Two-stage Budgeting with Bounded Rationality

Pretnar, Nick and Olivola, Christopher Y. and Montgomery, Alan

University of California Santa Barbara, Carnegie Mellon University Tepper School of Business

15 January 2021

Two-stage Budgeting with Bounded Rationality

Nick Pretnar[1,2*] Christopher Y. Olivola[2,3] Alan Montgomery[2,4]

[1]*University of California Santa Barbara, Laboratory for Aggregate Economics and Finance*
[2]*Carnegie Mellon University, Tepper School of Business*
[3]*Carnegie Mellon University, Department of Social and Decision Sciences*
[4]*Carnegie Mellon University, Machine Learning Department*

January 26, 2021[†]

Abstract

We construct a unifying theory of two-stage budgeting and bounded rationality with mental accounting features. Mental accounting and rational inattention induce behavioral wedges between first-stage and second-stage expenditure budgets. Because reviewing one's financial activities is cognitively costly, consumers might reassess only a subset of their spending budgets every period. Over- or under-spending affects future budgeting and expenditure decisions. We apply latent Bayesian inference to agent-level weekly expenditure data in order to structurally estimate the degree to which low-income consumers appear rationally constrained with respect to budgeting. Our findings provide insight into how consumers may respond to interventions that encourage more disciplined budgeting behavior, like push notifications in budgeting apps. If consumers are acutely aware of budget misses, they may adjust budgets upward to avoid the dis-utility of over-expenditure, driving savings rates and balances downward. In this manner, push notifications that warn consumers about budget thresholds could backfire and actually lead to budgeting behavior that reduces savings and wealth in the long-run.

Keywords: budgeting, mental accounting, bounded rationality, expenditure, savings
JEL Classification: D11, D12, D91

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

†We acknowledge financial support from the PNC Center for Financial Services Innovation at Carnegie Mellon University. Nick acknowledges support from the National Science Foundation Graduate Research Fellowship under Grant No. DGE1252522. We would like to thank Laurence Ales, William Bednar, Siddhartha Chib, Tony Cookson, Finn Kydland, Xiao Liu, Ryan Oprea, Ali Shourideh, Stephen Spear, Kannan Srinivasan, and Oleg Urminsky for valuable insights. Also, we would like to thank seminar participants at the 2017 INFORMS Marketing Science Conference in Los Angeles, the 2019 Association for Consumer Research Conference in Atlanta, the Certified Financial Planner Board of Standard's 2020 Academic Research Colloquium in Washington D.C., the 2020 NBER-NSF SBIES Virtual Conference, the 2020 INFER AC Virtual Conference, and seminars at the University of California Santa Barbara and Carnegie Mellon University, Tepper School of Business.

1 Introduction

We derive a tractable demand model with behavioral features, like mental accounting, which can be readily estimated using consumer spending data. To do this we build on the classical two-stage budgeting literature (Strotz 1957; Gorman 1959; Deaton and Muellbauer 1980b). In traditional two-stage budgeting consumers optimally form expenditure budgets for broad commodity groups every period, and liquidity is thus perfectly fungible. Here, we relax this implicit fungibility assumption. Instead, consumers may optimally set spending budgets for only a subset of consumption categories each period, similar to the sparse optimization framework of Gabaix (2014). The degree to which a consumer treats liquidity as fungible varies both between commodity groups and over time.

Typically, in two-stage budgeting models, decisions about both allocations to broad commodity groups and expenditure transactions are made effortlessly and simultaneously (Strotz 1957; Gorman 1959; Deaton and Muellbauer 1980b). These models contain no behavioral constraints that would cause consumers to fail to make optimal budget plans or even miss their budgeting goals. In addition to specifying a sparse optimization problem, we introduce explicit timing frictions where first-stage budgets are formed prior to the realization of a second-stage preference shock. Consumers thus make budgeting decisions with incomplete information regarding their future consumption needs. One can think of the model's timing frictions as producing something akin to a planner/doer decision problem, where the consumer plans his/her spending in the first-stage and buys things in the second-stage, as described in Shefrin and Thaler (1981) and Thaler (1999).

Mental accounting tendencies may arise from the timing frictions inherent in a planner/doer problem if consumers seek to adjust future spending after either over- or under-spending in the present relative to the budgets they previously set. We thus allow for a dynamic accounting mechanism whereby consumers keep track of second-stage expenditure relative to their first-stage budgets. They can then use that information to shape both future budgets and future expenditure. In equilibrium these mental accounting dynamics affect budgeting decisions in often complicated ways. Since consumers do not necessarily transfer funds between various first-stage budgets with perfect fungibility, whether or not over-spending in one category leads to upward or downward budget adjustments in other categories depends both on which budgets are being re-evaluated and the elasticities of substitution associated with the individual consumer's preferences. Heterogeneous cognitive constraints that lead to anchoring and narrow choice bracketing thus affect the degree to which a consumer exhibits fungible behavior consistent with unbounded rationality.

To validate our novel additions to traditional demand systems we provide an empirical application using anonymized, agent-level transaction data from a large North American bank serving underbanked, low-income consumers. Our empirical estimates indicate that consumer preferences and behavioral tendencies reside on a continuum. Some consumers exhibit behavior that appears near fully rational within the construct of our model, while others appear significantly bounded by planning frictions. We thus observe heterogeneity in consumer behavior that highlights the need for tractable demand models that allow for differences not only in preferences but also in the underlying structure of decision processes.

More broadly, our results speak to the limits of prescriptive solutions to encourage greater financial discipline for low-income and/or liquidity-constrained consumers. Specifically, nudges designed to induce higher savings by encouraging strict budgeting behavior may have unintended consequences. One of our findings is that the degree to which rationality constraints bind for various consumers is uncorrelated with other consumer-specific economic outcomes, such as savings rates and income. In a counterfactual simulation, we also find that relaxing rationality constraints leads to divergent behaviors that depend in complicated ways on consumers' preference primitives. Some consumers would indeed experience welfare improvements and higher savings rates when optimally re-evaluating their expenditure budgets every period, but many consumers would simply use more frequent budget reallocations to justify increasingly higher levels of spending. For this latter group, welfare would fall relative to a model environment where rationality constraints appear to bind more often. A small subset of this group even go into debt, and possibly go bankrupt, when allowed to adjust budgets too easily. We conclude by discussing how these results may inform the effective design of now widely-available financial planning and budgeting applications and softwares.

1.1 Outline

The paper proceeds as follows. In the next section we position our work in the broader literature of two-stage budgeting, bounded rationality, and mental accounting. We then build a model of two-stage budgeting that incorporates features from behavioral economics to allow for endogenous variation and heterogeneity in observed weekly spending patterns. After presenting and discussing the theoretical model, we describe our unique dataset of consumer-level weekly spending and income. Next, we develop a latent-inference estimation routine to uncover consumers' unobserved budgeting decisions and mental accounting state variables. Finally, we discuss the results of our estimation and their implications for inference with regards to mental accounting theory and financial-planning interventions.

2 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. Building on this result while exploiting the equilibrium property that relative prices and marginal rates of substitution equate, Deaton and Muellbauer (1980b) 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 (1980a), 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 empirical 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 set up in Gabaix (2014). This places our research into conversation with contemporary economic research, popularized in Sims (2003), where consumers make decisions under limited information or with limited cognitive resources to optimize. In our sparse max framework, consumers might only re-evaluate a subset of their first-stage budgets in any given period (due to the cognitive costs), thereby engaging in narrow choice bracketing (Kahneman and Lovallo 1993; Read, Loewenstein, and Rabin 1999; Rabin and Weizsäcker 2009; Felső and Soetevent 2014; Koch and Nafziger 2016, 2019). 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, Looney, and Kroft 2009; Bordalo, Gennaioli, and Shleifer 2014; Kőszegi and Szeidl 2013; Schwartzstein 2014; Caplin and Dean 2015; Bordalo, Gennaioli, and Shleifer 2020; Kőszegi and Matějka 2020). By allowing for such heterogeneity, we will use our model to show empirically that different consumers appear rationally constrained to varying degrees over different time intervals.

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: s/he lacks the intrinsic ability to think about how to optimally distribute resources amongst commodities each of which is substitutable with another to varying degrees. Thus, Kőszegi and Matějka (2020), using behavioral insights, predict the classical aggregation result in Hicks (1936), Leontief (1936), and Deaton and Muellbauer (1980b) — that separability is supported if the elemental commodities that comprise broad commodity-groupings are highly substitutable.

2.3 Mental Accounting

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 model in Shefrin and Thaler (1981). On the other hand, we depart significantly from much of the existing mental accounting literature in that we seek to explain behavioral tendencies not as resulting from reference-dependent utility but rather generated by bounds on cognitive attention. In the context of our model, a boundedly rational consumer may only have the cognitive attention to re-evaluate and, if needed, update a subset of his/her first-stage expenditure budgets in a given period. The degree to which consumers are boundedly rational, in our model, is unrelated to the degree to which they engage in mental accounting. Nonetheless, mental accounting 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 and his/her 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 his/her 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 certain 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 validate 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 empirical approach thus 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 (Mani et al. 2013; Schilbach, Schofield, and Mullainathan 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 behavioral theories at the frontier of present research.

3 Model

Time is discrete and indexed by t. Column vectors and matrices are denoted with a bold font. All variables and functions presented, except market prices $\boldsymbol{p}_t$ 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. Further, consumption expenditure and budgeting decisions will depend on the previous period's relative over- or under-expenditure.

We proceed by describing the preference and expenditure mechanisms, then consumers' budget updating choices. After characterizing equilibrium budgeting decisions under rationality frictions, we describe the model's unique ex-ante own-price and income elasticity predictions. We then analyze the complex interactions between mental accounting and optimal budget updating under bounded rationality.

3.1 Preferences and Expenditure

Let z_{it} be marginal period-t savings, m_i be the consumer's borrowing limit, and b_{it} available bank balances. Account balances evolve according to

$$b_{i,t+1} = r_t b_{it} + z_{it} \tag{1}$$

The borrowing limit is such that $m_i > -b_{it}$ always.[1] Consumers have preferences over a J-dimensional vector of real quantities of consumption $\boldsymbol{q}_{it}$ and their total period-t available resources for spending, $z_{it} + m_i + r_t b_{it}$.[2]

[1] In 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 his/her pre-paid card 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.

[2] In this formulation the consumer has preferences over money holdings. This condition of our model ensures that consumers never devote all of their available resources to expenditure, desiring instead to carry money 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, Fisher, and Serletis 1992; Walsh 2010).

We assume utility has a separable form that allows for zero expenditure:

$$u_{it}(\boldsymbol{q}_{it}, z_{it}) = \sum_{j=1}^{J} \alpha_{ij} \ln(q_{ijt} + 1) + \alpha_{i,J+1} \ln(z_{it} + m_i + r_t b_{it}) \tag{2}$$

This parameterization is chosen both for its classical intuition and identification purposes when it comes time to estimate the model. This will become readily apparent later on, as we infer budgeting behavior based on observing only expenditure. The flow utility from consumption, $\sum_{j=1}^{J} \alpha_{ij} \ln(q_{ijt} + 1)$, is the classic Stone-Geary representation where we also assume $\alpha_{ij} \in (0, 1)$ and $\sum_{j=1}^{J} \alpha_{ij} = 1$ (Geary 1950; Stone 1954). We set the Stone-Geary subsistence parameter to -1, thus bounding utility below at zero when $q_{ijt} = 0$. After characterizing the income and savings processes, we will return to the utility function to discuss the role of $\alpha_{i,J+1}$.

Let ℓ_{it} be period-t income. Let $x_{ijt} = p_{jt} q_{ijt}$ be expenditure on commodity group j. Note that marginal savings must be

$$z_{it} = \ell_{it} - \sum_{j=1}^{J} x_{ijt} \tag{3}$$

z_{it} can indeed be negative if consumers spend more than they earn in income. After substituting the expression for z_{it} into (2) it can be shown that the utility function is strictly increasing and strictly concave in q_{ijt} as long as m_i is sufficiently large.[3] Under these conditions, the utility function is well-behaved and yields unique equilibrium demand allocations for any given set of prices $\boldsymbol{p}_t$ and net liquid resources $m_i + \ell_{it} + r_t b_{it}$.

Having characterized savings and inflows, let us return now to the utility function. Later on it will become apparent that $\alpha_{i,J+1}$, which is proportional to the marginal propensity to save, will help govern a lot of the underlying mechanics of the model. Specifically,

[3] We require that any combination of $\alpha_{i,J+1}$ and m_i is such that maximization of u_{it} in $\boldsymbol{q}_{it}$ yields unique equilibrium outcomes. To ensure utility is strictly increasing in q_{ijt}, we must have

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

For strict concavity, we must have

$$\frac{\alpha_{i,J+1}}{\alpha_{ij}} p_{jt}^2 (q_{ijt} + 1)^2 < \left(\ell_{it} - \sum_{j=1}^{J} p_{jt} q_{ijt} + m_i + r_t b_{it} \right)^2$$

These conditions are readily apparent after twice differentiating u_{it} in q_{ijt} following substitution of $z_{it} = \ell_{it} - \sum_{j=1}^{J} p_{jt} q_{ijt}$.

note that $\alpha_{i,J+1}$ may be positive or negative. If $\alpha_{i,J+1} < 0$, the consumer on average borrows more than he saves. While this cannot be a permanent condition due to borrowing constraints, it is indeed possible to observe, in a finite sample of weeks, a consumer for whom $\alpha_{i,J+1} < 0$. As $\alpha_{i,J+1} \rightarrow_{+} 0$ from the right, preferences for saving and thus holding liquidity decline. Consumers for whom $\alpha_{i,J+1}$ is positive but close to zero can be thought of as engaging in hand-to-mouth behavior, consuming almost all of their income every period.

Expenditure in each commodity group is subject to a separate constraint rather than one single, perfectly-linear budget constraint. Coupled with ex-ante, category-specific expenditure uncertainty, this amounts to relaxation of the assumption that resources allocated toward first-stage budgets are perfectly fungible between commodity groups. In the standard two-stage budgeting model outlined in Deaton and Muellbauer (1980b), the realized expenditure share for group j is always exactly equal to the first stage budgeting share. However, we wish to introduce the potential for an intra-period discrepancy between planned and actual expenditure. For example, a consumer may plan a weekly budget on Sunday but be confronted with unexpected expenditure mid-week, thus overspending relative to the planned budget. In our model separate spending constraints along with timing frictions between budgeting and spending ensure that the ex-ante first-stage budgets never exactly equal second-stage realized expenditure, except in measure-zero occurrences for each j.

Period-t ex-post expenditure in commodity group j satisfies

$$x_{ijt} = \theta_{ijt}\ell_{it} + \gamma_i a_{ijt} + \zeta_{ijt} \tag{4}$$

θ_{ijt} is the share of income devoted to expenditure, a_{ijt} is the consumer's mental account balance which encodes the amount s/he over- or under-spent on category j in period $t-1$. γ_i captures how much a consumer's mental account balance influences his/her spending decision. γ_i is restricted to reside in the unit interval. As $\gamma_i \rightarrow 1$ the degree to which a consumer anchors his/her spending the amount over- or under-spent last period increases. At $\gamma = 0$ there is no mental-accounting anchoring effect. We let ζ_{ijt} be an *iid* idiosyncratic expenditure shock. We assume for each j, that ζ_{ijt} is orthogonal to $\zeta_{ij',t}$ where $j \neq j'$, and each shock is mean-zero normally distributed

$$\zeta_{ijt} \underset{iid}{\sim} \mathcal{N}(0, \sigma_{ij}^2)$$

This shock encodes price-independent unanticipated deviations from the spending plan, due to everything from weather-related variation to unexpected health shocks. In this sense, ζ_{ijt} is very similar to a taste shock, causing preferences to appear to vary period by period. A consumer's period-t ex-ante expected expenditure in commodity group j is

$$\mathbb{E}_{it}x_{ijt} = \theta_{ijt}\ell_{it} + \gamma_i a_{ijt} \tag{5}$$

where expectations are taken over $\boldsymbol{\zeta}_{it}$, not prices $\boldsymbol{p}_t$ which we assume are ex-ante known.

3.2 Mental Accounting

Denote the consumer's "mental account" by a_{ijt}. This is a state variable that encodes the amount the consumer over- or under-spent in the previous period, affecting his/her present expenditure. The consumer uses this mental account to discipline expenditure. If $a_{ijt} < 0$ then the consumer over-spent in period $t-1$, while if $a_{ijt} > 0$ he under-spent relative to expectations. The law of motion for mental account balances is

$$a_{ij,t+1} = \mathbb{E}_{it}x_{ijt} - x_{ijt} \tag{6}$$

$\boldsymbol{a}_{it}$ is $J+1$ dimensional where the $J+1$ component encodes how much a consumer over- or under-saved relative to expectations. This value is

$$a_{i,J+1,t+1} = -\sum_{j=1}^{J} a_{ij,t+1} \tag{7}$$

Note that if $a_{i,J+1,t} < 0$ then the consumer has under-saved relative to expectations and vice-versa for $a_{i,J+1,t+1} > 0$. By construction $\sum_{j=1}^{J+1} a_{ijt} = 0$ in every period, so that consumers' beliefs regarding leftover resources from the previous period do not shift their aggregated budget set. That is, their perception of how much they over- or under-spent last period in various commodity groups is exact.[4]

3.3 Expectations and Uncertainty

At the beginning of a period, consumers are uncertain about $\boldsymbol{\zeta}_{it}$ but not $\boldsymbol{p}_t$. Given the short period length we consider in the empirical application (one week), we argue that it is not unreasonable to suggest that consumers know the price levels for broad commodity aggregates prior to the week commencing. This is not the same as saying that consumers know the posted unit-prices of commodities within those aggregates. This assertion may seem strong, but we argue that it is not. Aggregate price levels for broad

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

commodity groups change very little from week-to-week. A consumer planning his next week's consumption expenditure on a Sunday, having most recently gone to the store on a Thursday, would likely expect to face the same nominal price level. If the price of a particular commodity within the groceries category, like beef for example, rises, he may substitute toward a less-expensive commodity like chicken, in order to preserve his grocery expenditure budget. If this is happening implicitly within the category, then such price-responsiveness would be unobservable to us anyway, given our data feature only store-level spending totals.

Our argument is supported by evidence in Hastings and Shapiro (2013) that consumers may substitute toward lower-quality products in the event of price increases. In such a situation, spending for a broad commodity aggregate may appear constant in the data, despite the fact that the quality of products being purchased has declined. They show that the link between quality and price sensitivity is best explained by a model, such as ours, in which consumers budget for specific commodity categories.

Thus, we argue that not incorporating price expectations is a reasonable approximating assumption in our model for two reasons. Since we do not observe specific commodity prices, only indices, we would expect that since such indices barely move over time, consumers would implicitly know their value. Further, even if they do not, but consumers engage in the type of budgeting we posit, ex-post responses to such unexpected price changes will not be identifiable anyway. We thus assume that broad price levels are ex-ante known at the weekly level, which is the period length we consider.

3.4 Choosing Expenditure Budgets

Consumers enter each period knowing the vector of budget weights ascribed to last period's income $\boldsymbol{\theta}_{i,t-1}$ and the following state variables $\boldsymbol{a}_{it}$, ℓ_{it}, b_{it}, r_t, and $\boldsymbol{p}_t$. Herein lies the planner/doer formulation in the style of Shefrin and Thaler (1981): the "planner" chooses his budgets ex-ante and the "doer" engages in expenditure ex-post. The doer's decision is exactly determined by the expenditure constraint in (4). The introduction of uncertainty regarding realized expenditure is what differentiates Thaler's formulation of consumer decisions from two-stage budgeting models. We incorporate an additional friction that allows expenditure to exhibit temporary persistence over a few periods.

Drawing on psychological evidence showing that consumers are cognitively constrained with regards to the number of choices they can consider at one time (Miller 1956; Simon 1957; Cowan 2000), we allow for consumers to possibly update only a subset of their budgets in any given period. That is, consumers may optimally re-evaluate $k_{it} \leq J$ of their budgets for the J different commodity groups. Imposing an integer constraint on the number of changes that can be made, rather than the magnitudes of the budget updates themselves, is supported by findings reported in Leslie, Gelmand, and Gallistel (2008) who argue that integer representations are innate within individual cognitive processes.[5] Further, information processing capacity in human memory is fairly limited and has been shown to impact consumption decisions (Malhotra 1982).[6]

Following from psychological evidence for the existence of cognitive-processing constraints, the budget re-evaluation process operates as follows. In any given period a consumer may re-evaluate his/her weighting variable θ_{ijt} for each good or leave it alone. Let ψ_{ij} denote the probability that a consumer re-evaluates the budget for expenditure in commodity category j, and let Γ_{ijt} be an indicator variable that equals 1 when a re-evaluation is made for the budget in j and 0 otherwise. The re-evaluation decision is Bernoulli distributed

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

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 equilibrium decision structure like that 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.

Regarding this process, a couple of things are worth noting. First, we do not assume consumers are choosing whether to re-evaluate a budget; instead, whether or not a re-evaluation occurs is an exogenous event that happens to the consumer. In certain periods, this feature allows for spending on certain categories to be more salient to the consumer. Thus, our model allows for the consumer to exhibit the kind of attentiveness bias described in Bordalo, Gennaioli, and Shleifer (2014), Kőszegi and Szeidl (2013), and Schwartzstein (2014). Although a consumer does not choose whether or not to re-evaluate a budget, conditional upon a re-evaluation being made ($\Gamma_{ijt} = 1$) the consumer optimally updates his/her category-j budget by choosing θ_{ijt} to maximize expected indirect utility. Otherwise s/he sets $\theta_{ijt} = \theta_{ij,t-1}$ leaving the planned budget weight for commodity group j alone. Due to this heterogeneous integer constraint, a consumer may change

[5] For example, findings in Miller (1956) and Simon (1957) suggest that individuals can consider at most seven choice alternatives at once. Meanwhile, Cowan (2000) says this number is closer to four.

[6] 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).

some budgets but not others. Generally speaking this is fine, just assume that if only one budget is changed then the implicit budget for savings changes as well. Since the expenditure system and total savings are perfectly collinear, it is sufficient to only specify how consumers alter their expenditure budgets.

Let $k_{it} = \sum_{j=1}^{J} \Gamma_{ijt}$, so that k_{it} describes the total number of expenditure budgets the agent will optimally adjust in period t. The vector $\boldsymbol{\theta}_{it}$ is J-dimensional. Denote $\boldsymbol{\vartheta}_{it}$ as the k_{it}-dimensional vector which holds, in cardinal order, the adjustable budgeting parameters of all commodities j for which $\Gamma_{ijt} = 1$. Note that $\boldsymbol{\vartheta}_{it}$ is a sub-vector of $\boldsymbol{\theta}_{it}$ corresponding to the non-zero indices of $\boldsymbol{\Gamma}_{it}$. Let $\boldsymbol{\vartheta}^*_{it}$ be the optimally-chosen analog of this vector. Denote $\mathbb{E}_{it} v_{it}(\boldsymbol{\theta}_{it})$ as expected indirect utility after dividing each component of the expenditure system (4) by p_{jt} and substituting into $u_{it}(\boldsymbol{q}_{it}, z_{it})$ along with z_{it}:[7]

$$\mathbb{E}_{it} v_{it}(\boldsymbol{\theta}_{it}) = \mathbb{E}_{it} \left\{ \sum_{j=1}^{J} \alpha_{ij} \ln \left(\frac{\theta_{ijt}\ell_{it} + \gamma_i a_{ijt} + \zeta_{ijt}}{p_{jt}} + 1 \right) \right.$$
$$\left. + \alpha_{i,J+1} \ln \left(\ell_{it} - \sum_{j=1}^{J} [\theta_{ijt}\ell_{it} + \gamma_i a_{ijt} + \zeta_{ijt}] + m_i + r_t b_{it} \right) \right\}$$

Let $\boldsymbol{\iota}_{it}(\boldsymbol{\Gamma}_{it})$ be a vector-valued integer function that maps the index of components of $\boldsymbol{\vartheta}_{it}$ back into the index of components for $\boldsymbol{\theta}_{it}$. This function is $\boldsymbol{\iota}_{it}(\boldsymbol{\Gamma}_{it}) : \mathbb{N}^J \to \mathbb{N}^{k_{it}}$ and valid for $k_{it} > 0$. The function outputs in cardinal order the index of the components of $\boldsymbol{\theta}_{it}$ to which we assign the components of $\boldsymbol{\vartheta}^*_{it}$. For example, suppose $J = 4$ and $\boldsymbol{\Gamma}_{it} = \big(0, 1, 0, 1\big)$. Then this implies $k_{it} = 2$ and $\boldsymbol{\iota}_{it}(\boldsymbol{\Gamma}_{it}) = \big(2, 4\big)$.

Let y index the components of $\boldsymbol{\iota}_{it}(\boldsymbol{\Gamma}_{it})$, so that ι_{iyt} denotes the y^{th} component of the vector $\boldsymbol{\iota}_{it}$.[8] If $k_{it} > 0$ then an optimal choice of ϑ^*_{iyt} must satisfy

$$\mathbb{E}_{it} \frac{\partial v_{it}(\boldsymbol{\theta}_{it})}{\partial \vartheta_{iyt}} = \mathbb{E}_{it} \left\{ \frac{\partial u_{it}}{\partial q_{ijt}} \frac{\partial q_{ijt}}{\partial \vartheta_{iyt}} + \frac{\partial u_{it}}{\partial z_{it}} \frac{\partial z_{it}}{\partial \vartheta_{iyt}} \right\} = 0, \quad \forall y > 0 \quad \text{and} \quad j = \iota_{iyt} \tag{8}$$

Due to the money-in-the-utility-function structure, (8) will always be satisfied in equilibrium as consumers equate the expected marginal utility of additional consumption with the expected marginal utility of additional liquidity.

We can invert (8) under our utility parameterization to arrive at an analytical expression for equilibrium values $\boldsymbol{\vartheta}_{it}$. For some y indexing $\boldsymbol{\iota}_{it}$, let $\boldsymbol{\vartheta}^*_{i,-yt}$ denote the vector of optimally chosen budget shares which does not include y. Note that this vector may be empty

[7] Again, recall that expectations are taken only over $\boldsymbol{\zeta}_{it}$.

[8] In our example above, $\iota_{i1t} = 2$ and $\iota_{i2t} = 4$.

when no changes are made in a given period. The optimal budget share $\vartheta^*_{iyt}(\boldsymbol{\vartheta}^*_{i,-y,t})$ for good ι_{iyt} can be implicitly expressed as:

$$\begin{aligned}\vartheta^*_{iyt}(\boldsymbol{\vartheta}^*_{i,-y,t}) = \mathbb{E}_{it}\Bigg\{ & \frac{\alpha_{i,\iota_{iyt}}\ell_{it} - \alpha_{i,\iota_{iyt}}\sum_{s\neq y}(\vartheta^*_{ist}\ell_{it} + \gamma_i a_{i,\iota_{ist},t} + \zeta_{i,\iota_{ist},t})}{\ell_{it}(\alpha_{i,\iota_{iyt}} + \alpha_{i,J+1})} \\ & - \frac{(\alpha_{i,\iota_{iyt}} + \alpha_{i,J+1})(\gamma_i a_{i,\iota_{iyt},t} + \zeta_{i,\iota_{iyt},t}) + \alpha_{i,J+1}p_{\iota_{iyt},t}}{\ell_{it}(\alpha_{i,\iota_{iyt}} + \alpha_{i,J+1})} \\ & - \frac{\alpha_{i,\iota_{iyt}}\Big(\sum_{j=1}^{J}(1-\Gamma_{ijt})(\theta_{ijt}\ell_{it} + \gamma_i a_{ijt} + \zeta_{ijt}) - m_i - r_t b_{it}\Big)}{\ell_{it}(\alpha_{i,\iota_{iyt}} + \alpha_{i,J+1})}\Bigg\}\end{aligned} \tag{9}$$

where $\theta_{ijt} = \theta_{ij,t-1}$ for $j \notin \boldsymbol{\iota}_{it}(\boldsymbol{\Gamma}_{it})$. For the algebra behind this expression, see Appendix A.1. Recall we assume that ℓ_{it} and $\boldsymbol{p}_t$ are ex-ante known. When estimating the structural model without explicitly observing budget-updating behavior, we exploit the independence of the components of $\boldsymbol{\zeta}_{it}$ to iteratively sample the latent shocks Γ_{ijt} and update the components of $\boldsymbol{\vartheta}^*_{it}$ accordingly, taking $\boldsymbol{\vartheta}^*_{i,-y,t}$ as given. This allows for estimation of the mental accounting and budgeting parameters without having to iterate over the component functions of $\boldsymbol{\vartheta}_{it}$ as we progress through a Markov Chain Monte Carlo estimation routine.[9]

3.5 Personal Mental Accounting Equilibrium

Let t_{i0} be the period in which an individual consumer enters the economy as an autonomous, decision-making agent. Given a sequence of prices $\{\boldsymbol{p}_t, r_t\}_{t\geq t_{i0}}$, a sequence of income values $\{\ell_{it}\}_{t\geq t_{i0}}$, a sequence of idiosyncratic expenditure shocks $\{\boldsymbol{\zeta}_{it}\}_{t\geq t_{i0}}$, a sequence of idiosyncratic cognitive shocks $\{\boldsymbol{\Gamma}_{it}\}_{t\geq t_{i0}}$, and initial values for the budget weights, mental account balances, and bank balances $\{\boldsymbol{\theta}_{i,t_{i0}}, \boldsymbol{a}_{i,t_{i0}}, b_{i,t_{i0}}\}$, a personal mental accounting equilibrium consists of:

i. *Sequences of policies:* $\{\boldsymbol{q}_{it}, z_{it}, \boldsymbol{\theta}_{it}\}_{t\geq t_{i0}}$.

ii. *Sequences of balances:* $\{\boldsymbol{a}_{it}, b_{it}\}_{t\geq t_{i0}}$.

such that in each period

a. *Given* $\boldsymbol{\Gamma}_{it}$, $\boldsymbol{a}_{it}$, *and* b_{it}, $\boldsymbol{\theta}_{it}$ *satisfies the sparse max indirect-utility maximization program.*

b. *Given* $\boldsymbol{\theta}_{it}$, ℓ_{it}, $\boldsymbol{a}_{it}$, *and* $\boldsymbol{\zeta}_{it}$, $\boldsymbol{x}_{it}$ *satisfies (4).*

[9]The details of our estimation strategy are left to Online Technical Appendix A, which is available at the lead author's website: `https://www.npretnar.com/research`.

c. *Given $\boldsymbol{x}_{it}$ and ℓ_{it}, z_{it} satisfies (3).*

d. *Mental account balances $\boldsymbol{a}_{it}$ are updated according to (6) and (7).*

e. *Bank balances b_{it} evolve according to (1).*

4 Highlighted Properties of the Model

In this section 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.

Finally, we also explored how model-predicted price and income elasticities respond to narrow choice bracketing, but for space reasons have relegated these results, which turn out to be empirically un-interesting, to Online Technical Appendix C. In that appendix, we show, theoretically, that in the presence of timing frictions the composition of the narrow choice bracket for optimization can affect own- and cross-price elasticities. However, we find empirically that such variation is rather insignificant. We also demonstrate that preferences for liquidity holdings as governed by the parameter $\alpha_{i,J+1}$ can induce Giffen-like demand behavior amongst hand-to-mouth consumers. While interesting in theory, no such consumer in our sample exhibits Giffen-like demand behavior. For these reasons, we focus our primary theoretical and empirical explanations in the main text around the unique behavioral aspects of our formulation (e.g., narrow choice bracketing and mental accounting).

4.1 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. 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 thus 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 his/her 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, where the consumer takes prices as given, choosing a vector of real consumption $\boldsymbol{q}_{it}$ and savings z_{it} to maximize utility.

In this section we focus on our model's unique non-equivalence of choosing budgets ex-ante versus choosing consumption ex-post which is a consequence of the way timing frictions enter the sparse max optimization problem. Ex-post, after realizing $\boldsymbol{\zeta}_{it}$, there is no information lost when plugging (3) into (1) to get back the standard budget constraint without mental accounting frictions:

$$\sum_{j=1}^{J} x_{ijt} + b_{i,t+1} \leq r_t b_{it} + \ell_{it} \qquad \text{with} \qquad b_{i,t+1} > -m_i$$

Note that this substitution cannot be accomplished prior to realization of $\boldsymbol{\zeta}_{it}$ due to Jensen's inequality. Ex-ante, expected expenditure depends on the budgeting decision of the consumer and his/her expectations over $\boldsymbol{\zeta}_{it}$. When choosing budgets, the consumer internalizes how balances b_{it} will evolve so that expected expenditure depends on expected balances by way of the budgeting decision.

Proposition 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 no cognitive frictions ($\Gamma_{ijt} = 1$ for all j).

All proofs for propositions and their corollaries are featured in detail in Appendix A.2. The intuition behind the proof of this proposition is that ex-ante there exists an expected level of consumption $\mathbb{E}_{it} q_{ijt}$ that also exactly solves the integral in (8), 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 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 in (8). In this case, even if ζ_{ijt} is known, 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 (1980b) only when there are no budget-timing or cognition frictions. Uncertainty introduced by rationality frictions causes consumers to deviate from the optimal consumption allocations implied under classical two-stage budgeting.

4.2 Budgetary Bracketing in Response to Spending Misses $\boldsymbol{a}_{it}$

Here, we formalize intuition by examining how equilibrium budgets co-vary under different degrees of narrow bracketing. 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 $\boldsymbol{\vartheta}^*_{it}$ is an implicit function component-wise, the total derivative $\frac{\mathrm{d}\vartheta^*_{iyt}}{\mathrm{d}a_{i,\iota_{iyt},t}}$ includes information regarding how other budgets $\boldsymbol{\vartheta}^*_{i,-y,t}$ 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^*_{iyt}. 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 $j \notin \boldsymbol{\iota}_{it}$ 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.

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

$$\frac{\mathrm{d}\vartheta^*_{iyt}}{\mathrm{d}a_{ijt}} = -\frac{\gamma_i}{\ell_{it}}\mathbf{1}\{j = \iota_{iyt}\} - \frac{\gamma_i \alpha_{i,\iota_{iyt}}}{\alpha_{i,\iota_{iyt}} + \alpha_{i,J+1}}\mathbf{1}\{j \neq \iota_{iyt}\} - \sum_{s \neq y} \frac{\alpha_{i,\iota_{iyt}}}{\alpha_{i,\iota_{iyt}} + \alpha_{i,J+1}} \frac{\mathrm{d}\vartheta^*_{ist}}{\mathrm{d}a_{ijt}} \qquad \forall j \in \boldsymbol{\iota}_{it} \tag{10}$$

where $\mathbf{1}\{\cdot\}$ is the indicator function that $= 1$ if the condition it evaluates is true. (10) can be verified by implicitly differentiating (9) in mental account balances for commodities inside the bracket. For any given $j \in \boldsymbol{\iota}_{it}$, (10) constitutes a linear system in $\frac{\mathrm{d}\vartheta^*_{iyt}}{\mathrm{d}a_{ijt}}$, where y

indexes the components of $\boldsymbol{\iota}_{it}$. This system is of full-rank and admits a unique solution as long as certain conditions described in Assumption 1 hold.

Assumption 1: 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 2: Under Assumption 1, without loss of generality, let $\boldsymbol{\iota}_{it} = (1, 2, 3, \ldots)$ and suppose $\boldsymbol{\iota}_{it}$ is of dimension $J' \leq J$. Consider the total responsiveness of the components of $\boldsymbol{\vartheta}^*_{it}$ to a_{iyt} where $y \in \boldsymbol{\iota}_{it}$.

i. Higher a_{iyt} leads to lower ϑ^*_{iyt}, i.e. $\frac{\mathrm{d}\vartheta^*_{iyt}}{\mathrm{d}a_{iyt}} = -\frac{\gamma_i}{\ell_{it}}$.

ii. For all $s \in \boldsymbol{\iota}_{it}$ where $s \neq y$, $\frac{\mathrm{d}\vartheta^*_{ist}}{\mathrm{d}a_{iyt}} = 0$.

Proposition 2 characterizes both own- and cross-category responsivenesses described by (10). For categories inside the bracket, both losses and gains in category $j = \iota_{iyt}$ due to over- or under-spending, respectively, are fully integrated into the optimal budget update for that same category. In case (i) of Proposition 2, we can specifically see that $\frac{\mathrm{d}\vartheta^*_{iyt}}{\mathrm{d}a_{ijt}} = -\frac{\gamma_i}{\ell_{it}}$ for $j = \iota_{iyt}$. A budget for a specific category increases in constant proportion, depending on income, for every $1 of over-spending carried forward. For a category inside the bracket, consumers thus 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 $\boldsymbol{\iota}_{it}$, 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 2. Thus, if the inside bracket is broad 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 ι_{iyt} to a_{ijt}, where $j \notin \boldsymbol{\iota}_{it}$, has an implicit expression that depends on how other categories inside the bracket also respond to over- or under-spending outside the bracket:

$$\frac{\mathrm{d}\vartheta^*_{iyt}}{\mathrm{d}a_{ijt}} = -\frac{\gamma_i \, \alpha_{i,\iota_{iyt}}}{\ell_{it}(\alpha_{i,\iota_{iyt}} + \alpha_{i,J+1})} - \sum_{s \neq y} \frac{\alpha_{i,\iota_{iyt}}}{\alpha_{i,\iota_{iyt}} + \alpha_{i,J+1}} \frac{\mathrm{d}\vartheta^*_{ist}}{\mathrm{d}a_{ijt}} \qquad \forall j \notin \boldsymbol{\iota}_{it} \tag{11}$$

Corollary 1: Under Assumption 1 for categories outside the bracket where $j \notin \boldsymbol{\iota}_{it}$, $\frac{\mathrm{d}\vartheta^*_{iyt}}{\mathrm{d}a_{ijt}}$ is independent of α_{ij}.

Corollary 1 is apparent by directly inspecting (11). Basically, 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.

Corollary 2: Under Assumption 1, if $J - k_{it} \geq 2$, so that at least 2 categories are outside the bracket, then for both $j, j' \notin \boldsymbol{\iota}_{it}$, $\frac{\mathrm{d}\vartheta^*_{iyt}}{\mathrm{d}a_{ijt}} = \frac{\mathrm{d}\vartheta^*_{iyt}}{\mathrm{d}a_{ij',t}}$.

Corollary 2 is also readily apparent by directly inspecting (11). Note that in (11), j indexes none of the $\alpha_{i,\iota_{iyt}}$, which all correspond to preference weights for commodities inside the budgeting bracket. Thus, the 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.

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

Proposition 3 demonstrates that budget variation cannot be determined a priori. Generally, the cross-category responsiveness of inside categories to a_{ijt} outside the bracket is rather complex. This result builds on Corollaries 1 and 2, 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 $\boldsymbol{\alpha}_i$ and $\alpha_{i,J+1}$, as well as the elements of $\boldsymbol{\iota}_{it}$ — i.e. the budgets being re-evaluated.

Returning to the case where $k_{it} < J$, for intuition we present several figures showing how the components of $\boldsymbol{\vartheta}^*_{it}$ together respond to variation in a_{ijt} where $j \notin \boldsymbol{\iota}_{it}$, under different parameterizations and different compositions of the inside bracket. In these plots the total derivatives $\frac{\mathrm{d}\vartheta^*_{iyt}}{\mathrm{d}a_{ijt}}$ 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.

Figure 1 shows how $\frac{\mathrm{d}\vartheta^*_{iyt}}{\mathrm{d}a_{ijt}}$, for some $j \notin \boldsymbol{\iota}_{it}$, varies in $\boldsymbol{\alpha}_i$ and $\alpha_{i,J+1}$. Here, we consider an environment where $k_{it} = 3$, and $\boldsymbol{\iota}_{it} = (1, 2, 3)$, examining the responsiveness of inside categories to variation in the outside account, a_{i4t}. We let good $j = 3$ be associated with the largest expenditure share, while good $j = 1$ is associated with the smallest. In Figure 1 the utility weights $\boldsymbol{\alpha}_i$ are increasing from left to right. The responsiveness is non-monotonic in the savings propensity $\alpha_{i,J+1}$, which is evident in Figures 1a and 1b. When $\frac{\mathrm{d}\vartheta^*_{iyt}}{\mathrm{d}a_{ijt}} < 0$ 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 $\frac{\mathrm{d}\vartheta^*_{iyt}}{\mathrm{d}a_{ijt}} > 0$ the opposite is 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 $\alpha_{i,J+1} \to_+ 0$ from the right, hand-to-mouth consumers appear to respond to under-spending in j by transferring excess funds to categories with smaller utility weights. Since $\frac{\mathrm{d}\vartheta^*_{iyt}}{\mathrm{d}a_{ijt}}$ 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 expansion paths in available liquidity. This is a by-product of the non-fungibilities induced by narrow bracketing.

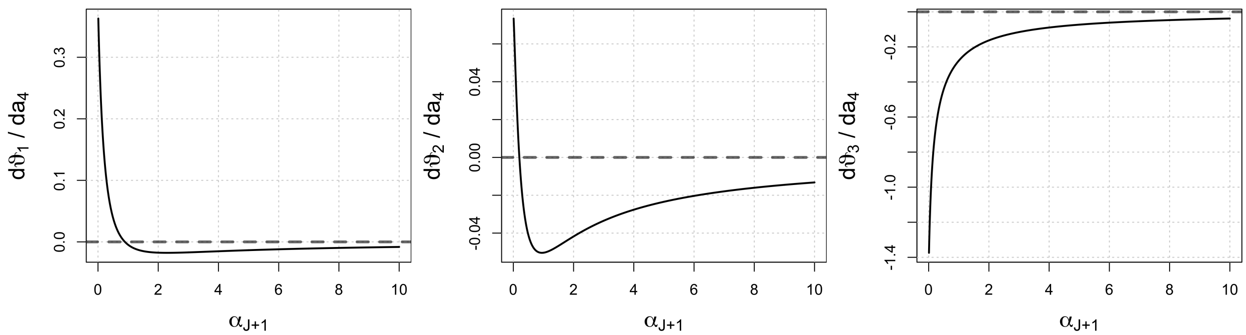

(a) $\alpha_{i1} = 0.1$ (b) $\alpha_{i2} = 0.15$ (c) $\alpha_{i3} = 0.4$

Figure 1: ($k_{it} = 3$; $\boldsymbol{\iota}_{it} = (1, 2, 3)$) From left to right we show $\frac{\mathrm{d}\vartheta^*_{iyt}}{\mathrm{d}a_{i4t}}$ 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 $\Gamma_{i4t} = 0$.

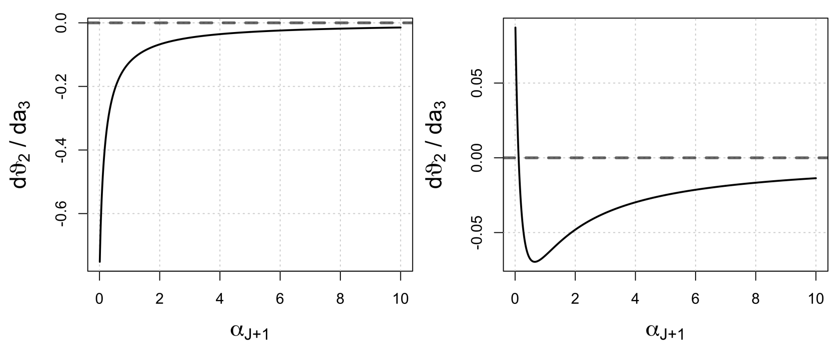

(a) $j = 2$; $\boldsymbol{\iota}_{it} = (1, 2)$ (b) $j = 2$; $\boldsymbol{\iota}'_{it} = (2, 4)$

Figure 2: ($k_{it} = 2$) From left to right, $\frac{\mathrm{d}\vartheta^*_{iyt}}{\mathrm{d}a_{i3t}}$ where budget responsiveness for $j = 2$ under composition $\boldsymbol{\iota}_{it} = (1, 2)$ is presented in panel (a) and under composition $\boldsymbol{\iota}'_{it} = (2, 4)$ 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 2. Between panels (a) and (b) of the figure, we keep $k_{it} = 2$ but vary the composition of the inside bracket from $\boldsymbol{\iota}_{it} = (1, 2)$ in Figure 2a to $\boldsymbol{\iota}'_{it} = (2, 4)$ in Figure 2b. In both panels we show the responsiveness of $j = 2$. 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 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.

5 Weekly Expenditure Data

One goal of our analysis is to estimate consumer demand parameters while simultaneously considering the endogenous link between spending and budgeting decisions. Moreover, we want to accomplish this using only expenditure data, where control variables to account for unobserved demand shifts for high frequency, weekly expenditure are not available to the econometrician. Further, we want to estimate the model despite lacking data about individual spending budgets. Incorporating widely-established bounded rationality and mental accounting features into a classical two-stage budgeting model offers two important advantages: first, it allows us to place structure on the underlying budgeting decisions in a way that is consistent with theory; second, it allows for spending to be responsive to an unobserved state variable that encodes consumers' budgeting misses. With such model components, we can estimate individual price and income elasticities as well as other behavioral demand features, such as the degree to which consumers appear rationally bounded, using only consumer-level spending data for broad commodity-group aggregates. In this section, we describe our unique consumer-level expenditure dataset and some of its features.

Our dataset is an anonymized sample of low-income pre-paid debit-card users from a large North American bank. It includes weekly totals of consumption expenditures, income, and running balances. Expenditure is categorized by 4-digit Visa merchant category classification codes (MCC).[10] 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.[11] The explicit categories we examine, and their 2-digit MCCs are Groceries (54), Gasoline (55), and Food Away from

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

[11]Note 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.

Home (58).[12] 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. Thus, the data panel is unbalanced.

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 28 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. That is, these consumers deposit 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. 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.

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
Expenditure							
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 $\sum_i \sum_{t_i} = 71,859$

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

[12] Our 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.

agent-level averages.[13] 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 United States household income in 2015 was \$56,516 according to the U.S. Census bureau (Proctor, Semega, and Kollar 2015). 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.[14]

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").

Figure 3 presents the time series of observed expenditures for a selected consumer whose income is closest to the sample median. We observe 18 weeks of expenditure for this consumer. 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, that can predict 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 mental accounting and taste shocks, persistence is a function of bounded rationality and mental accounting, and changes to average spending and trends are associated with budget updates. Fitting our model to data thus provides quantitative inferences as to how consumers make high-frequency consumption expenditure decisions and the tradeoffs they face in doing so.

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

[14] We do not observe the household size in our dataset.

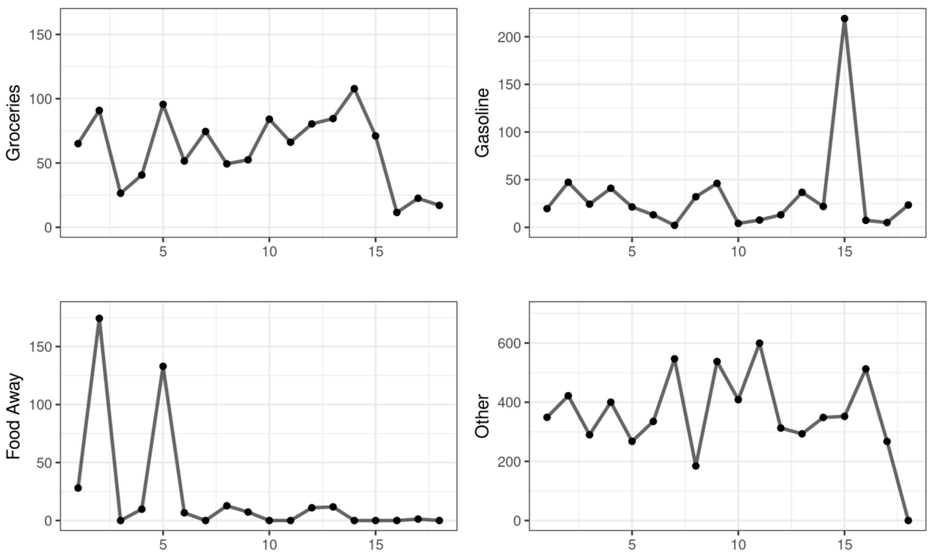

Figure 3: These are the time series of spending for a selected consumer whose average weekly income is closest to the sample median. Mean weekly expenditure in each category is as follows: groceries \$60.67, auto/gasoline \$32.50, food away from home \$21.98, and other expenditure \$357.00. This consumer earns an average of \$460.05 per week and maintains an average weekly card balance of \$115.21 after spending.

6 Structural Model Estimation

Few economics papers that grapple with behavioral phenomena in a decision-theoretic structural manner attempt to estimate such models using field data. A notable exception is Hastings and Shapiro (2018), though the authors' modeling and estimation approach is different than ours. Specifically, Hastings and Shapiro (2018) take a non-parametric approach to identifying and estimating the degree to which low-income consumers appear to be constrained by mental accounting behavior. Given one aim of this paper is to infer budgeting decisions without observing actual budgets, only expenditure, we require a parametric model and estimation approach 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 unobserved, our approach requires us to require that the equilibrium condition in (9) and the equilibrium laws of motion of mental account balances in (6) and(7) 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, $\boldsymbol{\Gamma}_{it}$. 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). Our prior assumptions, the likelihood function, an identification discussion, and a thorough description of the sampling algorithm are all featured in detail in Online Technical Appendix B. Constrained by page limitations, we focus the paper on both the new theoretical insights that result from our model formulation as well as their empirical counterparts which result from our estimation routine.

Our structural estimation procedure targets the posterior distribution of model parameters conditional upon observed data. We must estimate budget weights, mental accounts, and cognitive frictions as latent time series variables, $\{\boldsymbol{\theta}_{it}, \boldsymbol{a}_{it}, \boldsymbol{\Gamma}_{it}\}_{t=1}^{T_i}$ for each i. Changes to budgets are conditionally identified by level shifts in average spending. The unknown initial mental account balance $\boldsymbol{a}_{i1}$ is conditionally identified by the first-period deviation from average spending. The budgets themselves are identified conditional upon the other parameters and the parametric structure of the model itself.

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.[15] In this section we first explore the posterior distribution estimates from several versions of the theoretical model with different behavioral features turned on or off in order to understand how our modeling assumptions contribute to overall model fit. We then engage in a quantitative discussion regarding the implications of our estimates for behavioral features such as budgeting, anchoring, and choice bracketing.

6.1 Posterior Distribution Estimates

We estimate several different model specifications using our MCMC estimation algorithm. We then assess and compare the model fit of each in order to understand the degree to which budgeting frictions and mental accounting features affect predictive inference. First, we estimate the full model described in Section 3. Measures of fitness for this model are in the first row of Table 2, where we allow ψ_{ij}, which governs the marginal probability of a budget update, to be interior to the unit interval and the anchoring-effect of $\gamma_i = 1$. Then we turn off different model features one at a time and re-run the estima-

[15]The exact dimension of the parameter space is 4,664,212 with 4,598,976 of these parameters being the values of latent, time-dependent mental accounting variables.

tion under different parameterizations. Specifically, our five different model estimations are run under the following assumptions:

i. Heterogeneous rationality $\psi_{ij} \in (0,1)$ and full anchoring, $\gamma_i = 1$ (baseline model).

ii. No anchoring, $\psi_{ij} \in (0,1)$ and $\gamma_i = 0$.

iii. Constant budget weights, $\psi_{ij} = 0$ and $\gamma_i = 1$.

iv. Constant budget weights with no anchoring, $\psi_{ij} = 0$ and $\gamma_i = 0$.

v. No budget-updating frictions, $\psi_{ij} = 1$ and $\gamma_i = 1$ (full rationality model).

The last model is designed to best mimic a classical two-stage budgeting problem with ex-ante preference uncertainty and constant expenditure share preference weights.

Table 2: Model Performance and Comparisons

Model	MAE / $\widehat{\ell}_i$ [a]	lppd [b]	lppd % Baseline [c]	*MH Acc. Rates* Mean	Median	Mode
$\psi_{ij} \in (0,1)$, $\gamma_i = 1$	0.236	-208,208,458	—	0.159	0.196	0.231
$\psi_{ij} \in (0,1)$, $\gamma_i = 0$	0.131	-208,209,461	-4.818×10^{-4}	0.220	0.225	0.230
$\psi_{ij} = 0$, $\gamma_i = 1$	3,904.927	-208,208,979	-2.504×10^{-4}	0.060	0.017	0.008
$\psi_{ij} = 0$, $\gamma_i = 0$	180,077.100	-208,209,690	-6.368×10^{-4}	0.178	0.194	0.215
$\psi_{ij} = 1$, $\gamma_i = 1$	0.148	-208,209,783	-5.920×10^{-4}	0.199	0.230	0.231

[a] MAE / $\widehat{\ell}_i = \frac{1}{IJ} \sum_i \sum_j \frac{1}{T_i} \frac{\sum_t |\widehat{x}_{ijt} - x_{ijt}|}{\widehat{\ell}_i}$ where the predictive mean is $\widehat{x}_{ijt} = \frac{1}{N} \sum_n \widetilde{x}_{ijtn}$ with atomic MCMC predictions $\widetilde{x}_{ijtn}$, indexed by $n \in \{1, \ldots, N\}$.

[b] lppd is the log-pointwise predictive density which we take from Gelman et al. (2013). This number is $\text{lppd} = \sum_i \sum_t \sum_j \ln \left(\frac{1}{N} \sum_n \phi\left(\frac{\widetilde{x}_{ijtn} - x_{ijt}}{\sigma_{ijn}} \right) \right)$. Bigger numbers (the closer to 0 in our case) indicate better model fits.

[c] Since the baseline estimation of the full model with $\psi_{ij} \in (0,1)$ and $\gamma_i = 1$ has the largest lppd, this column represents the percentage loss of information in other models' predictive power relative to the baseline. Negative numbers indicate worse fit than baseline, positive better. Note that all values are negative, so that the full baseline model fits best.

In each estimation, we initialize the MCMC integration scheme with random generates from the prior distributions, except for $\boldsymbol{a}_{i1}$ which we set to $\mathbf{0}$ to start. We then proceed to iterate through a chain of length 100,000. That is, we operate on the agent-level MH within Gibbs blocks 100,000 times and the global blocks 100,000/100 = 1,000 times. We set a burn-in period of 30,000, keeping every 100^{th} agent-level draw and every global draw thereafter, giving us a total of 701 draws after burn-in and trimming. The sampler operates with agent-level blocks parallelized over a 128-core computer, requiring just under 48 hours to completion. For the internal Gaussian quadrature routine on (9), we use $4^4 = 256$ quadrature points.

In Table 2 we present measures of model performance and fitness.[16] In the full baseline estimation, ψ_{ij} is fully sampled and allowed to vary in $(0,1)$ and agents' expenditure time series are serially dependent via a_{ijt} since $\gamma_i = 1$. Our two main summary statistics comparing model fit are 1) mean absolute predictive error (MAE) as a fraction of each agent's average income $\widehat{\ell}_i$ and 2) a log-pointwise predictive density (lppd) information criterion as described in Gelman et al. (2013). As a measure of model convergence and performance we also present the mean, median, and modal MH acceptance rates for the agent-specific MH within Gibbs sampler. Model fit is notably poor when budget weights are assumed to be constant, $\psi_{ij} = 0$. When budget updating frictions are lifted ($\psi_{ij} = 1$ with $\gamma_i = 1$) and serial dependence is shut off ($\gamma_i = 0$ with $\psi_{ij} \in (0,1)$), MAE as a fraction of average income is lower than in the full model, yet the lppd information criterion suggests the full model is still a better fit. For this reason the full baseline model is our preferred specification because of this, so we refer to its parameter estimates throughout our empirical and counterfactual analyses.

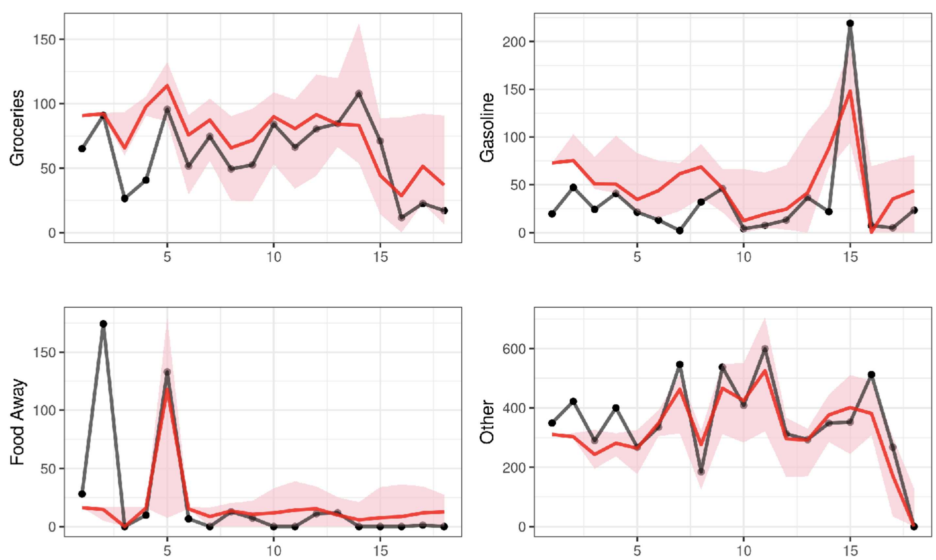

Figure 4: Black lines represent actual weekly spending for the median income agent, while red lines represent the baseline model's predicted means with 95% confidence regions in pink.

[16]In addition we present a grid of plots for the various agent-level parameter means and global parameter draws from the MCMC routine in Online Technical Appendix C to visualize the sampler's autocorrelation in n. Given our burn-in period length and trimming, autocorrelation is minimal and the sampler appears to have converged.

Table 3: Posterior Summary Statistics for Agent-Level Means, $\psi_{ij} \in (0,1)$ & $\gamma_i = 1$

What	Parameter	5%	1st Qu.	Median	Mean	3rd Qu.	95%	S.D.
Borrowing Limit	m_i	2,153.80	3,524.10	4,784.50	5,311.60	6,449.40	10,212.74	2,686.20
Long-run Expenditure Share	$\widehat{\alpha}_{i1}{}^{a}$	0.029	0.059	0.092	0.106	0.135	0.228	0.064
	$\widehat{\alpha}_{i2}$	0.037	0.080	0.130	0.146	0.192	0.313	0.088
	$\widehat{\alpha}_{i3}$	0.039	0.079	0.125	0.150	0.197	0.340	0.097
	$\widehat{\alpha}_{i4}$	0.034	0.506	0.605	0.598	0.706	0.822	0.148
Standard Deviation for Mental Accounts	$\widehat{\sigma}_{a_{i1}}$	0.006	0.009	0.012	0.013	0.015	0.024	0.007
	$\widehat{\sigma}_{a_{i2}}$	0.007	0.010	0.013	0.015	0.018	0.028	0.009
	$\widehat{\sigma}_{a_{i3}}$	0.008	0.012	0.016	0.018	0.022	0.034	0.011
	$\widehat{\sigma}_{a_{i4}}$	0.019	0.025	0.031	0.034	0.038	0.055	0.023
Likelihood Standard Deviation	$\widehat{\sigma}_{i1}$	10.413	17.618	24.948	29.422	35.938	58.297	24.353
	$\widehat{\sigma}_{i2}$	10.703	19.863	30.062	45.888	49.646	134.167	54.355
	$\widehat{\sigma}_{i3}$	15.890	30.065	46.923	63.459	72.629	169.564	63.851
	$\widehat{\sigma}_{i4}$	53.272	99.579	153.511	245.444	226.409	492.647	2,030.330
Shape Parameter for Likelihood Variance Prior	$\widehat{\tau}_{i1}$	0.021	0.124	0.163	0.750	0.303	3.660	2.546
	$\widehat{\tau}_{i2}$	0.020	0.115	0.160	0.997	0.298	4.576	3.340
	$\widehat{\tau}_{i3}$	0.017	0.103	0.139	0.903	0.253	3.808	3.633
	$\widehat{\tau}_{i4}$	0.013	0.081	0.106	0.569	0.189	2.881	2.015
Scale Parameter for Likelihood Variance Prior	$\widehat{\beta}_{i1}$	1.016	1.112	1.169	1.741	1.303	4.594	2.512
	$\widehat{\beta}_{i2}$	1.011	1.102	1.165	1.988	1.293	5.614	3.307
	$\widehat{\beta}_{i3}$	1.010	1.087	1.146	1.900	1.264	4.780	3.624
	$\widehat{\beta}_{i4}$	0.999	1.064	1.14	1.567	1.202	3.870	2.007
Liquidity Preferences	$\widehat{\alpha}_{i,J+1}$	8.763	10.483	11.557	12.166	13.047	17.521	3.336
Mean of $\widehat{\alpha}_{i,J+1}$	$\widehat{\mu}_{\alpha_{i,J+1}}$	15.536	15.559	15.576	15.577	15.593	15.619	0.026
Probability of Budget Update	$\widehat{\psi}_{i1}$	0.293	0.531	0.704	0.666	0.832	0.913	0.197
	$\widehat{\psi}_{i2}$	0.280	0.508	0.674	0.645	0.812	0.902	0.197
	$\widehat{\psi}_{i3}$	0.250	0.448	0.599	0.589	0.750	0.879	0.197
	$\widehat{\psi}_{i4}$	0.152	0.347	0.540	0.535	0.727	0.882	0.232
Budgeted Share of Income Chosen by the Consumer	$\widehat{\theta}_{i1}{}^{b}$	0.028	0.057	0.092	0.202	0.142	0.270	1.968
	$\widehat{\theta}_{i2}$	0.037	0.079	0.128	0.326	0.200	0.355	4.362
	$\widehat{\theta}_{i3}$	0.037	0.081	0.131	0.448	0.217	0.424	6.893
	$\widehat{\theta}_{i4}$	0.376	0.564	0.681	2.158	0.800	1.092	40.164
Indicator that Represents Cognitive Frictions w.r.t. Budget Updates	$\widehat{\Gamma}_{i1}{}^{c}$	0.304	0.175	0.307	0.348	0.488	0.943	0.205
	$\widehat{\Gamma}_{i2}$	0.291	0.196	0.340	0.369	0.513	0.935	0.205
	$\widehat{\Gamma}_{i3}$	0.260	0.260	0.417	0.427	0.575	0.911	0.205
	$\widehat{\Gamma}_{i4}$	0.156	0.283	0.476	0.484	0.680	0.919	0.241
Mental Account Balances	$\widehat{a}_{i1}$	−10.406	−4.310	−1.700	−0.305	1.410	10.685	36.213
	$\widehat{a}_{i2}$	−14.962	−4.045	−0.807	0.420	3.055	19.626	33.026
	$\widehat{a}_{i3}$	−26.813	−8.967	−3.824	−3.187	1.166	16.576	62.525
	$\widehat{a}_{i4}$	−83.045	−27.907	−8.817	−158.691	6.692	68.265	8,556.193
Global Mean of $\widehat{\mu}_{\alpha_{i,J+1}}$	$\widehat{\widehat{\mu}}_{\mu}{}^{d}$	9.677	12.351	15.269	15.260	18.036	20.959	3.593
Global S.D. of $\widehat{\mu}_{\alpha_{i,J+1}}$	$\widehat{\widehat{\sigma}}_{\mu}$	0.368	0.537	0.659	0.623	0.708	0.784	0.127

[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 baseline model where $\Gamma_{ijt} = 1$ if $|\theta_{ijt} - \theta_{ij,t-1}| > 0$.

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

To visually illustrate how well the baseline model fits the data, Figure 4 presents the time series of expenditures (black lines) along with predictive means (red lines) from the full model estimation with $\psi_{ij} \in (0,1)$ and $\gamma_i = 1$ for the median income consumer. Summary statistics over agent-level means for parameter estimates are presented in Table 3. For 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.

One thing stands out in Table 3 amongst the mental accounting parameters: the variances of the posterior distributions of agent-level mental account balance means $\widehat{\boldsymbol{a}}_i$ are large. Note that some of the likelihood variance estimates $\widehat{\sigma}^2_{ij}$ are also large. These high values are not surprising given the spikiness of the various time series. Remember, $a_{ij,t+1,n} = -\zeta_{ijtn}$, so high variance in ζ_{ijt} will lead to high variance in $\widehat{a}_{ij}$ as well. This variability propagates through the model via optimal budget updates ϑ^*_{iytn}. Lacking explicit budgeting data, this posterior variability is to be expected. Further, estimates at the fifth and ninety-fifth percentiles are reasonable, so high likelihood variance should not be cause for suspicion with respect to the validity of our results.

6.2 Budget-updating Behavior

Our model estimates provide empirical inference regarding the degree to which consumers face cognitive constraints with respect to budgeting. Here, we focus on the model's predicted number of total budget updates k_{it} each period.

By definition a budget update has occurred if $\theta_{ijt} \neq \theta_{ijt}$. In such a case, $\Gamma_{ijt} = 1$ regardless of the magnitude of the difference $|\theta_{ijt} - \theta_{ij,t-1}|$. This means that we may infer $\Gamma_{ijt} = 1$ even if the magnitude of the budget-share differences is very small. Since neither budget shares nor the shift indicators are known, model-estimated incremental changes of small magnitudes might be attributable to noise. In our baseline estimations we allow $\Gamma_{ijt} = 1$ any time $\theta_{ijt} \neq \theta_{ijt}$. We then consider several thresholds $\underline{\epsilon} > 0$ where we encode a switch (i.e., $\Gamma_{ijt} = 1$) when the threshold is exceeded — $|\theta_{ijt} - \theta_{ij,t-1}| > \underline{\epsilon}$.

Figure 5 provides visual summaries of the estimated total number of budget switches $k_{itn} = \sum_{j=1}^{J} \Gamma_{ijtn}$ under different updating thresholds. $\underline{\epsilon}$ roughly represents the minimum fraction (in units of income) the consumer would ever consider adjusting his budget on a week-by-week basis. For example, if a consumer makes on average \$100 per week, then $\underline{\epsilon} = 0.01$ means that the consumer would never actively adjust his budget by less than $\$1 = 0.01 \cdot \$100$. In panel (a) we present the density of agent-level averages, $\widehat{k}_i$. In the baseline model we estimate that consumers make an average of 2.48 budget updates

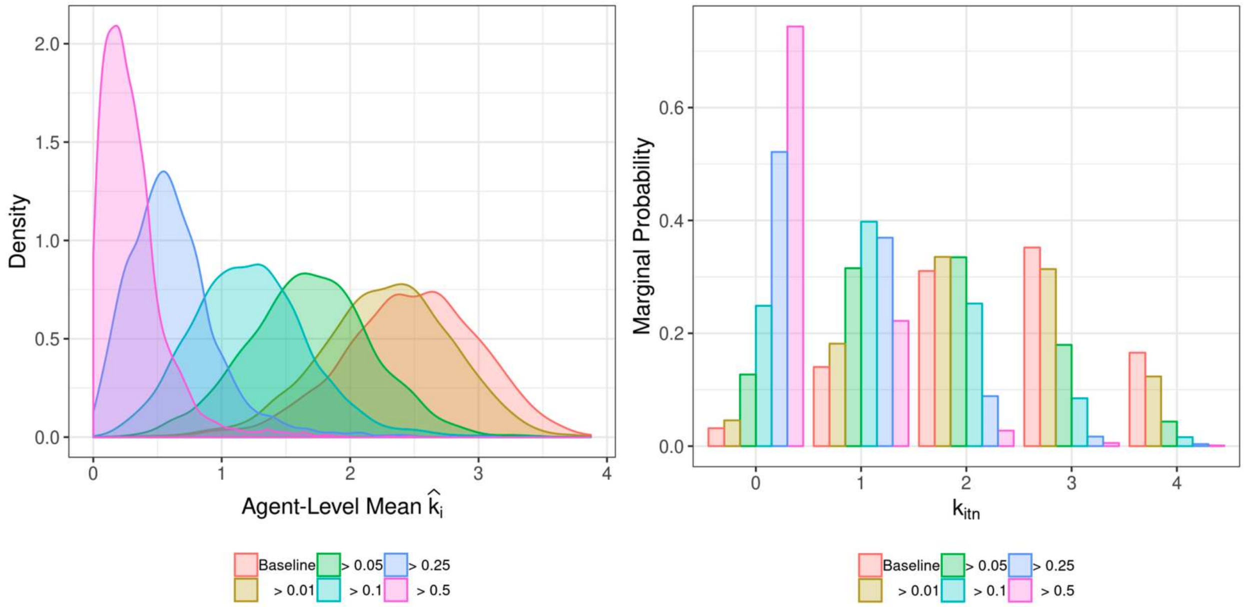

(a) (b)

Figure 5: In panel (a) we present the posterior density of $\widehat{k}_i = \frac{1}{N\,T_i}\sum_n \sum_t \sum_j \Gamma_{ijtn}$, i.e. the number of budget changes each period 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. In other scenarios we allow for the fact that model-estimated incremental changes may just represent noise around some true budget value. In these scenarios $\Gamma_{ijtn} = 1$ only if $|\theta_{ijtn} - \theta_{ij,t-1,n}| > \underline{\epsilon}$ where $\underline{\epsilon} \in \{0.01, 0.05, 0.1, 0.25, 0.5\}$.

every week. This number falls when we consider different minimum thresholds — 2.29 with a 1% threshold, 1.70 with a 5% threshold, 1.22 with a 10% threshold, 0.61 with a 25% threshold, and 0.30 with a 50% threshold. In panel (b) we present the marginal distribution of total budget updates across agents, time, and epochs. Even with only a 5% threshold (green), less consumers appear fully rational choosing $k_{it} = J = 4$ less than 5% of the time. With no threshold full rationality is exhibited just over 15% of the time.

Table 4 presents the fractions of budget updates that satisfy the given tolerance thresholds relative to the baseline threshold of zero. In the top half of the table, we compare the different thresholds to the baseline model ($\psi_{ij} \in (0, 1)$ and $\gamma_i = 1$). In the bottom half of the table, we compare the different thresholds to the model where consumers are assumed to be fully rational, so that $\psi_{ij} = 1$ and thus $\Gamma_{ijt} = 1$ always. We find very little difference between the two models regarding the fraction of budget-share updates that surpass the pre-defined tolerance thresholds. This attests to the stability of our results: regardless of whether consumers are assumed boundedly or fully rational, we estimate that first-order behavior is on average the same. This result lends support to our assumption that k_{it} is orthogonal to the budget-updating decision process in the boundedly-rational model. This is because, otherwise, we would expect more budget updates to exceed higher tolerances in the boundedly-rational model than the fully-rational model, where consumers are assumed to always make incremental adjustments. If k_{it} were to be endogenously linked to the magnitude of budget adjustments, then in such a model where k_{it} is allowed to be $< J$, sticky-budgeting behavior over multiple periods would lead to adjustments of high magnitude when they do occur, however infrequently, but that is not what we observe.

Broadly, our results indicate that few consumers update all of their budgets every period, providing evidence for 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. Thus, while agents do not consistently update their budgets every week, neither do they neglect these budgets over long extended periods.

Table 4: Share of Budget Updates by Magnitude of Change, Conditional on $\Gamma_{ijtn} = 1$

	(Baseline) Conditional on $\gamma_i = 1$ and $\psi_{ij} \in (0, 1)$				
	> 0.01	> 0.05	> 0.1	> 0.25	> 0.5
Groceries	0.888	0.551	0.310	0.083	0.028
Auto/Gas	0.908	0.615	0.385	0.135	0.052
Food Away	0.923	0.701	0.505	0.214	0.079
Other	0.984	0.921	0.842	0.628	0.368
Total	0.923	0.684	0.493	0.247	0.121
	(Full Rationality) Conditional on $\gamma_i = 1$ and $\psi_{ij} = 1$[a]				
	> 0.01	> 0.05	> 0.1	> 0.25	> 0.5
Groceries	0.874	0.532	0.300	0.079	0.025
Auto/Gas	0.893	0.591	0.366	0.132	0.052
Food Away	0.897	0.672	0.489	0.223	0.092
Other	0.985	0.927	0.859	0.669	0.430
Total	0.912	0.681	0.503	0.276	0.150

[a] In this simulation's baseline $\theta_{ijt} \neq \theta_{ij,t-1}$ always, since the budget weights are assumed to be optimally updated every period. The values presented in this part of the table represent the fraction of updates that exceed the posted thresholds, $\underline{\epsilon}$.

6.3 The Relationship Between Budgets and Mental Accounts

How do the sign and magnitude of estimated budget adjustments $\Delta\theta_{ijtn} = \vartheta^*_{iytn} - \theta_{ij,t-1,n}$, for $j = \iota_{iyt}$, depend on the underlying mental account balances $\boldsymbol{a}_{itn}$? In this section we consider some empirical analogs to Propositions 2 and 3 and Corollaries 1 and 2 described in Section 4.2. Recall, Proposition 2 states that for commodity groups contained in the inside bracket, $j \in \boldsymbol{\iota}_{it}$, losses and gains due to over- or under-spending in previous periods are fully incorporated into the given commodity group's updated budget share. Yet, Corollary 2 and Proposition 3 demonstrate that the composition of the outside bracket, $j \notin \boldsymbol{\iota}_{it}$, and a consumer's liquidity preferences, $\alpha_{i,J+1}$, will together determine both the sign and magnitude of the *total* change, $\Delta\theta_{ijtn}$. In this section we focus on the fractions of consumers who, on average, on average, appear to decrease their budget weights ($\Delta\theta_{ijtn} < 0$) whenever optimal updates are made ($\Gamma_{ijtn} = 1$), conditional upon the signs of the components of $\boldsymbol{a}_{itn}$ and the composition of $\boldsymbol{\iota}_{itn}$.

In Table 5, rows one and two present the fractions of consumers who, on average, adjust downward ($\Delta\theta_{ijtn} < 0$) after having over- or under-spent respectively in the same category associated with the optimal budget adjustment.[17] In row three, we present the fraction of consumers who, on average, appear to make larger absolute budget changes after over-spending relative to under-spending. This last data point is a measure of the proportion of consumers who tend to exhibit loss aversion with respect to correcting budgets in response to over- or under-expenditure. The definition of loss aversion we thus employ is one that occurs at the intensive margin (differences in the signs of budget updates conditional upon the signs of a_{ijt}), not the extensive margin (whether or not budget updates occur). We expect consumers' budgets to be more sensitive to $a_{ijtn} < 0$ if they are loss averse, and indeed we find that, across every consumption category, a slight majority of consumers in our sample seem to be loss averse. Note, though, that there is substantial heterogeneity, both across consumers and consumption categories, for all of these statistics.

Nonetheless, consumers, overall, tend to adjust their budgets downward both after over- and under-spending, though this tendency is slightly more pronounced after under-spending than over-spending. Note the distinction between the magnitude of the budget adjustment and the probability that a downward, as opposed to upward, adjustment occurs. Loss averse consumers will make budget adjustments of larger *magnitude*

[17]The reader should be aware that the statistics presented in Table 5 do not directly correspond to the total derivatives discussed in Section 4.2. This is because, in the context of the data-generating process we observe, all components of $\boldsymbol{a}_{itn}$ are varying simultaneously, so that within-category and cross-category effects from those categories for which $j \notin \boldsymbol{\iota}_{itn}$ both simultaneously impact the sign and magnitude of atomic estimates for $\Delta\theta_{ijtn}$.

after over-spending. Consumers in general, however, are more *likely* to downward adjust after under-spending, but such adjustments are of less magnitude than those after over-spending. Thus, a consumer may both exhibit loss aversion, according to our definition, but also be less likely (probabilistically) to make such downward adjustments.

Table 5: Budgeting & Mental Accounting Tendencies, Conditional on $\Gamma_{ijtn} = 1$

Conditional on the sign of Within-category a_{ijtn} [a]					
	Groceries	Auto/Gasoline	Food Away	Other	Total [e]
Share Adjust Down, $a_{ijtn} < 0$ [b]	0.561	0.590	0.535	0.466	0.472
Share Adjust Down, $a_{ijtn} > 0$ [c]	0.571	0.599	0.537	0.481	0.500
Share Loss Averse [d]	0.513	0.522	0.510	0.504	0.555

[a] Specifically, we examine the sign of $\Delta\theta_{ijtn}$ conditional on the sign of a_{ijtn}.

[b] The agent i statistic is $\frac{\sum_n \sum_t \Delta\theta_{ijtn}\mathbf{1}\{\Gamma_{ijtn}=1, a_{ijtn}<0\}}{\sum_n \sum_t \mathbf{1}\{\Gamma_{ijtn}=1, a_{ijtn}<0\}} < 0 \quad \text{for} \quad \iota_{iyt} = j$.

[c] Here, just flip the sign of a_{ijtn}: $\frac{\sum_n \sum_t \Delta\theta_{ijtn}\mathbf{1}\{\Gamma_{ijtn}=1, a_{ijtn}>0\}}{\sum_n \sum_t \mathbf{1}\{\Gamma_{ijtn}=1, a_{ijtn}>0\}} < 0 \quad \text{for} \quad \iota_{iyt} = j$.

[d] For each agent i, the underlying statistic is $\left|\frac{\sum_n \sum_t \Delta\theta_{ijtn}\mathbf{1}\{\Gamma_{ijtn}=1, a_{ijtn}>0\}}{\sum_n \sum_t \mathbf{1}\{\Gamma_{ijtn}=1, a_{ijtn}>0\}}\right| - \left|\frac{\sum_n \sum_t \Delta\theta_{ijtn}\mathbf{1}\{\Gamma_{ijtn}=1, a_{ijtn}<0\}}{\sum_n \sum_t \mathbf{1}\{\Gamma_{ijtn}=1, a_{ijtn}<0\}}\right| < 0 \quad \text{for} \quad \iota_{iyt} = j$.

[e] For the "Total" column, we must additionally average over j.

We previously noted that optimal budgets covary negatively with their own mental account balances, which is described in Proposition 2. Note, however, that Proposition 3 demonstrates that the cross-category variation can be ambiguous. This helps reconcile the results we observe in Table 5. Row one, which describes downward adjustments in response to over-spending, demonstrates, for example, that consumers adjust their budgets downward (following over-spending) approximately 59% of the time for "Auto/Gasoline" purchases. This suggests that cross-category variation dominates within-category variation for automotive-related purchases following over-spending, since Proposition 2 suggests we would otherwise expect an upward adjustment. The opposite is true for the catch-all "Other" category, which also is associated, on average, with larger basket shares α_{ij}. By inspecting the expression for $\frac{d\vartheta^*_{iyt}}{da_{ijt}}$ in (10), it is clear that if α_{ij} is large for the "Other" category, the within-category responsiveness will dominate potential ambiguities from cross-category responsiveness.

Recall, by Proposition 3, that budget updates respond to over- or under-spending in outside categories in ways that depend on the underlying preference parameters of the agent. Table 6 shows how cross-category responsiveness for outside categories, $j \notin \boldsymbol{\iota}_{itn}$, impacts the sign of $\Delta\theta_{ijtn}$. Like in Table 5 we consider the sign of budget adjustments in category j conditional on the sign of $a_{ij',tn}$ where $j' \neq j$ and $j' \notin \boldsymbol{\iota}_{itn}$. Downward adjustments are more common across consumers, across commodity groups, and over time for all inside commodity groups except the “Other” category. For “Auto/Gasoline” and “Food Away” we observe a slight, but consistently higher propensity to adjust downward in response to under-spending in all outside categories relative to over-spending. One can see this by noting the “Auto/Gasoline” and “Food Away” in the bottom half of the table all have slightly larger values than those in the top half. This runs counter to what mental accounting theory under broad bracketing would predict, providing model support that consumers are indeed narrow bracketers. Under broad bracketing we would expect more frequent upward adjustments of optimally-updated budgets conditional on under-spending in the outside bracket. In this scenario a consumer would implicitly move what he/her perceives to be excess funds left-over in outside categories to form larger budgets for inside categories. Instead, our results suggest that consumers are more likely to move inside-category budgets downward despite excess perceived funds, a tendency which is reflective of budgeting discipline that results from narrow-bracketing behavior. While “Groceries” and “Other” have more mixed outside-responsiveness differentials conditional on the signs of $a_{ij',tn}$, we only observe a tendency toward broad bracketing mental accounting behavior for the cross-category responsiveness of “Groceries” budgets to “Food Away” and the “Other” budget to “Auto/Gasoline.” Upward adjustments for “Groceries” are thus relatively more frequent after under-spending in the “Food Away” category, which reflects a relative tendency for consumers to persistently integrate gains from going out to eat back into their budget for eating at home. This slight differential thus may be indicative of consumers engaging in persistent bouts of frugality over multiple weeks by sacrificing eating out (arguably more of a luxury) for eating in. With regards to the relationship between “Other” and “Auto/Gasoline,” there is not an obvious or clear story that comes to mind. In general the results here support the notion that consumers are narrow bracketers with respect to mental accounting.

We draw several conclusions regarding mental accounting behavior and bounded rationality from these exercises. First, substantial heterogeneity is observed with respect to whether or not consumers exhibit loss aversion. Approximately 55.5% of consumers — a slight majority — exhibit consistent loss-averse behavior, though this varies across consumers and consumption categories. Second, mental accounting behavior is substantially idiosyncratic, which is not surprising given the observed heterogeneity in agent-level weekly expenditure patterns. Finally, we provide evidence that consumers engage in narrow choice bracketing with respect to mental accounting. That is, they do not appear to use mental accounting to justify over-spending in what category by implicitly re-appropriating funds left-over after under-spending in another category.

Table 6: Budget Updates Conditional on the Outside Bracket, $j' \notin \iota_{itn}$

	Share Adjust Down Conditional on $a_{ij',tn} < 0$			
Outside $a_{ij',tn}$	Groceries	Auto/Gasoline	Food Away	Other
Groceries	—	0.555	0.500	0.479
Auto/Gasoline	0.538	—	0.502	0.469
Food Away	0.554	0.570	—	0.481
Other	0.546	0.570	0.511	—
	Share Adjust Down Conditional on $a_{ij',tn} > 0$			
Outside $a_{ij',tn}$	Groceries	Auto/Gasoline	Food Away	Other
Groceries	—	0.560	0.511	0.476
Auto/Gas	0.545	—	0.524	0.476
Food Away	0.547	0.583	—	0.473
Other	0.549	0.577	0.523	—

6.4 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}. We find the following: 1) budget-updating behavior is generally random and uncorrelated with particular, broad characteristics of the consumer, such as his/her earnings potential and 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 $\widehat{k}_i$ 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 thus show that consumers are substantially heterogeneous with how they engage in budgeting behavior, while providing weak evidence that k_{it} may be endogenous. A 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 should validate our exercises here on datasets that record explicit budgeting behavior.

Figure 6 shows contour plots of the joint distributions of average income, average savings rates, and $\widehat{k}_i$. We find little correlation between the consumer-unit summary statistics: average weekly income appears only slightly positively correlated with average weekly budget updates (Pearson's $\rho = 0.104$), while the average savings rate appears uncorrelated with $\widehat{k}_i$ (Pearson's $\rho = -0.0125$). 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}, appear to be orthogonal to an agent's average income level $\widehat{\ell}_i$, which is exogenous, and propensity to save. We thus conclude that a consumer's average behavior with regards to budget attentiveness does not appear to be associated with either their earnings potential or average savings behavior.

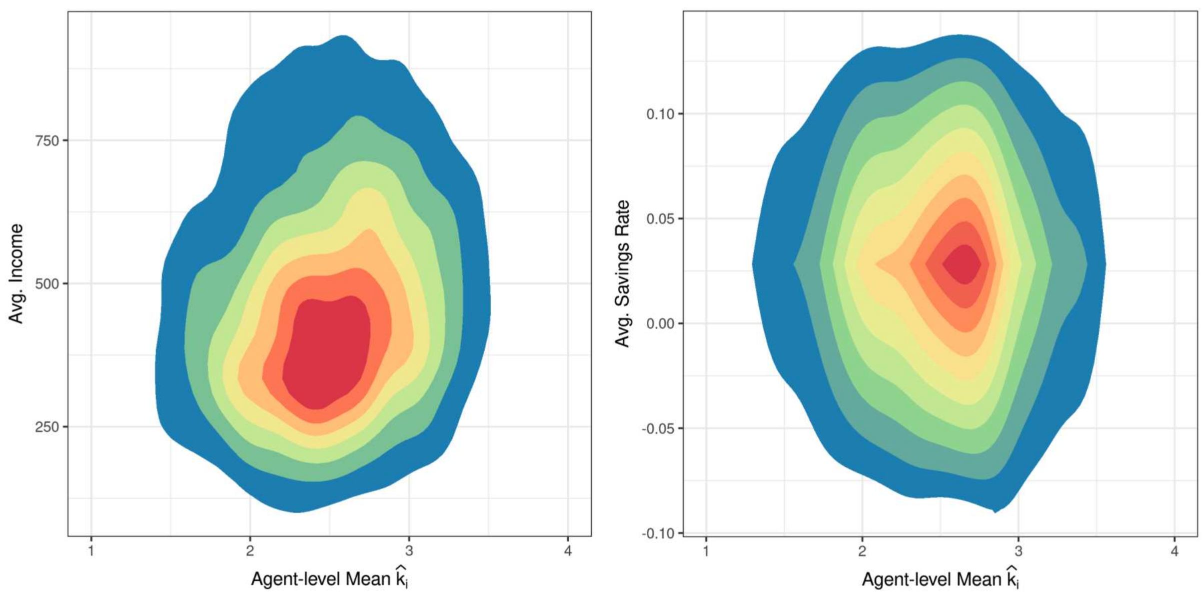

(a) (b)

Figure 6: This figure characterizes the joint densities of agent-level average income (a) and savings (b) and the agent-level average number of budget updates, $\widehat{k}_i$, where such averages are taken over posterior epochs and time.

However, within each consumer unit budget-updating behavior can still be strongly associated with fluctuations in personal income over time. Further, we find that most consumers (69.5% of the sample) 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 Pearson's coefficients for the correlation between ℓ_{it} and posterior $\widehat{k}_{it}$ as well as z_{it}/ℓ_{it} and posterior $\widehat{k}_{it}$ over time within the consumer unit. The density estimates across consumers for these coefficients are presented in Figure 7. 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 left panel. However, for more consumers periods with greater rationality are also associated with higher relative savings rates, which can be seen by noting the density in the right panel is skewed positive.

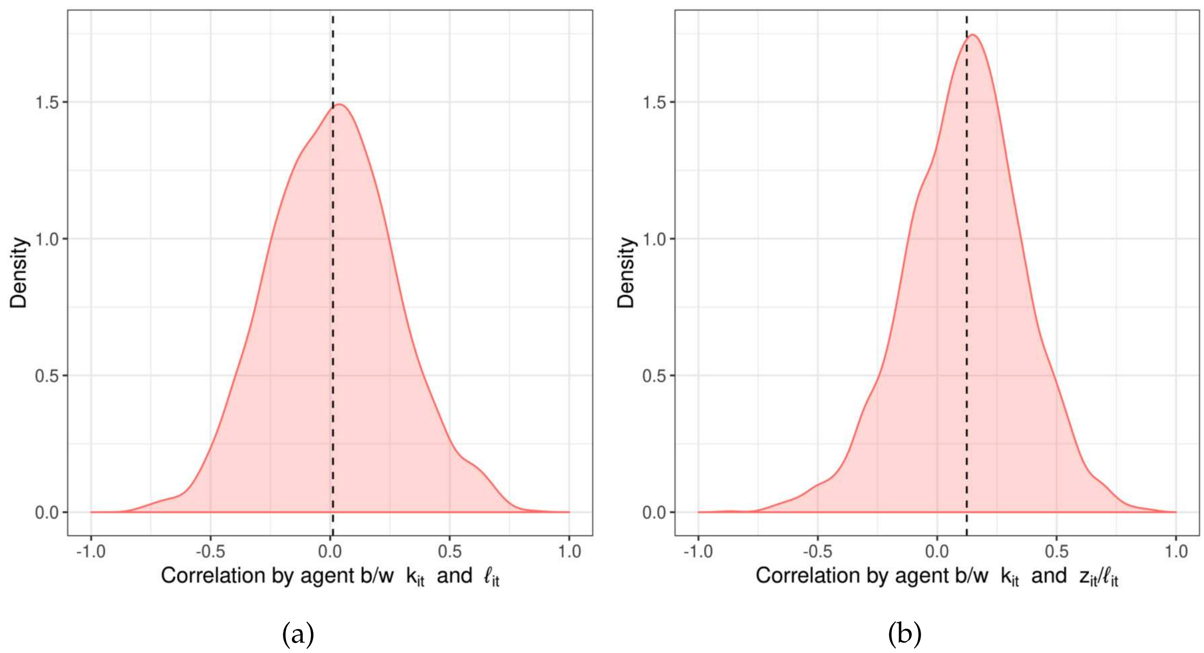

Figure 7: Here we present kernel density estimates of Pearson's coefficients across agents for correlations between $\widehat{k}_{it}$ and ℓ_{it} (a) and $\widehat{k}_{it}$ and z_{it}/ℓ_{it} (b). The median correlation (dashed black lines) in panel (a) is 0.012. In panel (b) it is 0.124.

To assess whether the relationship between budget-updating and income may be nonlinear and thus inadequately explained by Pearson's correlation coefficients, we check if there is any difference between $\widehat{k}_{it}$ in periods when ℓ_{it} is 10% higher or lower than $\ell_{i,t-1}$. Figure 8a plots the conditional kernel density estimates of $\widehat{k}_i$ under different income-shock scenarios against the baseline. Figure 8b plots the densities of differences of $\widehat{k}_i$ relative to the baseline estimate. Specifically, in Figure 8b, negative values mean that the baseline $\widehat{k}_i$ is *less* than the conditional estimate, while positive values mean the baseline estimate is *greater* than the conditional estimate. We find that 60.6% of consumers in our sample tend to make more budget updates in periods when income rises by 10%, while 38.3% of consumers tend to make more updates in periods when income falls by 10%. However, the magnitude of these differences is slight (median difference of -0.029 after 10% positive shocks versus a median of 0.038 after 10% negative shocks). We conclude that while $\widehat{k}_i$ is slightly correlated with income shocks, the relationship does not appear significant.

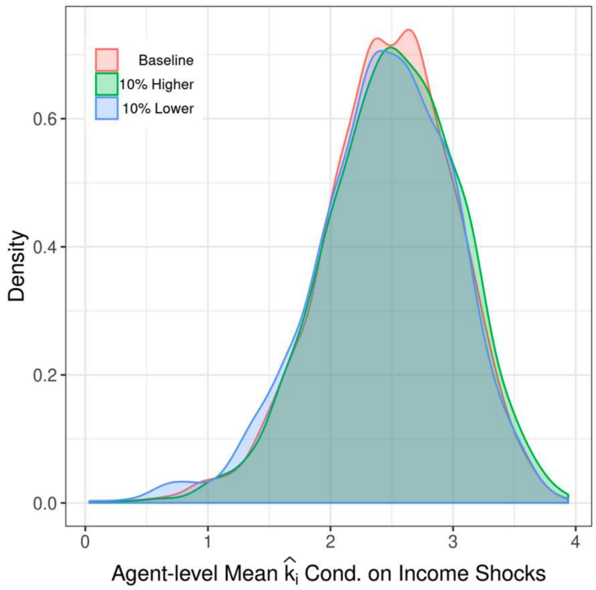

(a)

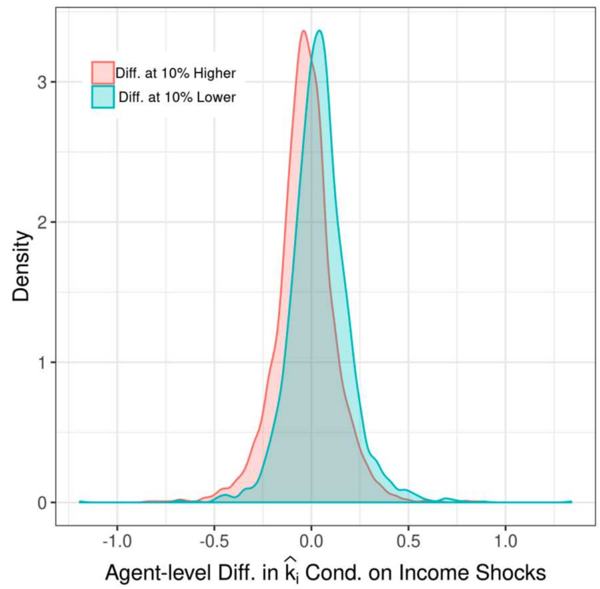

(b)

Figure 8: Panel (a) demonstrates that the estimated distribution of $\widehat{k}_i$, conditional on the consumer realizing a 10% positive or negative income shock, appears unchanged from the baseline estimate. Panel (b) demonstrates 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.

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 $\widehat{k}_{it}$ and $\ell_{i,t\pm s}$, where $s \in \{1,2,3\}$. In Figure 9 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 panel (b) where the densities are centered around zero. In panel (a) we see that there appears to be no systematic relationship between $\widehat{k}_{it}$ and income two and three periods in the past. However, $\widehat{k}_{it}$ and $\ell_{i,t-1}$ are negatively correlated for approximately 62.9% of consumers (red).

Together, the results presented in Figures 8 and 9 provide 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, though, 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

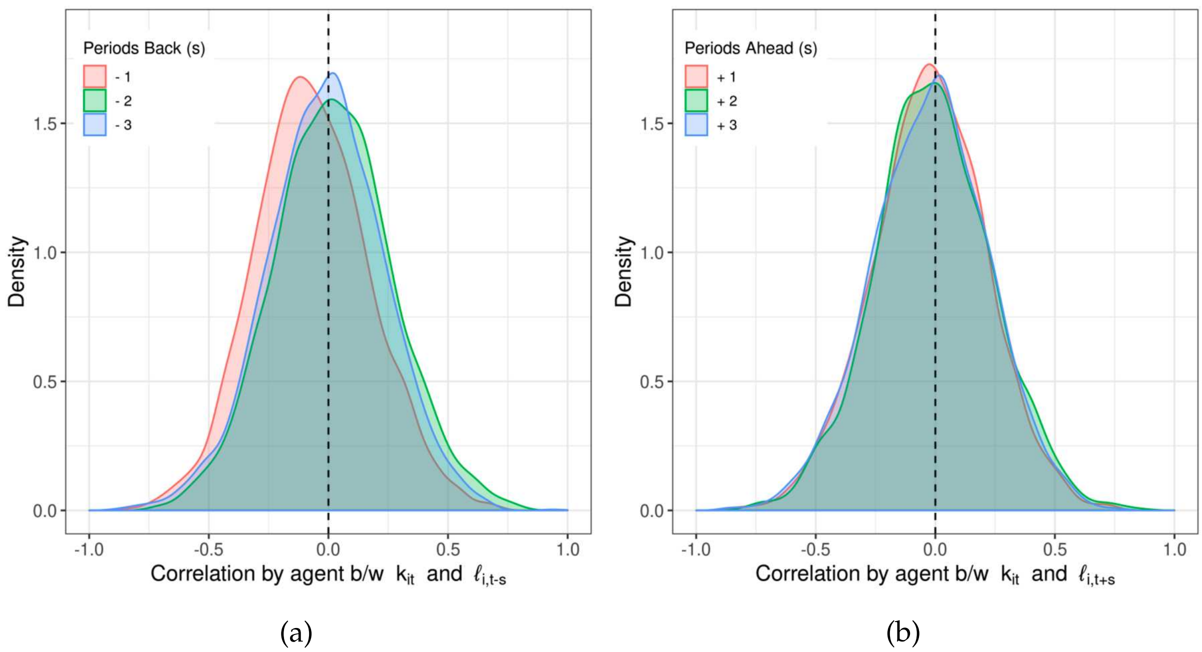

Figure 9: Panel (a) presents the densities of Pearson's coefficients across agents measuring the correlation between period-t $\widehat{k}_{it}$ and $\ell - i, t - s$. Panel (b) presents correlation coefficients measuring the relationship between $\widehat{k}_{it}$ and $\ell_{i,t+s}$.

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 appears to be 'no,' as we find no significant relationships between a consumers 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 relative 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.

6.5 The Relationship Between Savings and Mental Accounting

In this section we explore how savings in period t depends on whether or not the agent is estimated to have previously over- or under-saved. 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 $a_{i,J+1,t} > 0$. We find that after periods in which consumers over-spend on aggregate, they build up account balances by increasing savings rates, while after under-spending savings rates decline, and consumers are more likely to splurge. This is true for 87.1% of consumers in our sample. For the other 12.9% of consumers, savings rates after over-spending are on average lower than savings rates after under-spending. These consumers thus persistently, over multiple periods, engage in over- or under-spending relative to pre-determined budgets and do not appear to use mental accounting techniques to regulate their overall expenditure. Further, we assess whether the persistence of over- or under-saving is associated with the degree to which consumers are estimated to be boundedly rational, finding no significant relationship between estimates of $\widehat{k}_i$ and differences in savings behavior.

Figure 10 presents several plots that characterize relationships between average savings rates $\widehat{z_{it}/\ell_{it}}$ and posterior-estimated average over- or under-spending from the previous period, $\widehat{a}_{i,J+1,t}$. In panels (a) and (b) we show the joint densities of mean savings rates for periods following over- and under-spending, respectively. After over-spending savings rates are skewed positive, while after under-spending they are skewed negative, suggesting that many consumers in our sample appear to use mental accounting to regulate their period-by-period spending and bank account balances. This is evident by the fact that savings as a fraction of income is on average lower (higher), and indeed mostly negative (positive), after under(over-)-spending, so that consumers appear to spend down (replenish) positive (negative) mental account balances. In panel (c) we plot the density of the difference in agent-level average $\widehat{z_{it}/\ell_{it}}$ in periods following over-spending ($\widehat{a}_{i,J+1,t} > 0$) versus under-spending ($\widehat{a}_{i,J+1,t} < 0$). This distribution skews positive since most consumers have positive z_{it}/ℓ_{it} when $\widehat{a}_{i,J+1,t} > 0$ and negative z_{it}/ℓ_{it} when $\widehat{a}_{i,J+1,t} < 0$. Finally, we consider the relationship between the conditional difference in savings rates and $\widehat{k}_i$ finding no statistically significant correlation between the two. Thus, mental accounting consumers are no more likely than those whose savings behavior exhibits stickiness inconsistent with mental accounting to be more or less boundedly rational.

For most consumers $\widehat{z_{it}/\ell_{it}}$ and $\widehat{a}_{i,J+1,t}$ are positively correlated. In Figure 11 we select two consumers from our sample who differ in their savings behavior as it relates to mental accounting. In both panels black lines represent total absolute savings z_{it} while red lines are posterior average $\widehat{a}_{i,J+1,t}$ with a corresponding 95% confidence region in pink. Recall, when $\widehat{a}_{i,J+1,t} > 0$, the consumer under-saved (over-spent) in the previous period. Panel (a) features data from the median income consumer unit in our sample, while panel (b) features data from a consumer who responds to under-saving by continuing to negatively

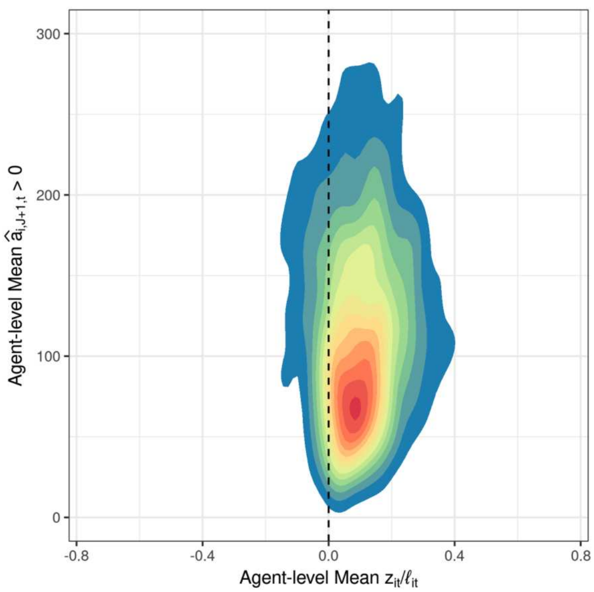

(a) Savings Rates After Over-spending

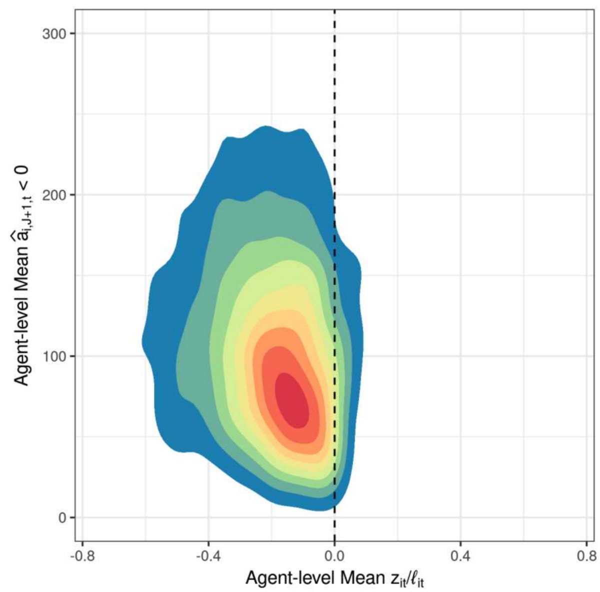

(b) Savings Rates After Under-spending

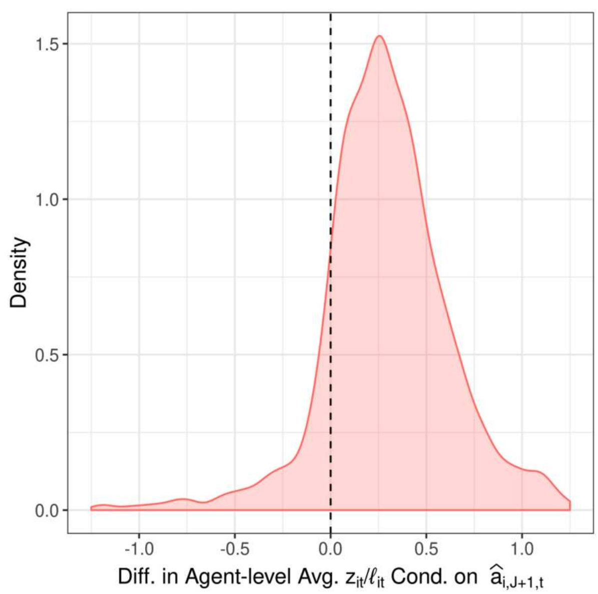

(c) Density of Differences Across Agents

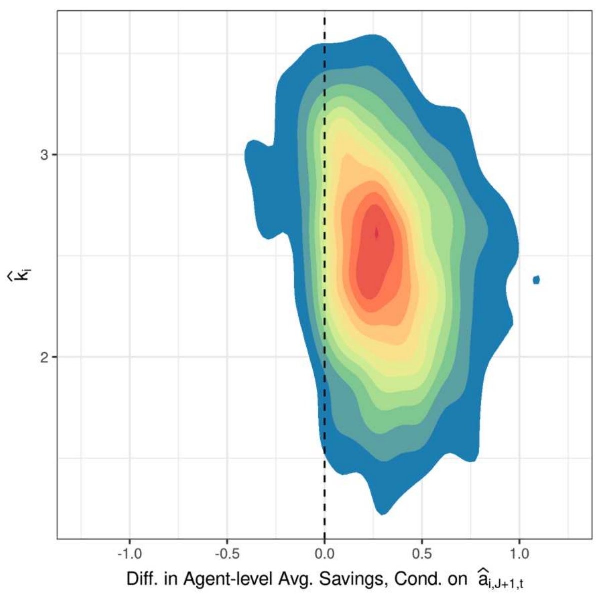

(d) Savings Differentials and Budget-updating

Figure 10: Here we plot several kernel density estimates to demonstrate how savings rates respond to relative over- or under-spending. Panels (a) and (b) show the joint densities, across agents, of average savings rates z_{it}/ℓ_{it} conditional on $\widehat{a}_{i,J+1,t} > 0$ and $\widehat{a}_{i,J+1,t} < 0$, respectively. The vertical axis in panel (a) features the absolute value of $|\widehat{a}_{i,J+1,t}|$. Panel (c) plots the density of the difference between conditional savings rates when $\widehat{a}_{i,J+1,t} > 0$ and savings rates when $\widehat{a}_{i,J+1,t} < 0$. Panel (d) compares this same difference with the agent-level average number of budget updates, $\widehat{k}_i$.

save over subsequent periods. Meanwhile, the median income consumer behaves according to standard mental accounting theory: s/he increases savings after under-saving, thus replenishing depleted balances. For the consumer in panel (b), z_{it} and $\widehat{a}_{i,J+1,t}$ appear to co-move negatively, suggesting a period of under-saving is followed by successive periods of under-saving, though each of lesser magnitude, until finally after a few periods the persistent cycle of under-saving is broken. The two examples thus illustrate the degree to which consumers in our sample are heterogeneous with respect to how they use mental accounting to regulate total spending and savings.

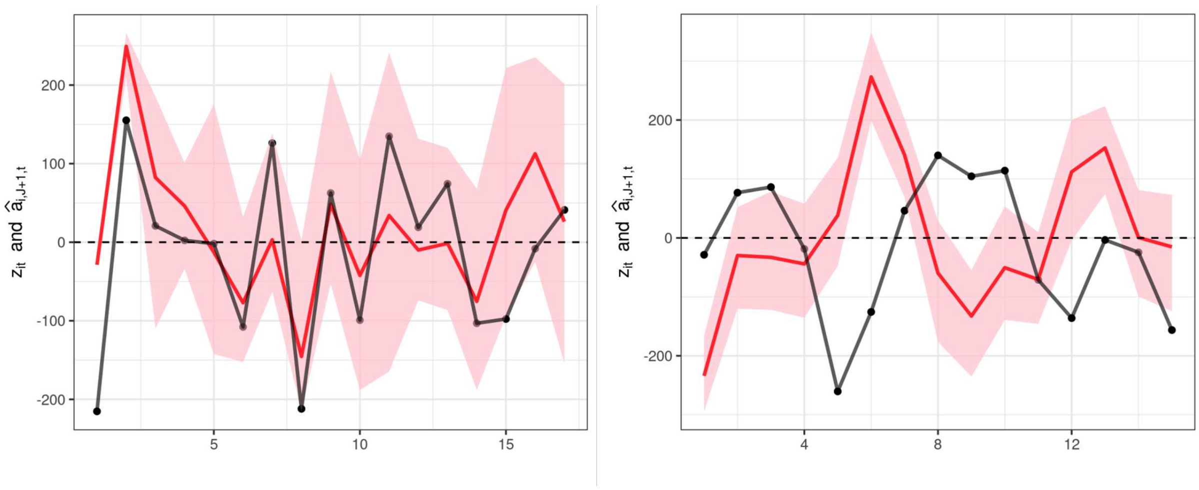

(a) Median Income Consumer

(b) Consumer with Negative Co-movement

Figure 11: Black lines represent z_{it}, while red lines represent posterior average $\widehat{a}_{i,J+1,t}$. The 95% confidence region is in pink. In panel (a) the median income consumer saves 33.2% more after over-spending than after under-spending. This is close to the median consumer in our sample who saves 29.2% more on average, where the first and third quartiles of this distribution are 11.1% and 50.7%, respectively. The consumer in panel (b) saves on average 23.4% less after over-spending than after under-spending.

6.6 Counterfactual Welfare Under Relaxed Rationality Constraints

To understand how consumer welfare is affected by relaxing the sparse max constraint on budget updates, we counterfactually simulate consumer expenditure using the posterior distribution of parameters from the full baseline estimation, except we force $\Gamma_{ijt} = 1$ for all i, j, and t combinations. This is equivalent to setting the primitive preference parameter $\psi_{ij} = 1$, for all i and j. 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. Note 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 main estimation 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 full baseline 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 predictive outcomes under coarse variation in ψ_{ij}.

In the counterfactual set up, each period every agent chooses ϑ^*_{it} by maximizing indirect utility subject to the relaxed constraint that $k_{it} = J$ always. Broadly speaking, under this counterfactual the average consumer in our sample experiences a 1.8% reduction in posterior average predicted flow utility when $\psi_{ij} = 1$ versus when $\psi_{ij} \in (0, 1)$. For 70.4% of consumers, engaging in weekly budget updates for every commodity category does not lead to welfare improvements relative to the posterior predictive baseline.

However, as we have thus far emphasized, there is substantial heterogeneity with respect to how different individual consumers' behaviors change when rationality constraints are relaxed. In fact we find that those who are most rationally constrained are most vulnerable to experiencing adverse outcomes due to interventions designed to increase financial attentiveness by sending $k_{it} \to J$.

Let $\widetilde{u}_i$ be the average (over atomic epochs and across time) counterfactual posterior flow utility of agent i. Let $\widehat{u}_i$ be the average (again, over atomic epochs and across time) predicted posterior flow utility of agent i. The ratio $\widetilde{u}_i/\widehat{u}_i$ is our primary measure of counterfactual variation in individual welfare due to relaxing the rationality constraint. If $\widetilde{u}_i/\widehat{u}_i > 1$, the consumer is better off when fully rational. If $\widetilde{u}_i/\widehat{u}_i < 1$, the consumer is worse off when fully rational. Amongst those who are worse off, the most vulnerable have $\widetilde{u}_i/\widehat{u}_i < 0$ due to $\widetilde{u}_i < 0$. Note that average utility is negative if, on average, savings plus account balances approach the borrowing limit, i.e. $z_{it} + r_t b_{it} \to -m_i$.[18] We classify this last group of consumers as bankrupt, and they constitute a subset of those who are worse off: just over 0.5% of the sample goes bankrupt which is approximately 1% of those who are worse off.

Figure 12 presents kernel density estimates of $\widehat{k}_i$ conditional on consumers being "Better Off," "Worse Off," and/or going "Bankrupt" due to the relaxation of the rationality constraint. On the whole there is no statistically significant difference between the distributions of consumers who are "Better Off" and "Worse Off," though "Bankrupt" con-

[18]Posterior predicted average utility, $\widehat{u}_i$, is always greater than zero for all consumers since borrowing limits are never reached in reality.

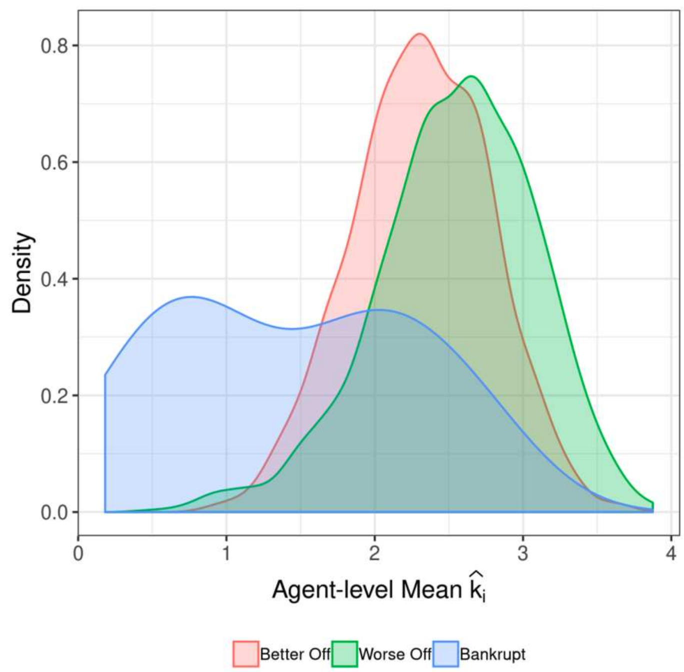

Figure 12: This figure features the density of $\widehat{k}_i$ conditional on the counterfactual type of consumers. Those who are "Better Off" derive higher utility from being more budget attentive, while those who are "Worse Off" derive lower utility. Consumers who go "Bankrupt" constitute a small subset of those who are "Worse Off."

sumers have significantly (statistically speaking) different underlying behavioral profiles than either of the broad groups. Of those who go bankrupt, they have on average lower $\widehat{k}_i$ in the baseline estimation. Thus, they appear to be more likely to use sticky budgets to regulate their spending.

Now consider these results in the context of financial-planning and budgeting apps, such as Mint, YNAB (You Need a Budget), EveryDollar, Honeydue, and Personal Capital. With such apps, users can opt-in to receive push notifications to their smart phones telling them when their spending in a given period for a specific spending category is approaching some pre-set budget. If the app allows consumers to readily change their budgets without incurring a penalty, consumers may opt to raise their budgets in the middle of a period in order to avoid being continually pestered by push notifications. At the extreme, this behavior would lead to the bankruptcies we observe, but can still make other consumers who do not go bankrupt worse off. The push notifications, presumably designed to nudge consumers toward financial discipline, could thus have the opposite effect. A consumer who raises his budget in response to push notifications now anchors future spending around a higher pre-set target while income remains the same. This then causes savings to fall, and the process may repeat itself over future periods until the borrowing limit is reached.

Our results thus provide evidence that consumers who use sticky budgets as heuristic rules of thumb to regulate spending patterns are most vulnerable to adverse outcomes when such budgets can be easily changed. This counterfactual result should be informative to financial-planning app designers. Apps should limit the ability of consumers to change budgets on the fly, either by placing restrictions on how often or when budgets can be changed or by imposing some kind of cost. Even simply making it difficult for a consumer to change a budget by burying such features deep in an app could mitigate potential backfires.

7 Conclusion

We have developed a structural model of two-stage budgeting with bounded rationality and mental accounting features. The model generates reference dependence and loss aversion with respect to expenditure budgets using a standard, quasi-concave, monotone, and continuously-differentiable utility function. By incorporating bounded rationality and mental accounting into the classical two-stage budgeting model, we can endogenously explain idiosyncratic, short-term variation in agent-level consumption expenditure patterns. Further, by allowing for narrow choice bracketing and reference dependence, we show that our model with all the features motivated from behavioral economics generates empirical estimated of behavioral inferences with respect to budgeting, choice bracketing, and anchoring. These results should encourage future work that seeks to unify well-established classical theories with contemporary behavioral ones in order to structurally explain empirical phenomena in consumer decision making.

An important result in our paper 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. Instead, most consumers are 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-none' theories that assume either full-rationality or little-to-no rationality. More broadly, we hope to see the debates between 'rationalists' and 'behavioralists' shift from focusing on whether consumers are rational to instead asking how rational they are, what variability there is in the level of rationality, and what could be driving this variability, as also suggested in Benjamin, Brown, and Shapiro (2013). Whereas previous work has identified within-person variance in the extent to which behaviors follow standard rational (vs. behavioral) models of decision making, such as Olivola and Wang (2016), our work highlights the importance of considering the variability of behavior across consumers.

Finally, we see several avenues for future work. 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 economists looking to study consumer spending in more detail. The work we present here is likely thus the first in a long series of forthcoming papers that may use agent-level, high-frequency data to structurally gain better insight into consumer behavior. Specifically, 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 validate our latent inferences. 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. There thus exists broad potential for many new insights to be gleaned from these treasure troves, and we hope that we have inspired economists to explore such agent-level spending data in more detail.

A Mathematical Appendix

A.1 Inversion of Optimal Budgeting Equation

Suppose there exists at least one j for which $\Gamma_{ijt} = 1$. For some y indexing $\boldsymbol{\iota}_{it}$, let $\boldsymbol{\vartheta}^*_{i,-yt}$ denote the vector of optimally chosen budget shares which does not include y. Note that this vector may be empty. The optimal choice of budget share $\vartheta^*_{iyt}(\boldsymbol{\vartheta}^*_{i,-y,t})$ for good ι_{iyt} can be expressed

$$\begin{aligned}\vartheta^*_{iyt}(\boldsymbol{\vartheta}^*_{i,-y,t}) = \mathbb{E}_{it}\Bigg\{ & \frac{\alpha_{i,\iota_{iyt}}\ell_{it} - \alpha_{i,\iota_{iyt}}\sum_{j\in\boldsymbol{\iota}_{it}(\boldsymbol{\Gamma}_{it})\, j\neq\iota_{iyt}}(\vartheta^*_{ijt}\ell_{it} + \gamma_i a_{ijt} + \zeta_{ijt})}{\ell_{it}(\alpha_{i,\iota_{iyt}} + \alpha_{i,J+1})} \\ & - \frac{(\alpha_{i,\iota_{iyt}} + \alpha_{i,J+1})(\gamma_i a_{i,\iota_{iyt},t} + \zeta_{i,\iota_{iyt},t}) + \alpha_{i,J+1}p_{\iota_{iyt},t}}{\ell_{it}(\alpha_{i,\iota_{iyt}} + \alpha_{i,J+1})} \\ & - \frac{\alpha_{i,\iota_{iyt}}\left(\sum_{j=1}^{J}(1-\Gamma_{ijt})(\theta_{ijt}\ell_{it} + \gamma_i a_{ijt} + \zeta_{ijt}) - m_i - r_t b_{it}\right)}{\ell_{it}(\alpha_{i,\iota_{iyt}} + \alpha_{i,J+1})}\Bigg\}\end{aligned} \tag{A.1}$$

where $\theta_{ijt} = \theta_{ij,t-1}$ for $j \notin \boldsymbol{\iota}_{it}(\boldsymbol{\Gamma}_{it})$, and expectations are over $\boldsymbol{\zeta}_{it}$.

To show this, first, we will make a slight transformation to the indirect utility function which will greatly aid with computation of first order conditions. For each additively separable utility component we can write $\ln(q_{ijt}+1) = \ln(x_{ijt}/p_{jt}+1) = \ln(x_{ijt}+p_{jt}) - \ln(p_{jt})$. We can thus transform indirect utility to permit direct substitution of x_{ijt} for each j without needing to divide out p_{jt}. This yields the expected indirect utility function

$$\begin{aligned}\mathbb{E}_{it}v_{it}(\boldsymbol{\theta}_{it}) = \mathbb{E}_{it}\Bigg\{ & \sum_{j=1}^{J}\alpha_{ij}\ln(\theta_{ijt}\ell_{it} + \gamma_i a_{ijt} + \zeta_{ijt} + p_{jt}) \\ & + \alpha_{i,J+1}\ln\left(\ell_{it} - \sum_{j=1}^{J}[\theta_{ijt}\ell_{it} + \gamma_i a_{ijt} + \zeta_{ijt}] + m_i + r_t b_{it}\right)\Bigg\} - \sum_{j=1}^{J}\ln(p_{jt})\end{aligned} \tag{A.2}$$

where $-\sum_{j=1}^{J}\ln(p_{jt})$ can be safely left outside of the expectation since we assume price levels entering the week are known to the consumer. Now under this parameterization, the first-order condition with respect to ϑ_{iyt} is

$$\mathbb{E}_t \left\{ \frac{\alpha_{i,\iota_{iyt}} \ell_{it}}{\vartheta_{iyt}\ell_{it} + \gamma_i a_{i,\iota_{iyt},t} + \zeta_{i,\iota_{iyt},t} + p_{\iota_{iyt},t}} \right\}$$
$$= \mathbb{E}_t \left\{ \frac{\ell_{it}\alpha_{i,J+1}}{\ell_{it} - \sum_{y'=1}^{k_t}(\vartheta_{iy',t}\ell_{it} + \gamma_i a_{i,\iota_{iy',t},t} + \zeta_{i,\iota_{iy',t},t}) - \sum_{j=1}^{J}(1-\Gamma_{ijt})(\theta_{ijt}\ell_{it} + \gamma_i a_{ijt} + \zeta_{ijt}) + m_i + r_t b_{it}} \right\} \tag{A.3}$$

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

$$\begin{aligned} &\alpha_{i,\iota_{iyt}}\ell_{it} - \alpha_{i,\iota_{iyt}}\ell_{it}\vartheta_{iyt} - \alpha_{i,\iota_{iyt}} \sum_{\substack{j \in \iota_{it}(\Gamma_{it}) \\ j \neq \iota_{iyt}}} (\vartheta_{ijt}\ell_{it} + \gamma_i a_{ijt} + \zeta_{ijt}) \\ &- \alpha_{i,\iota_{iyt}}(\gamma_i a_{i,\iota_{iyt},t} + \zeta_{i,\iota_{iyt},t}) - \alpha_{i,\iota_{iyt}} \sum_{j=1}^{J}(1-\Gamma_{ijt})(\theta_{ijt}\ell_{it} + \gamma_i a_{ijt} + \zeta_{ijt}) + \alpha_{i,\iota_{iyt}}(m_i + r_t b_{it}) \\ &= \alpha_{i,J+1}\vartheta_{iyt}\ell_{it} + \alpha_{i,J+1}(\gamma_i a_{i,\iota_{iyt},t} + \zeta_{i,\iota_{iyt},t}) + \alpha_{i,J+1}p_{\iota_{iyt},t} \end{aligned} \tag{A.4}$$

Isolate ϑ_{iyt} and take expectations to get (A.1).

A.2 Proofs

Proposition 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 no cognitive frictions ($\Gamma_{ijt} = 1$ for all j).

Proof. For the following proof, fix ℓ_{it}, $\boldsymbol{a}_{it}$, and $\boldsymbol{p}_{it}$.

($\Leftarrow$) With perfect foresight, by definition $\boldsymbol{\zeta}_{it} = \boldsymbol{0}$ 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 $\boldsymbol{\zeta}_{it} = \boldsymbol{0}$. Further, note that since z_{it} can be written as a function of q_{it} the equilibrium condition in (8) in the main text, where consumers are choosing ϑ^*_{iyt} such that $j = \iota_{iyt}$, can be written

$$\frac{\partial u_{it}}{\partial q_{ijt}}\frac{\partial q_{ijt}}{\partial \vartheta_{iyt}} + \frac{\partial u_{it}}{\partial z_{it}}\frac{\partial z_{it}}{\partial \vartheta_{iyt}} \tag{A.5}$$

$$= \frac{\partial u_{it}}{\partial q_{ijt}}\frac{\partial q_{ijt}}{\partial \vartheta_{iyt}} + \frac{\partial u_{it}}{\partial z_{it}}\frac{\partial z_{it}}{\partial q_{ijt}}\frac{\partial q_{ijt}}{\partial \vartheta_{iyt}} \tag{A.6}$$

$$= \frac{\partial u_{it}}{\partial q_{ijt}} + \frac{\partial u_{it}}{\partial z_{it}}\frac{\partial z_{it}}{\partial q_{ijt}} = 0 \tag{A.7}$$

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 (A.7). The choices are thus equivalent.

($\Rightarrow$) For this logical direction, we engage in proof by contrapositive. Suppose now that either consumers do not face perfect foresight, so that they do not know $\boldsymbol{\zeta}_{it}$ ex-ante, or there are meaningful cognitive frictions, that is $\sum_{j=1}^{J} \Gamma_{ijt} < J$.

First, let us consider the case where $\boldsymbol{\zeta}_{it}$ 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 $\mathbb{E}_{it} q_{ijt} = \widetilde{q}_{ijt} = \frac{\vartheta^*_{iyt}\ell_{it} + \gamma_i a_{ijt}}{p_{jt}}$. Now replace $\frac{\vartheta^*_{iyt}\ell_{it} + \gamma_i a_{ijt}}{p_{jt}}$ with $\widetilde{q}_{ijt}$ in the expected indirect utility function. Note that for each j optimal expected consumption $\widetilde{q}^*_{ijt}$ must exactly solve

$$\begin{aligned} &\mathbb{E}_{it}\left\{\frac{\partial u_{it}}{\partial \widetilde{q}_{ijt}} + \frac{\partial u_{it}}{\partial z_{it}}\frac{\partial z_{it}}{\partial \widetilde{q}_{ijt}}\right\} \\ &= \mathbb{E}_{it}\left\{\frac{\alpha_{ij}}{\widetilde{q}_{ijt} + \frac{\zeta_{ijt}}{p_{jt}} + 1} - \frac{p_{jt}\alpha_{i,J+1}}{\ell_{it} - \sum_{j=1}^{J}[p_{jt}\widetilde{q}_{ijt} + \zeta_{ijt}] + m_i + r_t b_{it}}\right\} = 0 \end{aligned} \tag{A.8}$$

Ex-post $q^*_{ijt}(\zeta_{ijt})$ satisfies (A.7) by construction. Since the random variable ζ_{ijt} enters (A.8), $\widetilde{q}^*_{ijt} \neq q^*_{ijt}$ except for a measure-zero realization of ζ_{ijt}. Since $\widetilde{q}^*_{ijt}$ 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 $\boldsymbol{\zeta}_{it}$ 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

$\theta_{ij,t-1}$,

$$\frac{\partial u_{it}(\theta_{ij,t-1})}{\partial \vartheta_{iyt}} + \frac{\partial u_{it}}{\partial z_{it}} \frac{\partial z_{it}(\theta_{ij,t-1})}{\partial \vartheta_{iyt}} \neq 0 \tag{A.9}$$

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 $\widetilde{q}_{ijt} = \frac{\theta_{ij,t-1}\ell_{it}+\gamma_i a_{ijt}}{p_{jt}}$, but this value will not force the left hand side of (A.9) to equal zero. Since $\theta_{ij,t-1}$ is fixed, there exists q^*_{ijt} that satisfies (A.7) and $q^*_{ijt} \neq \widetilde{q}_{ijt}$. The proof is complete. □

Proposition 2: Suppose at least one optimal budget update occurs so that $k_{it} > 0$. Assume $\ell_{it} > 0$, $\gamma_i > 0$, $\alpha_{ij} > 0$, $\forall j$, and $\alpha_{i,J+1} > 0$. Without loss of generality, let $\boldsymbol{\iota}_{it} = (1, 2, 3, \ldots)$ and suppose $\boldsymbol{\iota}_{it}$ is of dimension $J' \leq J$. Consider the total responsiveness of the components of $\boldsymbol{\vartheta}^*_{it}$ to a_{iyt} where $y \in \boldsymbol{\iota}_{it}$.

i. Higher a_{iyt} leads to lower ϑ^*_{iyt}, i.e. $\frac{\mathrm{d}\vartheta^*_{iyt}}{\mathrm{d}a_{iyt}} = -\frac{\gamma_i}{\ell_{it}}$.

ii. For all $s \in \boldsymbol{\iota}_{it}$ where $s \neq y$, $\frac{\mathrm{d}\vartheta^*_{ist}}{\mathrm{d}a_{iyt}} = 0$.

Proof. Consider the system of implicit component-wise total derivatives in equation (12) of the main text an suppose, without loss of generality, that $y = 1$:

$$\begin{aligned}
\frac{\mathrm{d}\vartheta^*_{i1t}}{\mathrm{d}a_{i1t}} &= -\frac{\gamma_i}{\ell_{it}} - \sum_{s\neq 1} \Gamma_{ist} \frac{\alpha_{i1}}{\alpha_{i1} + \alpha_{i,J+1}} \frac{\mathrm{d}\vartheta^*_{ist}}{\mathrm{d}a_{i1t}} \\
\frac{\mathrm{d}\vartheta^*_{i2t}}{\mathrm{d}a_{i1t}} &= -\frac{\gamma_i}{\ell_{it}} \frac{\alpha_{i2}}{\alpha_{i2} + \alpha_{i,J+1}} - \sum_{s\neq 2} \Gamma_{ist} \frac{\alpha_{i2}}{\alpha_{i2} + \alpha_{i,J+1}} \frac{\mathrm{d}\vartheta^*_{ist}}{\mathrm{d}a_{i1t}} \\
&\vdots \\
\frac{\mathrm{d}\vartheta^*_{iJ',t}}{\mathrm{d}a_{i1t}} &= -\frac{\gamma_i}{\ell_{it}} \frac{\alpha_{iJ'}}{\alpha_{iJ'} + \alpha_{i,J+1}} - \sum_{s\neq J'} \Gamma_{ist} \frac{\alpha_{iJ'}}{\alpha_{iJ'} + \alpha_{i,J+1}} \frac{\mathrm{d}\vartheta^*_{ist}}{\mathrm{d}a_{i1t}}
\end{aligned} \tag{A.10}$$

Note that (A.10) is linear, so after some algebra, we can write this system in canonical form:

$$\boldsymbol{A} \cdot \frac{\mathbf{d}\boldsymbol{\vartheta}^*_{it}}{\mathbf{d}a_{i1t}} = \boldsymbol{b} \tag{A.11}$$

where

$$
\boldsymbol{A} = \begin{pmatrix} 1 & \frac{\alpha_{i1}}{\alpha_{i1}+\alpha_{i,J+1}} & \cdots & \cdots & \frac{\alpha_{i1}}{\alpha_{i1}+\alpha_{i,J+1}} \\ \frac{\alpha_{i2}}{\alpha_{i2}+\alpha_{i,J+1}} & 1 & \frac{\alpha_{i2}}{\alpha_{i2}+\alpha_{i,J+1}} & \cdots & \frac{\alpha_{i2}}{\alpha_{i2}+\alpha_{i,J+1}} \\ \vdots & \cdots & \ddots & \cdots & \vdots \\ \vdots & \cdots & \cdots & \ddots & \vdots \\ \frac{\alpha_{iJ'}}{\alpha_{iJ'}+\alpha_{i,J+1}} & \frac{\alpha_{iJ'}}{\alpha_{iJ'}+\alpha_{i,J+1}} & \cdots & \frac{\alpha_{iJ'}}{\alpha_{iJ'}+\alpha_{i,J+1}} & 1 \end{pmatrix} \tag{A.12}
$$

$$
\frac{\mathbf{d}\boldsymbol{\vartheta}^*_{it}}{\mathbf{d}a_{i1t}} = \begin{pmatrix} \frac{\mathrm{d}\vartheta^*_{i1t}}{\mathrm{d}a_{i1t}} \\ \frac{\mathrm{d}\vartheta^*_{i2t}}{\mathrm{d}a_{i1t}} \\ \vdots \\ \frac{\mathrm{d}\vartheta^*_{iJ',t}}{\mathrm{d}a_{i1t}} \end{pmatrix} \tag{A.13}
$$

$$
\boldsymbol{b} = \begin{pmatrix} -\frac{\gamma_i}{\ell_{it}} \\ -\frac{\gamma_i}{\ell_{it}} \frac{\alpha_{i2}}{\alpha_{i2}+\alpha_{i,J+1}} \\ \vdots \\ -\frac{\gamma_i}{\ell_{it}} \frac{\alpha_{iJ'}}{\alpha_{iJ'}+\alpha_{i,J+1}} \end{pmatrix} \tag{A.14}
$$

Conjecture that a solution to (A.11) is

$$
\widetilde{\frac{\mathbf{d}\boldsymbol{\vartheta}^*_{it}}{\mathbf{d}a_{i1t}}} = \begin{pmatrix} -\frac{\gamma_i}{\ell_{it}} \\ 0 \\ \vdots \\ 0 \end{pmatrix} \tag{A.15}
$$

Note that (A.15) is indeed a solution, and it is the only solution since $\boldsymbol{A}$ has full rank J' under the assumption $\alpha_{i,J+1} > 0$. It is clear that

$$
\boldsymbol{A} \cdot \widetilde{\frac{\mathbf{d}\boldsymbol{\vartheta}^*_{it}}{\mathbf{d}a_{i1t}}} = \boldsymbol{b} \tag{A.16}
$$

by inspecting the row-wise dot products on the left hand side of (A.16). This result holds for all values of $y \in \boldsymbol{\iota}_{it}$ indexing a_{iyt}, so that a simple re-arrangement and re-definition of the indices for $y \neq 1$ will yield the same outcome. □

Corollary 1: Suppose at least one optimal budget update occurs so that $k_{it} > 0$. Assume

$\ell_{it} > 0$, $\gamma_i > 0$, $\alpha_{ij} > 0, \forall j$, and $\alpha_{i,J+1} > 0$. For categories outside the bracket where $j \notin \boldsymbol{\iota}_{it}$, $\frac{\mathrm{d}\vartheta^*_{iyt}}{\mathrm{d}a_{ijt}}$ is independent of α_{ij}.

Proof. Note that the only channel through which $\frac{\mathrm{d}\vartheta^*_{iyt}}{\mathrm{d}a_{ijt}}$ could possibly depend on α_{ij} is $\frac{\mathrm{d}\vartheta^*_{ist}}{\mathrm{d}a_{ijt}}$, $s \neq y$. But, it can be verified by inspecting equation (11) in the main text that all values of $\frac{\mathrm{d}\vartheta^*_{iyt}}{\mathrm{d}a_{ijt}}$ only depend on $\alpha_{ij'}$ where $j' \in \boldsymbol{\iota}_{it}$. □

Corollary 2: Assume $\ell_{it} > 0$, $\gamma_i > 0$, $\alpha_{ij} > 0, \forall j$, and $\alpha_{i,J+1} > 0$. If $J - k_{it} \geq 2$, so that at least 2 categories are outside the bracket, then for both $j, j' \notin \boldsymbol{\iota}_{it}$, $\frac{\mathrm{d}\vartheta^*_{iyt}}{\mathrm{d}a_{ijt}} = \frac{\mathrm{d}\vartheta^*_{iyt}}{\mathrm{d}a_{ij',t}}$.

Proof. By Corollary 1, for all $j, j' \notin \boldsymbol{\iota}_{it}$, $\frac{\mathrm{d}\vartheta^*_{iyt}}{\mathrm{d}a_{ijt}}$ depends implicitly on $\frac{\mathrm{d}\vartheta^*_{iyt}}{\mathrm{d}a_{ij't}}$ and $\alpha_{i,\iota_{ist}}$ where s indexes all components of $\boldsymbol{\iota}_{it}$. By independence, we can redefine the indices for the outside categories and the value of $\frac{\mathrm{d}\vartheta^*_{iyt}}{\mathrm{d}a_{ijt}}$ 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. □

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

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 $\boldsymbol{\alpha}_i = (0.1, 0.15, 0.4, 0.35)^\top$ and consider two values for $\alpha_{i,J+1}$ — $\alpha^1_{i,J+1} = 0.1$ and $\alpha^2_{i,J+1} = 5$. $\boldsymbol{\iota}_{it} = (1, 2, 3)$ so that $j = 4$ is the category for which $\Gamma_{ijt} = 0$. The relevant system of equations is

$$\begin{aligned}
\frac{\mathrm{d}\vartheta^*_{i1t}}{\mathrm{d}a_{i4t}} &= -\frac{\alpha_{i1}}{\alpha_{i1} + \alpha_{i,J+1}} - \sum_{s\neq 1} \Gamma_{ist} \frac{\alpha_{i1}}{\alpha_{i1} + \alpha_{i,J+1}} \frac{\mathrm{d}\vartheta^*_{ist}}{\mathrm{d}a_{i4t}} \\
\frac{\mathrm{d}\vartheta^*_{i2t}}{\mathrm{d}a_{i4t}} &= -\frac{\alpha_{i2}}{\alpha_{i2} + \alpha_{i,J+1}} - \sum_{s\neq 2} \Gamma_{ist} \frac{\alpha_{i2}}{\alpha_{i2} + \alpha_{i,J+1}} \frac{\mathrm{d}\vartheta^*_{ist}}{\mathrm{d}a_{i4t}} \\
\frac{\mathrm{d}\vartheta^*_{i3t}}{\mathrm{d}a_{i4t}} &= -\frac{\alpha_{i2}}{\alpha_{i2} + \alpha_{i,J+1}} - \sum_{s\neq 3} \Gamma_{ist} \frac{\alpha_{i3}}{\alpha_{i3} + \alpha_{i,J+1}} \frac{\mathrm{d}\vartheta^*_{ist}}{\mathrm{d}a_{i4t}}
\end{aligned} \tag{A.17}$$

After some algebra, we get the canonical linear system

$$\boldsymbol{A} \cdot \frac{\mathbf{d}\boldsymbol{\vartheta}^*_{it}}{\mathbf{d}a_{i4t}} = \boldsymbol{b} \tag{A.18}$$

$$\text{where} \quad \boldsymbol{A} = \begin{pmatrix} 1 & \frac{\alpha_{i1}}{\alpha_{i1}+\alpha_{i,J+1}} & \frac{\alpha_{i1}}{\alpha_{i1}+\alpha_{i,J+1}} \\ \frac{\alpha_{i2}}{\alpha_{i2}+\alpha_{i,J+1}} & 1 & \frac{\alpha_{i2}}{\alpha_{i2}+\alpha_{i,J+1}} \\ \frac{\alpha_{i3}}{\alpha_{i3}+\alpha_{i,J+1}} & \frac{\alpha_{i3}}{\alpha_{i3}+\alpha_{i,J+1}} & 1 \end{pmatrix} \tag{A.19}$$

$$\boldsymbol{b} = \begin{pmatrix} \frac{\alpha_{i1}}{\alpha_{i1}+\alpha_{i,J+1}} \\ \frac{\alpha_{i2}}{\alpha_{i2}+\alpha_{i,J+1}} \\ \frac{\alpha_{i3}}{\alpha_{i3}+\alpha_{i,J+1}} \end{pmatrix} \tag{A.20}$$

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

$$\frac{\mathbf{d}\boldsymbol{\vartheta}^*_{it}}{\mathbf{d}a_{i4t}} = \begin{pmatrix} 0.227 \\ 0.033 \\ -0.933 \end{pmatrix} \tag{A.21}$$

With $\alpha^1_{i,J+1} = 5$ we have

$$\frac{\mathbf{d}\boldsymbol{\vartheta}^*_{it}}{\mathbf{d}a_{i4t}} = \begin{pmatrix} -0.014 \\ -0.023 \\ -0.073 \end{pmatrix} \tag{A.22}$$

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}. □

References

Aiyagari, S. Rao. 1994. "Uninsured Idiosyncratic Risk and Aggregate Saving". *The Quarterly Journal of Economics* 109 (3): 659–684. (Cit. on p. 6).

Barnett, William, Douglas Fisher, and Apostolos Serletis. 1992. "Consumer Theory and the Demand for Money". *Journal of Economic Literature* 30 (4): 2086–2119. (Cit. on p. 6).

Barten, Anton. 1964. "Consumer Demand Functions under Conditions of Almost Additive Preferences". *Econometrica* 32 (1/2): 1–38. (Cit. on p. 3).

— . 1967. "Evidence on the Slutsky Conditions for Demand Equations". *The Review of Economics and Statistics* 49 (1): 77–84. (Cit. on p. 3).

— . 1977. "The Systems of Consumer Demand Functions Approach: A Review". *Econometrica* 45 (1): 23–50. (Cit. on p. 3).

Benjamin, Daniel J., Sebastian A. Brown, and Jesse M. Shapiro. 2013. "Who is 'behavioral'? Cognitive ability and anomalous preferences". *Journal of the European Economic Association* 11 (6): 1231–1255. (Cit. on p. 45).

Bewley, Truman. 1986. "Stationary monetary equilibrium with a continuum of independently fluctuating consumers". In *Contributions to mathematical economics in honor of Gérard Debreu*, ed. by Werner Hildenbrand. North Holland. (Cit. on p. 6).

Bordalo, Pedro, Nicola Gennaioli, and Andrei Shleifer. 2020. "Memory, Attention, and Choice". *The Quarterly Journal of Economics*, no. 135: 1399–1442. (Cit. on p. 4).

— . 2014. "Salience and Consumer Choice". *Journal of Political Economy* 121 (5): 803–843. (Cit. on pp. 4, 11).

Brock, William. 1974. "Money and Growth: The Case of Long Run Perfect Foresight". *International Economic Review* 15 (3): 750–777. (Cit. on p. 6).

Calvo, Guillermo. 1983. "Staggered Price in a Utility-Maximizing Framework". *Journal of Monetary Economics* 12:383–398. (Cit. on p. 6).

Caplin, Andrew, and Mark Dean. 2015. "Revealed Preference, Rational Inattention, and Costly Information Acquisition". *American Economic Review* 105 (7): 2183–2203. (Cit. on p. 4).

Chetty, Raj, Adam Looney, and Kory Kroft. 2009. "Salience and Taxation: Theory and Evidence". *American Economic Review* 99 (4): 1145–1177. (Cit. on p. 4).

Chib, Siddhartha. 1998. "Estimation and comparison of multiple change-point models". *Journal of Econometrics* 86:221–241. (Cit. on p. 25).

Cowan, Nelson. 2000. “The magical number 4 in short-term memory: A reconsidration of mental storage capacity”. *Behavioral and Brain Sciences* 24:87–185. (Cit. on pp. 10, 11).

Deaton, Angus, and John Muellbauer. 1980a. “An Almost Ideal Demand System”. *The American Economic Review* 70 (3): 312–326. (Cit. on p. 3).

— . 1980b. “Economics and Consumer Behavior”. Chap. 5, 119–147. Cambridge University Press: New York. (Cit. on pp. 1, 3, 4, 8, 16).

Farhi, Emmanuel, and Xavier Gabaix. 2020. “Optimal Taxation with Behavioral Agents”. *American Economic Review* 110 (1): 298–336. (Cit. on pp. 5, 9).

Feenstra, Robert. 1986. “Functional Equivalence Between Liquidity Costs and the Utility of Money”. *Journal of Monetary Economics* 17:271–291. (Cit. on p. 6).

Feinberg, Richard. 1986. “Credit Cards as Spending Facilitating Stimuli: A Conditioning Interpretation”. *Journal of Consumer Research* 13 (3): 348–356. (Cit. on p. 5).

Felső, Flóra, and Adriaan Soetevent. 2014. “Broad and Narrow Bracketing in Gift Certificate Spending”. *European Economic Review* 66:284–302. (Cit. on p. 4).

Gabaix, Xavier. 2014. “A Sparsity-Based Model of Bounded Rationality”. *The Quarterly Journal of Economics* 129 (4): 1661–1710. (Cit. on pp. 1, 4, 11).

Gathergood, John, et al. 2019. “How Do Individuals Repay Their Debt? The Balance-Matching Heuristic”. *American Economic Review* 109 (3): 844–875. (Cit. on p. 6).

Geary, Roy. 1950. “A Note on ‘A Constant-Utility Index of the Cost of Living’”. *Review of Economic Studies* 18 (2): 65–66. (Cit. on pp. 3, 7).

Gelman, Andrew, et al. 2013. *Bayesian Data Analysis*. Third. Chap. 7. Taylor & Francis, CRC Press: Boca Raton, FL. (Cit. on pp. 26, 27).

Gelman, Michael, et al. 2014. “Harnessing naturally occurring data to measure the response of spending to income”. *Science* 345:212–215. (Cit. on p. 5).

Geman, Stuart, and Donald Geman. 1984. “Stochastic Relaxation, Gibbs Distributions, and Bayesian Restoration of Images”. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 6 (6): 721–741. (Cit. on p. 25).

Gorman, William. 1959. “Separable Utility and Aggregation”. *Econometrica* 27 (3): 469–481. (Cit. on p. 1).

Hastings, Justine, and Jesse Shapiro. 2013. “Fungibility and Consumer Choice: Evidence from Commodity Price Shocks”. *The Quarterly Journal of Economics* 128 (4): 1449–1498. (Cit. on p. 10).

— . 2018. "How Are SNAP Benefits Spent? Evidence from a Retail Panel". *American Economic Review* 108 (12): 3493–3540. (Cit. on pp. 5, 24).

Hastings, Wilfred. 1970. "Monte Carlo Sampling Methods Using Markov Chains and Their Applications". *Biometrika* 57 (1): 97–109. (Cit. on p. 25).

Hicks, John. 1936. *Value and Capital*. Oxford: Oxford University Press. (Cit. on pp. 3, 4).

Houthakker, Hendrik. 1960. "Additive Preferences". *Econometrica* 28 (2): 244–257. (Cit. on p. 3).

Huggett, Mark. 1993. "The risk-free rate in heterogeneous-agent incomplete-insurance economies". *Journal of Economic Dynamics and Control* 17:953–969. (Cit. on p. 6).

— . 1996. "Wealth distribution in life-cycle economies". *Journal of Monetary Economics* 38:469–494. (Cit. on p. 6).

Kahneman, Daniel, and Dan Lovallo. 1993. "Timid Choices and Bold Forecasts: A Cognitive Perspective on Risk Taking". *Management Science* 39 (1): 17–31. (Cit. on p. 4).

Kőszegi, Botond, and Filip Matějka. 2020. "Choice Simplification: A Theory of Mental Budgeting and Naive Diversification". *The Quarterly Journal of Economics* 135 (2): 1153–1207. (Cit. on pp. 4, 5).

Kőszegi, Botond, and Adam Szeidl. 2013. "A Model of Focusing in Economic Choice". *The Quarterly Journal of Economics* 128 (1): 53–104. (Cit. on pp. 4, 11).

Koch, Alexander, and Julia Nafziger. 2019. "Correlates of Narrow Bracketing". *The Scandinavian Journal of Economics* 121 (4): 1441–1472. (Cit. on p. 4).

— . 2016. "Goals and bracketing under mental accounting". *Journal of Economic Theory* 162 (C): 305–351. (Cit. on p. 4).

Koop, Gary, and Simon M. Potter. 2007. "Estimation and Forecasting in Models with Multiple Breaks". *The Review of Economic Studies* 74 (3): 763–789. (Cit. on p. 25).

Leontief, Wassily. 1936. "Composite Commodities and the Problem of Index Numbers". *Econometrica* 4 (1): 39–59. (Cit. on pp. 3, 4).

Leslie, Alan, Rochel Gelmand, and C.R. Gallistel. 2008. "The generative basis of natural number concepts". *Trends in Cognitive Sciences* 12 (6): 213–218. (Cit. on p. 11).

Malhotra, Naresh. 1982. "Information Load and Consumer Decision Making". *Journal of Consumer Research* 8 (4): 419–430. (Cit. on p. 11).

Mani, Anandi, et al. 2013. "Poverty Impedes Cognitive Function". *Science* 341:976–980. (Cit. on p. 5).

McCulloch, Robert, and Ruey Tsay. 1993. "Bayesian Inference and Prediction for Mean and Variance Shifts in Autoregressive Time Series". *Journal of the American Statistical Association* 88 (423): 968–978. (Cit. on p. 25).

Miller, George A. 1956. "The magical number seven, plus or minus two: Some limits on our capacity for processing information". *Psychological Review* 63 (2): 81–97. (Cit. on pp. 10, 11).

Mullainathan, Sendhil. 2002. "A Memory-Based Model of Bounded Rationality". *The Quarterly Journal of Economics* 107 (3): 735–774. (Cit. on p. 5).

Obstfeld, Maurice, and Kenneth Rogoff. 1983. "Speculative Hyperinflations in Maximizing Models: Can We Rule Them Out?" *Journal of Political Economy* 91 (4): 675–687. (Cit. on p. 6).

Olivola, Christopher Y., and Stephanie W. Wang. 2016. "Patience auctions: The impact of time vs. money bidding on elicited discount rates". *Experimental Economics* 19 (4): 864–885. (Cit. on p. 46).

Oprea, Ryan. 2020. "What Makes a Rule Complex?" *American Economic Review* 110 (12): 3913–3951. (Cit. on p. 11).

Prelec, Drazen, and George Loewenstein. 1998. "The Red and the Black: Mental Accounting of Savings and Debt". *Marketing Science* 17 (1): 4–28. (Cit. on p. 5).

Proctor, Bernadette, Jessica Semega, and Melissa Kollar. 2015. *Income and Poverty in the United States: 2015*. `https://www.census.gov/library/publications/2016/demo/p60-256.html`. Accessed: 2019-05-06. (Cit. on p. 23).

Rabin, Matthew, and Georg Weizsäcker. 2009. "Narrow Bracketing and Dominated Choices". *American Economic Review* 99 (4): 1508–1543. (Cit. on p. 4).

Read, Daniel, George Loewenstein, and Matthew Rabin. 1999. "Choice Bracketing". *Journal of Risk and Uncertainty* 19 (1): 171–197. (Cit. on p. 4).

Schilbach, Frank, Heather Schofield, and Sendhil Mullainathan. 2016. "The psychological lives of the poor". *American Economic Review: Papers & Proceedings* 106 (5): 435–440. (Cit. on p. 5).

Schwartzstein, Joshua. 2014. "Selective Attention and Learning". *Journal of the European Economic Association* 12 (6): 1423–1452. (Cit. on pp. 4, 11).

Shefrin, Hersh, and Richard Thaler. 1981. "An Economic Theory of Self-Control". *Journal of Political Economy* 89 (2): 392–406. (Cit. on pp. 1, 4, 10).

Silverman, Bernard. 1986. *Density Estimation for Statistics and Data Analysis*. 48. London: Chapman / Hall. (Cit. on p. 30).

Simon, Herbert. 1957. *Models of Man*. John Wiley & Sons: New York. (Cit. on pp. 10, 11).

Sims, Christopher. 2003. "Implications of rational inattention". *Journal of Monetary Economics* 50 (3): 665–690. (Cit. on p. 4).

Stone, Richard. 1954. "Linear Expenditure Systems and Demand Analysis: An Application to the Pattern of British Demand". *The Economic Journal* 64 (255): 511–527. (Cit. on pp. 3, 7).

Strotz, Robert. 1957. "The Empirical Implications of a Utility Tree". *Econometrica* 25 (1): 269–280. (Cit. on p. 1).

Thaler, Richard. 1985. "Mental Accounting and Consumer Choice". *Marketing Science* 4 (3): 199–214. (Cit. on p. 5).

— . 1999. "Mental Accounting Matters". *Journal of Behavioral Decision Making* 12 (3): 183–206. (Cit. on p. 1).

Thaler, Richard, et al. 1997. "The Effect of Myopia and Loss Aversion on Risk Taking: An Experimental Test". *The Quarterly Journal of Economics* 112 (2): 647–661. (Cit. on p. 5).

Theil, Henri. 1965. "The Information Approach to Demand Analysis". *Econometrica* 33 (1): 67–87. (Cit. on p. 3).

— . 1976. *Theory and Measurement of Consumer Demand*. Vol. 2. Elsevier: New York. (Cit. on p. 3).

Walsh, Carl. 2010. *Monetary Theory and Policy*. 3rd ed. MIT Press: Cambridge. (Cit. on p. 6).