Shefrin and Thaler \(1981\)](#) and [Shefrin \(2020\)](#), as well as the dual-self, budget commitment model of [Galperti \(2019\)](#) and the multi-self, mental accounting model of [Lian \(2021\)](#). We depart from much of the existing mental accounting literature in that we seek to explain behavioral tendencies not as resulting from reference-dependent utility or price adaptation, such as the model in [Baucells and Hwang \(2017\)](#), but rather generated by bounds on cognitive atten-

³As an aside, we should note that we are not the only recent working paper to find heterogeneity in the degree to which consumers exhibit rationality in their consumption decisions. We will later show that the rationality heterogeneity implied by our model is broadly consistent with recent findings on consumption responsiveness discussed in [Andre et al. \(2025\)](#) and the theoretical prediction in [Lian \(2023\)](#) that mistakes in future consumption planning generate high marginal propensities to consume.

tion and numeracy. Thus, our approach is most similar to [Lian \(2021\)](#) who considers how narrow thinking can lead to outcomes that resemble narrow choice bracketing, though we build our model with an explicit aim toward estimation on high-frequency expenditure data as opposed to theoretical exposition. In our framework a boundedly rational consumer may only have the cognitive attention to re-evaluate and, if needed, update a subset of their first-stage expenditure budgets in a given period. In our model, mental accounting is just a book-keeping mechanism which may inform budget responsiveness: over-spending in period t could lead to reduced or increased budgets in period $t + 1$ depending both on the consumer’s disutility of over-spending (via reduced available liquidity) and their marginal propensity to save. The so-called ex-ante planner may take into consideration relative over- or under-expenditure in multiple different commodity groups when making a narrow decision regarding their budget for one particular category, only. In this way, we introduce the accounting mechanism from [Thaler \(1985\)](#) that dynamically keeps track of over- or under-spending, into a model environment where the breadth of choice bracketing determines the degree to which over- or under-spending actually affects future budgets.

This interaction governing the degree to which rationality (i.e., budget attentiveness) is bounded and over- or under-spending impact future budgets is what we interpret as mental accounting. As [Farhi and Gabaix \(2020\)](#) point out, however, the literature contains no widely agreed-upon definition as to what exactly constitutes mental accounting. For example, Thaler uses the term to describe different behaviors, such as keeping track both of spending on certain consumption items and, separately, spending using specific (different) liquidity sources ([Thaler, 1985](#); [Thaler et al., 1997](#)). In this paper, we abstract from method of payment mental accounting dependencies as in [Feinberg \(1986\)](#), [Prelec and Loewenstein \(1998\)](#), and [Mullainathan \(2002\)](#). This is not a problematic abstraction for us because the data we use to discipline the model are from pre-paid debit card users who are underbanked and most likely to use their card as their dominant liquidity source. Since low-income consumers are more likely to be liquidity-constrained and liquidity-constrained households are budget-sensitive and often engage in non-fungible spending behavior for basic necessities like food ([Gelman et al., 2014](#); [Hastings and Shapiro, 2018](#)), studying the spending patterns of low-income households is a natural application for a model with endogenous mental accounting and budgeting features. Our quantitative approach follows after recent theoretical work in [Kőszegi and Matějka \(2020\)](#), who show that category-specific budgets are useful for consumers with high attention costs, many of whom are more likely to be low-income ([Schilbach et al., 2016](#)). Given our field data, we thus contribute to the mental accounting and budgeting literature by examining the degree to which low-income consumers, via their spending behavior, appear to exhibit traits consistent with contemporary behavioral theories at the frontier of present research.

3 Model

We construct the model in order that it may be tractably estimated with our dataset featuring only expenditure data,⁴ while still featuring as many behavioral features as we can possibly allow for. Specifically, in our data we observe individual, weekly-level expenditures from which we want to infer latent budgets. Since latent inference of underlying decisions is required, this necessitates a highly stylized model that can easily be mapped into a latent-inference algorithm to probabilistically recover latent budgeting choices on which expenditure observations depend.

Time is discrete and indexed by t . Column vectors and matrices are denoted with a bold typeface. All variables and functions presented, except market prices p_t and gross interest rates r_t , are agent-specific, with consumer units indexed by i . Consumers make decisions in a sequentially dynamic environment, choosing savings so as to satisfy utility over liquidity holdings.⁵ A single period's decision process is divided into two sub-stages. At the beginning of each period consumers choose/plan their budgets subject to some restrictions to maximize their utility and end the period by realizing unanticipated random shocks to preferences and prices on which their actual expenditure depends.

In this sense our decision-theoretic setup has a planner/doer flavor a la [Shefrin and Thaler \(1981\)](#). At the beginning of each period in the first stage of the decision problem, the planner chooses optimal budgets subject to various cognitive constraints, which we will describe below; meanwhile, in the second stage, the doer reacts to some shock process that causes prices and preferences to deviate from their ex-ante expected values. The doer thus engages in expenditure which deviates from the budget/plan formed by the planner in the first stage: the doer may over- or under-spend, depending on the nature of the shock realization. Then, in the next period when the consumer is again making a planning decision, they do so internalizing the degree to which the doer over- or under-spent in the previous period. The amount of over- or under-spending by consumption category is thus the mental accounting state variable on which budgeting decisions depend.

Behavioral frictions are incorporated into the model along several dimensions. First, consumers are assumed to face cognitive constraints regarding their budget evaluations: that is, in any given period, they may only have the cognitive capacity to re-evaluate some subset of their consumption budgets. Second, consumers imperfectly choose budgets so as to optimally set ex-ante expenditure expectations, though expenditure itself (either due to unanticipated price and/or utility shocks) will only ever exactly adhere to the ex-ante budget in measure-zero outcomes. Third, ex-ante budget re-evaluation is also assumed to be subject to a numeracy constraint: consumers

⁴We describe the dataset in more detail in [§4](#).

⁵Consumers have preferences for carrying forward liquidity (balances) and, via this mechanism, indirectly smooth consumption. In this way consumers do not solve for today's consumption as a function of a forecast of future consumption. This setup is strictly undertaken for tractability purposes: as will become apparent, the balance of liquidity carried forward acts as a sufficient statistic for summarizing the dynamics of latent decisions, which in our setup depend on this balance rather than on agents' rational forecasts of optimal (and to the econometrician, unknown) future consumption.

may not necessarily make the effort to change a budget if the proposed, optimal budget is not significantly different from the budget already in place.

The equilibrium concept we deploy is that of the *personal equilibrium*, where decisions must be consistent with boundedly rational (i.e., constrained) optimization of some decision-theoretic, concave programming problem.⁶ In this sense we do not use the term “equilibrium” to describe a set of statistics that summarize market outcomes such as prices.⁷ Rather, we use the term to describe a set of decision-theoretic outcomes that summarize the deliberations undertaken in some kind of personal decision process where agents are subject to behavioral constraints and behave as if (at times) they are actually multiple, rational agents each engaging in different kinds of decisions. Specifically, our model formulation and the equilibrium concept associated with it will resemble the planner/doer formulation of [Shefrin and Thaler \(1981\)](#) in which a consumer can be thought to consist of two “selves” as in [Galperti \(2019\)](#) and [Lian \(2021\)](#) — 1) a planner who establishes ex-ante budgetary constraints; 2) a doer who engages in consumption expenditure subject to those constraints, which may or may not be associated with some kind of stochastic shock process. The planner can be thought of as forward-looking, forming budgets while thinking about the future, while the doer myopically reacts to various exogenous stimuli leading to expenditure deviations from the planner’s budgets.

We proceed in §3.1 by first describing the decision process of the myopic doer, who essentially just responds to various exogenous stimuli. In this section we also describe ex-post realized consumption preferences and outline how actual liquidity balances evolve. In §3.2 we move on to describe the ingredients of the ex-ante problem faced by the planner. Note that each period the planner acts *before* the doer, anticipating the doer’s actions in that period subject to some set of behavioral constraints. The planner is thus the “first stage” actor in a classic two-stage budgeting model without behavioral frictions, while the doer is the “second stage” spender ([Deaton and Muellbauer, 1980a](#)). Unlike the doer, the planner is not myopic: they make plans in period t by acting on information with regards to the doer’s behavior in period $t - 1$, as summarized by a mental accounting state variable that keeps track of the doer’s over- or under-expenditure relative to last period’s budget. In §3.3 we characterize the decision problem faced by the first-stage planner when engaging in budget updates, including how numeracy constraints affect budgets. Finally, we conclude the model outline in §3.4 where we formally define the *personal mental accounting equilibrium*.

3.1 The Doer

In the spirit of [Shefrin and Thaler \(1981\)](#), the “doer” in our model is the aspect of the consumer that makes consumption expenditure decisions subject to exogenous and random deviations around

⁶The term *personal equilibrium* is commonly deployed in behavioral economics to characterize decision-theoretic outcomes subject to various kinds of rationality constraints (see, for example, [Kőszegi \(2010\)](#) and [Freeman \(2017\)](#), along with discussions in [Thaler \(2016\)](#)).

⁷This is the common usage of the word in aggregate economics, for example.

a given budget. The doer realizes utility from consumption, while the planner forms budgets by trying to maximize expected utility subject to various behavioral constraints.

Utility. Consumers have direct utility defined over a J -dimensional vector of real quantities of consumption q_{it} and unallocated liquidity for period- t of $\bar{z}_{it}$. For tractability purposes we choose a strongly separable utility parameterization:⁸

$$u_i(q_{it}, \bar{z}_{it}) = \sum_{j=1}^J \alpha_{ij} \ln(q_{ijt} + 1) + \alpha_{i,J+1} \ln(\bar{z}_{it}) \quad (1)$$

Flow utility from consumption has a Stone-Geary representation that also allows for zero consumption in a given period (Geary, 1950; Stone, 1954). Further, the consumer has preferences over available liquidity for expenditure, $\bar{z}_{it}$; expenditure decisions today will indirectly affect tomorrow's available liquidity, so that a consumer who spends heavily today trades off consumption utility today for disutility from having less available liquidity in the future.⁹ We assume $\alpha_{ij} \in (0, 1)$ and $\sum_{j=1}^J \alpha_{ij} = 1$. After characterizing the income and savings processes, we will return to the utility function to discuss the role of $\alpha_{i,J+1}$.

Unallocated Liquidity from Borrowing, Savings, and Balances. The sum of marginal period- t savings z_{it} (flow of savings), the consumer's borrowing limit m_i , and available bank balances b_{it} is unallocated liquidity, $\bar{z}_{it} = z_{it} + m_i + r_t b_{it}$. Account balances evolve according to $b_{i,t+1} = r_t b_{it} + z_{it}$. The borrowing limit is such that $m_i > -b_{it}$ always, and returns are non-negative.¹⁰

Income and Savings. Let ℓ_{it} be period- t income and $x_{ijt} = p_{jt} q_{ijt}$ be expenditure on commodity group j . Note that marginal savings is $z_{it} = \ell_{it} - \sum_{j=1}^J x_{ijt}$. Savings can be negative ($z_{it} < 0$) if consumers spend more than they earn in income. After substituting the expression for z_{it} into $\bar{z}_{it}$ and then (1) it can be shown that the utility function is strictly increasing and strictly concave in q_{ijt} , thus leading to unique allocations, as long as m_i is sufficiently large, and $\alpha_{i,J+1}$ is strictly positive.¹¹

Expenditure Shocks. We model ex-post expenditure shocks that hit the consumer in the second-stage of the decision process as follows. Let ω_{ijt} be the budget inherited from the planner's first-stage constrained decision problem. We will discuss the structure of this object in more detail in §3.2 below. The doer takes this object as given, as if it is a fixed parameter. The doer realizes ζ_{ijt} ,

⁸Strong separability ensures we can analytically solve for the planner's equilibrium budgeting decisions, which will become apparent in §3.3.

⁹This condition of our model ensures that consumers never devote all of their available resources to expenditure, desiring instead to carry liquidity forward in time. Since the dynamics of the model are sequential rather than in terms of continuation values, if consumers care about balance-matching, as Gathergood et al. (2019) have shown, we expect that they also care about balance holdings and how their behavior contributes to such holdings. For other "money in the utility function" models readers should refer to the monetary economics literature (Brock, 1974; Calvo, 1983; Obstfeld and Rogoff, 1983; Feenstra, 1986; Barnett et al., 1992; Walsh, 2010).

¹⁰In a consumption/savings model with borrowing, such as those featured in Bewley (1986), Huggett (1993), Aiyagari (1994), and Huggett (1996), m_i is a limit to draw upon all available credit. The interpretation of m_i in the context of our model is more narrow. Rather, m_i represents the amount a consumer could feasibly over-draft their account. The mechanism is only needed for the rare occasion that over-drafting is observed, and it is not a critical component of the model. Indeed, in practice m_i is never binding.

¹¹We present a proof of this in Lemma B.1 in Supplemental Appendix B.

an *iid* idiosyncratic expenditure shock, which captures how much ex-post expenditure deviates from the ex-ante budget:

$$x_{ijt} = \omega_{ijt} + \zeta_{ijt}, \quad (2)$$

Given ω_{ijt} , (2) exactly characterizes expenditure, so that consumption is $q_{ijt} = (\omega_{ijt} + \zeta_{ijt})/p_{jt}$. Since the vector ζ_{it} is J -dimensional, realized indirect utility, v_{it} can be written strictly as a function of the expenditure shocks, prices, the budgeting parameters, period- t income, and balances:

$$\begin{aligned} v_{it}(\zeta_{it}; \mathbf{p}_t, \boldsymbol{\omega}_{it}, \ell_{it}, b_{it}) &= \sum_{j=1}^J \alpha_{ij} \ln \left(\frac{\omega_{ijt} + \zeta_{ijt}}{p_{jt}} + 1 \right) \\ &+ \alpha_{i,J+1} \ln \left(\ell_{it} - \sum_{j=1}^J [\omega_{ijt} + \zeta_{ijt}] + m_i + r_t b_{it} \right) \end{aligned} \quad (3)$$

The sense in which the doer is myopic can be seen by noting that, given income, balances, budgets, and prices, realized utility depends entirely on the realization of ζ_{it} . Thus, from the perspective of the doer, utility is random.

But, what is ζ_{it} ? This shock encodes unanticipated deviations from the spending plan, due to everything from weather-related variation to unexpected health shocks (i.e., taste shocks) to unexpected price shocks. Let ζ_{it}^u be a J -dimensional vector of the components of ζ_{it} corresponding to the preference shocks, and let ζ_{it}^p be a J -dimensional vector of the components of ζ_{it} which correspond to the price shocks. We assume orthogonality conditions such that $\zeta_{ijt} \perp \zeta_{ij't}$ and $\zeta_{ijt}^u \perp \zeta_{ij't}^p$ where $j' \neq j$. This means that across categories, expenditure shocks are uncorrelated, and within a particular category, preference shocks are uncorrelated with price shocks.

Finally, set $\zeta_{it} = \zeta_{it}^u + \zeta_{it}^p$ and let the j^{th} component of ζ_{it} be mean-zero normally distributed — $\zeta_{ijt} \sim_{iid} \mathcal{N}(0, \sigma_{ij}^2)$. Note that when it comes to our quantitative strategy, we seek to infer budgeting behavior from expenditure data, but we do not observe prices or quantities. It is thus not possible to separately identify the degree to which the shock, ζ_{it} , results from preference heterogeneity versus unexpected price variation. For the purposes of our quantitative theory, though, this does not matter, because it does not matter *how* consumer budgets are violated (e.g., from preference shocks versus price shocks) but rather how they *respond* to such violations by updating their budgets.

3.2 The Planner

While the doer takes budgets, $\boldsymbol{\omega}_{it}$, as given, the planner must first decide what these budgets shall be conditional upon various behavioral and cognitive processing constraints. We now outline the planner's ex-ante decision process around which the doer's myopic expenditure decisions vary.

Mental Accounting. We start our description of the budgeting problem faced by the planner by first describing how the mental accounting state variable evolves. Denote the consumer's "mental account" by a_{ijt} . This is a state variable that encodes the amount the consumer over- or under-

spent on category j in the previous period, $t - 1$, relative to their budget. We loosely follow the mental accounting model in [Thaler \(1985\)](#), where past over- or under-expenditure in a particular category affects an individual's willingness to engage in future expenditure. The law of motion for mental account balances is

$$a_{ijt} = \omega_{ij,t-1} - x_{ij,t-1} = -\zeta_{ij,t-1} \quad (4)$$

The mental account balance for consumption category j is exactly equal to the negative of that category's idiosyncratic expenditure shock in the previous period. If $a_{ijt} < 0$ then the consumer over-spent in period $t - 1$, while if $a_{ijt} > 0$ they under-spent relative to their budget. The vector $\mathbf{a}_{it}$ is $J + 1$ dimensional where the $J + 1$ component, $a_{i,J+1,t}$, encodes how much a consumer over- or under-saved relative to their budget. Note that if $a_{i,J+1,t} < 0$ then the consumer has under-spent (over-saved) relative to their budget and vice-versa for $a_{i,J+1,t+1} > 0$. We assume that $\sum_{j=1}^{J+1} a_{ijt} = 0$ in every period, so that consumers' beliefs regarding leftover resources from the previous period do not shift their total resource constraint. That is, their perception of how much they over- or under-spent last period in various commodity groups is exact.¹²

Budgeting Function. We defined the category j budget in period t as ω_{ijt} . This object is a function of the consumer's mental account balance carried forward from the previous period, as well as their income, according to:

$$\omega_{ijt} = \theta_{ijt}\ell_{it} + \gamma_i a_{ijt} \quad (5)$$

In (5) γ_i captures how much a consumer's mental account balance influences their budget. γ_i is restricted to reside in the unit interval. As $\gamma_i \rightarrow 1$ the degree to which a consumer anchors their spending to the amount over- or under-spent last period increases. At $\gamma_i = 0$ there is no mental-accounting anchoring effect.

Assuming γ_i is a fixed, agent-specific parameter and that consumers take their income flows as given, then in (5) the only object that consumers can ex-ante control is θ_{ijt} , which is the quasi-share of income, ℓ_{it} , the consumer intends to spend on category j consumption in period t . This object is a quasi-share because it need not be interior to the unit interval: for example, if a consumer enters period t with a large mental account balance, a_{ijt} , due to under-spending in the previous period, the consumer may choose to set $\theta_{ijt} < 0$, thus reducing their period- t budget, ω_{ijt} , such that it is less than the effective mental account balance (i.e., $\omega_{ijt} < \gamma_i a_{ijt}$). Indeed, we assume that when making optimal budgeting decisions, the planner affects ω_{ijt} by setting θ_{ijt} , given values for ℓ_{it} and a_{ijt} which they know at the beginning of the period.

Cognitive Constraints on Budget Changes. While consumers choose θ_{ijt} when making an optimal budgeting decision, they do not necessarily make such decisions for every consumption category j in every period t because they face cognitive processing constraints. Let ψ_{ij} be the probability a consumer re-evaluates their budget-weighting variable θ_{ijt} in the first stage of period t . Let $\underline{\Gamma}_{ijt}$

¹²The mental accounting model in [Farhi and Gabaix \(2020\)](#) features the same implicit assumption.

be an indicator variable that equals 1 when a re-evaluation is made for the budget in j and 0 otherwise. Γ_{ijt} is Bernoulli distributed

$$\Gamma_{ijt} \underset{iid}{\sim} \text{Bernoulli}(\psi_{ij}) \quad (6)$$

with Γ_{ijt} assumed to be orthogonal to $\Gamma_{ij't}$ for all $j' \neq j$. This independence assumption induces narrow choice bracketing and results in a sparse-max budgeting structure, like the decision process analyzed in [Gabaix \(2014\)](#). In the context of our problem, a consumer choosing how to allocate funds across multiple different consumption budgets might only optimally re-evaluate one or two of those budgets in any given period, leaving the remainder fixed. In this way $\underline{k}_{it} = \sum_{j=1}^J \Gamma_{ijt}$. Note that just because a budget was re-evaluated does not necessarily mean that the new, candidate budget will be accepted by the consumer. This is because the absolute difference between the new candidate budget and the old budget must exceed some exogenously imposed numeracy threshold, which we discuss more below.

That in some periods $\underline{k}_{it} < J$ (strictly) is a modeling choice which draws on psychological evidence that consumers are cognitively constrained with regards to the amount of information they can process.¹³ Constrained consumers may possibly re-evaluate and subsequently update only a subset, $0 \leq \underline{k}_{it} \leq J$, of their budgets in any given period, where $\underline{k}_{it}$ describes the total number of expenditure budgets the agent optimally adjusts in a period.^{14,15} This integer constraint is exogenous and probabilistic: our premise is that consumers do not choose how rational they want to be.

Information Set. At the beginning of every period the planner realizes the vector of budget-evaluation indicators, Γ_{it} . Consumers remain uncertain about ζ_{it} until the second stage of the period when the doer is in action. The planner also enters each period knowing the vector of budget weights, $\theta_{i,t-1}$, ascribed to last period's income and the mental account balances, a_{it} . From the consumer's perspective $\{\theta_{i,t-1}, a_{it}\}$ comprises the entirety of the set of period- t endogenous state variables. Additionally, consumers know the following exogenous state variables prior to setting their expenditure budgets — income ℓ_{it} , bank balances b_{it} , rates of return on such balances r_t , and the vector of average commodity-group aggregate price levels $\bar{p}_t$. Note that $\bar{p}_t$ is different from the actual price the consumer faces at the check-out counter. Rather, this price level can be

¹³This is synonymous with the number of choices they can consider at one time ([Miller, 1956](#); [Simon, 1957](#); [Cowan, 2000](#)).

¹⁴In some periods they may update the full set of budgets, and in some periods they may update none of their budgets.

¹⁵Imposing this integer constraint on the number of changes that can be made, rather than the magnitudes of the budget updates themselves, is supported by findings in [Leslie et al. \(2008\)](#) who argue that integer representations are innate within individual cognitive processes. For example, findings in [Miller \(1956\)](#) and [Simon \(1957\)](#) suggest that individuals can consider at most seven “chunks” of information at once. Meanwhile, [Cowan \(2000\)](#) says this number is closer to four. Further, information processing capacity in human memory is fairly limited and has been shown to impact consumption decisions ([Malhotra, 1982](#)). There is also a relationship between our work and recent work showing that agents exhibit aversion to complex decision making and prefer decisions wrought from more simplified rules ([Oprea, 2020](#)). See also [Abeler et al. \(2025\)](#), who provide experimental evidence that incentive complexity causes agents to deviate from optimal behavior in ways consistent with bounded rationality.

thought of as an index, like the Consumer Price Index (CPI) or Personal Consumption Expenditure (PCE) index, which encode the relative value of baskets of various goods.

Numeracy Constraints. Prior to characterizing the optimality conditions associated with the expenditure budgets, we briefly describe how the numeracy constraints work, though this will become more obvious upon reading ahead. First, in any period and for any consumption category for which $\underline{\Gamma}_{ijt} = 1$, the consumer will choose a new candidate budget so as to maximize utility. This candidate consumption-category j budget will only be accepted as a new, optimal budget, if it is *significantly different enough* (i.e., satisfies the numeracy constraint) from the budget entering the period, $\theta_{ijt,t-1}$. If the planner rejects a budget because it does not satisfy the numeracy condition, then we treat such a condition as the planner having not changed the budget at all, thus setting $\theta_{ijt} = \theta_{ijt,t-1}$. We introduce the new state variable Γ_{ijt} (sans underline) which is equal to unity *if and only if* both $\underline{\Gamma}_{ijt} = 1$ and the candidate optimal budget satisfies the numeracy constraint. Thus, it must always be the case that $\Gamma_{ijt} \leq \underline{\Gamma}_{ijt}$. Further, letting $k_{ijt} = \sum_{j=1}^J \Gamma_{ijt}$ count the number of actually implemented budget changes which occur, it is necessarily so that $k_{ijt} \leq \underline{k}_{ijt}$ always, with equality only happening if all proposed budgets satisfy the numeracy constraints.

3.3 Choosing Expenditure Budgets

In each period prior to realizing ζ_{it} , the planner must set expenditure budgets subject to the constraints imposed by the behavioral frictions. Meanwhile, the doer's decision is exactly determined by the expenditure constraint in (2).¹⁶

Optimal Candidate Budgets. Recall that the vector θ_{it} is J -dimensional. Denote ϑ_{it} as the $\underline{k}_{it}$ -dimensional vector which holds, in cardinal order, the adjustable budgeting parameters of all commodities j for which $\underline{\Gamma}_{ijt} = 1$. These are the new, candidate budgets that must be evaluated against the numeracy constraint.¹⁷ Let $\mathbb{E}_{it}^{\zeta}$ be the expectation operator over the random variable ζ_{it} . After substituting (5) into (3) and replacing p_{jt} with the price index $\bar{p}_{jt}$, we can write expected indirect utility as

$$\begin{aligned} \mathbb{E}_{it}^{\zeta} v_{it}(\theta_{it}; \zeta_{it}) = \mathbb{E}_{it} \left\{ \sum_{j=1}^J \alpha_{ij} \ln \left(\frac{\theta_{ijt} \ell_{it} + \gamma_i a_{ijt} + \zeta_{ijt}}{\bar{p}_{jt}} + 1 \right) \right. \\ \left. + \alpha_{i,J+1} \ln \left(\ell_{it} - \sum_{j=1}^J [\theta_{ijt} \ell_{it} + \gamma_i a_{ijt} + \zeta_{ijt}] + m_i + r_t b_{it} \right) \right\} \end{aligned} \quad (7)$$

In (7) we suppress the explicit listing of arguments for known (from the planner's perspective) state variables in the definition of $v_{it}(\cdot)$. Note that (7) is the objective function the planner seeks to evaluate, subject to the integer constraints imposed by $\underline{\Gamma}_{it}$.

¹⁶In this sense our formulation aligns with the theory of [Ilut and Valchev \(2023\)](#) that decisions can be thought of as either "automatic" or "analytical." Here, the planner is the analytical/boundedly rational thinker, while spending is something that just happens automatically, though still constrained by the planner's budget and the distribution of the shock process.

¹⁷Note that ϑ_{it} is a sub-vector of θ_{it} corresponding to the non-zero indices of $\underline{\Gamma}_{it}$.

Let $\iota_{it}(\underline{\Gamma}_{it})$ be a vector-valued integer mapping that maps the index of components of ϑ_{it} back into the index of components of θ_{it} . This mapping is $\iota_{it}(\underline{\Gamma}_{it}) : \mathbb{N}^J \rightarrow \mathbb{N}^{k_{it}}$ and valid for $k_{it} > 0$. Let ϑ_{it}^* be the optimally-chosen analog of ϑ_{it} , prior to numeracy-constraint evaluation. The mapping $\iota_{it}(\underline{\Gamma}_{it})$ outputs in cardinal order the index of the components of θ_{it} to which we can assign the components of ϑ_{it}^* .¹⁸ Let y index the components of $\iota_{it}(\underline{\Gamma}_{it})$, so that ι_{iyt} denotes the y^{th} component of the vector ι_{it} . Finally, let θ_{it}^* be the planner's optimally-chosen set of *candidate* budget shares (prior to numeracy-constraint evaluation), which also includes budgets in categories for which $\underline{\Gamma}_{ijt} = 0$. Then, these objects are

$$\begin{aligned} \vartheta_{it}^* = \operatorname{argmax}_{\vartheta_{it}} \quad & \mathbb{E}_{it}^{\zeta} v_{it}(\vartheta_{it}; \theta_{it} \setminus \vartheta_{it}, \zeta_{it}) \\ \text{subject to} \quad & \theta_{ijt}^* = \begin{cases} \theta_{ij,t-1} & \underline{\Gamma}_{ijt} = 0 \\ \vartheta_{iyt}^* & \underline{\Gamma}_{ijt} = 1 \ \& \ j = \iota_{iyt} \end{cases} \end{aligned} \quad (8)$$

In (8) the function $v_{it}(\vartheta_{it}; \theta_{it} \setminus \vartheta_{it}, \zeta_{it})$ has been re-written so that the set of elements of θ_{it} for which $\underline{\Gamma}_{ijt} = 0$ is $\theta_{it} \setminus \vartheta_{it}$ and is taken as given by the planner (i.e., these budgets are not adjusted as part of the optimization problem). The solution to (8) is unique for large enough values of $\alpha_{i,J+1} > 0$, following from Lemma B.1 in Supplemental Appendix B and the fact that q_{ijt} is strictly increasing in θ_{ijt} .¹⁹ Clearly, if $k_{it} = 0$ because $\underline{\Gamma}_{ijt} = 0, \forall j$, then we set $\theta_{it} = \theta_{i,t-1}$, and the budget shares do not change. However, note that as long as $a_{it} \neq a_{i,t-1}$, the total budget ω_{it} (i.e., ex-ante expected expenditure under the assumption that ζ_{it} is mean zero) *will* change (i.e., $\omega_{it} \neq \omega_{i,t-1}$), since ω_{it} depends on a_{it} .

For $k_{it} > 0$ a candidate optimal choice of ϑ_{iyt}^* must satisfy

$$\mathbb{E}_{it}^{\zeta} \left\{ \frac{\partial v_{it}}{\partial \vartheta_{iyt}} \right\} = \mathbb{E}_{it}^{\zeta} \left\{ \frac{\partial u_i}{\partial q_{ijt}} \frac{\partial q_{ijt}}{\partial \vartheta_{iyt}} + \frac{\partial u_i}{\partial \bar{z}_{it}} \frac{\partial \bar{z}_{it}}{\partial \vartheta_{iyt}} \right\} = 0, \quad \forall y > 0 \quad \text{and} \quad j = \iota_{iyt}, \quad (9)$$

which is equivalent to the problem's representation in (8). Due to the money-in-the-utility-function structure, any allocation that solves the consumer's problem will ensure (9) is satisfied, as consumers equate the expected marginal utility of additional consumption with the expected marginal utility of additional liquidity. Further, note that (9) is an equivalent, component-wise representation of the optimization problem specified (8): when partitioning θ_{it} into ϑ_{it} and $\theta_{it} \setminus \vartheta_{it}$, the optimization problem can be represented in an unconstrained form so that each component of the vector ϑ_{it}^* must simply satisfy (9).

We can invert (9) under our utility parameterization to arrive at an analytical expression for ϑ_{it}^* . For some y indexing ι_{it} , let $\vartheta_{i,-yt}^*$ denote the vector of optimally chosen budget candidates which does not include y . Note that this vector may be empty when no changes are made in a

¹⁸For example, suppose $J = 4$ and $\underline{\Gamma}_{it} = (0, 1, 0, 1)$. Then this implies $k_{it} = 2$ and $\iota_{it}(\underline{\Gamma}_{it}) = (2, 4)$.

¹⁹Note that $\frac{\partial q_{ijt}}{\partial \theta_{ijt}} = \ell_{it}/p_{jt} > 0$, always.

given period. The optimal budget share $\vartheta_{iyt}^*(\vartheta_{i,-y,t}^*)$ for good ι_{iyt} can be expressed as:

$$\begin{aligned} \vartheta_{iyt}^*(\vartheta_{i,-y,t}^*) = \mathbb{E}_{it} \left\{ \frac{\alpha_{i,\iota_{iyt}} \ell_{it} - \alpha_{i,\iota_{iyt}} \sum_{s \neq y} (\vartheta_{ist}^* \ell_{it} + \gamma_i a_{i,\iota_{ist},t} + \zeta_{i,\iota_{ist},t})}{\ell_{it} (\alpha_{i,\iota_{iyt}} + \alpha_{i,J+1})} \right. \\ \left. - \frac{(\alpha_{i,\iota_{iyt}} + \alpha_{i,J+1}) (\gamma_i a_{i,\iota_{iyt},t} + \zeta_{i,\iota_{iyt},t}) + \alpha_{i,J+1} \bar{p}_{\iota_{iyt},t}}{\ell_{it} (\alpha_{i,\iota_{iyt}} + \alpha_{i,J+1})} \right. \\ \left. - \frac{\alpha_{i,\iota_{iyt}} \left(\sum_{j \notin \iota_{it}(\Gamma_{it})} (\theta_{ijt} \ell_{it} + \gamma_i a_{ijt} + \zeta_{ijt}) - m_i - r_t b_{it} \right)}{\ell_{it} (\alpha_{i,\iota_{iyt}} + \alpha_{i,J+1})} \right\} \end{aligned} \quad (10)$$

where $\theta_{ijt} = \theta_{ij,t-1}$ for $j \notin \iota_{it}(\Gamma_{it})$.²⁰ Recall we assume that ℓ_{it} and $\bar{p}_t$ are ex-ante known. When implementing the model quantitatively without explicitly observing budget-updating behavior, we exploit the independence of the components of ζ_{it} to iteratively sample the latent shocks Γ_{ijt} and update the components of ϑ_{it}^* accordingly, taking $\vartheta_{i,-y,t}^*$ as given. This allows for estimation of the mental accounting and budgeting parameters without having to iterate over the component functions of ϑ_{it} as we progress through a Markov Chain Monte Carlo estimation routine.

Numeracy Constraints. From our perspective as modelers, numeracy constraints must be evaluated *after* candidate budgets are computed, since one must have the candidate budget in hand to compare against the previous budget and then evaluate whether or not the numeracy constraint is satisfied. In terms of consumers' actual behaviors though, one can simply assume that the entire process characterized in §3.3 happens *all at once*: consumers solve (8) and evaluate the numeracy constraints we describe in this section simultaneously to arrive at period t 's optimal budget shares.

We incorporate numeracy constraints into the optimal budgeting problem for practical issues: if a budget change is really small, then it will have very little impact. We call the margin whereby a possible budget re-evaluation could occur (described above) the “extensive margin.” Let $\Gamma_{ijt} = 1$ (no underline) when both the extensive margin and the numeracy margin (intensive) for commodity group j are satisfied. Since re-evaluation occurs before the numeracy constraint is assessed, $\Gamma_{ijt} \leq \underline{\Gamma}_{ijt}$ always. Let k_{it} be the total number of budget updates inclusive of the numeracy-evaluation condition, so that $k_{it} \leq \underline{k}_{it}$ always, as well.²¹ We consider two possible types of numeracy constraints, which are evaluated after the initial draws of $\underline{\Gamma}_{ijt}$:

- i. The candidate budget will be accepted and $\Gamma_{ijt} = 1$ if the constraint ϵ on the minimum percent-change of the budget in question is satisfied: $|\vartheta_{iyt}^* - \theta_{i,\iota_{iyt},t-1}| \geq \epsilon$.
- ii. The candidate budget will be accepted and $\Gamma_{ijt} = 1$ if the constraint ϵ_ℓ on the absolute dollar-value of the budget change in question is satisfied:²² $\ell_{it} |\vartheta_{iyt}^* - \theta_{i,\iota_{iyt},t-1}| \geq \epsilon_\ell$.

²⁰For the algebra behind this expression, see Supplemental Appendix A.

²¹It should be apparent that we use underlines to describe indicators for the extensive margin prior to the intensive-margin numeracy evaluation taking place.

²²The ℓ -subscript on ϵ_ℓ indicates that the constraint parameter has the same units as income (dollars).

Note that both ϵ and ϵ_ℓ will be exogenously imposed by the econometrician.²³

In a slight abuse of notation, we will let $\iota_{it}(\Gamma_{it}) \subseteq \iota_{it}(\underline{\Gamma}_{it})$ be the possible subset of components of the extensive-margin index list for which the numeracy constraints are also satisfied.²⁴ From here on, assume any mention of ι_{it} without its dependency argument is such that $\iota_{it} \equiv \iota_{it}(\Gamma_{it})$. Thus, in our analyses we focus only on budget changes for which both the extensive margin and the numeracy constraints are met.

Note that, in practice, imposing the numeracy constraint amounts to assuming the consumer first receives a draw of $\underline{\Gamma}_{ijt} = 1$. Next, the consumer solves for the candidate budget ϑ_{iyt}^* in (10). Finally, if such a budget satisfies the numeracy constraint, then $\Gamma_{ijt} = \underline{\Gamma}_{ijt} = 1$, and the new budget is accepted. Otherwise $\Gamma_{ijt} = 0$ and $\theta_{ijt} = \theta_{ijt-1}$.

3.4 Personal Mental Accounting Equilibrium

Figure 1 summarizes the flow of the decision process within a period, including the within-period timing of the different shock processes. First, the candidate re-evaluation vector $\underline{\Gamma}_{it}$ is realized. Next, optimal candidate budgets are computed if a re-evaluation is warranted. In the third and final step of the planner's decision process, the candidates are evaluated against the numeracy constraints; if no numeracy constraint is present, then we skip to the doer's problem. After the planner sets budgets, the vector of expenditure shocks, ζ_{it} , is realized $\forall j$. Period- t expenditure takes into account the mental account balance, a_{ijt} , which was set at the end of period $t - 1$. Next, after all expenditure is realized, new mental account balances are set and then carried forward into $t + 1$. The process then repeats itself, starting with the realization of the budget re-evaluation shock in period $t + 1$.




Figure 1: This directed graph presents the timing of the decision process within the period.

In Supplemental Appendix B we provide lemmas which characterize the conditions under

²³Prior to estimation we must assume that either the consumer is constrained to make certain percentage-based updates or absolute dollar-based updates.

²⁴Note that $\iota_{it}(\Gamma_{it}) : \mathbb{N}^J \rightarrow \mathbb{N}^{k_{it}}$, instead of $\mathbb{N}^{k_{it}}$.

which optimal budgets are unique and for which we can solve for such outcomes. Note that uniqueness holds almost always, as shown by Lemma B.2, so it does not threaten the validity of our estimation. We now formally define a “personal mental accounting equilibrium.”

Definition (Personal Mental Accounting Equilibrium): Assume Lemma B.2 holds. Let t_{i0} be the period in which an individual consumer enters the economy as an autonomous, decision-making agent. Given a numeracy threshold level (either ϵ or ϵ_ℓ), a sequence of prices $\{\bar{p}_t, r_t\}_{t \geq t_{i0}}$, a sequence of positive income values $\{\ell_{it}\}_{t \geq t_{i0}}$, a sequence of idiosyncratic expenditure shocks $\{\zeta_{it}\}_{t \geq t_{i0}}$, a sequence of idiosyncratic cognitive shocks $\{\underline{\Gamma}_{it}\}_{t \geq t_{i0}}$, and initial values for the budget weights, mental account balances, and bank balances $\{\theta_{i,t_{i0}}, \mathbf{a}_{i,t_{i0}}, b_{i,t_{i0}}\}$, a personal mental accounting equilibrium consists of:

- i. Sequences of policies: $\{q_{it}, z_{it}, \theta_{it}\}_{t \geq t_{i0}}$.
- ii. Sequences of balances: $\{a_{it}, b_{it}\}_{t \geq t_{i0}}$.

such that in each period

- a. Given $\underline{\Gamma}_{it}$, a_{it} , and b_{it} , θ_{it} satisfies the sparse-max indirect-utility maximization program in (8) and corresponding numeracy constraints.
- b. Given θ_{it} , ℓ_{it} , and a_{it} , ω_{it} is exactly (5).
- c. Given θ_{it} , ℓ_{it} , ω_{it} , and ζ_{it} , x_{it} satisfies the expenditure criterion in (2).
- d. Given x_{it} and ℓ_{it} , z_{it} satisfies the marginal savings rule.
- e. Mental account balances, a_{it} , are updated according to (4) with $a_{i,J+1,t} = -\sum_{j=1}^J a_{ijt}$.
- f. Bank balances evolve according to $b_{i,t+1} = r_t b_{it} z_{it}$.

4 Weekly Expenditure Data

To discipline our theoretical framework with observed expenditure outcomes, we turn to a novel dataset of individual, consumer-level weekly-aggregated expenditure. The goal of our endeavor is to estimate consumer demand parameters while simultaneously considering the endogenous link between spending and budgeting decisions. Our challenge is that we only observe expenditure variation, not budgets.²⁵ Given model structure, we thus must latently infer how budgets and

²⁵An ideal dataset that tracks spending, prices, and income, along with consumers’ budgets, does not actually exist. Financial-planning and budgeting apps (e.g., Mint, Empower, YNAB (You Need a Budget), EveryDollar, Honeydue, and Personal Capital) record budgeting goals, sometimes by category, but such apps still do not record transaction-level prices. This means that categorical aggregation of expenditure would still need to happen in order to match spending to aggregated price indices. Further, users opt in to using such apps, meaning that conclusions drawn from their data are likely plagued by selection bias. Our expenditure data comes from enterprise banking data, which is not affected by such a bias. We can thus infer budgeting behavior from all types of consumers, including those who otherwise would not otherwise choose to use expenditure monitoring services.

mental accounts evolve in order to explore what the model implies about the degree to which observed expenditure patterns are consistent with bounded rationality.

Our dataset is an anonymized sample of low-income, pre-paid debit-card users from a large North American bank. We observe and classify users by card number, though multiple card numbers may be associated with a single household. Our most basic unit of observation is a “swipe” (i.e., card utilization) and/or a “load” (i.e., when income is added to a card, either by an employer’s direct deposit or some other kind of bank transfer). From swipes and loads, we sum up expenditure by category, income, and running balances to get weekly aggregates.

Expenditure is categorized by 4-digit Visa merchant category classification codes (MCC).²⁶ To classify expenditure categories we consider only the first two digits of the MCCs, which give us three explicit consumption categories and one commodity group that collects all other expenditure.²⁷ The explicit categories we examine, and their 2-digit MCCs are Groceries (54), Gasoline (55), and Food Away from Home (58).²⁸ The timeline of observations spans from September 2013 to January 2016 with most agents appearing as active users over only the latter part of this time period.

After data cleaning and selection, we have a total of $I = 2,509$ customers with a total of 71,859 weekly observations. This gives an average of approximately 29 weeks per customer. Importantly, as we alluded to above, we restrict our sample of consumers to those who appear to use their pre-paid debit card as their only banking product. These consumers are associated with deposits of regular, weekly income to their pre-paid card and use the card for expenditures on a regular basis. On one hand, restricting the analysis to pre-paid debit card users limits inferences to only a small subset of the population, namely the underbanked with limited access to other banking products, such as checking accounts and credit cards.²⁹ On the other hand given data limitations, doing so greatly increases the probability that the income and expenditure profiles we observe represent complete consumption profiles for those agents who use this particular product. Each consumer profile contains at least 16 consecutive weeks of regular income. This ensures that we are observing a regular income process for the consumer. Finally, we exclude from our sample consumers for whom we observe at least four consecutive weeks of zero expenditure in one or more of the aforementioned categories.

Summary statistics for our entire sample are presented in Table 1, where the units are agent-

²⁶Within the card industry this code is common across all card processors and is also known as the Standard Industry Code (SIC).

²⁷Note that the fourth category representing all remaining expenditure includes both classifiable transactions that fall outside the 2-digit codes above and cash withdrawals from the account. This means that some possibly classifiable expenditure may be mis-categorized. Unfortunately, this cannot be reconciled without making some very strong assumptions with regards to how consumers spend their money.

²⁸Our goal with category selection is to look at purchasing behavior in categories whose demand has also been widely analyzed in the literature. For the current analysis categories have been exogenously defined in order to discipline inference.

²⁹We have verified that these customers do not have other accounts at the bank from which our data originates. Although it is possible that these individuals have accounts at other banks that we cannot see, we believe this to be unlikely given the nature of the selected customer base.

Table 1: Summary Statistics Over Agent-level Means

	5%	1st Qu.	Median	Mean	3rd Qu.	95%	S.D.
Weeks	16	20	25	28.640	35	52	11.131
Income	207.096	338.855	460.052	510.732	620.131	981.994	258.288
Balances	6.216	89.953	204.734	500.744	438.464	1,415.661	1,868.651
<i>Expenditure</i>							
Groceries	11.472	22.282	34.715	42.018	54.195	94.548	30.622
Auto/Gas	13.393	29.191	48.430	60.941	80.198	148.186	46.064
Food Away	14.614	30.179	51.965	68.724	86.684	177.296	61.656
Other	99.211	194.847	283.406	326.124	415.037	687.897	195.551

Number of Consumers $I = 2,509$; Total Consumer Weeks $\sum_i \sum_{t_i} = 71,859$

level averages.³⁰ The median consumer in our sample earns \$460.05 per week after taxes, which aggregates to \$23,922.60 per year in nominal dollars. For comparison, consider that median U.S. household income in 2015 was \$56,516 according to the U.S. Census bureau. 96.8% of our sample fall below 2015 median household income. The 2015 poverty thresholds for one, two, three, and four person households where the head was under the age of 65 were \$12,331, \$15,952, \$18,871, and \$24,257; 7.8%, 18.3%, 30.8%, and 51.4% of our sample respectively fall below these corresponding poverty thresholds.³¹

We do not observe purchase prices at the transaction-level, so we turn to the Consumer Price Index (CPI) for price indices. Price units are \$1982-84 dollars. All price indices are published at the monthly level, so we use simple linear interpolation to get weekly, aggregate price estimates. For the prices of groceries and food away from home, we use the U.S. city average food and beverage price index (series I.D. "CUUR0000SAF"). For the price of gasoline and other miscellaneous automotive maintenance expenses, we use the U.S. city average transportation price index (series I.D. "CUUR0000SAT"). All other expenditure will be deflated using the U.S. city average all items index (series I.D. "CUUR0000SA0").

5 Quantitative Implementation

Our quantitative strategy serves to discipline the model's free parameters so that its implications are grounded in observed expenditure patterns. The model thus provides a coherent theoretical lens through which to interpret the data. The value of the exercise lies in the model's ability to generate quantitative predictions about latent objects (i.e., budgets, mental accounts, attentiveness) that are consistent with observed expenditure and that can be used to conduct structured quantitative analysis, including counterfactuals. Given one aim of this paper is to infer budgeting decisions without observing actual budgets, only expenditure, we desire a parametric model to

³⁰For example, in Table 1 "Median" income corresponds to the median agent-level average income and so on.

³¹We do not observe the household size in our dataset.

avoid placing further burdens on the data. We have so far described such a model that places structure on the agents' decisions. We can then exploit our assumed structure to estimate equilibrium budgeting decisions as parameters. We do this by turning to the literature in Bayesian statistics that deals with detecting latent change points of observed time series. Since our time series of budgeting decisions, ϑ_{it}^* , are themselves endogenous and unobserved, we require that the equilibrium condition in (10) and the equilibrium laws of motion of mental account balances in (4) all exactly hold conditional upon the time-independent preference parameters as well as estimated latent change points, which are identical to the vector of unobserved budget-updating indicators, Γ_{it} .³² This is accomplished with an algorithm that borrows from the endogenous change-point inference techniques described in McCulloch and Tsay (1993), Chib (1998), and Koop and Potter (2007) nested within a hierarchical Metropolis-Hastings (MH) within Gibbs Markov Chain Monte Carlo (MCMC) sampling algorithm (Hastings, 1970; Geman and Geman, 1984).

The key challenge of our quantitative implementation is that we must estimate budget weights, mental accounts, and cognitive frictions as latent time series variables, $\{\theta_{it}, a_{it}, \Gamma_{it}\}_{t=1}^{T_i}$ for each i using only variation in observed expenditure. Due to the vast degree of observed heterogeneity across agents, most of our estimation routine operates only on agent-level parameters. Because of the need to sample latent time series of consumer decisions, the parameter space is large, containing over four million parameters.³³ Within the model's structure, budget changes correspond to level shifts in average spending, and the initial mental account balance a_{i1} is pinned down by first-period deviations from average spending. We emphasize that these identifications are conditional on the model: they reflect what the theory, combined with data, implies about latent budgeting behavior.³⁴

The remainder of this section describes how the theoretical model is made operational. We first explain our estimation routine in rigorous detail, providing intuition as to how budgeting choices are conditionally identified. We then assess model fitness to confirm that the parameterized theory can reproduce observed expenditure patterns — a necessary step before the model can serve as a usable laboratory for quantitative analysis. Additional MCMC fitness measures and diagnostic statistics are presented in Supplemental Appendix D. With a disciplined, working model in hand, we then use it to explore the quantitative implications of our estimates for behavioral features such as budgeting, anchoring, and choice bracketing.

³²We count budget updates only if they satisfy both the extensive and intensive (numeracy) margins.

³³The exact dimension of the parameter space in our preferred model is 4,664,212 with 4,598,976 of these parameters being the values of latent, time-dependent budgeting and mental accounting variables.

³⁴Alternative theoretical structures could attribute the same data patterns to different mechanisms. Our contribution is not to definitively resolve this attribution but to show what a coherent model of two-stage budgeting with behavioral frictions implies quantitatively when disciplined by observed expenditure.

5.1 Description of Estimator

5.1.1 Overview

The estimation procedure translates the theoretical structure developed in Section ?? into a parameterized analytical instrument that can be directly interrogated. We estimate multiple different models, varying assumptions regarding numeracy constraints and mental accounting behavior. For each estimation we initialize the MCMC integration scheme described here with random generates from the prior distributions, except for a_{i1} which we set to $\mathbf{0}$ to start. We then proceed to iterate through a chain of length 400,000. That is, we operate on the agent-level MH within Gibbs blocks 400,000 times and the global blocks $400,000/400 = 1,000$ times. We set a burn-in period of 120,000, keeping every 400^{th} agent-level draw and every global draw thereafter, giving us a total of 700 draws after burn-in and trimming. The sampler operates with agent-level blocks parallelized over a 128-core computer, requiring just under 100 hours to completion for each of the 13 different estimation routines.³⁵

5.1.2 Likelihood Function and Prior Distributional Assumptions

In this section we describe the consumer-specific likelihood function associated with observing given combinations of expenditure, income, and balances. We then describe several prior distributional assumptions on agent-specific parameters and global hyper-parameters.

Under the assumption that the expenditure errors, ζ_{ijt} , are *iid* and Gaussian normal with agent- and commodity-specific variances, σ_{ij}^2 , let $\phi(\cdot)$ denote the standard normal probability density function and $\Phi(\cdot)$ denote the associated cumulative distribution function. To accommodate the zero lower-bound on expenditure, we assume that the likelihood of realizing agent i 's data under the given structural parameterization takes the following form:

$$\begin{aligned} \mathcal{L}_i \left(\underbrace{\{\mathbf{x}_{it}, \ell_{it}\}_t}_{\text{Data}} \mid \underbrace{\{\{\boldsymbol{\theta}_{it}, \mathbf{a}_{it}\}_t, \boldsymbol{\sigma}_i^2, \gamma_i\}}_{\text{Structural Parameters}} \right) \\ = \prod_t \prod_j \underbrace{\left(\frac{1}{\sigma_{ij}} \phi \left(\frac{x_{ijt} - \theta_{ijt} \ell_{it} - \gamma_i a_{ijt}}{\sigma_{ij}} \right) \right)^{\mathbf{1}(x_{ijt} > 0)} \left(1 - \Phi \left(\frac{\theta_{ijt} \ell_{it} + \gamma_i a_{ijt}}{\sigma_{ij}} \right) \right)^{\mathbf{1}(x_{ijt} \leq 0)}}_{\mathcal{L}_{ijt}(x_{ijt}, \ell_{it} \mid \theta_{ijt}, a_{ijt}, \sigma_{ij}^2, \gamma_i) :=} \end{aligned}$$

In a slight abuse of notation we denote the agent i , commodity group j , and period t likelihood of observing x_{ijt} and ℓ_{it} , $\mathcal{L}_{ijt}(x_{ijt}, \ell_{it} \mid \theta_{ijt}, a_{ijt}, \sigma_{ij}^2, \gamma_i)$.³⁶

Inference on the agent's budget weights, θ_{ijt} , and mental account balances, a_{ijt} , operates on the inverted ex-ante expected indirect utility first-order condition. Note that the entire time series of

³⁵There are 6 separate routines with different numeracy thresholds and 7 separate routines, including the full model, with different assumptions regarding the degree to which consumers are rationally constrained and/or engage in mental accounting. We discuss the various, separate estimators further below.

³⁶This notation will come in handy when discussing the sampling steps for latent budgeting parameters.

mental account balances, a_{it} , and by extension budget weights, θ_{it} , depends on the initial balances, a_{i1} , which are unknown parameters. Therefore, we give these initial balances, a_{i1} , a normal prior structure:

$$a_{ij1} \sim \mathcal{N}(0, \sigma_{a_{ij1}}^2)$$

$\sigma_{a_{ij1}}^2$ is fixed for each good and each agent. We set this value to correspond to the expenditure variance for good j over the whole time series of observations divided by average agent- i income, $\widehat{\ell}_i$, specifically $\sigma_{a_{ij1}}^2 = \text{var}(x_{ijt}) / \widehat{\ell}_i$. Given a_{i1} , initial values of θ_{ij1} are then set every iteration of the sampler by passing a_{i1} and data for $t = 1$ through the integral in the expression for ϑ_{iyt}^* defined in (10). Finally, amongst the budgeting variables, cognitive shocks Γ_{ijt} occur with probability ψ_{ij} , which we give a flat uniform prior:

$$\psi_{ij} \sim U[0, 1]$$

In our various estimation exercises we fix either the mental accounting weight to $\gamma_i = 1$ (base-line estimation) or $\gamma_i = 0$ (estimation with no serial dependence). For the utility weights we require $\sum_{j=1}^J \alpha_{ij} = 1$ with $0 \leq \alpha_{ij} \leq 1$ for all $j \leq J$. Following [Geary \(1950\)](#), [Stone \(1954\)](#), and [Deaton and Muellbauer \(1980a\)](#), these values encode the shares of total expenditure devoted toward category j consumption. We fix these values as the mean expenditure shares taken over the entire time series we observe for each consumer: $\widehat{\alpha}_{ij} = \frac{1}{T_i} \sum_t \left(\frac{x_{ijt}}{\sum_j x_{ijt}} \right)$.

Note that m_i and $\alpha_{i,J+1}$ cannot be separately identified. We sample $\alpha_{i,J+1}$ explicitly while fixing the consumption borrowing limit, m_i , to be 20% of an agent's annualized income.³⁷ From simulations we have inferred that $\alpha_{i,J+1}$ is weakly identified given m_i , so we impose a three-tier hierarchical prior structure with empirical priors for the global hierarchical mean and variance:

$$\begin{aligned} \alpha_{i,J+1} &\sim \mathcal{N}(\mu_{i,\alpha_{i,J+1}}, 10^2) \quad \text{with} \quad \mu_{i,\alpha_{i,J+1}} \sim \mathcal{N}(\bar{\mu}_\mu, \bar{\sigma}_\mu^2) \\ \text{where } \bar{\mu}_\mu &= \frac{1}{I} \sum_i \mu_{i,\alpha_{i,J+1}} \quad \& \quad \bar{\sigma}_\mu^2 = \frac{1}{I-1} \sum_i (\mu_{i,\alpha_{i,J+1}} - \bar{\mu}_\mu)^2 \end{aligned} \quad (11)$$

We allow the prior variance on the draw of the structural parameter $\alpha_{i,J+1}$ to be large and identical across all agents, inducing flatness for the initial draw. In our simulations, we found that the sampler achieves convergence more efficiently under this structure than one in which the prior variance on $\alpha_{i,J+1}$ is explicitly sampled for each agent.

The orthogonality assumption on error terms such that $\zeta_{ijt} \perp \zeta_{ij't}, j' \neq j$, allows us to impose

³⁷Annualized income is computed by multiplying agent i 's weekly average income, $\widehat{\ell}_i = \frac{1}{T_i} \sum_t \ell_{it}$, by 52.

separate, agent-specific, hierarchical independent gamma priors on the likelihood precisions, $\frac{1}{\sigma_{ij}^2}$.³⁸

$$\frac{1}{\sigma_{ij}^2} \sim \text{Gamma}\left(\tau_{ij}, \frac{1}{\beta_{ij}}\right) \quad \text{with} \quad \tau_{ij} \propto \frac{\beta_{ij}^{\tau_{ij}}}{\Gamma(\tau_{ij})\sigma_{ij}^{2(\tau_{ij}-1)}} \quad \& \quad \beta_{ij} \sim \text{Gamma}(1, 1)$$

The non-conjugate prior on τ_{ij} is discussed in [Fink \(1997\)](#). Its posterior can be easily sampled using a Metropolis-Hastings step. The rate parameter β_{ij} has a conjugate gamma posterior under a gamma prior.

5.1.3 Estimating Latent Budgeting Variables

Let n index the atomic posterior estimation draws (epochs) from the Metropolis-Hastings (MH) within Gibbs sampler. Let $\mathcal{P}_{in}$ be an atomic draw of parameters associated with agent i on iteration n , and let $\mathcal{H}_n$ be the set of global hierarchical parameters common to all agents.³⁹ Let $\mathcal{D}_i$ denote the set of agent i 's data. These sets (with vectors in bold) are:

$$\begin{aligned} \mathcal{P}_{in} &= \left\{ \sigma_{in}^2, \boldsymbol{\tau}_{in}, \boldsymbol{\beta}_{in}, \alpha_{i,J+1,n}, \mu_{i,\alpha_{i,J+1,n}}, \boldsymbol{\psi}_{in}, \{\boldsymbol{\theta}_{itn}, \mathbf{a}_{itn}, \boldsymbol{\Gamma}_{itn}\}_{t=1}^{T_i} \right\} \\ \mathcal{H}_n &= \{\bar{\mu}_{\mu_n}, \bar{\sigma}_{\mu_n}^2\} \\ \mathcal{D}_i &= \left\{ \{\mathbf{x}_{it}, \ell_{it}, b_{it}, \bar{\mathbf{p}}_t, r_t\}_{t=1}^{T_i} \right\} \end{aligned}$$

Note that we do not observe prices, but prices feature in the calculation of $\boldsymbol{\vartheta}_{it}^*$. Thus, we use CPI estimates, $\bar{\mathbf{p}}_t$. The sampler proceeds in several blocks. Here, we describe the algorithm used to make atomic draws of the budgeting parameters, $\{\boldsymbol{\theta}_{itn}, \mathbf{a}_{itn}, \boldsymbol{\Gamma}_{itn}\}_{t=1}^{T_i}$, and proceed to the description of the remaining Gibbs blocks in the next section. The algorithm for sampling the cognitive shocks, Γ_{ijt} , is based on the latent time-series level-shift identification algorithm from [McCulloch and Tsay \(1993\)](#). While ours is not a standard linear time series model, the spirit of the exercise is the same: to infer whether a change to the budget weight for commodity group j is made from period $t-1$ to t , exploit information provided by data observations in periods t to T_i .

Assume, for now, that $\gamma_i = 1$.⁴⁰ Take the non-budgeting parameters, $\mathcal{P}_{in} \setminus \{\{\boldsymbol{\theta}_{itn}, \mathbf{a}_{itn}, \boldsymbol{\Gamma}_{itn}\}_{t=1}^{T_i}\}$, as given, where $\setminus$ is the set exclusion operator. We want to estimate the posterior distribution of $\{\boldsymbol{\Gamma}_{it}\}_t$. Ignoring MCMC iteration indexing for now, consider the sequence $\{\mathbf{a}_{it}\}_{t>1}$, which, given $\{\boldsymbol{\theta}_{it}\}_t$ and data, is completely determined by the law of motion for mental account balances. For

³⁸We use the shape-scale parameterization of the gamma distribution such that if $(1/\sigma^2) \sim \text{Gamma}(\tau, 1/\beta)$ then $1/\sigma^2$ has probability density function

$$\frac{\beta^\tau}{\Gamma(\tau)\sigma^{2(\tau-1)}} \exp\left\{-\frac{\beta}{\sigma^2}\right\}$$

where $\Gamma(\tau)$ is the gamma function.

³⁹Recall, parameters fixed for the entirety of the MCMC sampling routine are $\gamma_i, m_i, \alpha_{ij}$ for all $j \in \{1, \dots, J\}$, and $\sigma_{a_{ij}}^2$.

⁴⁰This is the value we set in our primary baseline estimation and our other preferred estimation, which we discuss further in this section.

each j recursively substitute out $a_{ijt}, a_{ij,t-1}, \dots, a_{ij2}$ to get

$$\mathbb{E}_{it}^{\zeta} x_{ijt} = \sum_{s=1}^t \theta_{ijs} \ell_{is} - \sum_{s=1}^{t-1} x_{ijs} + a_{ij1}$$

Thus, for each t , $\mathbb{E}_{it}^{\zeta} x_{ijt}$ can be written solely as a function of a_{ij1} . Similarly, ϑ_{iyt}^* , where $j = i_{yt}$ is a function only of the vector of first period mental accounting state variables $\mathbf{a}_{i1}$.

Re-introducing MCMC indexing, n , take $\{\Gamma_{it,n-1}\}_{t=1}^{T_i}$ as given. We use an MH step to draw $\mathbf{a}_{i1n}$, stepping through from $t = 1$ to T_i to update $\theta_{ijtn} = \vartheta_{iytn}^*(\mathbf{a}_{i1n})$ whenever $\Gamma_{ijt,n-1} = 1$ and setting $\theta_{ijtn} = \theta_{ij,t-1,n}$ whenever $\Gamma_{ijt,n-1} = 0$. For each t , we step through, from $j = 1$ to J , and check $\Gamma_{i1t,n-1}, \Gamma_{i2t,n-1}, \dots, \Gamma_{ij,t,n-1}$ evaluating the integral in the expression for ϑ_{iyt}^* sequentially. Thus, on any given draw, updated values of ϑ_{iytn}^* depend implicitly on values $y' < y$ from the current iterate n and potentially $y'' > y$ from the previous iterate $n - 1$, as well as past values θ_{isn} , $s < t$:

$$\vartheta_{iytn}^*(\vartheta_{i1tn}^*, \dots, \vartheta_{i,y-1,tn}^*, \vartheta_{i,y+1,t,n-1}^*, \dots, \vartheta_{i,k_{it},t,n-1}^*, \theta_{i,t-1,n}, \dots, \theta_{i1n})$$

If numeracy constraints are also imposed, we must also check whether new, candidate budget updates satisfy either ϵ or ϵ_{ℓ} . With the candidate sequence $\{\theta_{itn}, \mathbf{a}_{itn}, \Gamma_{it,n-1}\}_{t=1}^{T_i}$, we can accept or reject the draw of $\mathbf{a}_{i1n}$ using a standard MH acceptance rule. If the draw is rejected, the entire sequence is reset so that $\theta_{ijtn} = \theta_{ij,t,n-1}$ and $a_{ijtn} = a_{ij,t,n-1}$ for all j and t .⁴¹ Thus, accepted behavioral parameter *sequences* sufficiently (relative to some probabilistic threshold) explain the underlying expenditure data. As this process is repeated over many different MCMC epochs, we arrive at an estimated distribution of behavioral parameters, including latent mental accounting and decision variables, that reconciles the different, observed agent-level data series.

Yet, the previous paragraph just describes how initial $\mathbf{a}_{i1n}$ is sampled, which itself in terms of statistical inference depends on *future* values of the other behavioral parameters tied to it: that is, values of future mental accounts and budgeting variables provide information regarding the initial value of $\mathbf{a}_{i1n}$, given data. Therefore, having performed the MH step for $\mathbf{a}_{i1n}$, we must now (conditional upon a value of $\mathbf{a}_{i1n}$ in hand) proceed to use Gibbs sampling to now jointly draw the sequences $\{\Gamma_{itn}\}_{t=1}^{T_i}$ and $\{\theta_{itn}\}_{t=1}^{T_i}$. For $t = 1$, we assume $\Gamma_{ij1n} = 1$, so that $k_{i1} = J$ and $\theta_{ij1n} = \vartheta_{ij1n}^*$ always (i.e., we always evaluate (10) to get the first period's budgets). For $t > 1$, θ_{ijtn} depends on past values $\theta_{ij,t-1,n}, \theta_{ij,t-2,n}, \dots, \theta_{ij,t-s,n}$ where s is the number of periods that have passed since the last budget update, $\Gamma_{ij,t-s,n} = 1$. Since $\theta_{ij,t-s,n}$ is a function of $\Gamma_{ij,t-s,n}$, then by extension so is θ_{ijtn} . Thus, any new draw of Γ_{ijtn} will affect all future realizations of the budget weights, $\{\theta_{ijsn}\}_{s=t}^{T_i}$. Further, since updating Γ_{ijtn} induces a change to the sequence $\{\theta_{ijsn}\}_{s=t}^{T_i}$, values of $\{a_{ijsn}\}_{s=t}^{T_i}$ are

⁴¹In practice, this MH step involves evaluating a joint draw of all agent-level parameters associated with a conditionally non-conjugate posterior. These parameters are $\mathbf{a}_{i1n}$, $\boldsymbol{\tau}_{in}$, and $\alpha_{i,J+1,n}$. We use the adaptive algorithm of [Atchade and Rosenthal \(2005\)](#) to achieve convergence towards a theoretically optimal acceptance rate of 0.234 ([Roberts et al., 1997](#)). All MH draws are agent-specific, so each agent's parameters are associated with a unique adaptation step and thus a unique acceptance rate.

also affected and must be updated to ensure the equilibrium mental account balance law of motion always holds.

Let $\mathcal{D}_{i,t:T_i} \subseteq \mathcal{D}_i$ that only contains data for observed periods t to T_i . Let $\mathcal{P}_{i,t:T_i,n} \subseteq \mathcal{P}_{in}$ similarly represent agent i 's parameter set containing budgeting parameters only for periods t to T_i , $\mathcal{P}_{i,t:T_i,n} = \mathcal{P}_{in} \setminus \{\{\theta_{isn}, \mathbf{a}_{isn}, \Gamma_{isn}\}_{s=1}^{t-1}\}$. Sampling $\{\Gamma_{ijtn}\}_{t>1}$ proceeds as follows. If $\Gamma_{ijt,n-1} = 0$, then set $\Gamma_{ijtn} = 1$ as a candidate draw and compute the sequence of budget weight policies $\{\theta_{ijsn}(\Gamma_{ijtn}, \Gamma_{ijt+1,n-1}, \dots, \Gamma_{ijs,n-1})\}_{s \geq t}$, re-computing ϑ_{iysn}^* and setting $\theta_{ijsn} = \vartheta_{iysn}^*$ for $j = i_{ys}$ whenever $\Gamma_{ijs,n-1} = 1$ and $s > t$. Future values of θ_{ijsn} will change because they depend on $\mathbf{a}_{ijsn}$, which in turn depends on past values of budget weights. After computing the new sequence of implied budget weights, evaluate the posterior probability of observing a budget update:

$$\begin{aligned} & \mathbb{P}(\Gamma_{ijtn} = 1 \mid \mathcal{D}_{i,t:T_i}, \mathcal{P}_{i,t:T_i,n} \setminus \{\Gamma_{ijtn}\}, \mathcal{H}_n) \\ &= \psi_{ijn} \prod_{s \geq t} \mathcal{L}_{ijs}(\vartheta_{ijsn}^*(\Gamma_{ijtn} = 1)) \Bigg/ \left[\psi_{ijn} \prod_{s \geq t} \mathcal{L}_{ijs}(\vartheta_{ijsn}^*(\Gamma_{ijtn} = 1)) \right. \\ & \quad \left. + (1 - \psi_{ijn}) \prod_{s \geq t} \mathcal{L}_{ijs}(\theta_{ijs,n-1}(\Gamma_{ijt,n-1} = 0)) \right] \end{aligned} \quad (12)$$

where $\mathcal{L}_{ijs}(\cdot)$ is the aforementioned marginal likelihood function with notational dependencies on data and other parameters suppressed. Note that $\mathcal{L}_{ijs}(\cdot)$ depends on Γ_{ijtn} only via ϑ_{ijtn}^* or $\theta_{ijs,n-1}$. This is the probability of simultaneously realizing Γ_{ijt} and θ_{ijt} , where θ_{ijt} is exactly determined by an equilibrium function of data and other parameters. In a slight abuse of notation we index $\theta_{ijs,n-1}(\Gamma_{ijt,n-1} = 0)$ with $n-1$ even though for any $t < s$, θ_{ijsn} and $\mathbf{a}_{ijsn}$ will each have been re-evaluated and updated previously to ensure that the equilibrium conditions hold *exactly* for the entire sequence at each draw of Γ_{ijtn} for each t . Under Gibbs sampling, with some probability we accept that $\Gamma_{ijtn} = 1$, or we reject it and set $\Gamma_{ijtn} = \Gamma_{ijt,n-1} = 0$, and the budget weight sequence reverts back to that which was previously computed as a result of following the same procedure to sample $\Gamma_{ij,t-1,n}$. Thus, rejecting $\Gamma_{ijtn} = 1$ also amounts to rejecting the entire future sequence $\{\vartheta_{ijsn}^*(\Gamma_{ijtn} = 1)\}_{s \geq t}$. This procedure is repeated, stepping forward in t . In this process θ_{ij1n} is computed once, θ_{ij2n} twice, and θ_{ijtn} t times, once at every draw of Γ_{ijsn} for all $s \leq t$. Further, this process is performed for each j , starting at $j = 1$ and on to $j = J$. Note that, despite the fact that ϑ_{iytn}^* depends on $\vartheta_{i,-ytn}^*$ (other optimal budget re-evaluations) and thus, by extension $\Gamma_{i,-jtn}$, we do not re-evaluate the entire vector ϑ_{itn}^* every time a change is made to Γ_{ijtn} for each particular j . Thus, other values of $\theta_{i,-jtn}$ are treated as parameters and taken as given when re-evaluating θ_{ijtn} . The relevant likelihood function is then only that for commodity-group j under the assumption that the components of ζ_{it} are independent.

5.1.4 Metropolis-Hastings within Gibbs Steps

In the previous subsection we described how the latent budgeting decisions and associated cognitive shocks $\{\theta_{itn}, \mathbf{a}_{itn}, \Gamma_{itn}\}_{t=1}^{T_i}$ are sampled. Given $\{\theta_{itn}, \mathbf{a}_{itn}, \Gamma_{itn}\}_{t=1}^{T_i}$, we now describe the

Metropolis-Hastings (MH) within Gibbs sampling steps for the budgeting and mental accounting parameters $\mathcal{P}_{in} \setminus \{\{\theta_{itn}, \mathbf{a}_{itn}, \Gamma_{itn}\}_{t=1}^{T_i}\}$. We exploit conjugacy where possible.

Here, we start with describing the agent-level draws then move on to describing the sampling statements for the global, hierarchical parameters. Note that the following agent-level parameters are associated with non-conjugate priors, $\{\mathbf{a}_{i1n}, \alpha_{i,J+1,n}, \boldsymbol{\tau}_{in}\}$, while the remaining agent-level parameters, $\{\sigma_{in}^2, \beta_{in}, \mu_{i,\alpha_{i,J+1,n}}, \boldsymbol{\psi}_{in}\}$, can be conditionally sampled directly. In practice, we sample $\{\mathbf{a}_{i1n}, \alpha_{i,J+1,n}, \boldsymbol{\tau}_{in}\}$ all together using a random-walk MH step with the adaptive algorithm described in [Atchade and Rosenthal \(2005\)](#).⁴²

Each agent is associated with an adaptive MH step to conditionally draw the column vector

$$\begin{aligned} & \begin{pmatrix} \mathbf{a}_{i1n} \\ \alpha_{i,J+1,n} \\ \ln \boldsymbol{\tau}_{in} \end{pmatrix} \left| \mathcal{D}_i, \mathcal{P}_{in} \setminus \{\mathbf{a}_{i1n}, \alpha_{i,J+1,n}, \boldsymbol{\tau}_{in}\}, \mathcal{H}_n \propto \mathcal{L}_i(\mathbf{a}_{i1n}, \alpha_{i,J+1,n}, \boldsymbol{\tau}_{in}) \right. \\ & \times \left(\underbrace{\prod_j \phi\left(\frac{a_{ij1n}}{\sigma_{a_{ij1n}}}\right)}_{\text{Priors on } a_{ij1n}} \underbrace{\frac{\beta_{ij,n-1}^{\tau_{ijn}}}{\Gamma(\tau_{ijn}) \sigma_{ij,n-1}^{2(\tau_{ijn}-1)}} \exp\left\{\frac{\beta_{ij,n-1}}{\sigma_{ij,n-1}^2}\right\}}_{\text{Posterior on } \tau_{ijn} \text{ given } \sigma_{ij,n-1}} \right) \\ & \times \underbrace{\phi\left(\frac{\alpha_{i,J+1,n} - \mu_{i,\alpha_{i,J+1,n}}}{10}\right)}_{\text{Prior on } \alpha_{i,J+1,n}} \times \underbrace{\tau_{ijn}}_{\text{Jacobian of Log Transform}} \end{aligned} \quad (13)$$

The adaptation routine is agent-specific, in that each agent's MH step is associated with its own acceptance rate. Let $\boldsymbol{\Lambda}_{in}$ be the column vector of parameters sampled in each agent's MH step, where $\ln \boldsymbol{\tau}_{in}$ is the vector of logged shape parameters. We draw this vector, $\boldsymbol{\Lambda}_{in}$, according to

$$\boldsymbol{\Lambda}_{in} = \begin{pmatrix} \mathbf{a}_{i1n} \\ \alpha_{i,J+1,n} \\ \ln \boldsymbol{\tau}_{in} \end{pmatrix} \sim \mathcal{MN}(\boldsymbol{\Lambda}_{i,n-1}, \chi_{i,n-1} \mathbb{I}_{2J+1}) \quad (14)$$

where $\mathcal{MN}$ denotes the multivariate normal distribution, parameterized in terms of the variance/covariance matrix, which is simply the identity matrix scaled by an adaptive scalar, $\chi_{i,n-1}$, a la [Atchade and Rosenthal \(2005\)](#). Let $\pi_i(\boldsymbol{\Lambda}_{in})$ be a collect-all expression for the prior densities of MH-sampled parameters on the right-hand side of (13).⁴³ Let acc_{in} be the algorithm's acceptance rate for agent i at iterate n , and let $\overline{acc} = 0.234$ be the optimally-targeted acceptance rate.⁴⁴ At

⁴²Through simulations we find that the adaptation algorithm of [Atchade and Rosenthal \(2005\)](#) performs better in our high-dimensional model with substantial non-linearities and heterogeneity than similar adaptive Metropolis routines described in [Haario et al. \(1998, 2001\)](#).

⁴³Note that we draw the log of τ_{ijn} in the Metropolis-Hastings step, so that we must apply the Jacobian of the transform $y = \ln \tau$, i.e. $\tau = e^y$ which has derivative $e^y = \tau$, to the expression for the density function.

⁴⁴See [Roberts et al. \(1997\)](#) for discussion of optimal acceptance rates for MH algorithms.

each iterate, χ_{in} is updated according to:

$$\chi_{in} = \chi_{i,n-1} + \frac{2.5}{n} \left(\min \left\{ 1, \frac{\mathcal{L}_i(\boldsymbol{\Lambda}_{in}) \pi_i(\boldsymbol{\Lambda}_{in})}{\mathcal{L}_i(\boldsymbol{\Lambda}_{i,n-1}) \pi_i(\boldsymbol{\Lambda}_{i,n-1})} \right\} - \overline{acc} \right) \quad (15)$$

We set minimal and maximal values for χ_{in} of 0.0001 and 1000 respectively.

Our prior specifications allow for conditionally conjugate draws of σ_{in}^2 , β_{in} , $\mu_{i,\alpha_{i,J+1},n}$, and ψ_{in} . Given the independence of ζ_{ijt} and $\zeta_{i'j',t}$ we can draw each of the J components of the aforementioned vectors separately. Let $\bar{\zeta}_{ijn}$ be the conditional sample mean of the expenditure error process under the current MCMC draw, and let $\overline{SSQ}_{ijn}$ be the sample sum of squared residuals for good j in the current draw, each defined as follows:

$$\begin{aligned} \bar{\zeta}_{ijn} &= \frac{1}{T_i} \sum_t (x_{ijt} - \theta_{ijt} \ell_{it} - a_{ijt}) \\ \overline{SSQ}_{ijn} &= \sum_t (x_{ijt} - \theta_{ijt} \ell_{it} - a_{ijt} - \bar{\zeta}_{ijn})^2 \end{aligned}$$

In practice we draw the precisions rather than the variances. The conditionally conjugate sampling statements are as follows:

$$\begin{aligned} \left(\frac{1}{\sigma_{ijn}^2} \right) &| \mathcal{D}_i, \mathcal{P}_{in} \setminus \{\sigma_{ijn}\}, \mathcal{H}_n \sim \text{Gamma} \left(\tau_{ijn} + \frac{T_i}{2}, \frac{\beta_{ij,n-1} \overline{SSQ}_{ijn} + 2}{2 \beta_{ij,n-1}} \right) \\ \beta_{ijn} &| \mathcal{D}_i, \mathcal{P}_{in} \setminus \{\beta_{ijn}\}, \mathcal{H}_n \sim \text{Gamma} \left(1 + \tau_{ijn}, \left(1 + \frac{1}{\sigma_{ijn}^2} \right)^{-1} \right) \\ \mu_{i,\alpha_{i,J+1},n} &| \mathcal{D}_i, \mathcal{P}_{in} \setminus \{\mu_{i,\alpha_{i,J+1},n}\}, \mathcal{H}_n \sim \mathcal{N} \left(\left(\frac{\mu_\mu}{\sigma_\mu^2} + \frac{\alpha_{i,J+1,n}}{100} \right) \left(\frac{1}{\frac{1}{\sigma_\mu^2} + \frac{1}{100}} \right), \frac{1}{\frac{1}{\sigma_\mu^2} + \frac{1}{100}} \right) \\ \psi_{ijn} &| \mathcal{D}_i, \mathcal{P}_{in} \setminus \psi_{ijn}, \mathcal{H}_n \sim \text{Beta} \left(1 + \sum_t \Gamma_{ijt}, T_i - \sum_t \Gamma_{ijt} \right) \end{aligned}$$

There are only two global parameters common to all agents — μ_μ and σ_μ^2 , the mean and variance associated with each agent's $\mu_{i,\alpha_{i,J+1},n}$ parameter. Having placed empirical priors on these parameters, updating them is a matter of simply re-evaluating (11). Given substantial autocorrelation in the agent-level draws due to the complex filtration process needed to uncover the latent mental accounting series, we sample the global hyperparameters only once every 400 times through the chain. This means that for each global parameter draw, we proceed through the agent-level sampling steps 400 times for each agent. The idea is to engage in the most computationally efficient sampling scheme without sacrificing inference.

5.1.5 Intuition Regarding Identification of Budget Weight Updates

Since our estimation is of the Bayesian learning variety, different parameters may be identified to varying degrees. The key is to assess whether or not prior distributions imposed by the modeler

appear different than their posterior estimates. Such a statistical assessment of identification for time-varying mental accounting parameters is difficult; θ_{it} and a_{it} are pre-determined by all the other parameters and so are not directly sampled, though they are still determined probabilistically in the sense that the probability of observing a sequence of values depends on the priors, likelihood, and data.

Suppose for sake of argument, θ_{it} and a_{it} were given proper prior distributions and sampled directly, not subject to the constraints on their values imposed by the equilibrium modeling conditions. That is, fix $\gamma_i = 1$ and consider the expenditure relation $x_{ijt} = \theta_{ijt}\ell_{it} + a_{ijt} + \zeta_{ijt}$ strictly in some reduced-form sense, where θ_{ijt} and a_{ijt} are time-varying slopes and intercepts. Then, for any given error variance σ_{ij}^2 , there exists a thick vein of possible pairs of (θ_{it}, a_{it}) that are all equally likely realizations. Under our model structure however, the likelihood of these realizations are restricted by theory. For any atomic parameter draw, excluding the budget weights in category j , where the draw is the set $\mathcal{P}_{in} \setminus \{\{\theta_{ijtn}\}_{t=1}^{T_i}\} \cup \mathcal{H}_n$, the equilibrium condition for ϑ_{iyt}^* and Lemma B.2 in Supplemental Appendix B together ensure that there exists a unique sequence $\{\theta_{ijtn}\}_{t=1}^{T_i}$ that conditionally satisfies the personal mental accounting equilibrium. Thus, if we can sample the parameters in $\mathcal{P}_{in}$ and $\mathcal{H}_n$, then we can conditionally identify a unique sequence of budget weights.

Identifying atomic realizations of the cognitive shocks driving budget updates, $\{\Gamma_{ijtn}\}_{t=1}^{T_i}$, relies on exploiting information regarding changes in the average levels of the expenditure time series, conditional on income, over multiple periods. This is discussed for more general time series level changes in McCulloch and Tsay (1993). Take σ_{ijn}^2 , the likelihood variance, and a_{ij1n} , the initial mental account balance, as given. Also, suppose all future budget-updating indicators, $\{\Gamma_{ijs,n-1}\}_{s>t}$, are taken as given from the previous iteration of the MCMC sampler. Recall further that ζ_{it} is mean zero. Suppose for illustration we observe $x_{ijt} \gg \theta_{ij,t-1,n}\ell_{it} + \sum_{s=1}^{t-1}(\theta_{ijsn}l_{is} - x_{ijs}) + a_{ij1n}$, but after this particular deviation of expenditure we observe $x_{ij,t+s'} \approx \sum_{s=t+s'}^S(\theta_{ij,t-1,n}l_{is} - x_{ijs}) + \sum_{s=1}^{t-1}(\theta_{ijsn}l_{is} - x_{ijs}) + a_{ij1n}$ for some $S \geq t$ and all s' such that $t + s' \leq S$.⁴⁵ That is, observations of future expenditure are only small, idiosyncratic deviations from expected expenditure. Then the Gibbs sampler would probabilistically conclude $\Gamma_{ijtn} = 0$, and no switch is necessary (i.e., $\theta_{ijtn} = \theta_{ij,t-1,n}$). Now suppose instead for most s' , $x_{ij,t+s'} \gg \sum_{s=t+s'}^S(\theta_{ij,t-1,n}l_{is} - x_{ijs}) + \sum_{s=1}^{t-1}(\theta_{ijsn}l_{is} - x_{ijs}) + a_{ij1n}$. This compounds and leads the period- t likelihoods to shift further and further out, so that each successive expenditure realization, $x_{ij,t+s'}$, appears less and less likely to the Gibbs sampler. The sampler thus infers that a permanent budget update has been made (i.e. $\Gamma_{ijtn} = 1$) and computes ϑ_{iytn}^* with $\iota_{iyt} = j$ according to the equilibrium condition. Aggregating over all MCMC draws, n , we can thus identify periods in the data that feature higher probabilities of optimal budget updates, conditional upon our modeling assumptions being true. Joint variation in expenditure, income, and balances, via both the spending constraints and the inversion formula, thus inform the probabilistic identification of cognitive shocks leading to budget

⁴⁵Note that it is possible that $s' = 0$ if $\Gamma_{ij,t+1,n-1} = 1$. Further, if S' is the next period at which $\Gamma_{ijs',n-1} = 1$, then $S = S' - 1$, so that it is possible that $S = t$.

updates.

5.2 Model Fitness

Having described the estimation apparatus, we now assess whether the disciplined model can reproduce observed expenditure patterns — a prerequisite for its use as a functioning quantitative framework. We estimate several different model specifications using the MCMC algorithm and compare the fit of each to understand the degree to which assumptions regarding numeracy, budgeting frictions, and mental accounting affect predictive inference. With respect to numeracy, we consider six different thresholds, depending on whether we assume numeracy is constrained along the relative (percentage) or absolute (dollar) margins:⁴⁶

- i. Along the relative margin, we estimate four models with thresholds $\epsilon \in \{0, 0.01, 0.05, 0.1\}$.
- ii. Along the absolute margin, we estimate two models with thresholds $\epsilon_\ell \in \{0.01, 1\}$.

In the baseline model where $\epsilon = \epsilon_\ell = 0$, we count estimated budget updates anytime the candidate is such that $\vartheta_{iyt}^* \neq \theta_{i,t_{iyt},t-1}$. Here, in the main text we present estimation results from both the baseline and \$1-threshold models, which is the best-fitting model according to the information criterion we employ.⁴⁷

In addition to estimating six versions of the full model with different assumptions regarding numeracy, we also estimate several different models where different behavioral features are turned off. Specifically, we estimate the baseline model ($\epsilon = \epsilon_\ell = 0$) and the \$1-threshold model ($\epsilon_\ell = 1$) under the following separate assumptions:

- i. Heterogeneous rationality $\psi_{ij} \in (0, 1)$ and full anchoring, $\gamma_i = 1$.
- ii. No anchoring, $\psi_{ij} \in (0, 1)$ and $\gamma_i = 0$.
- iii. Constant budget weights, $\Gamma_{ijt} = 0$ and $\gamma_i = 1$.
- iv. Constant budget weights with no anchoring, $\Gamma_{ijt} = 0$ and $\gamma_i = 0$.
- v. No budget-updating frictions, $\Gamma_{ijt} = 1$ and $\gamma_i = 1$.
- vi. Initial mental account balances are always zero, $a_{ij,t_{i0}} = 0$, where t_{i0} is the first observation in the sample for agent i .

Model (v) is designed to best mimic a classical two-stage budgeting problem with ex-ante preference uncertainty and constant expenditure share preference weights. Alternative models (iii) and

⁴⁶Recall that, for example, $\epsilon = 0.01$ implies that optimal budget updates are only recorded as such if the absolute difference between current and candidate expenditure shares is $|\vartheta_{iyt}^* - \theta_{i,t_{iyt},t-1}| > \epsilon$. For simulations where budget updates are bounded along the absolute margin, budget updates are only recorded as such if the absolute difference is $\ell_{it} |\vartheta_{iyt}^* - \theta_{i,t_{iyt},t-1}| > \epsilon_\ell$.

⁴⁷Including the \$1-dollar threshold model amongst our primary results is also a nod to the literature suggesting consumers exhibit so-called “left digit biases” (Lacetera et al., 2012; Sokolova et al., 2020)

(iv) are identical across numeracy constraints since we are assuming both the extensive and intensive optimization margins are non-binding always (i.e., both $\Gamma_{ijt} = 0$ and $\underline{\Gamma}_{ijt} = 0$). Model (vi) is included to assess how much our prior distributional assumptions for the initial mental account state variable drive the results (they do not).

In Supplemental Appendix D.1 we describe fitness summary statistics and present a table that compares all of the models by numeracy threshold and the various assumptions on their behavioral parameters. We use an information criterion, comparing predictive densities across models a la Gelman et al. (2013), to assess model fitness. Note that all fully-parameterized models ($\psi_{ij} \in (0, 1)$ and $\gamma_i = 1$) with non-zero numeracy constraints are marginally better fits than the baseline model. Among the fully-parameterized models, the \$1-threshold model achieves the largest information gain on the baseline model. This result conforms with theory and extensive evidence that individuals are often constrained by integer representations when making numerical decisions (Miller, 1956; Simon, 1957; Cowan, 2000; Leslie et al., 2008). Given these fitness results, we focus most of our discussion on results generated by two models: 1) the baseline model with no numeracy constraints and 2) the \$1-threshold model, each with all of the mental accounting features (i.e., $\psi_{ij} \in (0, 1)$ and $\gamma_i = 1$). Both models allow for anchoring and heterogeneous rationality.




Figure 2: Black lines represent actual weekly spending for the median income agent. Red lines represent the baseline model’s predicted means with 95% confidence regions in pink. The blue lines represent the \$1-threshold model’s predicted means with 95% confidence regions in light blue. Notice that predictions from both models are almost identical.

To visually illustrate how well both the baseline model and \$1-threshold model fit the data, Figure 2 presents the time series of expenditures (black lines) along with predictive means (red and blue lines) from the two fully-parameterized estimations with $\psi_{ij} \in (0, 1)$ and $\gamma_i = 1$ for a selected consumer whose average weekly income is closest to the sample median. Mean weekly

expenditure in each category is as follows: groceries \$60.67, auto/gasoline \$32.50, food away from home \$21.98, and other expenditure \$357.00. This consumer earns an average of \$460.05 per week and maintains an average weekly card balance of \$115.21 after spending. Notice that weekly expenditure patterns are spiky yet weakly persistent — a pattern we observe for many consumers in our sample. Consequently, explaining these patterns requires a demand model, formalized by theory, with features that can endogenously explain spending spikes, persistence, and trend decline or growth in individual consumption expenditure series. The model we have developed can predict all of these features: spending spikes result from both endogenous mental accounting and exogenous taste shocks; persistence is a function of budgetary bounded rationality and mental accounting; and changes to average spending and trends are associated with budget updates. That the model captures these salient patterns establishes it as an operational platform for structured inquiry — one that can be used to examine latent behavioral mechanisms and counterfactual scenarios that the expenditure data alone cannot reveal.

Summary statistics over agent-level means for parameter estimates are presented in Table 2 for both the baseline model sans numeracy threshold and the \$1-threshold model.⁴⁸ Amongst the mental accounting parameters, the variances of the posterior distributions of agent-level mental account balance means, $\hat{a}_i$, are large. Some of the likelihood variance estimates, $\hat{\sigma}_{ij}^2$, are also large. These high values result from behavior whereby a consumer goes from very few (or no) purchases one week to high expenditure values the next. Such high values are not surprising given the spikiness of the various time series. Remember, $a_{ij,t+1} = -\zeta_{ijt}$, so high variance in ζ_{ijt} will lead to high variance in $\hat{a}_{ij}$ as well. This variation propagates through the model via optimal budget updates, ϑ_{iyt}^* . Lacking explicit budgeting data, this posterior variability is to be expected. Further, though not presented here, estimates at the fifth and ninety-fifth percentiles are reasonable for both models.

5.3 Budget-updating Behavior

With the calibrated model in hand, we now interrogate its quantitative implications for consumer behavior, beginning with the frequency and nature of budget updates. Figure 3 provides visual summaries of the estimated total number of budget switches $k_{itn} = \sum_{j=1}^J \Gamma_{ijtn}$ under different numeracy thresholds. The baseline model implies an average of 2.48 budget updates every week. When the absolute threshold is set at $\epsilon_\ell = 0.01$, so that we encode budget updates only when the absolute difference in old versus new budgets exceeds one cent, the average number of updates is 2.51 each week. Estimates of $\hat{k}_i$ between the baseline model and the model with $\epsilon_\ell = 0.01$ are not significantly (statistically) different from each other. When the \$1-threshold model is considered, the implied average falls to 2.11 updates per week. These numbers fall when we consider different minimum relative thresholds — 1.45 when $\epsilon = 0.01$, 0.89 when $\epsilon = 0.05$, and 0.75 when $\epsilon = 0.1$. With a \$1 threshold (teal), full rationality ($k_{it} = J = 4$) obtains only 10% of the time. Even without

⁴⁸For parameters that are time dependent, we average both across MCMC draws and over time for each agent. For the global parameters, we simply take summary statistics over the outputted chain.

Table 2: Posterior Summary Statistics

Description	Parameter	Baseline Model $\epsilon = \epsilon_\ell = 0, \psi_{ij} \in (0, 1), \& \gamma_i = 1$		\$1-threshold Model $\epsilon_\ell = 1, \psi_{ij} \in (0, 1), \& \gamma_i = 1$	
		Mean	S.D. ^e	Mean	S.D. ^e
<i>Borrowing Limit</i>	m_i	5,311.618	2,686.200	5,311.618	2,686.200
<i>Long-run</i>	$\hat{\alpha}_{i1}^a$	0.106	0.064	0.106	0.064
<i>Expenditure</i>	$\hat{\alpha}_{i2}$	0.146	0.088	0.146	0.088
<i>Share</i>	$\hat{\alpha}_{i3}$	0.150	0.097	0.150	0.097
	$\hat{\alpha}_{i4}$	0.598	0.148	0.598	0.148
<i>Likelihood</i>	$\hat{\sigma}_{i1}$	29.346	24.325	34.604	200.396
<i>Standard</i>	$\hat{\sigma}_{i2}$	45.774	54.622	54.062	188.935
<i>Deviation</i>	$\hat{\sigma}_{i3}$	63.186	64.446	73.899	251.622
	$\hat{\sigma}_{i4}$	245.342	2,031.048	364.204	1,785.711
<i>Shape Parameter</i>	$\hat{\tau}_{i1}$	0.746	2.537	3.90×10^{10}	1.96×10^{13}
<i>for Likelihood</i>	$\hat{\tau}_{i2}$	0.978	3.305	2.18×10^{21}	1.09×10^{23}
<i>Variance</i>	$\hat{\tau}_{i3}$	0.903	3.625	0.353	1.997
<i>Prior</i>	$\hat{\tau}_{i4}$	0.566	2.009	1.29×10^{18}	6.44×10^{19}
<i>Scale Parameter</i>	$\hat{\beta}_{i1}$	1.737	2.501	4.657	167.197
<i>for Likelihood</i>	$\hat{\beta}_{i2}$	1.968	3.266	9.682	417.278
<i>Variance</i>	$\hat{\beta}_{i3}$	1.899	3.620	1.350	1.990
<i>Prior</i>	$\hat{\beta}_{i4}$	1.566	2.011	1,626.734	81,422.770
<i>Liquidity Preferences</i>	$\hat{\alpha}_{i,j+1}$	12.178	3.388	11.488	10.442
<i>Mean of $\hat{\alpha}_{i,j+1}$</i>	$\hat{\mu}_{\alpha_{i,j+1}}$	15.315	0.018	13.251	1.004
<i>Probability of</i>	$\hat{\psi}_{i1}$	0.667	0.194	0.545	0.207
<i>Budget</i>	$\hat{\psi}_{i2}$	0.644	0.197	0.545	0.201
<i>Update</i>	$\hat{\psi}_{i3}$	0.590	0.195	0.517	0.188
	$\hat{\psi}_{i4}$	0.536	0.231	0.515	0.227
<i>Budgeted Share</i>	$\hat{\theta}_{i1}^b$	0.204	2.005	0.114	0.108
<i>of Income</i>	$\hat{\theta}_{i2}$	0.330	4.495	0.156	0.146
<i>Chosen by</i>	$\hat{\theta}_{i3}$	0.451	6.954	0.171	0.168
<i>the Consumer</i>	$\hat{\theta}_{i4}$	2.188	41.125	0.708	0.295
<i>Cognitive and</i>	$\hat{\Gamma}_{i1}^c$	0.694	0.201	0.566	0.215
<i>Numeracy</i>	$\hat{\Gamma}_{i2}$	0.669	0.204	0.566	0.209
<i>Frictions w.r.t.</i>	$\hat{\Gamma}_{i3}$	0.613	0.202	0.537	0.195
<i>Budget Updates</i>	$\hat{\Gamma}_{i4}$	0.557	0.239	0.536	0.236
<i>Mental</i>	$\hat{a}_{i1}$	-0.182	35.937	3.663	305.530
<i>Account</i>	$\hat{a}_{i2}$	0.303	32.942	7.228	279.077
<i>Balances</i>	$\hat{a}_{i3}$	-2.942	65.058	6.223	372.457
	$\hat{a}_{i4}$	-158.191	8,548.589	211.587	2,594.666
<i>Standard</i>	$\hat{\sigma}_{a_{i1}}$	0.013	0.007	0.013	0.007
<i>Deviation</i>	$\hat{\sigma}_{a_{i2}}$	0.015	0.009	0.015	0.009
<i>for Mental</i>	$\hat{\sigma}_{a_{i3}}$	0.018	0.011	0.018	0.011
<i>Accounts</i>	$\hat{\sigma}_{a_{i4}}$	0.034	0.023	0.034	0.023
<i>Global Mean of $\hat{\mu}_{\alpha_{i,j+1}}$</i>	$\hat{\mu}_\mu^d$	14.844	3.395	13.074	7.117
<i>Global S.D. of $\hat{\mu}_{\alpha_{i,j+1}}$</i>	$\hat{\sigma}_\mu$	0.501	0.117	3.512	0.709

^a Second subscript corresponds to commodity group — groceries ($j = 1$), auto/gasoline ($j = 2$), food away from home ($j = 3$), and other expenditure ($j = 4$).

^b For time-dependent parameters, θ_{ijt} , Γ_{ijt} , and a_{ijt} , we average over both posterior draws and time for each agent and each commodity category. The statistics presented are summary statistics over these agent-level averages.

^c The agent-level means for the budget-updating indicators are taken from the \$1-threshold model where $\Gamma_{ijt} = 1$ if $\ell_{it} |\theta_{ijt} - \theta_{ij,t-1}| > \kappa_\ell = 1$.

^d $\hat{\mu}_\mu$ and $\hat{\sigma}_\mu$ are the only two global parameters. They are summary statistics over agent-level $\hat{\mu}_{\alpha_{i,j+1}}$. From simulations, we found that these parameters are needed to induce posterior shrinkage.

^e This is the standard deviation of the posterior distribution *not* the standard error of the posterior mean.




Figure 3: In panel (a) we present the posterior density of $\hat{k}_i = \frac{1}{NT_i} \sum_n \sum_t \sum_j \Gamma_{ijtn}$, i.e. the number of budget changes each period, which satisfy both attentiveness and numeracy constraints, averaged over time and across MCMC draws for each agent, using a smoothed Gaussian kernel with the Silverman rule-of-thumb for choosing the bandwidth (Silverman, 1986). In panel (b) we show the marginal posterior probability mass function for the atomic posterior draws (epochs) of k_{itn} . Colors denote different thresholds for judging when a budget update has occurred. In the baseline scenario we assume that any time $\theta_{ijtn} \neq \theta_{ij,t-1,n}$ a budget update has occurred, while the other scenarios correspond to those associated with numeracy thresholds $\epsilon \in \{0.1, 0.05, 0.01\}$ or $\epsilon_\ell \in \{1, 0.01\}$.

allowing for numeracy constraints, just over 80% of consumer-week combinations exhibit some degree of bounded rationality — a pattern consistent with the presence of financial-attentiveness frictions under the model’s assumptions.

Raising the numeracy threshold from zero to \$1 leads to approximately 0.34 fewer updates, on average, per week across our sample of consumers. When taking into consideration the relative thresholds, that value climbs to 0.95 fewer updates when $\epsilon = 0.01$, 1.54 fewer updates when $\epsilon = 0.05$, and 1.69 fewer updates when $\epsilon = 0.1$. Overall, total budget updates across consumers and over time drop 14.9% when the \$1 threshold is implemented. Under the relative thresholds, the model yields a 41.8% drop when $\epsilon = 0.01$, a 64% drop when $\epsilon = 0.05$, and a 70% drop when $\epsilon = 0.1$.

Broadly, the model implies that few consumers update all of their budgets every period, a pattern consistent with binding cognitive constraints and narrow bracketing. At the same time, these constraints do not appear to force agents to stick with budgets for long extended periods, since on average they still make a few updates every week. Thus, while agents do not consistently update their budgets every week, neither do they neglect these budgets over an extended amount of time.

5.4 Classifying Consumers By Budgetary Responsiveness

A further advantage of the model as a working quantitative apparatus is that it generates latent state variables that permit richer consumer classifications than expenditure data alone would allow. In this section we explore how expected total spending (or expected savings) and total realized spending (or realized savings) depend on total over- or under-expenditure from the previous period, $a_{i,J+1,t}$. Recall that over-saving in period $t - 1$ implies $a_{i,J+1,t} < 0$, not positive, so that over-spending in $t - 1$ is associated with $a_{i,J+1,t} > 0$. Consumers can respond to $a_{i,J+1,t}$ ex-ante by adjusting their budget shares and thus their expected *total* expenditure. While our modeling approach does not take an explicit stand on how consumers optimally choose their actual ex-post expenditure after having ex-ante set their optimal budgets, we can still examine the degree to which actual expenditure varies in the estimated state variable $a_{i,J+1,t}$.

We type-cast consumers based on the probability we observe a budget (ex-ante) or expenditure (ex-post) increase following over/under spending in the previous period (i.e., conditional upon the sign of $a_{i,J+1,t}$). Let $\hat{\mathbb{P}}_i(\cdot|\cdot)$ be a posterior-estimated conditional probability function. Let $\text{sgn}\{y\}$ be the sign function that describes whether or not the variable y is negative. The estimated probabilities we consider, depending on whether we are casting ex-ante or ex-post types, are $\hat{\mathbb{P}}_i\left(\text{sgn}\left\{\sum_j(\omega_{ijt} - \omega_{ij,t-1})\right\} \middle| \text{sgn}\{a_{i,J+1,t}\}\right)$ and $\hat{\mathbb{P}}_i\left(\text{sgn}\left\{\sum_j(x_{ijt} - x_{ij,t-1})\right\} \middle| \text{sgn}\{a_{i,J+1,t}\}\right)$. We can classify consumers into four broad types based on their average behavior.^{49,50}

- i. (*Budget prioritizers* — negative correlation patterns) Those who respond to over-spending in $t - 1$ by reducing expected (or realized) spending in t and under-spending in $t - 1$ by increasing expected (or realized) spending in t . These consumers use mental accounting as a correction mechanism in the original spirit described in [Thaler \(1985\)](#), with the goal of balancing out prior over- or under-spending by spending less or more in the next period, so that their budgets remain fairly stable and disciplined, overall. These are consumers for whom the posterior observed probabilities are such that

$$\begin{aligned} \hat{\mathbb{P}}_i\left(\sum_j(\omega_{ijt} - \omega_{ij,t-1}) < 0 \middle| a_{i,J+1,t} > 0\right) \\ = \frac{\sum_n \sum_t \mathbf{1}\{\sum_j(\tilde{\omega}_{ijtn} - \tilde{\omega}_{ij,t-1,n}) < 0\} \mathbf{1}\{\tilde{a}_{i,J+1,tn} > 0\}}{\sum_n \sum_t \mathbf{1}\{\tilde{a}_{i,J+1,tn} > 0\}} > 0.5 \end{aligned} \quad (16)$$

$$\begin{aligned} \text{and } \hat{\mathbb{P}}_i\left(\sum_j(\omega_{ijt} - \omega_{ij,t-1}) > 0 \middle| a_{i,J+1,t} < 0\right) \\ = \frac{\sum_n \sum_t \mathbf{1}\{\sum_j(\tilde{\omega}_{ijtn} - \tilde{\omega}_{ij,t-1,n}) > 0\} \mathbf{1}\{\tilde{a}_{i,J+1,tn} < 0\}}{\sum_n \sum_t \mathbf{1}\{\tilde{a}_{i,J+1,tn} < 0\}} > 0.5 \end{aligned} \quad (17)$$

⁴⁹In what follows “realized” refers to ex-post types inferred from observed spending (not unobserved budgeting) behavior.

⁵⁰To some degree we engage in similar exercises to [Andre et al. \(2025\)](#) by type-casting consumers based on their behavior. Where our endeavors depart from theirs, however, is that they simply look at how people spend and save out of transitory income shocks, whereas we examine behavioral responsiveness to shocks conditional upon estimated, latent budgeting patterns.

Together, the signs of (16) and (17) fully characterize a “budget prioritizer” (ex-ante type i). Replace $\tilde{\omega}$ ’s in (16) and (17) with actual x ’s (i.e., expenditure *data*) and you have the statistics that summarize a “spending prioritizer” (ex-post type i). Both (16) and (17) must hold for a consumer to be classified as a budget and/or spending prioritizer.

- ii. (*Spending followers* — positive correlation patterns) Those who respond to over-spending in $t - 1$ by increasing expected (or realized) spending in t and under-spending in $t - 1$ by reducing expected (or realized) spending in t . These consumers engage in persistent over-spending and under-spending over successive periods. When updating budgets, they often update them to be more in line with their prior spending, rather than the other way around. The ex-ante probabilities which must hold are $\hat{\mathbb{P}}_i\left(\sum_j(\omega_{ijt} - \omega_{ij,t-1}) > 0 \middle| a_{i,J+1,t} > 0\right) > 0.5$ and $\hat{\mathbb{P}}_i\left(\sum_j(\omega_{ijt} - \omega_{ij,t-1}) < 0 \middle| a_{i,J+1,t} < 0\right) > 0.5$ (and the same for ex-post).⁵¹
- iii. (*Spendthrifts*) Those who respond to both over- and under-spending in $t - 1$ by increasing expected (or realized) spending in t . These consumers use under-spending as an excuse to spend more in the next period (i.e., use their surplus, implicit savings) and over-spending as a reason to increase their future budgets rather than decrease their spending. The ex-ante probabilities which must hold are $\hat{\mathbb{P}}_i\left(\sum_j(\omega_{ijt} - \omega_{ij,t-1}) > 0 \middle| a_{i,J+1,t} > 0\right) > 0.5$ and $\hat{\mathbb{P}}_i\left(\sum_j(\omega_{ijt} - \omega_{ij,t-1}) > 0 \middle| a_{i,J+1,t} < 0\right) > 0.5$ (and the same for ex-post).
- iv. (*Frugal savers*) Those who respond to both over- and under-spending in $t - 1$ by decreasing expected (or realized) spending in t . These consumers use under-spending as a motivator to save more (i.e., continue spending less) in the next period and respond to over-spending by reducing their future budget.⁵² The ex-ante probabilities which must hold are $\hat{\mathbb{P}}_i\left(\sum_j(\omega_{ijt} - \omega_{ij,t-1}) < 0 \middle| a_{i,J+1,t} > 0\right) > 0.5$ and $\hat{\mathbb{P}}_i\left(\sum_j(\omega_{ijt} - \omega_{ij,t-1}) < 0 \middle| a_{i,J+1,t} < 0\right) > 0.5$ (and the same for ex-post).

We can also characterize the joint distribution of consumers across both how their ex-ante budgeting behavior and ex-post spending behavior relate to $a_{i,J+1,t}$. This more granular type casting yields $4^2 = 16$ different consumer types. In what follows we present the marginal distributions of consumers pertaining to their particular ex-ante and ex-post type, analyzing how income and savings vary across types to understand if any systematic relationships exists. Then, we repeat the exercise examining the joint distribution of ex-ante and ex-post types. Finally, we focus on the spending and savings patterns for the subset of consumers who are of different ex-ante and ex-post types in order to understand why, between budgeting and spending, type-inconsistent behavior occurs.

⁵¹The sample counterparts of all probabilities are computed in the manner as (16) and (17). We avoid the algebraic details here to save space.

⁵²Type (iii) and (iv) consumers fall on the so-called “spendthrift” and “tightwad” spectrum proposed in Rick et al. (2008). Spendthrifts lack in self-regulation, experiencing low average dis-utility with respect to over-spending, while tight wads are overly sensitive to over-spending.

Table 3: Marginal Distributions of Ex-ante and Ex-post Types ^a

<i>Expected Expenditure</i>				
	Type			
Model ^b	(i)	(ii)	(iii)	(iv)
Baseline	0.826	0.042	0.075	0.058
\$1-threshold	0.787	0.057	0.096	0.060
<i>Realized Expenditure</i>				
	Type			
Model ^b	(i)	(ii)	(iii)	(iv)
Baseline	0.110	0.305	0.425	0.160
\$1-threshold	0.104	0.318	0.468	0.110

^a Each row sums to unity.

^b These are the fully-parameterized models with the corresponding numeracy thresholds.

5.4.1 Quantifying Ex-ante and Ex-post Types

Table 3 presents the marginal distributions of types, conditional on the numeracy threshold and whether we care about ex-ante (top half) or ex-post (bottom half) behavior. Ex-ante, most consumers (> 75%) are classified as type (i) under the model, exhibiting aggregate budgeting behavior consistent with what mental accounting theory would predict. However, the model’s implications with respect to ex-post spending are more mixed. In fact, the model classifies a plurality of consumers as spendthrifts (iii), who are more likely to increase spending regardless of whether they previously over- or under-spent. The next most common group of spenders are type (ii), responding to over-spending by raising expenditure and under-spending by reducing expenditure. Type (ii)’s are thus habitual (serial) over- or under-spenders.

Table 4 shows the joint distributions of ex-ante and ex-post types. Not surprisingly, the model places a plurality of consumers as ex-ante type (i) and ex-post type (iii) (> 37%). Under this classification, many consumers exhibit behavior consistent with using mental accounting heuristics to inform their budgeting decisions but then having difficulty disciplining their actual spending to conform to their ex-ante budgets. This group exhibits patterns consistent with so-called “planning fallacy syndrome,” where they make plans with overly optimistic expectations for their future behavior but mostly fail to actually achieve their budgeting goals (Kahneman and Tversky, 1979; Ülkümen et al., 2008; Peetz and Buehler, 2009; Buehler et al., 2010). Next, > 22% of consumers are classified by the model as type (i) ex-ante and type (ii) ex-post, so that they respond to over-spending by reducing their budget according to standard mental accounting theory but then, nonetheless, increase spending anyway. However, what separates these consumers from

relatively un-disciplined spend-thrifts is that after under-spending they raise budgets but still maintain ex-post expenditure discipline. These consumers thus exhibit patterns consistent with “planning fallacy syndrome” only after over-spending.

Table 4: Joint Distributions of Types ^a

<i>Baseline</i> ^b				
Ex-post Type				
Ex-ante Type	(i)	(ii)	(iii)	(iv)
(i)	0.102	0.220	0.375	0.128
(ii)	0.004	0.017	0.012	0.008
(iii)	0.001	0.035	0.030	0.008
(iv)	0.003	0.032	0.007	0.016
<i>\$1 Threshold</i> ^b				
Ex-post Type				
Ex-ante Type	(i)	(ii)	(iii)	(iv)
(i)	0.093	0.223	0.382	0.089
(ii)	0.005	0.018	0.026	0.007
(iii)	0.003	0.043	0.046	0.004
(iv)	0.003	0.034	0.014	0.010

^a Each matrix sums to unity.

^b These are the fully-parameterized models with the corresponding numeracy thresholds.

5.5 Counterfactually Relaxing Rationality Constraints

To explore what the model implies about consumer welfare when the sparse-max constraint on budget updates is relaxed, we counterfactually simulate consumer expenditure using the posterior distribution of parameters, except we force $\Gamma_{ijt} = 1$ for all i, j , and t combinations.⁵³ We perform these simulations for both the baseline model with no numeracy constraints and the \$1-threshold model. Specifically, the reader can think of this counterfactual as what might happen if a consumer had a financial planner or budgeting application constantly nudging them to engage in optimal behavior, so that they always check their budgets for all categories every week, while still being subjected to the same sequence of expenditure/budgeting shocks, $\{\zeta_{it}\}_t$, we back out

⁵³This exercise is equivalent to setting the primitive preference parameter $\psi_{ij} = 1$, for all i and j .

of the estimator.⁵⁴ In the counterfactual setup each period every agent chooses ϑ_{it}^* by maximizing indirect utility subject to the relaxed constraint that $k_{it} = J$ always. When the \$1 threshold is present, budget updates will always occur in every category unless the numeracy constraint is not satisfied. These simulations isolate the effect of cognitive inattentiveness, independent of numeracy constraints, thus allowing us to understand how cognitive inattentiveness may impact individual economic outcomes and welfare.

For 68.4% of consumers (baseline) and 68.2% (\$1 threshold), the model implies that engaging in weekly budget updates for every commodity category would *not* lead to welfare improvements relative to the posterior predictive baseline. However, as we have so far emphasized, the model generates substantial heterogeneity with respect to how different individual consumers' behaviors change when rationality constraints are relaxed. In fact the model implies that those who are most constrained with respect to budgetary attentiveness are most vulnerable to experiencing adverse outcomes due to interventions designed to increase financial attentiveness by sending $k_{it} \rightarrow J$ every period. Thus, like the findings of Koch and Nafziger (2020) suggest, for some consumers narrow choice bracketing can be disciplinary.

While the model implies that most consumers (>65%) have lower overall utility relative to the predictive baseline when rationality constraints are relaxed, a small fraction engage in such extreme behavior that their savings plus account balances approach the borrowing limit, i.e. $z_{it} + r_t b_{it} \rightarrow -m_i$. We classify this last group of consumers as bankrupt. In the baseline model < 1% of all consumers go bankrupt, while this number jumps to 3.3% when the \$1 threshold is imposed. For these consumers setting $\Gamma_{ijt} = 1$ leads to over attentiveness in response to extreme expenditure shocks; thus, the degree to which consumers experience adverse outcomes when overly attentive is also a function of the variance of their personal shock processes.

Figure 4 presents kernel density estimates of predictive $\hat{k}_i$ (not the counterfactual values) conditional on consumers being “Better Off,” “Worse Off, Not Bankrupt,” and “Worse Off, Bankrupt” due to the relaxation of the rationality constraint. On the whole there is no statistically significant difference between the pre-counterfactual $\hat{k}_i$ -distributions of consumers who are “Better Off” and “Worse Off, Not Bankrupt” in either the baseline or \$1-threshold model. Yet, differences between both bankrupt consumers and all others are significant at the 1% level in both models. Bankrupt consumers thus have significantly different underlying behavioral profiles than either of the other groups. Of those who go bankrupt, they are associated with lower estimated $\hat{k}_i$ — an average of 1.25 updates per week (both models) compared to 2.48 (baseline) and 2.11 (\$1 threshold) updates across the whole sample — as well as on average higher shock variances. These consumers thus are relatively stickier budgeters than their peers, and the counterfactual results suggest that,

⁵⁴Note that this experiment is not the same as the estimation where we force $\psi_{ij} = 1$ and estimate the additional parameters separately. Here, we use the parameter estimates from the estimations where ψ_{ij} is a free parameter and $\gamma_i = 1$, counterfactually *only* assuming the distribution of Γ_{ijt} changes, not the other model primitives. Specifically, we compare the counterfactual predictive utility against the posterior predictive utility from the estimated model, as opposed to actual utility under the posterior parameterization. The reason for this is that we care about how model *predictions* change, given the parameter change. The exercise amounts to assessing the total variation in model-predicted outcomes under coarse variation in ψ_{ij} .




Figure 4: This figure features the density of $\hat{k}_i$ conditional on the counterfactual type of consumers. Those who are “Better Off” derive higher utility from being more budget attentive, while those who are “Worse Off” derive lower utility. Consumers who are “Worse Off” and go “Bankrupt” constitute a small subset of those who are “Worse Off,” overall.

within the model, their budget stickiness functions as budgeting discipline in conjunction with their more volatile spending patterns.

Of the counterfactually-bankrupt consumers, a plurality appear to be ex-ante type (i) *budgeters* who respond to over-spending by reducing their budgets and under-spending by increasing them. Most, however, are type (ii) *spenders*. The budgeting/spending marginal types and cross types for the “Better Off” and “Worse Off, No Bankrupt” groups are largely the same as the aggregate distributions of types. Thus, under the model’s classification, “Bankrupt” consumers are more likely to engage in full type-switching behavior between budgeting and spending, wanting to be true mental accounters when budgeting, though engaging in *anti*-mental accounting behavior when spending. Within the model, this joint budgeting/spending behavior causes excessive borrowing and has the effect of reducing bank balances and thus available liquidity over the long run.

Let us now illustrate the mechanics of how someone who is naturally less financially attentive than the average person but who also tries to use mental accounting to discipline budgets (ex-ante type (i)) but fails to do so (ex-post type (ii)) goes bankrupt. An ex-post type (ii) consumer responds to over-spending in period t by over-spending again in period $t + 1$. But, since this individual is ex-ante type (i), they reduce their expenditure budget from t to $t + 1$ in response to having over-spent. Having over-spent again in $t + 1$, the consumer wants to try to make up for this in $t + 2$, and so they again reduce the spending budget. But, type (ii) spending behavior dominates, so that the consumer over-spends again. After several periods the budget has been reduced so much

that as an anchoring point it becomes effectively meaningless. This is a direct consequence of the fact that the consumer’s budget is being re-evaluated every period and thus ratcheted downward in response to over-spending, ostensibly in an effort to correct the over-spending behavior. In an environment in which budgets are sticky between periods, there is time between budget re-evaluations for the consumer to realize a favorable expenditure shock that forces spending below the budget, thus breaking the over-spending cycle undertaken by their ex-post type-(ii) self.

Under the model’s assumptions, these results carry implications for the design of financial-planning and budgeting nudges (e.g. via consumer finance apps). If nudges are overly salient and consumers can readily change their budgets without incurring a penalty, the model implies that certain types of consumers may engage in overly attentive behavior that actually induces *more* frequent periods of over-spending as a direct consequence of the expenditure shock process they face. At the extreme, this behavior would lead to the model-implied bankruptcies described above, but can still make other consumers who do not go bankrupt worse off. Nudge, presumably designed to facilitate financial discipline, could thus have the opposite effect. Certain consumers may make hasty short-term decisions in response to the salient push-notification stimulus, which has a dynamic, adverse impact on long-run budgeting goals, just like the departures from full rationality described in [Bordalo et al. \(2022\)](#). The model thus generates the counterfactually testable hypothesis (within its structural context) that consumers who use sticky budgets as rules of thumb to regulate spending patterns are most vulnerable to adverse outcomes when such budgets can be easily changed. Future empirical work using data from (e.g.) financial-planning apps, where budgeting behavior and attentiveness can be directly observed, could directly test this model prediction in the data.

6 Conclusion

We return to the questions that motivated our research. First, through the lens of our model, observed expenditure patterns are consistent with budget rigidities (i.e., extended periods when budgets remain stable), which we interpret as a consequence of bounded rationality. Second, the model implies substantial heterogeneity in the degree to which consumers appear boundedly rational. Some consumers’ expenditure patterns are consistent with mostly fungible treatment of money, while others show signs of mental accounting. Our model endogenously generates idiosyncratic, short-term variation in agent-level consumption expenditure patterns. Third, by allowing for narrow choice bracketing and reference dependence, the model generates behavioral inferences with respect to budgeting and anchoring.

An important implication of our model is that most consumers are neither fully rational in the traditional sense (i.e., no cognitive constraints; regularly updating all of their budgets) nor fully behavioral (i.e., fully bounded; never updating any of their budgets). Rather, the model places most consumers somewhere in between: they update some, but not all, of their budgets every period. This idea, that consumers are rational to some extent, provides a more nuanced model

of consumer decision-making than ‘all-or-nothing’ theories that assume either full-rationality or little-to-no rationality. This finding is consistent with recent evidence in [Andre et al. \(2025\)](#), who show that near-rational “quick-fixing” heuristics in consumption carry small utility costs, suggesting that the degree of bounded rationality our model implies is economically plausible.

Finally, we see several avenues for future work. While our results are model-dependent, the proliferation of financial-planning apps to which consumers link all of their expenditure and savings accounts, from credit cards, checking accounts, and investment accounts, should prove useful to researchers looking to study consumer spending and budgeting in more detail. Future work should consider using data from financial-planning apps where budgeting and attentiveness, like say through app log-ins, can be explicitly measured in order to assess our latent inferences against directly observed behavior. Such apps could also prove useful for engaging in field experiments to explicitly test how individuals respond to different kinds of notifications and to then understand how their responses are related to other aspects of their financial behavior.

References

- Johannes Abeler and Felix Marklein. Fungibility, labels, and consumption. *Journal of the European Economic Association*, 15(1), 2017.
- Johannes Abeler, David Huffman, and Collin Raymond. Incentive complexity, bounded rationality, and effort provision. *American Economic Review*, 115(12):4404–4437, 2025.
- S. Rao Aiyagari. Uninsured idiosyncratic risk and aggregate saving. *The Quarterly Journal of Economics*, 109(3):659–684, 1994.
- Carlos Alos-Ferrer and Michele Garagnani. Part-time Bayesians: Incentives and behavioral heterogeneity in belief updating. *Management Science*, 69(9):5523–5542, 2023.
- Peter Andre, Joel P. Flynn, Georgios Nikolakoudis, and Karthik A. Sastry. Quick-fixing: Near-rationality in consumption and savings behavior. *NBER Working Paper No. 33464*, 2025.
- Yves Atchade and Jeffrey Rosenthal. On adaptive markov chain monte carlo algorithms. *Bernoulli*, 11(5):815–828, 2005.
- William Barnett, Douglas Fisher, and Apostolos Serletis. Consumer theory and the demand for money. *Journal of Economic Literature*, 30(4):2086–2119, 1992.
- Anton Barten. Consumer demand functions under conditions of almost additive preferences. *Econometrica*, 32(1/2):1–38, 1964.
- Anton Barten. Evidence on the slusky conditions for demand equations. *The Review of Economics and Statistics*, 49(1):77–84, 1967.
- Anton Barten. The systems of consumer demand functions approach: A review. *Econometrica*, 45(1):23–50, 1977.
- Manel Baucells and Woonam Hwang. A model of mental accounting and reference price adaptation. *Management Science*, 63(12), 2017.
- Truman Bewley. Stationary monetary equilibrium with a continuum of independently fluctuating consumers. In Werner Hildenbrand, editor, *Contributions to mathematical economics in honor of Gérard Debreu*. North Holland, 1986.
- Charles Blackorby and Robert Russell. Two-stage budgeting: An extension of gorman’s theorem. *Economic Theory*, 9(1):185–193, 1997.
- Pedro Bordalo, Nicola Gennaioli, and Andrei Shleifer. Salience and consumer choice. *Journal of Political Economy*, 121(5):803–843, 2014.
- Pedro Bordalo, Nicola Gennaioli, and Andrei Shleifer. Memory, attention, and choice. *The Quarterly Journal of Economics*, 135(3):1399–1442, 2020.

- Pedro Bordalo, Nicola Gennaioli, and Andrei Shleifer. Salience. *Annual Review of Economics* (forthcoming), 2022.
- William Brock. Money and growth: The case of long run perfect foresight. *International Economic Review*, 15(3):750–777, 1974.
- Roger Buehler, Dale Wesley Griffin, and Johanna Peetz. The planning fallacy. *Advances in Experimental Social Psychology*, 43:1–62, 2010.
- Guillermo Calvo. Staggered price in a utility-maximizing framework. *Journal of Monetary Economics*, 12:383–398, 1983.
- Andrew Caplin and Mark Dean. Revealed preference, rational inattention, and costly information acquisition. *American Economic Review*, 105(7):2183–2203, 2015.
- Raj Chetty, Adam Looney, and Kory Kroft. Salience and taxation: Theory and evidence. *American Economic Review*, 99(4):1145–1177, 2009.
- Siddhartha Chib. Estimation and comparison of multiple change-point models. *Journal of Econometrics*, 86:221–241, 1998.
- Syngjoo Choi, Shachar Kariv, Wieland Müller, and Dan Silverman. Who is (more) rational? *American Economic Review*, 104(6):1518–1550, 2014.
- Nelson Cowan. The magical number 4 in short-term memory: A reconsideration of mental storage capacity. *Behavioral and Brain Sciences*, 24:87–185, 2000.
- Angus Deaton and John Muellbauer. *Economics and Consumer Behavior*, chapter 5, pages 119–147. Cambridge University Press: New York, 1980a.
- Angus Deaton and John Muellbauer. An almost ideal demand system. *The American Economic Review*, 70(3):312–326, 1980b.
- Emmanuel Farhi and Xavier Gabaix. Optimal taxation with behavioral agents. *American Economic Review*, 110(1):298–336, 2020.
- Robert Feenstra. Functional equivalence between liquidity costs and the utility of money. *Journal of Monetary Economics*, 17:271–291, 1986.
- Richard Feinberg. Credit cards as spending facilitating stimuli: A conditioning interpretation. *Journal of Consumer Research*, 13(3):348–356, 1986.
- Flóra Felső and Adriaan Soetevent. Broad and narrow bracketing in gift certificate spending. *European Economic Review*, 66:284–302, 2014.
- Daniel Fink. A compendium of conjugate priors. *Mimeo: Environmental Statistics Group, Department of Biology, Montana State University*, 1997.

- David J Freeman. Preferred personal equilibrium and simple choices. *Journal of Economic Behavior & Organization*, 143:165–172, 2017.
- Xavier Gabaix. A sparsity-based model of bounded rationality. *The Quarterly Journal of Economics*, 129(4):1661–1710, 2014.
- Simone Galperti. A theory of personal budgeting. *Theoretical Economics*, 14:173–210, 2019.
- John Gathergood, Neale Mahoney, Neil Stewart, and J’org Weber. How do individuals repay their debt? the balance-matching heuristic. *American Economic Review*, 109(3):844–875, 2019.
- Roy Geary. A note on ‘a constant-utility index of the cost of living’. *Review of Economic Studies*, 18(2):65–66, 1950.
- Andrew Gelman, John Carlin, Hal Stern, David Dunson, Aki Vehtari, and Donald Rubin. *Bayesian Data Analysis*. Taylor & Francis, CRC Press: Boca Raton, FL, third edition, 2013.
- Michael Gelman, Shachar Kariv, Matthew Shapiro, Dan Silvermand, and Steven Tadelis. Harnessing naturally occurring data to measure the response of spending to income. *Science*, 345:212–215, 2014.
- Stuart Geman and Donald Geman. Stochastic relaxation, gibbs distributions, and bayesian restoration of images. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 6(6):721–741, 1984.
- William Gorman. Separable utility and aggregation. *Econometrica*, 27(3):469–481, 1959.
- Heikki Haario, Eero Saksman, and Johanna Tamminen. Adaptive proposal distribution for random walk metropolis algorithm. *Computational Statistics*, 1998.
- Heikki Haario, Eero Saksman, and Johanna Tamminen. An adaptive metropolis algorithm. *Bernoulli*, 7(2), 2001.
- Justine Hastings and Jesse Shapiro. How are snap benefits spent? evidence from a retail panel. *American Economic Review*, 108(12):3493–3540, 2018.
- Wilfred Hastings. Monte carlo sampling methods using markov chains and their applications. *Biometrika*, 57(1):97–109, 1970.
- John Hicks. *Value and Capital*. Oxford: Oxford University Press, 1936.
- Hendrik Houthakker. Additive preferences. *Econometrica*, 28(2):244–257, 1960.
- Mark Huggett. The risk-free rate in heterogeneous-agent incomplete-insurance economies. *Journal of Economic Dynamics and Control*, 17:953–969, 1993.
- Mark Huggett. Wealth distribution in life-cycle economies. *Journal of Monetary Economics*, 38:469–494, 1996.

- Cosmin Ilut and Rosen Valchev. Economic agents as imperfect problem solvers. *The Quarterly Journal of Economics*, 138(1):313–362, 2023.
- Daniel Kahneman and Dan Lovallo. Timid choices and bold forecasts: A cognitive perspective on risk taking. *Management Science*, 39(1):17–31, 1993.
- Daniel Kahneman and Amos Tversky. Intuitive prediction: Biases and corrective procedures. *TIMS Studies in Management Science*, 12:313–327, 1979.
- Botond Kőszegi and Filip Matějka. Choice simplification: A theory of mental budgeting and naive diversification. *The Quarterly Journal of Economics*, 135(2):1153–1207, 2020.
- Botond Kőszegi and Adam Szeidl. A model of focusing in economic choice. *The Quarterly Journal of Economics*, 128(1):53–104, 2013.
- Alexander Koch and Julia Nafziger. Goals and bracketing under mental accounting. *Journal of Economic Theory*, 162(C):305–351, 2016.
- Alexander Koch and Julia Nafziger. Correlates of narrow bracketing. *The Scandinavian Journal of Economics*, 121(4):1441–1472, 2019.
- Alexander Koch and Julia Nafziger. Motivational goal bracketing: An experiment. *Journal of Economic Theory*, 185:104949, 2020.
- Alexander Koch and Julia Nafziger. Motivational goal bracketing with non-rational goals. *Journal of Behavioral and Experimental Economics*, 94:101740, 2021.
- Gary Koop and Simon M. Potter. Estimation and forecasting in models with multiple breaks. *The Review of Economic Studies*, 74(3):763–789, 2007.
- Botond Kőszegi. Utility from anticipation and personal equilibrium. *Economic Theory*, 44(3):415–444, 2010.
- Nicola Lacetera, Devin G. Pope, and Justin R. Sydnor. Heuristic thinking and limited attention in the car market. *American Economic Review*, 102(5):2206–2236, 2012.
- Wassily Leontief. Composite commodities and the problem of index numbers. *Econometrica*, 4(1):39–59, 1936.
- Alan Leslie, Rochel Gelmand, and C.R. Gallistel. The generative basis of natural number concepts. *Trends in Cognitive Sciences*, 12(6):213–218, 2008.
- Chen Lian. A theory of narrow thinking. *Review of Economic Studies*, 88:2344–2374, 2021.
- Chen Lian. Mistakes in future consumption, high MPCs now. *American Economic Review: Insights*, 5(4):499–514, 2023.

- Bartosz Maćkowiak, Filip Matějka, and Mirko Wiederholt. Rational inattention: A review. *Journal of Economic Literature*, 61(1):226–273, 2023.
- Naresh Malhotra. Information load and consumer decision making. *Journal of Consumer Research*, 8(4):419–430, 1982.
- Robert McCulloch and Ruey Tsay. Bayesian inference and prediction for mean and variance shifts in autoregressive time series. *Journal of the American Statistical Association*, 88(423):968–978, 1993.
- George A. Miller. The magical number seven, plus or minus two: Some limits on our capacity for processing information. *Psychological Review*, 63(2):81–97, 1956.
- Sendhil Mullainathan. A memory-based model of bounded rationality. *The Quarterly Journal of Economics*, 107(3):735–774, 2002.
- Maurice Obstfeld and Kenneth Rogoff. Speculative hyperinflations in maximizing models: Can we rule them out? *Journal of Political Economy*, 91(4):675–687, 1983.
- Ryan Oprea. What makes a rule complex? *American Economic Review*, 110(12):3913–3951, 2020.
- Johanna Peetz and Roger Buehler. Is there a budget fallacy? the role of savings goals in the prediction of personal spending. *Personality and Social Psychology Bulletin*, 35(12):1579–1591, 2009.
- Drazen Prelec and George Loewenstein. The red and the black: Mental accounting of savings and debt. *Marketing Science*, 17(1):4–28, 1998.
- Matthew Rabin and Georg Weizsäcker. Narrow bracketing and dominated choices. *American Economic Review*, 99(4):1508–1543, 2009.
- Daniel Read, George Loewenstein, and Matthew Rabin. Choice bracketing. *Journal of Risk and Uncertainty*, 19(1):171–197, 1999.
- Daniel Read, Christopher Y. Olivola, and David J. Hardisty. The value of nothing: Asymmetric attention to opportunity costs drives intertemporal decision making. *Management Science*, 63(12), 2017.
- Scott I. Rick, Cynthia E. Cryder, and George Loewenstein. Tightwads and spendthrifts. *Journal of Consumer Research*, 34(6):767–782, 2008.
- Gareth Roberts, Andrew Gelman, and Walter Gilks. Weak convergence and optimal scaling of random walk metropolis algorithms. *The Annals of Applied Probability*, 7(1):110–120, 1997.
- Frank Schilbach, Heather Schofield, and Sendhil Mullainathan. The psychological lives of the poor. *American Economic Review: Papers & Proceedings*, 106(5):435–440, 2016.

- Joshua Schwartzstein. Selective attention and learning. *Journal of the European Economic Association*, 12(6):1423–1452, 2014.
- Hersh Shefrin. Unfinished business: a multicommodity intertemporal planner–doer framework. *Review of Behavioral Finance*, 2020.
- Hersh Shefrin and Richard Thaler. An economic theory of self-control. *Journal of Political Economy*, 89(2):392–406, 1981.
- Bernard Silverman. *Density Estimation for Statistics and Data Analysis*. London: Chapman and Hall, 1986.
- Herbert Simon. *Models of Man*. John Wiley & Sons: New York, 1957.
- Christopher Sims. Implications of rational inattention. *Journal of Monetary Economics*, 50(3):665–690, 2003.
- Tatiana Sokolova, Satheesh Seenivasan, and Manoj Thomas. The left-digit bias: When and why are consumers penny wise and pound foolish? *Journal of Marketing Research*, 57(4):771–788, 2020.
- Victor Stango and Jonathan Zinman. We are all behavioral, more or less: A taxonomy of consumer decision making. *The Review of Economic Studies*, 90(3):1470–1498, 2023.
- Richard Stone. Linear expenditure systems and demand analysis: An application to the pattern of british demand. *The Economic Journal*, 64(255):511–527, 1954.
- Abigail B. Sussman and Christopher Y. Olivola. Axe the tax: Taxes are disliked more than equivalent costs. *Journal of Marketing Research*, 48:S91–S101, 2011.
- Richard Thaler. Mental accounting and consumer choice. *Marketing Science*, 4(3):199–214, 1985.
- Richard Thaler. Mental accounting matters. *Journal of Behavioral Decision Making*, 12(3):183–206, 1999.
- Richard Thaler, Amos Tversky, Daniel Kahneman, and Alan Schwartz. The effect of myopia and loss aversion on risk taking: An experimental test. *The Quarterly Journal of Economics*, 112(2): 647–661, 1997.
- Richard H Thaler. Behavioral economics: Past, present, and future. *American economic review*, 106(7):1577–1600, 2016.
- Henri Theil. The information approach to demand analysis. *Econometrica*, 33(1):67–87, 1965.
- Henri Theil. *Theory and Measurement of Consumer Demand*, volume 2. Elsevier: New York, 1976.
- Gülden Ülkümen, Manoj Thomas, and Vicki G. Morwitz. Will i spend more in 12 months or a year? the effect of ease of estimation and confidence on budget estimates. *Journal of Consumer Research*, 35(2):245–256, 2008.

Carl Walsh. *Monetary Theory and Policy*. MIT Press: Cambridge, 3 edition, 2010.

Supplemental Appendices

(For Online Publication Only)

Mental Accounting Through Two-stage Budgeting Under Bounded Rationality

Nick Pretnar

*Laboratory for Aggregate Economics and Finance, UCSB
University of Austin Texas (UATX)*

Christopher Y. Olivola

Kelley School of Business, Indiana University

Alan Montgomery

Tepper School of Business, Carnegie Mellon University

A Deriving the Optimal Budget

Note that $\ln(q_{ijt} + 1) = \ln(x_{ijt}/p_{jt} + 1) = \ln(x_{ijt} + p_{jt}) - \ln(p_{jt})$, yielding the expected indirect utility function

$$\begin{aligned} \mathbb{E}_{it}^{\zeta} v_{it} = & \mathbb{E}_{it}^{\zeta} \left\{ \sum_{j=1}^J \alpha_{ij} \ln(\theta_{ijt} \ell_{it} + \gamma_i a_{ijt} + \zeta_{ijt} + p_{jt}) \right. \\ & \left. + \alpha_{i,J+1} \ln \left(\ell_{it} - \sum_{j=1}^J [\theta_{ijt} \ell_{it} + \gamma_i a_{ijt} + \zeta_{ijt}] + m_i + r_t b_{it} \right) \right\} - \sum_{j=1}^J \ln(p_{jt}) \end{aligned} \quad (\text{A.1})$$

The FOC with respect to ϑ_{iyt} is

$$\begin{aligned} & \mathbb{E}_{it}^{\zeta} \left\{ \frac{\alpha_{i,\iota_{iyt}} \ell_{it}}{\vartheta_{iyt} \ell_{it} + \gamma_i a_{i,\iota_{iyt},t} + \zeta_{i,\iota_{iyt},t} + p_{\iota_{iyt},t}} \right\} \\ & = \mathbb{E}_{it}^{\zeta} \left\{ \frac{\ell_{it} \alpha_{i,J+1}}{\ell_{it} - \sum_{y'=1}^k (\vartheta_{iy',t} \ell_{it} + \gamma_i a_{i,\iota_{iy'},t} + \zeta_{i,\iota_{iy'},t}) - \sum_{j=1}^J (1 - \Gamma_{ijt}) (\theta_{ijt} \ell_{it} + \gamma_i a_{ijt} + \zeta_{ijt}) + m_i + r_t b_{it}} \right\} \end{aligned} \quad (\text{A.2})$$

Ignore expectations for now, divide by ℓ_{it} which is assumed known, and multiply both sides by the terms in the denominator to get

$$\begin{aligned} & \alpha_{i,\iota_{iyt}} \ell_{it} - \alpha_{i,\iota_{iyt}} \ell_{it} \vartheta_{iyt} - \alpha_{i,\iota_{iyt}} \sum_{\substack{j \in \iota_{it}(\Gamma_{it}) \\ j \neq \iota_{iyt}}} (\vartheta_{ijt} \ell_{it} + \gamma_i a_{ijt} + \zeta_{ijt}) \\ & - \alpha_{i,\iota_{iyt}} (\gamma_i a_{i,\iota_{iyt},t} + \zeta_{i,\iota_{iyt},t}) - \alpha_{i,\iota_{iyt}} \sum_{j=1}^J (1 - \Gamma_{ijt}) (\theta_{ijt} \ell_{it} + \gamma_i a_{ijt} + \zeta_{ijt}) + \alpha_{i,\iota_{iyt}} (m_i + r_t b_{it}) \\ & = \alpha_{i,J+1} \vartheta_{iyt} \ell_{it} + \alpha_{i,J+1} (\gamma_i a_{i,\iota_{iyt},t} + \zeta_{i,\iota_{iyt},t}) + \alpha_{i,J+1} p_{\iota_{iyt},t} \end{aligned} \quad (\text{A.3})$$

Isolate ϑ_{iyt} and take expectations to get the optimality expression in the main text.

B Equilibrium Appendix

Lemma B.1: *The utility function is strictly increasing in each q_{ijt} as long as the borrowing constraint is such that*

$$m_i > \frac{\alpha_{i,J+1}}{\alpha_{ij}} p_{jt} (q_{ijt} + 1) + \sum_{j=1}^J p_{jt} q_{ijt} - \ell_{it} - r_t b_{it}$$

Further, if $\alpha_{i,J+1} > 0$ then the utility function is strictly concave in the vector $\mathbf{q}_{it}$.

Proof. Substitute z_{it} into the utility function and differentiate in q_{ijt} to get

$$\frac{\partial u_i}{\partial q_{ijt}} = \frac{\alpha_{ij}}{q_{ijt} + 1} - \frac{\alpha_{i,J+1} p_{jt}}{\ell_{it} - \sum_j p_{jt} q_{ijt} + m_i + r_t b_{it}} \quad (\text{B.1})$$

Re-arranging (B.1) and isolating m_i gives the condition for strict monotonicity. The Hessian matrix of u_i for optimization over $\mathbf{q}_{it}$ has the following diagonal terms

$$\frac{\partial^2 u_i}{\partial q_{ijt}^2} = -\alpha_{ij}(q_{ijt} + 1)^{-2} - \alpha_{i,J+1} p_{jt}^2 \left(\ell_{it} - \sum_j p_{jt} q_{ijt} + m_i + r_t b_{it} \right)^{-2} \quad (\text{B.2})$$

Assume $\alpha_{i,J+1} > 0$ and $\alpha_{ij} \in (0, 1)$ for all j . Then, it follows that (B.2) is always negative since it must be that

$$\left(\ell_{it} - \sum_j p_{jt} q_{ijt} + m_i + r_t b_{it} \right)^{-2} > 0$$

Note that all $q_{ijt} > 0$ are interior solutions. By the continuous differentiability of u_i in $\mathbf{q}_{it}$, the Hessian is symmetric and well-defined. Off-diagonal terms are

$$\frac{\partial^2 u_i}{\partial q_{ijt} \partial q_{ikt}} = -p_{jt} p_{kt} \alpha_{i,J+1} \left(\ell_{it} - \sum_j p_{jt} q_{ijt} + m_i + r_t b_{it} \right)^{-2} \quad (\text{B.3})$$

which are always negative as long as $\alpha_{i,J+1} > 0$. Under our assumptions, the trace of the Hessian is strictly negative always. Let $\mathbf{H}_{it}$ stand-in for the agent's Hessian and let $\mathbf{h}_{ijt}$ be the j^{th} column vector of the Hessian. Note that every element of both $\mathbf{h}_{ijt}$ and $\mathbf{H}_{it}$ is strictly negative. Therefore

$$\mathbf{h}_{ijt}^\top \cdot \mathbf{H}_{it} \cdot \mathbf{h}_{ijt} < 0, \quad \forall j \quad (\text{B.4})$$

which implies that $\mathbf{H}_{it}$ is negative definite and concave in $\mathbf{q}_{it}$. **QED**

Lemma B.2:

- i. Suppose $\underline{k}_{it} > 0$ and $\ell_{it} > 0$. ϑ_{it}^* is unique as long as the sufficient conditions underlying Lemma B.1 are satisfied. Further, if $k_{it} > 0$ also, then each ϑ_{iyt}^* is unique for all y for which the candidate budget, ϑ_{iyt}^* , satisfies the numeracy threshold.
- ii. Suppose $\underline{k}_{it} = 0$ or $\underline{k}_{it} > 0$ but $k_{it} = 0$. Then $\boldsymbol{\theta}_{it} = \boldsymbol{\theta}_{i,t-1}$, which is unique.

Proof. Start with part (i). Let $\underline{k}_{it} > 0$. Pick y . Clearly, v_{it} is increasing in ϑ_{iyt}^* by linearity of $q_{i,\ell_{iyt},t}$ in $\theta_{i,\ell_{iyt},t}$. Note that

$$\frac{\partial q_{ijt}}{\partial \vartheta_{iyt}} = \ell_{it} \quad \text{and} \quad \frac{\partial z_{it}}{\partial \vartheta_{iyt}} = -\ell_{it} \quad (\text{B.5})$$

Following (B.1), we get

$$\frac{\partial v_{it}}{\partial \vartheta_{iyt}} = \ell_{it} \mathbb{E}_{it} \left\{ \frac{\alpha_{i, \iota_{iyt}}}{q_{i, \iota_{iyt}, t} + 1} - \frac{\alpha_{i, J+1} p_{\iota_{iyt}, t}}{\ell_{it} - \sum_j p_{jt} q_{ijt} + m_i + r_t b_{it}} \right\} \quad (\text{B.6})$$

where each q_{ijt} is a function of either of ϑ_{iyt} or θ_{ijt} depending on whether j is a member of the narrow choice bracket. Expectations are taken over ζ_{it} . Note that, as long as Lemma B.1 is satisfied, the integrand (i.e., the object inside the expectation operator) is always positive since negative q_{ijt} cannot be observed. Assuming ζ_{ijt} is censored so that $\zeta_{ijt} \geq -\theta_{ijt} \ell_{it} - \gamma_i a_{ijt}$ always, then the integrand is strictly positive and $\frac{\partial v_{it}}{\partial \vartheta_{iyt}} > 0$.

Now consider the generalized second-order conditions, where the diagonal terms of the Hessian are

$$\frac{\partial^2 v_{it}}{\partial \vartheta_{iyt}^2} = \ell^2 \frac{\partial^2 u_i}{\partial q_{ijt}^2} \quad (\text{B.7})$$

For $\underline{k}_{it} \geq 2$ the off-diagonal terms are

$$\frac{\partial^2 v_{it}}{\partial \vartheta_{iyt} \partial \vartheta_{ist}} = \ell^2 \frac{\partial^2 u_i}{\partial q_{ijt} \partial q_{ikt}} \quad (\text{B.8})$$

Since $\ell \geq 0$ always, clearly (B.7) and (B.8) are each always non-positive, though they could be zero if $\ell_{it} = 0$. Under the assumption that $\ell_{it} > 0$, then if the conditions of Lemma B.1 are satisfied, the same argument applies with respect to assessing the negative definiteness of the Hessian.

Finally, if a potential equilibrium under $\underline{k}_{it} > 0$ is unique, then any refinement of the choice set such that $k_{it} < \underline{k}_{it}$ will also be unique, since uniqueness was proven thus far for a choice set of arbitrary dimension.

Proof of part (ii) is trivial.

QED

C Highlighted Properties of the Model

In this supplemental appendix we explore selected aspects of the model's behavioral features to understand how budgeting, anchoring, and choice bracketing impact equilibrium outcomes. First, we demonstrate that the timing frictions between first-stage and second-stage decisions, along with uncertainty surrounding expenditure shocks, lead to non-fungible behavior. We show that by relaxing cognitive frictions and uncertainty, we can obtain the classical two-stage budgeting model. Second, we explore how anchoring induced by mental accounting behavior impacts optimal first-stage budgets. Our results indicate that mental accounting coupled with strongly separable preferences together imply that losses due to over-spending are fully integrated into their own category's budget re-evaluation decision as long as cognitive constraints are relaxed. In the event the consumer is cognitively constrained and does not re-evaluate a particular budget, losses from over-spending are spread out across the budgets of multiple consumption categories.

C.1 Non-fungibility Under Budgeting Frictions

In our formulation the consumer's ex-ante budgeting decisions satisfy the first-order conditions of a sparse-max indirect utility optimization problem, where sparsity is induced by both the extensive margin re-evaluation condition and the numeracy constraints. By contrast ex-post expenditure is subject to stochastic deviations around expected expenditure. In this way the consumer anchors expenditure around a chosen budget. Through this channel, resources allocated toward first-stage budgets are not perfectly fungible nor transferable across second-stage expenditure.

The structure and timing of this decision process has implications relative to how we typically model a neo-classical consumption/savings problem where the consumer chooses their real consumption level and is not subject to budgeting frictions. Indeed, if the consumer was to have one-period-ahead perfect foresight, then our model would collapse into a standard consumption/savings problem with money in the utility function.

Proposition C.1: *Optimally choosing ex-ante budgets is equivalent to optimally choosing ex-post consumption if and only if consumers have perfect foresight over spending shocks ζ_{ijt} and face no cognitive or numeracy frictions ($\Gamma_{ijt} = 1$ for all j).*

Supplemental Appendix C.3, at the end of this section, contains all proofs for propositions and their corollaries. The intuition behind Proposition C.1 is that ex-ante there exists an expected level of consumption $\mathbb{E}_{it}q_{ijt}$ that also exactly solves the optimal budgeting condition, in the case where ζ_{ijt} is not known. This expected level of consumption is a function of the chosen budget weight, θ_{ijt} . Realizing this expected level of consumption as an ex-post actual consumption level is a measure-zero outcome as long as the measure associated with the distribution of ζ_{ijt} is absolutely continuous. However, even if ζ_{ijt} is known beforehand, if $\Gamma_{ijt} = 0$, so that consumers face cognitive frictions and numeracy constraints in making optimal budget updates, then the budget share is $\theta_{ijt} = \theta_{ij,t-1}$. It follows that the ex-ante value of consumption, constrained by a sub-optimal budget weight, will not solve the first-order condition for the optimal budget. In this latter case as well, the utility-maximizing value of q_{ijt} will not be equal to the quantity of consumption associated with the sub-optimal budget weight, except, again, in a measure-zero case. Proposition 1 thus shows that the two-stage budgeting model with mental accounting collapses into the standard two-stage budgeting model of Deaton and Muellbauer (1980a) *only* when there are no budget-timing, cognition, or numeracy frictions.

C.2 Budgetary Bracketing in Response to Spending Misses a_{it}

Here, we examine how equilibrium budgets co-vary under different degrees of narrow bracketing.⁵⁵ Both the number of categories being updated and the composition of those categories affects how budgeting updates will respond to over- or under-spending. Since ϑ_{it}^* is an implicit function

⁵⁵Assume throughout this section that all equilibrium candidates ϑ_{iyt}^* satisfy the exogenously-imposed numeracy conditions, so that $u_{it}(\Gamma_{it}) = u_{it}(\Gamma_{it}) \equiv u_{it}$.

component-wise, the total derivative $\frac{d\vartheta_{iyt}^*}{da_{i,\iota_{iyt},t}}$ includes information regarding how other budgets $\vartheta_{i,-y,t}^*$ respond to variation in the current mental account for commodity ι_{iyt} . Variation in budget ϑ_{iyt}^* on over- or under-spending in category $j \neq \iota_{iyt}$ will also inform equilibrium values of ϑ_{ijt}^* . For categories featuring optimal budget updates ($\Gamma_{ijt} = 1$), the responsiveness of ϑ_{iyt}^* to variation in a_{ijt} , where j is one of these optimally-updated budget categories, is systematic and predictable. Yet, over- or under-spending in outside categories $j \notin \iota_{it}$ induces ambiguous optimal budget responsiveness in the updated categories. In this section and the next, we will refer to categories for which $\Gamma_{ijt} = 1$ as “inside” the bracket and other categories for which $\Gamma_{ijt} = 0$ as “outside” the bracket.⁵⁶

For all commodities inside the bracket, the total responsiveness of optimal budget shares ϑ_{iyt}^* to over- or under-spending a_{ijt} is

$$\frac{d\vartheta_{iyt}^*}{da_{ijt}} = -\frac{\gamma_i}{\ell_{it}} \mathbf{1}\{j = \iota_{iyt}\} - \frac{\gamma_i \alpha_{i,\iota_{iyt}}}{\alpha_{i,\iota_{iyt}} + \alpha_{i,J+1}} \mathbf{1}\{j \neq \iota_{iyt}\} - \sum_{s \neq y} \frac{\alpha_{i,\iota_{iyt}}}{\alpha_{i,\iota_{iyt}} + \alpha_{i,J+1}} \frac{d\vartheta_{ist}^*}{da_{ijt}} \quad \forall j \in \iota_{it} \quad (\text{C.1})$$

where $\mathbf{1}\{\cdot\}$ is an indicator function.⁵⁷ For any given $j \in \iota_{it}$, (C.1) constitutes a linear system in $\frac{d\vartheta_{iyt}^*}{da_{ijt}}$, where y indexes the components of ι_{it} . This system is of full-rank and admits a unique solution as long as certain conditions described in Assumption 1 hold.

Assumption C.2: Assume at least one optimal budget update occurs so that $k_{it} > 0$. Assume further that $\ell_{it} > 0$, $\gamma_i > 0$, $\alpha_{ij} > 0$, $\forall j$, and $\alpha_{i,J+1} > 0$.

Proposition C.3: Under Assumption C.2, without loss of generality, let $\iota_{it} = (1, 2, 3, \dots)$ and suppose ι_{it} is of dimension $J' \leq J$. Consider the total responsiveness of the components of ϑ_{it}^* to a_{iyt} where $y \in \iota_{it}$.

- i. Higher a_{iyt} leads to lower ϑ_{iyt}^* , i.e. $\frac{d\vartheta_{iyt}^*}{da_{iyt}} = -\frac{\gamma_i}{\ell_{it}}$.
- ii. For all $s \in \iota_{it}$ where $s \neq y$, $\frac{d\vartheta_{ist}^*}{da_{iyt}} = 0$.

Proposition C.3 characterizes both own- and cross-category responsivenesses described by (C.1). For categories inside the bracket, both losses and gains due to over- or under-spending, respectively, are fully integrated into the optimal budget update for that *same category*. Since consumers have over spent when $a_{ijt} < 0$, case (i) of Proposition C.3 says that for a category inside the bracket, consumers respond to over-spending in that same category by increasing that category’s budget.

Now consider the remaining categories inside the bracket. Budget adjustments do not respond to cross-category over- or under-expenditure. For $s \neq y$ where s indexes a component of ι_{it} , simultaneous optimal budget updates to ϑ_{ist}^* are independent of the mental account balance in category $j = \iota_{iyt}$. This can be seen in case (ii) of Proposition C.3. Thus, if the inside bracket is broad

⁵⁶Again, since numeracy conditions are assumed to hold, throughout this section, $\Gamma_{ijt} = \underline{\Gamma}_{ijt} = 1$ for all j and t .

⁵⁷Expression (C.1) can be verified by implicitly differentiating the optimal budget condition in mental account balances for commodities inside the bracket.

enough then cross-category responsiveness is minimized. In fact if the inside bracket consists of all J categories, all over- or under-expenditure from the previous period is separately but fully integrated into the optimal budget updates for each category. For example, suppose a consumer is re-evaluating separate budgets for “Groceries” and “Gasoline.” The size of the inside bracket is thus $k_{it} = 2$. Suppose the consumer under-spent on his grocery budget by \$10 ($a_{ijt} = 10$) for the given week but over-spent on his gasoline budget by \$25 ($a_{ijt} = -25$). The updated budget for groceries will take into consideration under-expenditure in groceries but *will not* take into consideration over-expenditure on gasoline. This is a direct result of the linearity of the expenditure system in budget shares and mental account balances, so that when totally differentiating ϑ_{iyt}^* , variation due to over- or under-expenditure in other inside categories is fully captured by those categories’ optimal budget updates.

Over- and under-expenditure in categories outside the optimization bracket, however, *may still* affect inside category budgets. If $k_{it} < J$, strictly, then when updating a budget inside the bracket, consumers take into account how much they over- or under-spent in categories outside the bracket. Note that the total derivative describing the responsiveness of budgets in category i_{yt} to a_{ijt} , where $j \notin i_{it}$, has an implicit expression that depends on how other categories inside the bracket also respond to over- or under-spending outside the bracket:

$$\frac{d\vartheta_{iyt}^*}{da_{ijt}} = -\frac{\gamma_i \alpha_{i,i_{yt}}}{\ell_{it}(\alpha_{i,i_{yt}} + \alpha_{i,J+1})} - \sum_{s \neq y} \frac{\alpha_{i,i_{yt}}}{\alpha_{i,i_{yt}} + \alpha_{i,J+1}} \frac{d\vartheta_{ist}^*}{da_{ijt}} \quad \forall j \notin i_{it} \quad (\text{C.2})$$

It is apparent by directly inspecting (C.2) that optimal budgetary responsiveness to over- or under-spending only depends on liquidity preferences and the long-run expenditure weights for categories inside the bracket. Preference weights for categories outside the bracket play no direct role. Further, it can readily be shown that if there are at least two categories, j and j' , outside of the bracket, then $\frac{d\vartheta_{iyt}^*}{da_{ijt}} = \frac{d\vartheta_{iyt}^*}{da_{ij't}}$. Thus, the inside-bracket adjustment rate depends only on other optimal budget adjustments and the within-category spending preferences, regardless of which category outside the bracket is associated with the greater budget miss. Corollaries C.4 and C.5 formalize these points in Supplemental Appendix C.3.

Proposition C.6: *Under Assumption C.2, without loss of generality, let $i_{it} = (1, 2, 3, \dots)$ and suppose i_{it} is of dimension $J' < J$, strictly. Then both the sign and magnitude of the total responsiveness of the components of ϑ_{it}^* to a_{ijt} , where $j \notin i_{it}$, are ambiguous and depend on the underlying values of the utility parameters α_i and $\alpha_{i,J+1}$.*

Proposition C.6 demonstrates that the cross-category budgetary responsiveness of inside categories to outside categories cannot be determined a priori. This result builds on Corollaries C.4 and C.5 in Supplemental Appendix C.3 which show that only the values of $\alpha_{i,J+1}$ and α_{ij} for inside categories matter, but the magnitudes of these parameters determine the sign of the budget change. Indeed, there is no hard-and-fast rule as to which inside-bracket categories increase or decrease in response to increases in a_{ijt} . Some categories may see budget increases while others

see budget decreases. Further, this responsiveness will be different for different consumers in different periods since it depends heavily on both the underlying values of α_i and $\alpha_{i,J+1}$, as well as the elements of ι_{it} (i.e. the budgets being re-evaluated).




Figure C.1: ($k_{it} = 3$; $\iota_{it} = (1, 2, 3)$) From left to right we show $\frac{d\theta_{iyt}^*}{da_{i4t}}$ where $y \in (1, 2, 3)$ as a function of $\alpha_{i,J+1}$. Zero is denoted by a dashed line in each panel. In panel (a) the value of α_{ij} associated with good $j = y = 1$ is 0.1, in panel (b) it is 0.15, and in panel (c) it is 0.4. Notice that depending on α_{ij} and $\alpha_{i,J+1}$ optimal budget updates respond differently to over- or under-spending in the non-updated category for which $\Gamma_{i4t} = 0$.

Returning to the case where $k_{it} < J$, for intuition we present several figures showing how the components of ϑ_{it}^* together respond to variation in a_{ijt} where $j \notin \iota_{it}$, under different parameterizations and different compositions of the inside bracket.⁵⁸ Figure C.1 shows how $\frac{d\theta_{iyt}^*}{da_{ijt}}$, for some $j \notin \iota_{it}$, varies in α_i and $\alpha_{i,J+1}$. Here, we consider an environment where $k_{it} = 3$, and $\iota_{it} = (1, 2, 3)$, examining the responsiveness of inside categories to variation in the outside account, a_{i4t} . In Figure C.1 the utility weights α_i are increasing from left to right. The responsiveness is non-monotonic in the savings propensity $\alpha_{i,J+1}$, which is most-evident in Figure C.1b. When $\frac{d\theta_{iyt}^*}{da_{ijt}} < 0$ under-spending in j leads to a budget reduction in ι_{iyt} . Consumers thus implicitly use the fact that they believe they have more freedom to spend in j to justify a more frugal approach to ι_{iyt} . When $\frac{d\theta_{iyt}^*}{da_{ijt}} > 0$ the opposite occurs, which seems to be simultaneously associated with smaller values of $\alpha_{i,J+1}$ overall and with categories that have smaller values of $\alpha_{i,\iota_{iyt}}$.

The magnitude of $\alpha_{i,J+1}$ also plays an important role in determining budgetary responsiveness to over- or under-spending. As $\alpha_{i,J+1} \rightarrow_+ 0$ hand-to-mouth consumers appear to respond to under-spending in j by transferring excess funds to categories with smaller utility weights. Since $\frac{d\theta_{iyt}^*}{da_{ijt}}$ represents the responsiveness of budget shares, not absolute budgets, consumers thus increase budget shares more for categories with lower average expenditure shares, balancing out consumption across the various categories. In this manner, the preference representation under our decision-theoretic structure is endogenously non-homothetic, generating non-linear expan-

⁵⁸In these plots the total derivatives $\frac{d\theta_{iyt}^*}{da_{ijt}}$ are featured on the vertical axes while we vary liquidity preferences, $\alpha_{i,J+1}$, on the horizontal axes. We simulate several cases where $J = 4$, while fixing $\gamma_i = \ell_{it} = 1$. Further, we assign different values to α_{ij} , so that some commodities are associated with larger average consumption basket shares than others.

sion paths in available liquidity. This is a by-product of the non-fungibilities induced by narrow bracketing.




Figure C.2: ($k_{it} = 2$) From left to right, $\frac{d\theta_{iyt}^*}{da_{i3t}}$ where budget responsiveness for $j = 2$ under composition $\iota_{it} = (1, 2)$ is presented in panel (a) and under composition $\iota'_{it} = (2, 4)$ in panel (b). Again, zero is denoted by a dashed line.

Budget shares respond to outside over- or under-spending in ways that depend on the size and composition of the inside bracket. Consider, for example, the responsiveness of budgets for commodity group $j = 2$ to variation in outside category $j = 3$ presented in Figure C.2. Between panels (a) and (b) of the figure, we keep $k_{it} = 2$ but vary the composition of the inside bracket from $\iota_{it} = (1, 2)$ in Figure C.2a to $\iota'_{it} = (2, 4)$ in Figure C.2b. In both panels we show the responsiveness of $j = 2$ to variation in a_{i3t} , which is always outside the bracket. Note that the qualitative change in budgetary responsiveness depends on which of the other consumption categories is inside the bracket. In panel (a), when the inside bracket is $(1, 2)$, budget $j = 2$ responds to variation in a_{i3t} by falling for all values of $\alpha_{i,J+1} \in (0, 10)$. Meanwhile, in panel (b) when the inside bracket is $(2, 4)$, the responsiveness of budget $j = 2$ to variation in a_{i3t} is positive for near hand-to-mouth consumers ($\alpha_{i,J+1}$ near zero) and negative otherwise, though less negative as $\alpha_{i,J+1}$ increases.

From this exercise we conclude that a complex interaction between preferences and both the size and composition of the inside bracket determine the degree to which over- or under-spending in an outside category affects optimal budget weights.

C.3 Proofs of Propositions and Corollaries

Proposition C.1: *Proof.* For the following proof, fix ℓ_{it} , $\mathbf{a}_{it}$, and $\mathbf{p}_{it}$.

($\Leftarrow$) With perfect foresight, by definition $\zeta_{it} = \mathbf{0}$ since there is no uncertainty. If $\Gamma_{ijt} = 1$ for all j then consumers make optimal updates to each budget weight for each j , where it is sufficient to drop the expectations operator since $\zeta_{it} = \mathbf{0}$. Further, note that since $\mathbf{z}_{it}$ can be written as a function of q_{it} then the equilibrium condition can be written

$$\frac{\partial u_{it}}{\partial q_{ijt}} \frac{\partial q_{ijt}}{\partial \vartheta_{iyt}} + \frac{\partial u_{it}}{\partial z_{it}} \frac{\partial z_{it}}{\partial \vartheta_{iyt}} \quad (\text{C.3})$$

$$= \frac{\partial u_{it}}{\partial q_{ijt}} \frac{\partial q_{ijt}}{\partial \vartheta_{iyt}} + \frac{\partial u_{it}}{\partial z_{it}} \frac{\partial z_{it}}{\partial q_{ijt}} \frac{\partial q_{ijt}}{\partial \vartheta_{iyt}} \quad (\text{C.4})$$

$$= \frac{\partial u_{it}}{\partial q_{ijt}} + \frac{\partial u_{it}}{\partial z_{it}} \frac{\partial z_{it}}{\partial q_{ijt}} = 0 \quad (\text{C.5})$$

Now suppose rather than choosing ex-ante budget weights, consumers choose ex-post consumption. The optimal choice of q_{ijt} must satisfy the equilibrium condition in (C.5). The choices are thus equivalent.

($\Rightarrow$) For this logical direction, we engage in proof by contrapositive. It is sufficient to assume there are no numeracy constraints. Suppose now that either consumers do not face perfect foresight, so that they do not know ζ_{it} ex-ante, or there are meaningful cognitive frictions, that is $\sum_{j=1}^J \Gamma_{ijt} < J$.

First, let us consider the case where ζ_{it} is unknown ex-ante but $\sum_{j=1}^J \Gamma_{ijt} = J$, so that there are no cognitive frictions. Note that ex-ante budgets ϑ_{iyt}^* must satisfy the main equilibrium condition. Denote ex-ante expected consumption by $\mathbb{E}_{it} q_{ijt} = \tilde{q}_{ijt} = \frac{\vartheta_{iyt}^* \ell_{it} + \gamma_i a_{ijt}}{p_{jt}}$. Now replace $\frac{\vartheta_{iyt}^* \ell_{it} + \gamma_i a_{ijt}}{p_{jt}}$ with $\tilde{q}_{ijt}$ in the expected indirect utility function. Note that for each j optimal expected consumption $\tilde{q}_{ijt}^*$ must exactly solve

$$\begin{aligned} & \mathbb{E}_{it} \left\{ \frac{\partial u_{it}}{\partial \tilde{q}_{ijt}} + \frac{\partial u_{it}}{\partial z_{it}} \frac{\partial z_{it}}{\partial \tilde{q}_{ijt}} \right\} \\ &= \mathbb{E}_{it} \left\{ \frac{\alpha_{ij}}{\tilde{q}_{ijt} + \frac{\zeta_{ijt}}{p_{jt}} + 1} - \frac{p_{jt} \alpha_{i,J+1}}{\ell_{it} - \sum_{j=1}^J [p_{jt} \tilde{q}_{ijt} + \zeta_{ijt}] + m_i + r_t b_{it}} \right\} = 0 \end{aligned} \quad (\text{C.6})$$

Ex-post $q_{ijt}^*(\zeta_{ijt})$ satisfies (C.5) by construction. Since the random variable ζ_{ijt} enters (C.6), $\tilde{q}_{ijt}^* \neq q_{ijt}^*$ except for a measure-zero realization of ζ_{ijt} . Since $\tilde{q}_{ijt}^*$ is completely determined by ϑ_{iyt}^* , ℓ_{it} , a_{ijt} , and p_{jt} , it follows that ex-ante budget choices are not equivalent to choices of ex-post consumption if ζ_{ijt} is not known. Clearly this holds for all j .

Now suppose ζ_{it} is known but $\exists j$ such that $\Gamma_{ijt} = 0$. The consumer thus cannot choose an ex-ante budget for category j and sets his budget such that $\theta_{ijt} = \theta_{ij,t-1}$. There is no ex-ante uncertainty, so we can drop expectations. Except for measure zero values of $\theta_{ij,t-1}$,

$$\frac{\partial u_{it}(\theta_{ij,t-1})}{\partial \vartheta_{iyt}} + \frac{\partial u_{it}}{\partial z_{it}} \frac{\partial z_{it}(\theta_{ij,t-1})}{\partial \vartheta_{iyt}} \neq 0 \quad (\text{C.7})$$

where the dependencies are to note that the first order condition is evaluated at $\theta_{ij,t-1}$. Now the level of consumption associated with the budget share is $\tilde{q}_{ijt} = \frac{\theta_{ij,t-1} \ell_{it} + \gamma_i a_{ijt}}{p_{jt}}$, but this value will not force the left hand side of (C.7) to equal zero. Since $\theta_{ij,t-1}$ is fixed, there exists q_{ijt}^* that satisfies (C.5) and $q_{ijt}^* \neq \tilde{q}_{ijt}$. The proof is complete. QED

Proposition C.3: *Proof.* Consider the system of implicit component-wise total derivatives of ϑ_{iyt}^*

and suppose, without loss of generality, that $y = 1$:

$$\begin{aligned}
\frac{d\vartheta_{i1t}^*}{da_{i1t}} &= -\frac{\gamma_i}{\ell_{it}} - \sum_{s \neq 1} \Gamma_{ist} \frac{\alpha_{i1}}{\alpha_{i1} + \alpha_{i,J+1}} \frac{d\vartheta_{ist}^*}{da_{i1t}} \\
\frac{d\vartheta_{i2t}^*}{da_{i1t}} &= -\frac{\gamma_i}{\ell_{it}} \frac{\alpha_{i2}}{\alpha_{i2} + \alpha_{i,J+1}} - \sum_{s \neq 2} \Gamma_{ist} \frac{\alpha_{i2}}{\alpha_{i2} + \alpha_{i,J+1}} \frac{d\vartheta_{ist}^*}{da_{i1t}} \\
&\vdots \\
\frac{d\vartheta_{ij',t}^*}{da_{i1t}} &= -\frac{\gamma_i}{\ell_{it}} \frac{\alpha_{ij'}}{\alpha_{ij'} + \alpha_{i,J+1}} - \sum_{s \neq j'} \Gamma_{ist} \frac{\alpha_{ij'}}{\alpha_{ij'} + \alpha_{i,J+1}} \frac{d\vartheta_{ist}^*}{da_{i1t}}
\end{aligned} \tag{C.8}$$

Note that (C.8) is linear, so after some algebra, we can write this system in canonical form:

$$A \cdot \frac{d\vartheta_{it}^*}{da_{i1t}} = b \tag{C.9}$$

where

$$A = \begin{pmatrix} 1 & \frac{\alpha_{i1}}{\alpha_{i1} + \alpha_{i,J+1}} & \dots & \dots & \frac{\alpha_{i1}}{\alpha_{i1} + \alpha_{i,J+1}} \\ \frac{\alpha_{i2}}{\alpha_{i2} + \alpha_{i,J+1}} & 1 & \frac{\alpha_{i2}}{\alpha_{i2} + \alpha_{i,J+1}} & \dots & \frac{\alpha_{i2}}{\alpha_{i2} + \alpha_{i,J+1}} \\ \vdots & \dots & \ddots & \dots & \vdots \\ \vdots & \dots & \dots & \ddots & \vdots \\ \frac{\alpha_{ij'}}{\alpha_{ij'} + \alpha_{i,J+1}} & \frac{\alpha_{ij'}}{\alpha_{ij'} + \alpha_{i,J+1}} & \dots & \frac{\alpha_{ij'}}{\alpha_{ij'} + \alpha_{i,J+1}} & 1 \end{pmatrix} \tag{C.10}$$

$$\frac{d\vartheta_{it}^*}{da_{i1t}} = \begin{pmatrix} \frac{d\vartheta_{i1t}^*}{da_{i1t}} \\ \frac{d\vartheta_{i2t}^*}{da_{i1t}} \\ \vdots \\ \frac{d\vartheta_{ij',t}^*}{da_{i1t}} \end{pmatrix} \tag{C.11}$$

$$b = \begin{pmatrix} -\frac{\gamma_i}{\ell_{it}} \\ -\frac{\gamma_i}{\ell_{it}} \frac{\alpha_{i2}}{\alpha_{i2} + \alpha_{i,J+1}} \\ \vdots \\ -\frac{\gamma_i}{\ell_{it}} \frac{\alpha_{ij'}}{\alpha_{ij'} + \alpha_{i,J+1}} \end{pmatrix} \tag{C.12}$$

Conjecture that a solution to (C.9) is

$$\frac{\widetilde{\mathbf{d}\vartheta_{it}^*}}{\mathbf{d}a_{i1t}} = \begin{pmatrix} -\frac{\gamma_i}{\ell_{it}} \\ 0 \\ \vdots \\ 0 \end{pmatrix} \quad (\text{C.13})$$

Note that (C.13) is indeed a solution, and it is the only solution since A has full rank J' under the assumption $\alpha_{i,J+1} > 0$. It is clear that

$$A \cdot \frac{\widetilde{\mathbf{d}\vartheta_{it}^*}}{\mathbf{d}a_{i1t}} = \mathbf{b} \quad (\text{C.14})$$

by inspecting the row-wise dot products on the left hand side of (C.14). This result holds for all values of $y \in \iota_{it}$ indexing a_{iyt} , so that a simple re-arrangement and re-definition of the indices for $y \neq 1$ will yield the same outcome. *QED*

Corollary C.4: *Under Assumption C.2 for categories outside the bracket where $j \notin \iota_{it}$, $\frac{\mathbf{d}\vartheta_{iyt}^*}{\mathbf{d}a_{ijt}}$ is independent of α_{ij} .*

Proof. Note that the only channel through which $\frac{\mathbf{d}\vartheta_{iyt}^*}{\mathbf{d}a_{ijt}}$ could possibly depend on α_{ij} is $\frac{\mathbf{d}\vartheta_{ist}^*}{\mathbf{d}a_{ijt}}$, $s \neq y$. But, it can be verified by inspecting the equilibrium condition for ϑ_{iyt}^* that all values of $\frac{\mathbf{d}\vartheta_{iyt}^*}{\mathbf{d}a_{ijt}}$ only depend on $\alpha_{ij'}$ where $j' \in \iota_{it}$. *QED*

Corollary C.5: *Under Assumption C.2, if $J - k_{it} \geq 2$, so that at least 2 categories are outside the bracket, then for both $j, j' \notin \iota_{it}$, $\frac{\mathbf{d}\vartheta_{iyt}^*}{\mathbf{d}a_{ijt}} = \frac{\mathbf{d}\vartheta_{iyt}^*}{\mathbf{d}a_{ij't}}$.*

Proof. By Corollary C.4, for all $j, j' \notin \iota_{it}$, $\frac{\mathbf{d}\vartheta_{iyt}^*}{\mathbf{d}a_{ijt}}$ depends implicitly on $\frac{\mathbf{d}\vartheta_{ist}^*}{\mathbf{d}a_{ij't}}$ and $\alpha_{i,\iota_{ist}}$ where s indexes all components of ι_{it} . By independence, we can redefine the indices for the outside categories and the value of $\frac{\mathbf{d}\vartheta_{iyt}^*}{\mathbf{d}a_{ijt}}$ will be unchanged. Thus, the definition (index) of the outside category being considered has no bearing on the value of the derivative. This completes the proof. *QED*

Proposition C.6: *Proof.* We will prove this by considering two numerical parameterizations and showing that both the signs and magnitudes of the total derivatives are different under the different parameterizations. Without loss of generality, let $J = 4$ and $J' = 3$. Let $\gamma_i = \ell_{it} = 1$. Let $\alpha_i = (0.1, 0.15, 0.4, 0.35)^\top$ and consider two values for $\alpha_{i,J+1} - \alpha_{i,J+1}^1 = 0.1$ and $\alpha_{i,J+1}^2 = 5$.

$\iota_{it} = (1, 2, 3)$ so that $j = 4$ is the category for which $\Gamma_{ijt} = 0$. The relevant system of equations is

$$\begin{aligned}\frac{d\vartheta_{i1t}^*}{da_{i4t}} &= -\frac{\alpha_{i1}}{\alpha_{i1} + \alpha_{i,J+1}} - \sum_{s \neq 1} \Gamma_{ist} \frac{\alpha_{i1}}{\alpha_{i1} + \alpha_{i,J+1}} \frac{d\vartheta_{ist}^*}{da_{i4t}} \\ \frac{d\vartheta_{i2t}^*}{da_{i4t}} &= -\frac{\alpha_{i2}}{\alpha_{i2} + \alpha_{i,J+1}} - \sum_{s \neq 2} \Gamma_{ist} \frac{\alpha_{i2}}{\alpha_{i2} + \alpha_{i,J+1}} \frac{d\vartheta_{ist}^*}{da_{i4t}} \\ \frac{d\vartheta_{i3t}^*}{da_{i4t}} &= -\frac{\alpha_{i3}}{\alpha_{i3} + \alpha_{i,J+1}} - \sum_{s \neq 3} \Gamma_{ist} \frac{\alpha_{i3}}{\alpha_{i3} + \alpha_{i,J+1}} \frac{d\vartheta_{ist}^*}{da_{i4t}}\end{aligned}\tag{C.15}$$

After some algebra, we get the canonical linear system

$$A \cdot \frac{d\vartheta_{it}^*}{da_{i4t}} = b\tag{C.16}$$

$$\text{where } A = \begin{pmatrix} 1 & \frac{\alpha_{i1}}{\alpha_{i1} + \alpha_{i,J+1}} & \frac{\alpha_{i1}}{\alpha_{i1} + \alpha_{i,J+1}} \\ \frac{\alpha_{i2}}{\alpha_{i2} + \alpha_{i,J+1}} & 1 & \frac{\alpha_{i2}}{\alpha_{i2} + \alpha_{i,J+1}} \\ \frac{\alpha_{i3}}{\alpha_{i3} + \alpha_{i,J+1}} & \frac{\alpha_{i3}}{\alpha_{i3} + \alpha_{i,J+1}} & 1 \end{pmatrix}\tag{C.17}$$

$$b = \begin{pmatrix} \frac{\alpha_{i1}}{\alpha_{i1} + \alpha_{i,J+1}} \\ \frac{\alpha_{i2}}{\alpha_{i2} + \alpha_{i,J+1}} \\ \frac{\alpha_{i3}}{\alpha_{i3} + \alpha_{i,J+1}} \end{pmatrix}\tag{C.18}$$

Now consider two equilibrium solutions to this system under $\alpha_{i,J+1}^1$ and $\alpha_{i,J+1}^2$. Note that with $\alpha_{i,J+1}^1 = 0.1$ we have

$$\frac{d\vartheta_{it}^*}{da_{i4t}} = \begin{pmatrix} 0.227 \\ 0.033 \\ -0.933 \end{pmatrix}\tag{C.19}$$

With $\alpha_{i,J+1}^2 = 5$ we have

$$\frac{d\vartheta_{it}^*}{da_{i4t}} = \begin{pmatrix} -0.014 \\ -0.023 \\ -0.073 \end{pmatrix}\tag{C.20}$$

Clearly, depending on the value of $\alpha_{i,J+1}$, ϑ_{i1t}^* either goes up when a_{i4t} goes up or goes down when a_{i4t} goes up. The same holds for ϑ_{i2t}^* . **QED**

D Estimation Appendix

D.1 Model Fitness

In Table D.1 we present measures of model performance and fitness under different numeracy constraints and parameterizations. In the top part of the table we present estimates of the fully-parameterized models under different numeracy thresholds. Our two main summary statistics comparing model fit are 1) mean absolute predictive error (MAE) as a fraction of each agent's average income $\hat{\ell}_i$ and 2) a log-pointwise predictive density (lppd) information criterion as described in Gelman et al. (2013).⁵⁹ The mean absolute predictive error as a percentage of average income is

$$\frac{MAE}{\hat{\ell}_i} = \frac{1}{I \times J} \sum_i \sum_j \frac{1}{T_i} \frac{\sum_t |\hat{x}_{ijt} - x_{ijt}|}{\hat{\ell}_i} \quad (D.1)$$

where $\hat{x}_{ijt} = \frac{1}{N} \sum_n \tilde{x}_{ijtn}$,

with n indexing atomic MCMC predictions $\tilde{x}_{ijtn}$ and $\hat{x}_{ijt}$ being the average prediction for agent i 's expenditure on consumption category j in period t . The Gelman et al. (2013) log-pointwise predictive density is

$$\text{lppd} = \sum_i \sum_t \sum_j \ln \left[\frac{1}{N} \sum_n \phi \left(\frac{\tilde{x}_{ijtn} - x_{ijt}}{\sigma_{ijn}} \right) \right] \quad (D.2)$$

where the function $\phi(\cdot)$ is the normal density function following from our assumption that ζ_{ijt} is normally distributed and for each consumer independent across j and t . The object in (D.2) inside of the natural log function is the average of the predictive density over the MCMC samples, indexed by n , with $\tilde{x}_{ijtn}$ being the actual, atomic predictions of the model at every iteration n of the MCMC sampler. As a measure of model convergence and performance, in Table D.1 we also present the mean, median, and modal MH acceptance rates for the agent-specific MH within Gibbs sampler.

D.2 MCMC Mixing Properties

For a visual assessment of model performance and MCMC mixing properties, we plot the entire chain of posterior draws, averaging each draw across agents, for parameters that do not change over time. Specifically, if τ_{in} represents the n^{th} draw after burnin and trimming of parameter τ for agent i , then at each draw we compute $\hat{\tau}_n = \frac{1}{I} \sum_i \tau_{in}$. These chains are presented in Figure D.1 for the fully-parameterized baseline model with no numeracy constraints and Figure D.2 for the

⁵⁹Not surprisingly given the non-linearity and complexity of each model's mathematical structure, all of our models have highly-dispersed posterior distributions. However, in the case of the two models we primarily focus on (baseline and \$1-threshold under full parameterizations), this dispersion does not appear to adversely impact the predictive fit.

Table D.1: Model Performance and Comparisons

Model	Numeracy Threshold	MAE / $\widehat{\ell}_i$	lppd	lppd % Baseline ^a	MH Acc. Rates		
					Mean	Median	Mode
$\psi_{ij} \in (0, 1)$ & $\gamma_i = 1$	$\epsilon = \epsilon_\ell = 0$	0.235	-208,208,945	—	0.320	0.389	0.465
	$\epsilon_\ell = 1$	0.256	-208,208,025	4.417×10^{-4}	0.275	0.291	0.438
	$\epsilon_\ell = 0.01$	0.222	-208,208,539	1.950×10^{-4}	0.324	0.377	0.439
	$\epsilon = 0.01$	0.248	-208,208,239	3.389×10^{-4}	0.177	0.116	0.057
	$\epsilon = 0.05$	0.386	-208,208,882	3.023×10^{-5}	0.157	0.091	0.057
	$\epsilon = 0.1$	0.389	-208,208,576	1.772×10^{-4}	0.209	0.143	0.061
$\psi_{ij} \in (0, 1)$ & $\gamma_i = 0$	$\epsilon = \epsilon_\ell = 0$	0.131	-208,209,661	-3.440×10^{-4}	0.452	0.461	0.465
$\psi_{ij} = 0$ & $\gamma_i = 1$	$\epsilon \in [0, 1]$ or $\epsilon_\ell \geq 0$ ^b	1.726×10^{60}	-208,208,827	5.680×10^{-5}	0.123	0.034	0.017
$\psi_{ij} = 0$ & $\gamma_i = 0$	$\epsilon \in [0, 1]$ or $\epsilon_\ell \geq 0$ ^b	3,597.9	-208,209,434	-2.350×10^{-4}	0.377	0.421	0.446
$\psi_{ij} = 1$ & $\gamma_i = 1$	$\epsilon = \epsilon_\ell = 0$	0.148	-208,209,728	-3.761×10^{-4}	0.402	0.465	0.465
$\psi_{ij} \in (0, 1)$ & $\gamma_i = 0$	$\epsilon_\ell = 1$	0.158	-208,208,549	-2.516×10^{-4} ^c	0.342	0.386	0.411
$\psi_{ij} = 1$ & $\gamma_i = 1$	$\epsilon_\ell = 1$	0.448	-208,208,420	-1.896×10^{-4} ^c	0.239	0.246	0.060
$a_{ij,t_{i0}} = 0$	$\epsilon = \epsilon_\ell = 0$	0.478	-208,207,541	6.741×10^{-4}	0.281	0.311	0.362

^a This column represents the percentage gain in information for a model's predictive power relative to the baseline, where the consumer faces no numeracy constraints and all model features are turned on. Negative numbers indicate worse fit than baseline, positive better.

^b Note that whenever $\psi_{ij} = 0$, numeracy constraints do not come into play, so ϵ and/or ϵ_ℓ are free variables.

^c The lppd % Baseline statistics under alternative parameterizations with $\epsilon_\ell = 1$ are measured against the fully parameterized model, $\psi_{ij} \in (0, 1)$ and $\gamma_i = 1$, with $\epsilon_\ell = 1$, not $\epsilon = \epsilon_\ell = 0$.

\$1-threshold model. Note that the hyperparameters (τ_{ij}) associated with the likelihood variance exhibit no autocorrelation in the model without numeracy conditions, compared to substantial autocorrelation in the \$1-threshold model. It is for this reason that we technically prefer the fully-parameterized model, however, the results in the \$1-threshold model should not be completely discounted, as the sampling properties for behavioral parameters like ψ_{ij} , appear robust.




Figure D.1: ($\epsilon = \epsilon_\ell = 0$) On the horizontal axis 0 corresponds to the first MCMC draw after burn-in. Recall the indexing of goods: groceries ($j = 1$), auto/gasoline ($j = 2$), food away from home ($j = 3$), and other expenditure ($j = 4$).




Figure D.2: ($\epsilon_\ell = 1$) On the horizontal axis 0 corresponds to the first MCMC draw after burn-in. Recall the indexing of goods: groceries ($j = 1$), auto/gasoline ($j = 2$), food away from home ($j = 3$), and other expenditure ($j = 4$).

E Additional Behavioral Quantitative Results

E.1 Relations Between Bounded Rationality, Income, and Savings

In this section we analyze the joint relationships between posterior budget updates k_{it} , income ℓ_{it} , and savings z_{it} . The posterior estimates imply the following: 1) budget-updating behavior is generally random and uncorrelated with broad characteristics of the consumer, such as his/her average earnings and/or savings behavior; 2) while, on aggregate, income and savings rates tell us very little about how a consumer will systematically engage in budgeting, within a consumer unit income and savings rates over time are associated with variation in k_{it} , though relations are idiosyncratic, and these tendencies are subject to substantial heterogeneity across consumer units; 3) most consumers have slightly higher savings in periods with more budget updates, though this finding is also rather idiosyncratic; 4) budget updates occur at slightly higher rates when consumers experience positive versus negative income shocks in the same period, though the aggregate distribution of $\hat{k}_i$ conditional upon income shocks is unchanged; 5) in periods *after* high income shocks, most consumers actually adjust their budgets *less* often. The results outlined in this section indicate that, within the model, consumers are substantially heterogeneous with how they engage in budgeting behavior, while providing a cursory (though very weak) indication that k_{it} may be endogenous.⁶⁰

Figure E.1 shows contour plots of the joint distributions of average income and average savings rates against $\hat{k}_i$. The posterior estimates exhibit little correlation between the consumer-unit summary statistics: average weekly income appears only slightly positively correlated with average weekly budget updates (Pearson's $\rho = 0.108$ for the baseline model and $\rho = 0.263$ for the \$1-threshold model), while the average savings rate appears uncorrelated with $\hat{k}_i$ (Pearson's $\rho = -0.008$ for the baseline model and $\rho = 0.029$ for the \$1-threshold model). Notice that for the \$1-threshold model the correlation between average income and average weekly budget updates increases slightly but is still less than 0.3. Since we are integrating both over the entire posterior distribution of draws k_{itn} and time, this exercise demonstrates that estimates of cognitive frictions, Γ_{it} , may be orthogonal to an agent's average income level $\hat{\ell}_i$ and propensity to save. Within the model, a consumer's average behavior with regards to budget attentiveness does not appear to be strongly associated with either their earnings potential or average savings behavior.

However, within each consumer unit budget-updating behavior can still be strongly associated with fluctuations in personal income over time. Further, the posterior estimates indicate that most consumers (68.1% of the sample in the baseline and 69.5% with a \$1 threshold) tend to engage in relatively higher savings (for them) in periods featuring greater budget attentiveness. To demonstrate consumer heterogeneity along these dimensions, we compute agent-specific Pear-

⁶⁰A previously-discussed limitation of our exercise is that we could account for such endogeneity by allowing agents to choose when they want to pay attention to their budgets, thus implicitly selecting k_{it} , though our data limitations prevent this. Since we do not observe budgets or mental accounts, we must estimate them first in order to estimate k_{it} . These data limitations thus require us to assume k_{it} is a structural residual. Future work could assess these implications against datasets that record explicit budgeting behavior.




Figure E.1: This figure characterizes the joint densities of agent-level average income in (a) and (c) and average savings rates in (b) and (d) with the agent-level average number of budget updates, $\hat{k}_i$, where such averages are taken over posterior epochs and time. (a) and (b) contain estimates from the Baseline model, and (c) and (d) contain estimates from the \$1-Threshold model.

son's coefficients for the correlation between ℓ_{it} and posterior $\hat{k}_{it}$ as well as z_{it}/ℓ_{it} and posterior $\hat{k}_{it}$ over time within the consumer unit. The density estimates across consumers for these coefficients are presented in Figure E.2. Consumers appear equally likely to have either higher or lower relative income in periods of greater rationality, which can be seen by noting the symmetry of the densities in panels (a) and (c). However, for more consumers, periods with greater rationality are also associated with higher relative savings rates, which can be seen by noting that the densities in panels (b) and (d) are skewed positive.

To assess whether the relationship between budget-updating and income may be non-linear and thus inadequately explained by Pearson's correlation coefficients, we check if there is any difference between $\hat{k}_{it}$ in periods when ℓ_{it} is 10% higher or lower than $\ell_{i,t-1}$. Figures E.3a and




Figure E.2: Here we present kernel density estimates of Pearson's coefficients across agents for correlations between $\hat{k}_{it}$ and ℓ_{it} in (a) and (c) and $\hat{k}_{it}$ and z_{it}/ℓ_{it} in (b) and (d). The median correlations (dashed black lines) are 0.006 in panel (a), 0.121 in panel (b), 0.013 in panel (c), and 0.119 in panel (d).

[E.3c](#) plot the conditional kernel density estimates of $\hat{k}_{it}$ under different income-shock scenarios against the full estimate. Figures [E.3b](#) and [E.3d](#) show the densities of differences of $\hat{k}_{it}$ relative to the baseline estimate. Specifically, in Figures [E.3b](#) and [E.3d](#), negative values mean that the full $\hat{k}_{it}$ is *less* than the conditional estimate, while positive values mean the full estimate is *greater* than the conditional estimate. We find that 59.8% (baseline) and 59.4% (\$1 threshold) of consumers in our sample tend to make more budget updates in periods when income rises by 10%, while 38.8% (baseline) and 36.5% (\$1 threshold) of consumers tend to make more updates in periods when income falls by 10%. However, the magnitude of these differences is slight — median differences of -0.028 (baseline) and -0.029 (\$1 threshold) after a 10% positive shock versus median differences of 0.034 (baseline) and 0.046 (\$1 threshold) after 10% negative shocks. We conclude that while $\hat{k}_{it}$




Figure E.3: Panels (a) and (c) demonstrate that the estimated distribution of $\hat{k}_{it}$, conditional on the consumer realizing a 10% positive or negative income shock, appears unchanged from the full estimate. Panels (b) and (d) demonstrate that within a consumer unit, we observe enough heterogeneity such that more consumers are likely to make *more* adjustments (red) after a high shock and *fewer* adjustments (green) after a low shock.

is slightly correlated with income shocks, the relationship does not appear significant.

To address the possibility that a consumer's budget-updating behavior is made in response to past income or in anticipation of future income, we consider the correlation between $\hat{k}_{it}$ and $\ell_{i,t\pm s}$, where $s \in \{1, 2, 3\}$. In Figure E.4 we plot the conditional densities across agents of Pearson's coefficients measuring the correlation between budget updates and past and future income. We find no systematic relationship between budget updates and anticipated income, as seen in panels (b) and (d) where the densities are all centered around zero. In panels (a) and (c) we see that there appears to be no systematic relationship between $\hat{k}_{it}$ and income two and three periods in the past. However, $\hat{k}_{it}$ and $\ell_{i,t-1}$ (red) are negatively correlated for approximately 62.0% (baseline)




Figure E.4: Panels (a) and (c) present the densities of Pearson's coefficients across agents measuring the correlation between period- t $\hat{k}_{it}$ and $\ell_{i,t-s}$. Panels (b) and (d) present correlation coefficients measuring the relationship between $\hat{k}_{it}$ and $\ell_{i,t+s}$.

and 61.2% (\$1 threshold) of consumers.

Together, the results presented in Figures E.3 and E.4 provide possible cursory, yet weak, evidence that budget-updating behavior is correlated with exogenous income. Specifically, consumers appear to be less budget attentive in the immediate period *after* experiencing a high income shock, though more budget attentive during the period in which the high shock occurs. Again, however, this relationship is weakly systematic, as there is substantial heterogeneity with respect to how budgeting and income are related.

Finally, we can also check whether the correlations between individual budget-updating behavior, income, and savings over time are at all systematically associated with the average income distribution. That is, do consumers who are more likely to engage in budget updates after high (low) income shocks earn more on average? Do consumers who are more likely to save at higher

rates in periods of many (few) budget updates earn more on average? For both of these questions, the answer again is emphatically ‘no,’ as we find no significant relationships between a consumer’s estimated Pearson’s coefficient from the previous exercise and either their average weekly income or savings. Those who tend to save more in periods of greater rationality do not save more than others on average. Those who are more rational in periods when they have relatively higher income do not earn more than others on average. Thus, how people adjust their behavior period-by-period in response to income shocks and the relaxation of rationality constraints is also rather idiosyncratic and uncorrelated with average earnings.

E.2 Summary Statistics of Selected Agent Characteristics by Type

Table E.1: Median Agent-level Averages by Types ^a

		<i>Ex-ante Type</i>			
Statistic	Model ^b	(i)	(ii)	(iii)	(iv)
$\widehat{k}_{it}$	Baseline	2.514	2.300	2.323	2.514
	\$1-threshold	2.163	1.990	1.931	2.207
$\widehat{\ell}_{it}$	Baseline	460.52	489.08	417.71	490.30
	\$1-threshold	465.25	442.26	431.06	433.95
$\widehat{z_{it}/\ell_{it}}$	Baseline	0.002	0.002	0.001	2.233×10^{-4}
	\$1-threshold	0.002	5.739×10^{-5}	0.005	1.292×10^{-5}
		<i>Ex-post Type</i>			
Statistic	Model ^b	(i)	(ii)	(iii)	(iv)
$\widehat{k}_{it}$	Baseline	2.738	2.419	2.500	2.451
	\$1-threshold	2.398	2.063	2.071	2.267
$\widehat{\ell}_{it}$	Baseline	485.56	449.19	444.00	504.35
	\$1-threshold	471.71	447.86	454.82	502.85
$\widehat{z_{it}/\ell_{it}}$	Baseline	0.002	0.002	0.001	0.005
	\$1-threshold	0.002	0.003	0.001	0.001

^a We first take averages over agent-level outcomes then take the median of the cross-sectional averages.

^b These are the fully-parameterized models with the corresponding numeracy thresholds.