[redacted]

A Deriving the Optimal Budget

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

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

[redacted]

The FOC with respect to ϑ_{iyt} is

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

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

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

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

B Equilibrium Appendix

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

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

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

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

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

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

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

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

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

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

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

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

$$\boldsymbol{h}_{ijt}^\top \cdot \boldsymbol{H}_{it} \cdot \boldsymbol{h}_{ijt} < 0, \qquad \forall j \tag{B.4}$$

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

Lemma B.2.

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

ii. Suppose $\underline{k}_{it} = 0$ or $\underline{k}_{it} > 0$ but $k_{it} = 0$. Then $\boldsymbol{\theta}_{it} = \boldsymbol{\theta}_{i,t-1}$, which is unique.

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

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

Following (B.1), we get

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

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

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

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

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

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

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

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

Proof of part (ii) is trivial. ***QED***

C Highlighted Properties of the Model

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

C.1 Non-fungibility Under Budgeting Frictions

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

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

Proposition C.1. *Optimally choosing ex-ante budgets is equivalent to optimally choos-*

ing ex-post consumption if and only if consumers have perfect foresight over spending shocks ζ_{ijt} and face no cognitive or numeracy frictions ($\Gamma_{ijt} = 1$ for all j).

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

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

Here, we examine how equilibrium budgets co-vary under different degrees of narrow bracketing.[1] 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^*_{ijt}. For categories featuring optimal budget updates ($\Gamma_{ijt} = 1$), the responsiveness of ϑ^*_{iyt} to variation in a_{ijt}, where j is one of these optimally-updated

[1] Assume throughout this section that all equilibrium candidates ϑ^*_{iyt} satisfy the exogenously-imposed numeracy conditions, so that $\boldsymbol{\iota}_{it}(\underline{\boldsymbol{\Gamma}}_{it}) = \boldsymbol{\iota}_{it}(\boldsymbol{\Gamma}_{it}) \equiv \boldsymbol{\iota}_{it}$.

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.[2]

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

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

where $\mathbf{1}\{\cdot\}$ is an indicator function.[3] For any given $j \in \boldsymbol{\iota}_{it}$, (C.1) 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 C.2. *Assume at least one optimal budget update occurs so that $k_{it} > 0$. Assume further that $\ell_{it} > 0$, $\gamma_i > 0$, $\alpha_{ij} > 0$, $\forall j$, and $\alpha_{i,J+1} > 0$.*

Proposition C.3. *Under Assumption C.2, without loss of generality, let $\boldsymbol{\iota}_{it} = (1, 2, 3, \ldots)$ 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 C.3 characterizes both own- and cross-category responsivenesses described by (C.1). For categories inside the bracket, both losses and gains due to over- or under-spending, respectively, are fully integrated into the optimal budget update for that *same category*. Since consumers have over spent when $a_{ijt} < 0$, case

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

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

(i) of Proposition C.3 says that for a category inside the bracket, consumers respond to over-spending in that same category by increasing that category's budget.

Now consider the remaining categories inside the bracket. Budget adjustments do not respond to cross-category over- or under-expenditure. For $s \neq y$ where s indexes a component of $\boldsymbol{\iota}_{it}$, 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 C.3. 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:

$$\frac{\mathrm{d}\vartheta^*_{iyt}}{\mathrm{d}a_{ijt}} = -\frac{\gamma_i\, \alpha_{i,\iota_{iyt}}}{\ell_{it}(\alpha_{i,\iota_{iyt}} + \alpha_{i,J+1})} - \sum_{s \neq y} \frac{\alpha_{i,\iota_{iyt}}}{\alpha_{i,\iota_{iyt}} + \alpha_{i,J+1}} \frac{\mathrm{d}\vartheta^*_{ist}}{\mathrm{d}a_{ijt}} \qquad \forall j \notin \boldsymbol{\iota}_{it} \tag{C.2}$$

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

Proposition C.6. *Under Assumption C.2, without loss of generality, let $\boldsymbol{\iota}_{it} = (1, 2, 3, \ldots)$ 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 C.6 demonstrates that the cross-category budgetary responsiveness of inside categories to outside categories cannot be determined a priori. This result builds on Corollaries C.4 and C.5 in Supplemental Appendix C.3 which show that only the values of $\alpha_{i,J+1}$ and α_{ij} for inside categories matter, but the magnitudes of these parameters determine the sign of the budget change. Indeed, there is no hard-and-fast rule as to which inside-bracket categories increase or decrease in response to increases in a_{ijt}. Some categories may see budget increases while others see budget decreases. Further, this responsiveness will be different for different consumers in different periods since it depends heavily on both the underlying values of $\boldsymbol{\alpha}_i$ and $\alpha_{i,J+1}$, as well as the elements of $\boldsymbol{\iota}_{it}$ (i.e. the budgets being reevaluated).

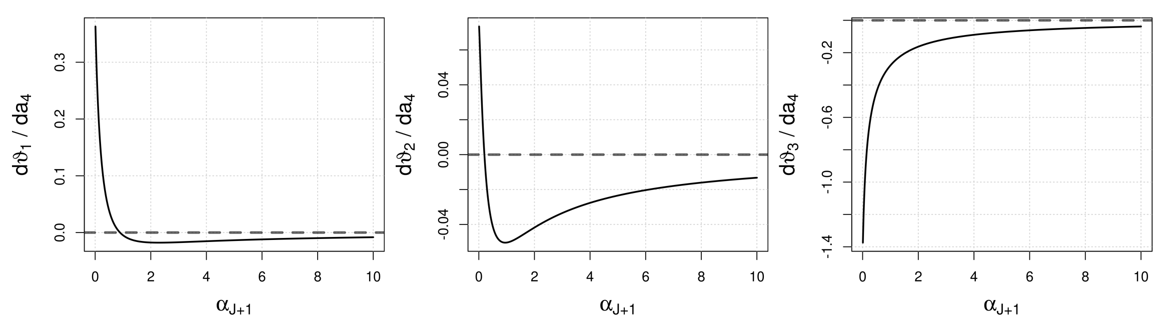

(a) $\alpha_{i1} = 0.1$ (b) $\alpha_{i2} = 0.15$ (c) $\alpha_{i3} = 0.4$

Figure C.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 $\Gamma_{i4t} = 0$.

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.[4] Figure C.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}. In Figure C.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 most-evident in Figure C.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 occurs, which seems to be simultaneously associated with smaller values of $\alpha_{i,J+1}$ overall and with categories that have smaller values of $\alpha_{i,\iota_{iyt}}$.

The magnitude of $\alpha_{i,J+1}$ also plays an important role in determining budgetary responsiveness to over- or under-spending. As $\alpha_{i,J+1} \rightarrow_+ 0$ hand-to-mouth consumers appear to respond to under-spending in j by transferring excess funds to categories with smaller utility weights. Since $\frac{\mathrm{d}\vartheta^*_{iyt}}{\mathrm{d}a_{ijt}}$ 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.

Budget shares respond to outside over- or under-spending in ways that depend on the size and composition of the inside bracket. Consider, for example, the responsiveness of budgets for commodity group $j = 2$ to variation in outside category $j = 3$ presented in Figure C.2. Between panels (a) and (b) of the figure, we keep $k_{it} = 2$ but vary the composition of the inside bracket from $\boldsymbol{\iota}_{it} = (1, 2)$ in Figure C.2a to $\boldsymbol{\iota}'_{it} = (2, 4)$ in Figure C.2b. In both panels we show the responsiveness of $j = 2$ to variation in a_{i3t}, which is always outside the bracket. Note

[4]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.

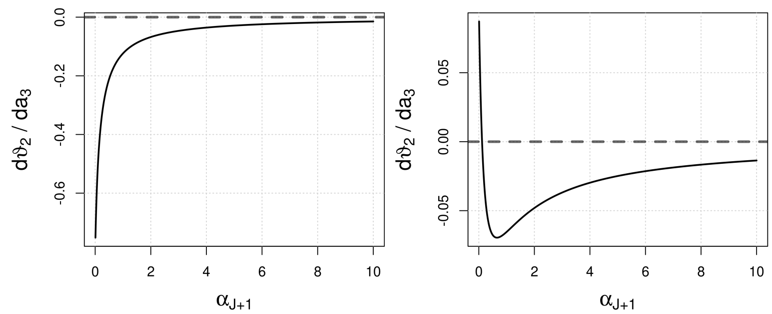

(a) $j = 2$; $\boldsymbol{\iota}_{it} = (1, 2)$ (b) $j = 2$; $\boldsymbol{\iota}'_{it} = (2, 4)$

Figure C.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.

that the qualitative change in budgetary responsiveness depends on which of the other consumption categories is inside the bracket. In panel (a), when the inside bracket is $(1, 2)$, budget $j = 2$ responds to variation in a_{i3t} by falling for all values of $\alpha_{i,J+1} \in (0, 10)$. Meanwhile, in panel (b) when the inside bracket is $(2, 4)$, the responsiveness of budget $j = 2$ to variation in a_{i3t} is positive for near hand-to-mouth consumers ($\alpha_{i,J+1}$ near zero) and negative otherwise, though less negative as $\alpha_{i,J+1}$ increases.

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

C.3 Proofs of Propositions and Corollaries

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

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

$$\frac{\partial u_{it}}{\partial q_{ijt}}\frac{\partial q_{ijt}}{\partial \vartheta_{iyt}}+\frac{\partial u_{it}}{\partial z_{it}}\frac{\partial z_{it}}{\partial \vartheta_{iyt}} \tag{C.3}$$

$$=\frac{\partial u_{it}}{\partial q_{ijt}}\frac{\partial q_{ijt}}{\partial \vartheta_{iyt}}+\frac{\partial u_{it}}{\partial z_{it}}\frac{\partial z_{it}}{\partial q_{ijt}}\frac{\partial q_{ijt}}{\partial \vartheta_{iyt}} \tag{C.4}$$

$$=\frac{\partial u_{it}}{\partial q_{ijt}}+\frac{\partial u_{it}}{\partial z_{it}}\frac{\partial z_{it}}{\partial q_{ijt}}=0 \tag{C.5}$$

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

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

First, let us consider the case where ζ_{it} is unknown ex-ante but $\sum_{j=1}^{J}\Gamma_{ijt}=J$, so that there are no cognitive frictions. Note that ex-ante budgets ϑ^*_{iyt} must satisfy the main equilibrium condition. Denote ex-ante expected consumption by $\mathbb{E}_{it}q_{ijt}=\widetilde{q}_{ijt}=\frac{\vartheta^*_{iyt}\ell_{it}+\gamma_i a_{ijt}}{p_{jt}}$. 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

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

Ex-post $q^*_{ijt}(\zeta_{ijt})$ satisfies (C.5) by construction. Since the random variable ζ_{ijt} enters (C.6), $\widetilde{q}^*_{ijt}\neq q^*_{ijt}$ 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 ζ_{it} is known but $\exists j$ such that $\Gamma_{ijt}=0$. The consumer thus cannot choose an ex-ante budget for category j and sets his budget such that $\theta_{ijt}=\theta_{ij,t-1}$.

There is no ex-ante uncertainty, so we can drop expectations. Except for measure zero values of $\theta_{ij,t-1}$,

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

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

Proposition C.3. *Proof.* Consider the system of implicit component-wise total derivatives of ϑ^*_{iyt} and suppose, without loss of generality, that $y = 1$:

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

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

$$\boldsymbol{A} \cdot \frac{\mathbf{d}\boldsymbol{\vartheta}^*_{it}}{\mathbf{d}a_{i1t}} = \boldsymbol{b} \tag{C.9}$$

where

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

$$
\frac{\mathbf{d}\boldsymbol{\vartheta}^*_{it}}{\mathbf{d}a_{i1t}} = \begin{pmatrix} \frac{\mathrm{d}\vartheta^*_{i1t}}{\mathrm{d}a_{i1t}} \\ \frac{\mathrm{d}\vartheta^*_{i2t}}{\mathrm{d}a_{i1t}} \\ \vdots \\ \frac{\mathrm{d}\vartheta^*_{iJ',t}}{\mathrm{d}a_{i1t}} \end{pmatrix} \tag{C.11}
$$

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

Conjecture that a solution to (C.9) is

$$
\widetilde{\frac{\mathbf{d}\boldsymbol{\vartheta}^*_{it}}{\mathbf{d}a_{i1t}}} = \begin{pmatrix} -\frac{\gamma_i}{\ell_{it}} \\ 0 \\ \vdots \\ 0 \end{pmatrix} \tag{C.13}
$$

Note that (C.13) 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

$$
\boldsymbol{A} \cdot \widetilde{\frac{\mathbf{d}\boldsymbol{\vartheta}^*_{it}}{\mathbf{d}a_{i1t}}} = \boldsymbol{b} \tag{C.14}
$$

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

Corollary C.4. *Under Assumption C.2 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 the equilibrium condition for ϑ^*_{iyt} that all values of $\frac{\mathrm{d}\vartheta^*_{iyt}}{\mathrm{d}a_{ijt}}$ only depend on $\alpha_{ij'}$ where $j' \in \boldsymbol{\iota}_{it}$. ***QED***

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

Proof. By Corollary C.4, 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. ***QED***

Proposition C.6. *Proof.* We will prove this by considering two numerical parameterizations and showing that both the signs and magnitudes of the total derivatives are different under the different parameterizations. Without loss of generality, let $J = 4$ and $J' = 3$. Let $\gamma_i = \ell_{it} = 1$. Let $\boldsymbol{\alpha}_i = (0.1, 0.15, 0.4, 0.35)^\top$ 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

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

After some algebra, we get the canonical linear system

$$\boldsymbol{A} \cdot \frac{\mathbf{d}\boldsymbol{\vartheta}^*_{it}}{\mathbf{d}a_{i4t}} = \boldsymbol{b} \tag{C.16}$$

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

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

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

$$\frac{\mathbf{d}\boldsymbol{\vartheta}^*_{it}}{\mathbf{d}a_{i4t}} = \begin{pmatrix} 0.227 \\ 0.033 \\ -0.933 \end{pmatrix} \tag{C.19}$$

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

$$\frac{\mathbf{d}\boldsymbol{\vartheta}^*_{it}}{\mathbf{d}a_{i4t}} = \begin{pmatrix} -0.014 \\ -0.023 \\ -0.073 \end{pmatrix} \tag{C.20}$$

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

D Estimation Appendix

D.1 Description of Estimator

D.1.1 Overview

We estimate multiple different models, varying assumptions regarding numeracy constraints and mental accounting behavior. For each estimation we initialize the MCMC integration scheme described here with random generates from the prior distributions, except for $\boldsymbol{a}_{i1}$ which we set to $\mathbf{0}$ to start. We then proceed to iterate through a chain of length 400,000. That is, we operate on the agent-level MH within Gibbs blocks 400,000 times and the global blocks 400,000/400 = 1,000 times. We set a burn-in period of 120,000, keeping every 400^{th} agent-level draw and every global draw thereafter, giving us a total of 700 draws after burn-in and trimming. The sampler operates with agent-level blocks parallelized over a 128-core computer, requiring just under 100 hours to completion for each of the 12 different estimation routines.[5]

D.1.2 Likelihood Function and Prior Distributional Assumptions

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

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

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

parameterization takes the following form:

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

In a slight abuse of notation we denote the agent i, commodity group j, and period t likelihood of observing x_{ijt} and ℓ_{it}, $\mathcal{L}_{ijt}(x_{ijt},\ell_{it} \mid \theta_{ijt}, a_{ijt}, \sigma^2_{ij}, \gamma_i)$.[6]

Inference on the agent's budget weights, θ_{ijt}, and mental account balances, a_{ijt}, operates on the inverted ex-ante expected indirect utility first-order condition. Note that the entire time series of mental account balances, $\boldsymbol{a}_{it}$, and by extension budget weights, $\boldsymbol{\theta}_{it}$, depends on the initial balances, $\boldsymbol{a}_{i1}$, which are unknown parameters. Therefore, we give these initial balances, $\boldsymbol{a}_{i1}$, a normal prior structure:

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

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

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

In our various estimation exercises we fix either the mental accounting weight to $\gamma_i = 1$ (baseline estimation) or $\gamma_i = 0$ (estimation with no serial dependence). For the utility weights we require $\sum_{j=1}^{J} \alpha_{ij} = 1$ with $0 \le \alpha_{ij} \le 1$ for all $j \le J$. Following Geary (1950), Stone (1954), and Deaton and Muellbauer (1980), these values encode the shares of total expenditure devoted toward category j consump-

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

tion. We fix these values as the mean expenditure shares taken over the entire time series we observe for each consumer: $\widehat{\alpha}_{ij} = \frac{1}{T_i}\sum_t \left(\frac{x_{ijt}}{\sum_j x_{ijt}}\right)$.

Note that m_i and $\alpha_{i,J+1}$ cannot be separately identified. We sample $\alpha_{i,J+1}$ explicitly while fixing the consumption borrowing limit, m_i, to be 20% of an agent's annualized income.[7] From simulations we have inferred that $\alpha_{i,J+1}$ is weakly identified given m_i, so we impose a three-tier hierarchical prior structure with empirical priors for the global hierarchical mean and variance:

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

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

The orthogonality assumption on error terms such that $\zeta_{ijt} \perp \zeta_{ij',t}$, $j' \neq j$, allows us to impose separate, agent-specific, hierarchical independent gamma priors on the likelihood precisions, $\frac{1}{\sigma^2_{ij}}$:[8]

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

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

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

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

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

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

D.1.3 Estimating Latent Budgeting Variables

Let n index the atomic posterior estimation draws (epochs) from the Metropolis-Hastings (MH) within Gibbs sampler. Let $\mathcal{P}_{in}$ be an atomic draw of parameters associated with agent i on iteration n, and let $\mathcal{H}_n$ be the set of global hierarchical parameters common to all agents.[9] Let $\mathcal{D}_i$ denote the set of agent i's data. These sets (with vectors in bold) are:

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

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

Assume, for now, that $\gamma_i = 1$.[10] Take the non-budgeting parameters, $\mathcal{P}_{in} \setminus \left\{\{\boldsymbol{\theta}_{itn}, \boldsymbol{a}_{itn}, \boldsymbol{\Gamma}_{itn}\}_{t=1}^{T_i}\right\}$, as given, where $\setminus$ is the set exclusion operator. We want to estimate the posterior distribution of $\{\boldsymbol{\Gamma}_{it}\}_t$. Ignoring MCMC iteration indexing for now, consider the sequence $\{\boldsymbol{a}_{it}\}_{t>1}$, which, given $\{\boldsymbol{\theta}_{it}\}_t$ and data, is completely determined by the law of motion for mental account balances. For each j recursively substitute out $a_{ijt}, a_{ij,t-1}, \ldots, a_{ij2}$ to get

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

[9] Recall, parameters fixed for the entirety of the MCMC sampling routine are γ_i, m_i, α_{ij} for all $j \in \{1, \ldots, J\}$, and $\boldsymbol{\sigma}^2_{a_{i1}}$.

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

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

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

$$\vartheta^*_{iytn}(\vartheta^*_{i1tn}, \ldots, \vartheta^*_{i,y-1,tn}, \vartheta^*_{i,y+1,t,n-1}, \ldots, \vartheta^*_{i,k_{it},t,n-1}, \boldsymbol{\theta}_{i,t-1,n}, \ldots, \boldsymbol{\theta}_{i1n})$$

If numeracy constraints are also imposed, we must also check whether new, candidate budget updates satisfy either ϵ or ϵ_ℓ. With the candidate sequence $\{\boldsymbol{\theta}_{itn}, \boldsymbol{a}_{itn}, \boldsymbol{\Gamma}_{it,n-1}\}_{t=1}^{T_i}$, we can accept or reject the draw of $\boldsymbol{a}_{i1n}$ using a standard MH acceptance rule. If the draw is rejected, the entire sequence is reset so that $\theta_{ijtn} = \theta_{ijt,n-1}$ and $a_{ijtn} = a_{ijt,n-1}$ for all j and t.[11]

Having performed the MH step for $\boldsymbol{a}_{i1n}$, we proceed to use Gibbs sampling to now jointly draw the sequences $\{\boldsymbol{\Gamma}_{itn}\}_{t=1}^{T_i}$ and $\{\boldsymbol{\theta}_{itn}\}_{t=1}^{T_i}$. For $t = 1$, we assume $\Gamma_{ij1n} = 1$, so that $k_{i1} = J$ and $\theta_{ij1n} = \vartheta^*_{ij1n}$ always (i.e., we always evaluate (**??**) to get the first period's budgets). For $t > 1$, θ_{ijtn} depends on past values $\theta_{ij,t-1,n}$, $\theta_{ij,t-2,n}$, $\ldots$, $\theta_{ij,t-s,n}$ where s is the number of periods that have passed since the last budget update, $\Gamma_{ij,t-s,n} = 1$. Since $\theta_{ij,t-s,n}$ is a function of $\Gamma_{ij,t-s,n}$, then by extension so is θ_{ijtn}. Thus, any new draw of Γ_{ijtn} will affect all future realizations of the budget weights, $\{\theta_{ijsn}\}_{s=t}^{T_i}$. Further, since updating Γ_{ijtn} induces a change to the sequence $\{\theta_{ijsn}\}_{s=t}^{T_i}$, values of $\{a_{ijsn}\}_{s=t}^{T_i}$ are also affected and must be updated to ensure the equilibrium mental account balance law of motion always holds.

Let $\mathcal{D}_{i,t:T_i} \subseteq \mathcal{D}_i$ that only contains data for observed periods t to T_i. Let $\mathcal{P}_{i,t:T_i,n} \subseteq$

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

$\mathcal{P}_{in}$ similarly represent agent i's parameter set containing budgeting parameters only for periods t to T_i, $\mathcal{P}_{i,t:T_i,n} = \mathcal{P}_{in} \setminus \big\{\{\boldsymbol{\theta}_{isn}, \boldsymbol{a}_{isn}, \boldsymbol{\Gamma}_{isn}\}_{s=1}^{t-1}\big\}$. Sampling $\{\Gamma_{ijtn}\}_{t>1}$ proceeds as follows. If $\Gamma_{ijt,n-1} = 0$, then set $\Gamma_{ijtn} = 1$ as a candidate draw and compute the sequence of budget weight policies $\{\theta_{ijsn}(\Gamma_{ijtn}, \Gamma_{ij,t+1,n-1}, \ldots, \Gamma_{ijs,n-1})\}_{s \geq t}$, re-computing ϑ^*_{iysn} and setting $\theta_{ijsn} = \vartheta^*_{iysn}$ for $j = \iota_{iys}$ whenever $\Gamma_{ijs,n-1} = 1$ and $s > t$. Future values of θ_{ijsn} will change because they depend on a_{ijsn}, which in turn depends on past values of budget weights. After computing the new sequence of implied budget weights, evaluate the posterior probability of observing a budget update:

$$
\begin{aligned}
&\mathbb{P}\Big(\Gamma_{ijtn} = 1 \mid \mathcal{D}_{i,t:T_i}, \mathcal{P}_{i,t:T_i,n} \setminus \{\Gamma_{ijtn}\}, \mathcal{H}_n\Big) \\
&= \psi_{ijn} \prod_{s \geq t} \mathcal{L}_{ijs}\Big(\vartheta^*_{ijsn}(\Gamma_{ijtn} = 1)\Big) \Big/ \Big[\psi_{ijn} \prod_{s \geq t} \mathcal{L}_{ijs}\Big(\vartheta^*_{ijsn}(\Gamma_{ijtn} = 1)\Big) \\
&\qquad + (1 - \psi_{ijn}) \prod_{s \geq t} \mathcal{L}_{ijs}\Big(\theta_{ijs,n-1}(\Gamma_{ijt,n-1} = 0)\Big)\Big]
\end{aligned}
\tag{D.2}
$$

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

D.1.4 Metropolis-Hastings within Gibbs Steps

In the previous subsection we described how the latent budgeting decisions and associated cognitive shocks $\{\boldsymbol{\theta}_{itn}, \boldsymbol{a}_{itn}, \boldsymbol{\Gamma}_{itn}\}_{t=1}^{T_i}$ are sampled. Given $\{\boldsymbol{\theta}_{itn}, \boldsymbol{a}_{itn}, \boldsymbol{\Gamma}_{itn}\}_{t=1}^{T_i}$, we now describe the Metropolis-Hastings (MH) within Gibbs sampling steps for the budgeting and mental accounting parameters $\mathcal{P}_{in} \setminus \{\{\boldsymbol{\theta}_{itn}, \boldsymbol{a}_{itn}, \boldsymbol{\Gamma}_{itn}\}_{t=1}^{T_i}\}$. We exploit conjugacy where possible.

Here, we start with describing the agent-level draws then move on to describing the sampling statements for the global, hierarchical parameters. Note that the following agent-level parameters are associated with non-conjugate priors, $\{\boldsymbol{a}_{i1n}, \alpha_{i,J+1,n}, \boldsymbol{\tau}_{in}\}$, while the remaining agent-level parameters, $\{\boldsymbol{\sigma}^2_{in}, \boldsymbol{\beta}_{in}, \mu_{i,\alpha_{i,J+1},,n}, \boldsymbol{\psi}_{in}\}$, can be conditionally sampled directly. In practice, we sample $\{\boldsymbol{a}_{i1n}, \alpha_{i,J+1,n}, \boldsymbol{\tau}_{in}\}$ all together using a random-walk MH step with the adaptive algorithm described in Atchade and Rosenthal (2005).[12]

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

$$
\begin{aligned}
&\left.\begin{pmatrix} \boldsymbol{a}_{i1n} \\ \alpha_{i,J+1,n} \\ \ln \boldsymbol{\tau}_{in} \end{pmatrix} \right| \mathcal{D}_i, \mathcal{P}_{in} \setminus \{\boldsymbol{a}_{i1n}, \alpha_{i,J+1,n}, \boldsymbol{\tau}_{in}\}, \mathcal{H}_n \propto \mathcal{L}_i(\boldsymbol{a}_{i1n}, \alpha_{i,J+1,n}, \boldsymbol{\tau}_{in}) \\
&\quad \times \left(\prod_j \underbrace{\phi\left(\frac{a_{ij1n}}{\sigma_{a_{ij1n}}}\right)}_{\text{Priors on } a_{ij1n}} \underbrace{\frac{\beta_{ij,n-1}^{\tau_{ijn}}}{\Gamma(\tau_{ijn})\sigma_{ij,n-1}^{2(\tau_{ijn}-1)}} \exp\left\{\frac{\beta_{ij,n-1}}{\sigma^2_{ijn}}\right\}}_{\text{Posterior on } \tau_{ijn} \text{ given } \sigma_{ij,n-1}} \right) \\
&\quad \times \underbrace{\phi\left(\frac{\alpha_{i,J+1,n} - \mu_{i,\alpha_{i,J+1},n-1}}{10}\right)}_{\text{Prior on } \alpha_{i,J+1,n}} \times \underbrace{\tau_{ijn}}_{\text{Jacobian of Log Transform}}
\end{aligned}
\tag{D.3}
$$

The adaptation routine is agent-specific, in that each agent's MH step is associated

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

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

$$\boldsymbol{\Lambda}_{in} = \begin{pmatrix} \boldsymbol{a}_{i1n} \\ \alpha_{i,J+1,n} \\ \ln \boldsymbol{\tau}_{in} \end{pmatrix} \sim \mathcal{MN}(\boldsymbol{\Lambda}_{i,n-1}, \chi_{i,n-1}\, \mathbb{I}_{2J+1}) \tag{D.4}$$

where $\mathcal{MN}$ denotes the multivariate normal distribution, parameterized in terms of the variance/covariance matrix, which is simply the identity matrix scaled by an adaptive scalar, $\chi_{i,n-1}$, a la Atchade and Rosenthal (2005). Let $\pi_i(\boldsymbol{\Lambda}_{in})$ be a collect-all expression for the prior densities of MH-sampled parameters on the right-hand side of (D.3).[13] Let acc_{in} be the algorithm's acceptance rate for agent i at iterate n, and let $\overline{acc} = 0.234$ be the optimally-targeted acceptance rate.[14] At each iterate, χ_{in} is updated according to:

$$\chi_{in} = \chi_{i,n-1} + \frac{2.5}{n}\left(\min\left\{1, \frac{\mathcal{L}_i(\boldsymbol{\Lambda}_{in})\pi_i(\boldsymbol{\Lambda}_{in})}{\mathcal{L}_i(\boldsymbol{\Lambda}_{i,n-1})\pi_i(\boldsymbol{\Lambda}_{i,n-1})}\right\} - \overline{acc}\right) \tag{D.5}$$

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

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

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

In practice we draw the precisions rather than the variances. The conditionally

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

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

conjugate sampling statements are as follows:

$$\left(\frac{1}{\sigma^2_{ijn}}\right) \mid \mathcal{D}_i, \mathcal{P}_{in} \setminus \{\sigma_{ijn}\}, \mathcal{H}_n \sim \text{Gamma}\left(\tau_{ijn} + \frac{T_i}{2}, \frac{\beta_{ij,n-1}\overline{SSQ}_{ijn} + 2}{2\,\beta_{ij,n-1}}\right)$$

$$\beta_{ijn} \mid \mathcal{D}_i, \mathcal{P}_{in} \setminus \{\beta_{ijn}\}, \mathcal{H}_n \sim \text{Gamma}\left(1 + \tau_{ijn}, \left(1 + \frac{1}{\sigma^2_{ijn}}\right)^{-1}\right)$$

$$\mu_{i,\alpha_{i,J+1},n} \mid \mathcal{D}_i, \mathcal{P}_{in} \setminus \{\mu_{i,\alpha_{i,J+1}.n}\}, \mathcal{H}_n \sim \mathcal{N}\left(\left(\frac{\mu_\mu}{\sigma^2_\mu} + \frac{\alpha_{i,J+1,n}}{100}\right)\left(\frac{1}{\frac{1}{\sigma^2_\mu} + \frac{1}{100}}\right), \frac{1}{\frac{1}{\sigma^2_\mu} + \frac{1}{100}}\right)$$

$$\psi_{ijn} \mid \mathcal{D}_i, \mathcal{P}_{in} \setminus \psi_{ijn}, \mathcal{H}_n \sim \text{Beta}\left(1 + \sum_t \Gamma_{ijtn}, T_i - \sum_t \Gamma_{ijtn}\right)$$

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

D.1.5 Intuition Regarding Identification of Budget Weight Updates

Since our estimation is of the Bayesian learning variety, different parameters may be identified to varying degrees. The key is to assess whether or not prior distributions imposed by the modeler appear different than their posterior estimates. Such a statistical assessment of identification for time-varying mental accounting parameters is difficult. $\boldsymbol{\theta}_{it}$ and $\boldsymbol{a}_{it}$ are pre-determined by all the other parameters and so are not directly sampled, though they are still determined probabilistically in the sense that the probability of observing a sequence of values depends on the priors, likelihood, and data.

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

Identifying atomic realizations of the cognitive shocks driving budget updates, $\{\Gamma_{ijtn}\}_{t=1}^{T_i}$, relies on exploiting information regarding changes in the average levels of the expenditure time series, conditional on income, over multiple periods. This is discussed for more general time series level changes in McCulloch and Tsay (1993). Take σ^2_{ijn}, the likelihood variance, and a_{ij1n}, the initial mental account balance, as given. Also, suppose all future budget-updating indicators, $\{\Gamma_{ijs,n-1}\}_{s>t}$, are taken as given from the previous iteration of the MCMC sampler. Recall further that ζ_{it} is mean zero. Suppose for illustration we observe $x_{ijt} >> \theta_{ij,t-1,n}\ell_{it} + \sum_{s=1}^{t-1}(\theta_{ijsn}l_{is} - x_{ijs}) + a_{ij1n}$, but after this particular deviation of expenditure we observe $x_{ij,t+s'} \approx \sum_{s=t+s'}^{S}(\theta_{ij,t-1,n}l_{is} - x_{ijs}) + \sum_{s=1}^{t-1}(\theta_{ijsn}l_{i,s} - x_{ijs}) + a_{ij1n}$ for some $S \geq t$ and all s' such that $t + s' \leq S$.[15] That is, observations of future expenditure are only small, idiosyncratic deviations from expected expenditure. Then the Gibbs sampler would probabilistically conclude $\Gamma_{ijtn} = 0$, and no switch is necessary (i.e., $\theta_{ijtn} = \theta_{ij,t-1,n}$). Now suppose instead for most s', $x_{ij,t+s'} >> \sum_{s=t+s'}^{S}(\theta_{ij,t-1,n}l_{is} - x_{ijs}) + \sum_{s=1}^{t-1}(\theta_{ijsn}l_{is} - x_{ijs}) + a_{ij1n}$. This compounds and leads the period-t likelihoods to shift further and further out, so that each successive expenditure realization, $x_{ij,t+s'}$, appears less and less likely to the Gibbs sampler. The sampler thus infers that a permanent budget update has been made (i.e. $\Gamma_{ijtn} = 1$) and computes ϑ^*_{iytn} with $\iota_{iyt} = j$ according to the equilibrium condition. Aggregating over all MCMC draws, n, we can thus identify periods in the data that feature higher probabilities of optimal budget updates, conditional upon our modeling assumptions being true. Joint variation in expenditure, income, and

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

balances, via both the spending constraints and the inversion formula, thus inform the probabilistic identification of cognitive shocks leading to budget updates.

D.2 Model Fitness

In Table D.1 we present measures of model performance and fitness under different numeracy constraints and parameterizations. In the top part of the table we present estimates of the fully-parameterized models under different numeracy thresholds. Our two main summary statistics comparing model fit are 1) mean absolute predictive error (MAE) as a fraction of each agent's average income $\widehat{\ell}_i$ and 2) a log-pointwise predictive density (lppd) information criterion as described in Gelman et al. (2013).[16] The mean absolute predictive error as a percentage of average income is

$$\frac{MAE}{\widehat{\ell}_i} = \frac{1}{I \times J} \sum_i \sum_j \frac{1}{T_i} \frac{\sum_t |\widehat{x}_{ijt} - x_{ijt}|}{\widehat{\ell}_i} \tag{D.6}$$
$$\text{where} \quad \widehat{x}_{ijt} = \frac{1}{N} \sum_n \widetilde{x}_{ijtn},$$

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

$$\text{lppd} = \sum_i \sum_t \sum_j \ln \left[\frac{1}{N} \sum_n \phi \left(\frac{\widetilde{x}_{ijtn} - x_{ijt}}{\sigma_{ijn}} \right) \right] \tag{D.7}$$

where the function $\phi(\cdot)$ is the normal density function following from our assumption that ζ_{ijt} is normally distributed and for each consumer independent across j and t. The object in (D.7) inside of the natural log function is the average of the predictive density over the MCMC samples, indexed by n, with $\widetilde{x}_{ijtn}$ being the actual, atomic predictions of the model at every iteration n of the MCMC sampler. As a measure of model convergence and performance, in Table D.1 we also present the

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

mean, median, and modal MH acceptance rates for the agent-specific MH within Gibbs sampler.

Table D.1: Model Performance and Comparisons

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

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

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

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

D.3 MCMC Mixing Properties

For a visual assessment of model performance and MCMC mixing properties, we plot the entire chain of posterior draws, averaging each draw across agents, for parameters that do not change over time. Specifically, if $\daleth_{in}$ represents the n^{th} draw after burnin and trimming of parameter $\daleth$ for agent i, then at each draw we compute $\widehat{\daleth}_n = \frac{1}{I}\sum_i \daleth_{in}$. These chains are presented in Figure D.1 for the fully-parameterized baseline model with no numeracy constraints and Figure D.2 for the \$1-threshold model. Note that the hyperparameters (τ_{ij}) associated with the likelihood variance exhibit no autocorrelation in the model without numeracy conditions, compared to substantial autocorrelation in the \$1-threshold model. It is for this reason that we technically prefer the fully-parameterized model, however, the results in the \$1-threshold model should not be completely discounted, as the sampling properties for behavioral parameters like ψ_{ij}, appear robust.

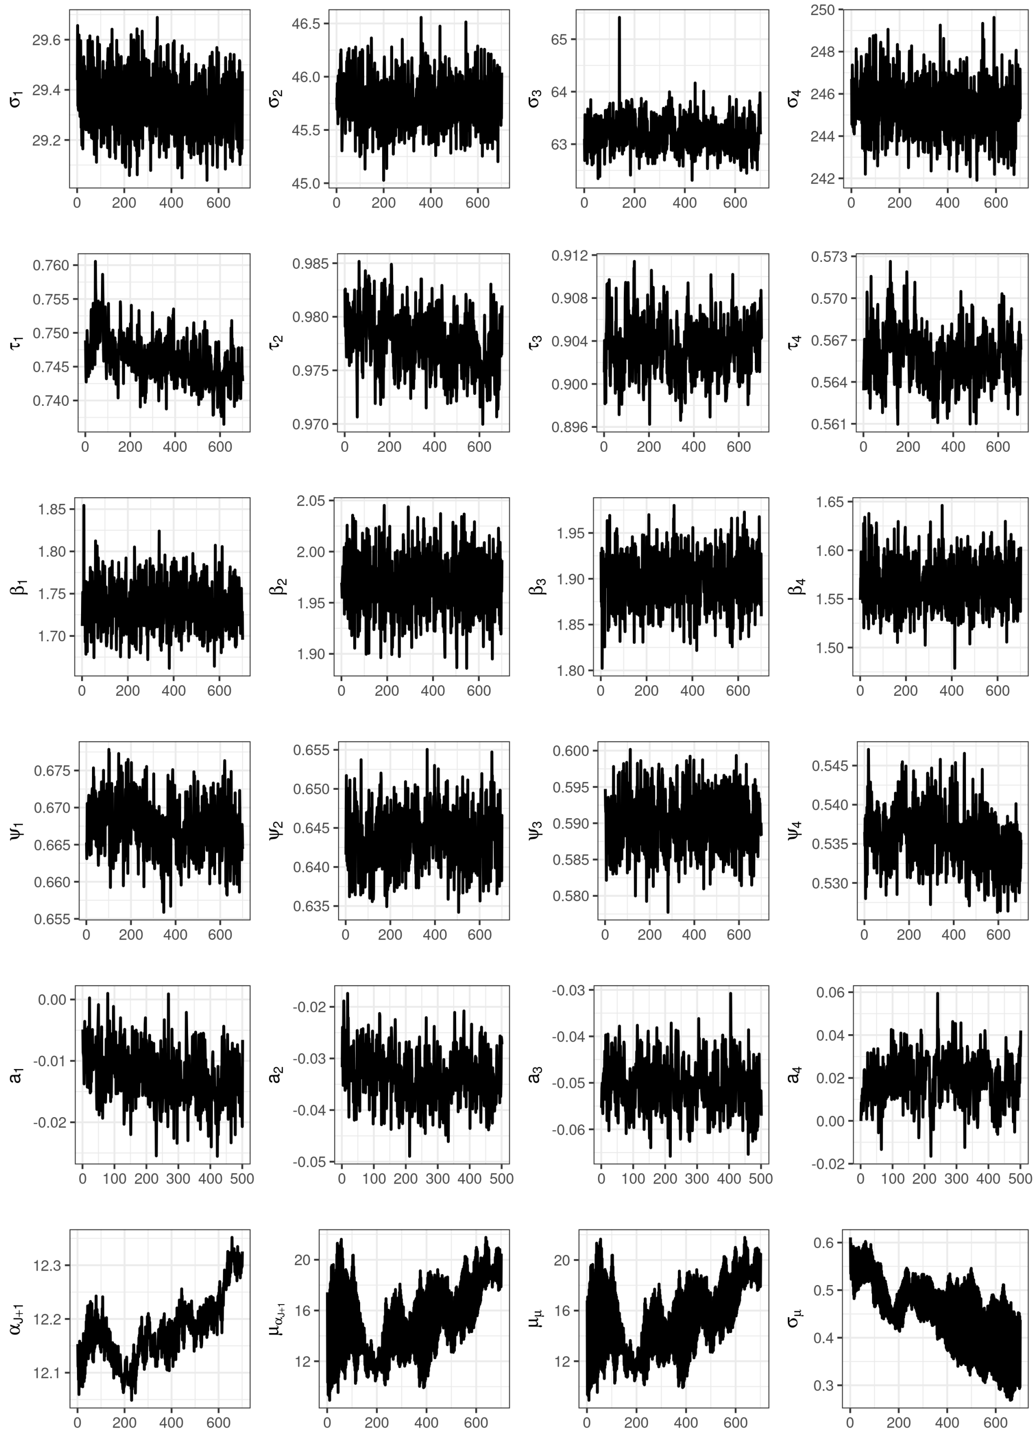

Figure D.1: ($\epsilon = \epsilon_\ell = 0$) On the horizontal axis 0 corresponds to the first MCMC draw after burn-in. Recall the indexing of goods: groceries ($j = 1$), auto/gasoline ($j = 2$), food away from home ($j = 3$), and other expenditure ($j = 4$).

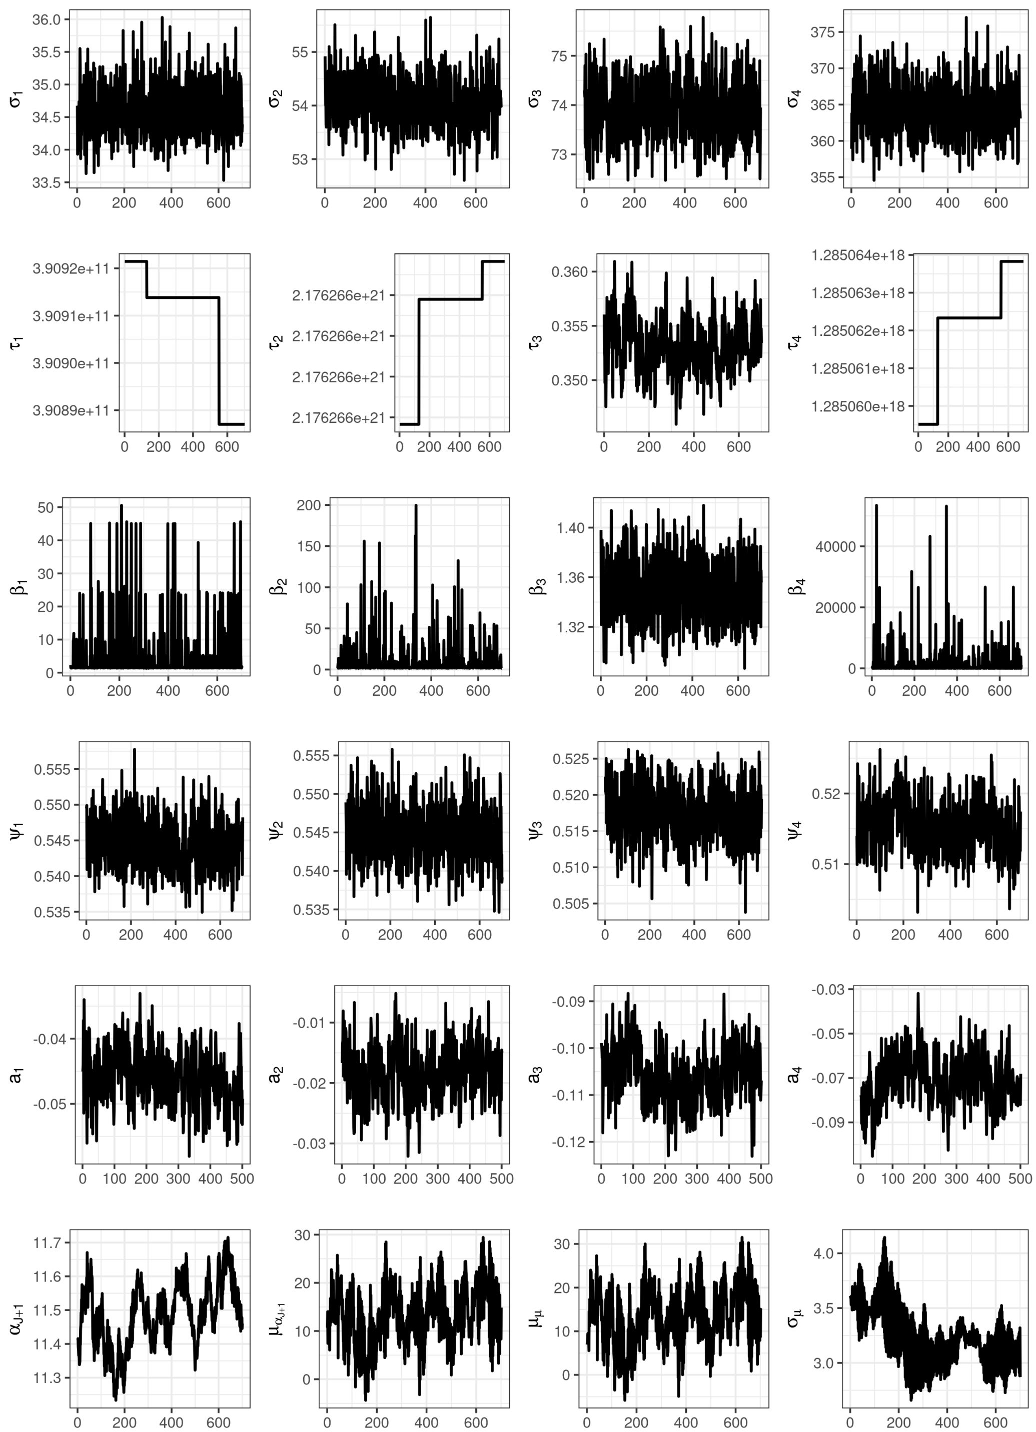

Figure D.2: ($\epsilon_\ell = 1$) On the horizontal axis 0 corresponds to the first MCMC draw after burn-in. Recall the indexing of goods: groceries ($j = 1$), auto/gasoline ($j = 2$), food away from home ($j = 3$), and other expenditure ($j = 4$).

E Additional Behavioral Quantitative Results

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

In this section we analyze the joint relationships between posterior budget updates k_{it}, income ℓ_{it}, and savings z_{it}. We find the following: 1) budget-updating behavior is generally random and uncorrelated with broad characteristics of the consumer, such as his/her average earnings and/or savings behavior; 2) while, on aggregate, income and savings rates tell us very little about how a consumer will systematically engage in budgeting, within a consumer unit income and savings rates over time are associated with variation in k_{it}, though relations are idiosyncratic, and these tendencies are subject to substantial heterogeneity across consumer units; 3) most consumers have slightly higher savings in periods with more budget updates, though this finding is also rather idiosyncratic; 4) budget updates occur at slightly higher rates when consumers experience positive versus negative income shocks in the same period, though the aggregate distribution of $\widehat{k}_i$ 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 cursory (though very weak) evidence that k_{it} may be endogenous.[17]

Figure E.1 shows contour plots of the joint distributions of average income and average savings rates against $\widehat{k}_i$. 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 $\rho = 0.108$ for the baseline model and $\rho = 0.263$ for the \$1-threshold model), while the average savings rate appears uncorrelated with $\widehat{k}_i$ (Pearson's $\rho = -0.008$ for the baseline model and $\rho = 0.029$ for the \$1-threshold model). Notice that for the \$1-threshold model the correlation between average income and average weekly budget updates in-

[17] 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.

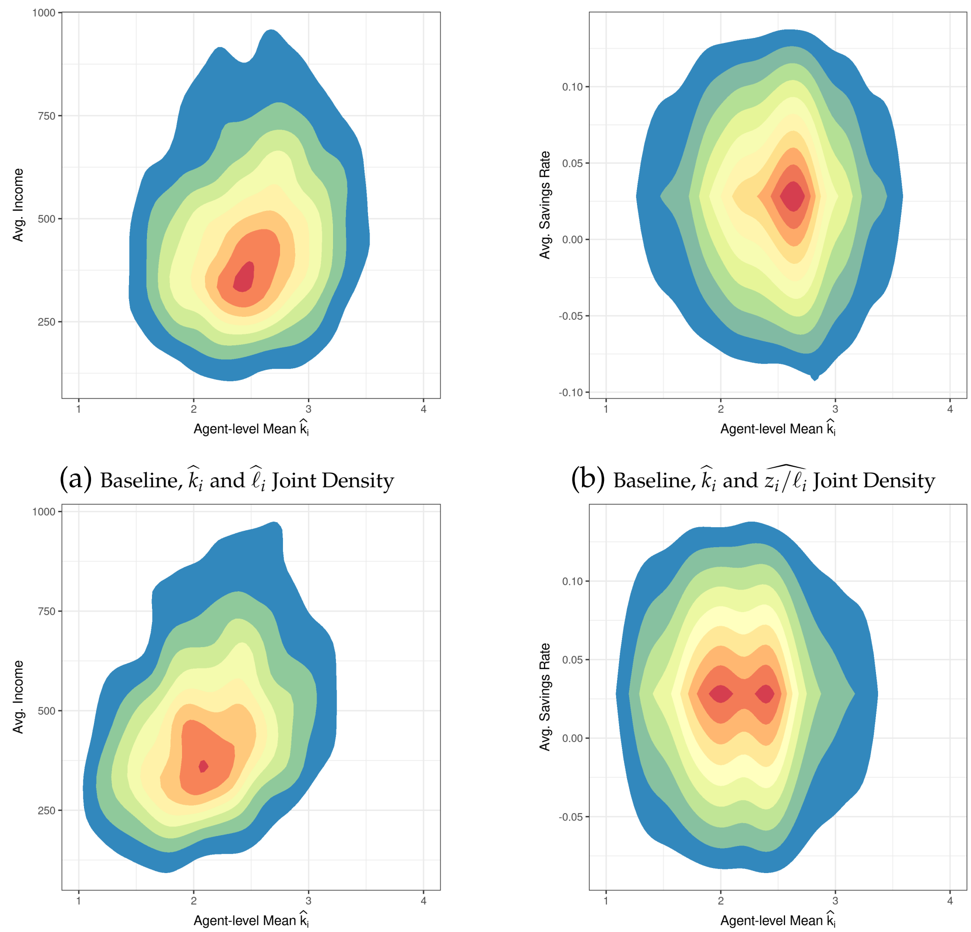

(a) Baseline, $\widehat{k}_i$ and $\widehat{\ell}_i$ Joint Density

(b) Baseline, $\widehat{k}_i$ and $\widehat{z_i/\ell_i}$ Joint Density

(c) \$1 Threshold, $\widehat{k}_i$ and $\widehat{\ell}_i$ Joint Density

(d) \$1 Threshold, $\widehat{k}_i$ and $\widehat{z_i/\ell_i}$ Joint Density

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

creases slightly but is still less than 0.3. Since we are integrating both over the entire posterior distribution of draws k_{itn} and time, this exercise demonstrates that estimates of cognitive frictions, Γ_{it}, may be orthogonal to an agent's average income level $\widehat{\ell}_i$ and propensity to save. We thus conclude that a consumer's average behavior with regards to budget attentiveness does not appear to be strongly associated with either their earnings potential or average savings behavior.

However, within each consumer unit budget-updating behavior can still be strongly associated with fluctuations in personal income over time. Further, we

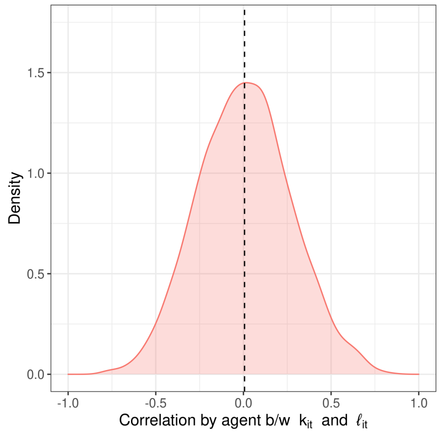

(a) Baseline, $\mathrm{Cor}(\widehat{k}_{it}, \ell_{it})$

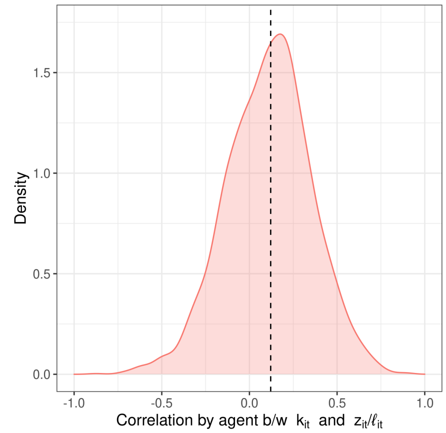

(b) Baseline, $\mathrm{Cor}(\widehat{k}_{it}, z_{it}/\ell_{it})$

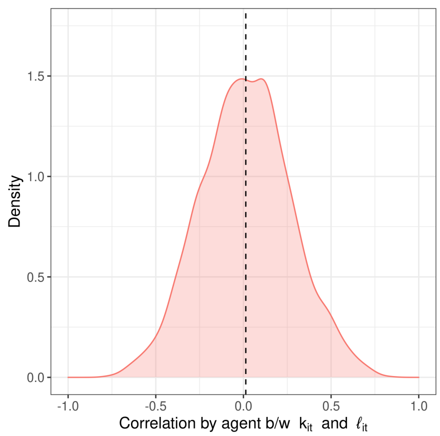

(c) \$1 Dollar, $\mathrm{Cor}(\widehat{k}_{it}, \ell_{it})$

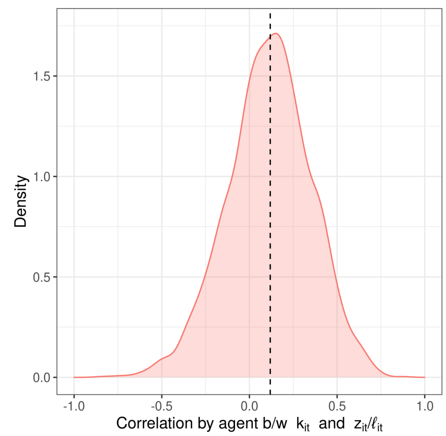

(d) \$1 Threshold, $\mathrm{Cor}(\widehat{k}_{it}, z_{it}/\ell_{it})$

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

find that most consumers (68.1% of the sample in the baseline and 69.5% with a \$1 threshold) tend to engage in relatively higher savings (for them) in periods featuring greater budget attentiveness. To demonstrate consumer heterogeneity along these dimensions, we compute agent-specific 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 E.2. Consumers appear equally likely to have either higher or lower relative income in periods of greater rationality, which can be seen by noting the symmetry of the densities in panels (a) and (c). However, for more consumers, periods with greater rationality are also associated with higher relative savings rates, which can be seen by noting that the densities in panels (b) and (d) are skewed positive.

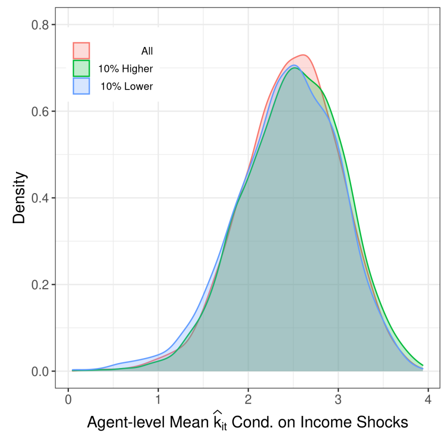

(a) Baseline, $\widehat{k}_{it}$ Conditional Density

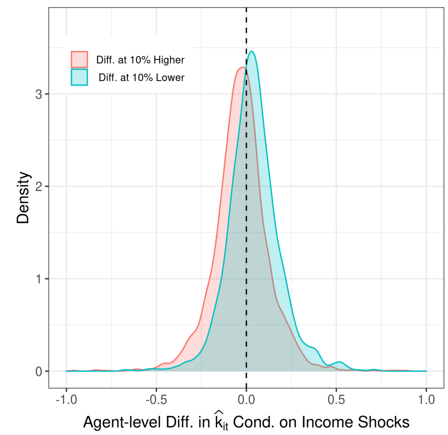

(b) Baseline, $\widehat{k}_{it}$ Conditional Density

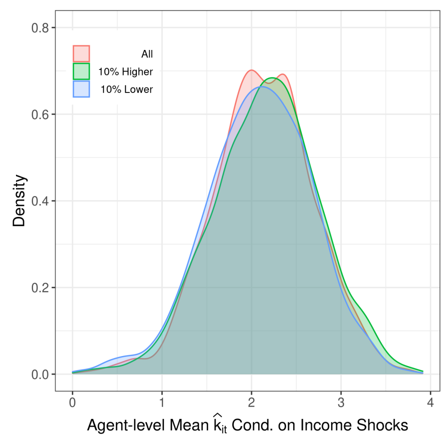

(c) \$1 Threshold, $\widehat{k}_{it}$ Conditional Density

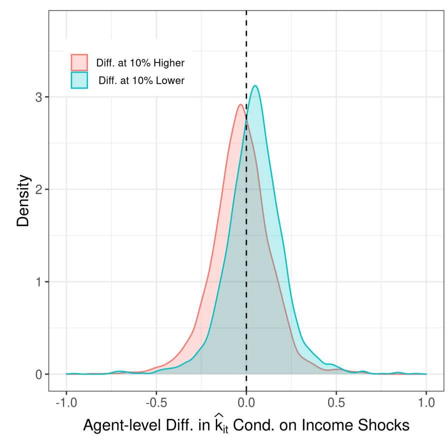

(d) \$1 Threshold, $\widehat{k}_{it}$ Conditional Density

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

To assess whether the relationship between budget-updating and income may be non-linear and thus inadequately explained by Pearson's correlation coefficients, we check if there is any difference between $\widehat{k}_{it}$ in periods when ℓ_{it} is 10%

higher or lower than $\ell_{i,t-1}$. Figures E.3a and E.3c plot the conditional kernel density estimates of $\widehat{k}_{it}$ under different income-shock scenarios against the full estimate. Figures E.3b and E.3d show the densities of differences of $\widehat{k}_{it}$ relative to the baseline estimate. Specifically, in Figures E.3b and E.3d, negative values mean that the full $\widehat{k}_{it}$ is *less* than the conditional estimate, while positive values mean the full estimate is *greater* than the conditional estimate. We find that 59.8% (baseline) and 59.4% ($1 threshold) of consumers in our sample tend to make more budget updates in periods when income rises by 10%, while 38.8% (baseline) and 36.5% ($1 threshold) of consumers tend to make more updates in periods when income falls by 10%. However, the magnitude of these differences is slight — median differences of -0.028 (baseline) and -0.029 ($1 threshold) after a 10% positive shock versus median differences of 0.034 (baseline) and 0.046 ($1 threshold) after 10% negative shocks. We conclude that while $\widehat{k}_{it}$ is slightly correlated with income shocks, the relationship does not appear significant.

To address the possibility that a consumer's budget-updating behavior is made in response to past income or in anticipation of future income, we consider the correlation between $\widehat{k}_{it}$ and $\ell_{i,t\pm s}$, where $s \in \{1, 2, 3\}$. In Figure E.4 we plot the conditional densities across agents of Pearson's coefficients measuring the correlation between budget updates and past and future income. We find no systematic relationship between budget updates and anticipated income, as seen in panels (b) and (d) where the densities are all centered around zero. In panels (a) and (c) we see that there appears to be no systematic relationship between $\widehat{k}_{it}$ and income two and three periods in the past. However, $\widehat{k}_{it}$ and $\ell_{i,t-1}$ (red) are negatively correlated for approximately 62.0% (baseline) and 61.2% ($1 threshold) of consumers.

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

Finally, we can also check whether the correlations between individual budget-updating behavior, income, and savings over time are at all systematically associated with the average income distribution. That is, do consumers who are more

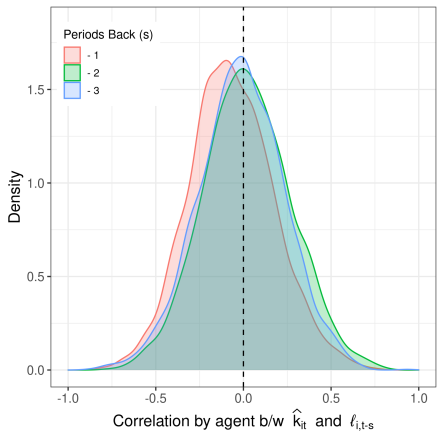

(a) Baseline, $\widehat{k}_{it}$ and Past Income

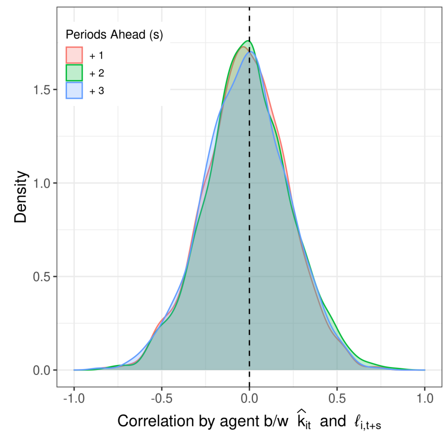

(b) Baseline, $\widehat{k}_{it}$ and Future Income

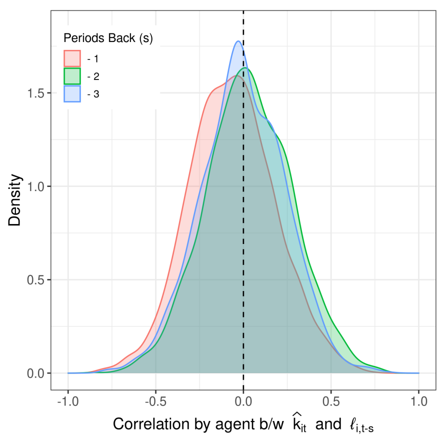

(c) \$1 Threshold, $\widehat{k}_{it}$ and Past Income

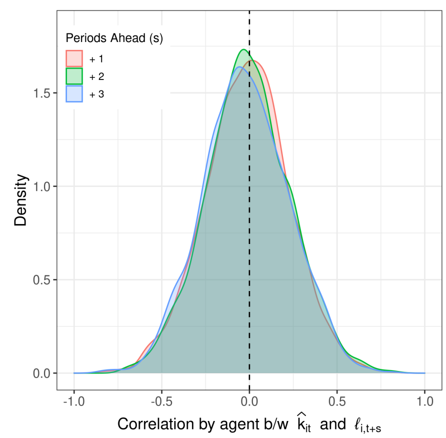

(d) \$1 Threshold, $\widehat{k}_{it}$ and Future Income

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

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

E.2 Summary Statistics of Selected Agent Characteristics by Type

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

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

[a] We first take averages over agent-level outcomes then take the median of the cross-sectional averages.
[b] These are the fully-parameterized models with the corresponding numeracy thresholds.

References

Yves Atchade and Jeffrey Rosenthal. On adaptive markov chain monte carlo algorithms. *Bernoulli*, 11(5):815–828, 2005.

Angus Deaton and John Muellbauer. *Economics and Consumer Behavior*, chapter 5, pages 119–147. Cambridge University Press: New York, 1980.

Daniel Fink. A compendium of conjugate priors. *Mimeo: Environmental Statistics Group, Department of Biology, Montana State University*, 1997.

Roy Geary. A note on 'a constant-utility index of the cost of living'. *Review of Economic Studies*, 18(2):65–66, 1950.

Andrew Gelman, John Carlin, Hal Stern, David Dunson, Aki Vehtari, and Donald Rubin. *Bayesian Data Analysis*. Taylor & Francis, CRC Press: Boca Raton, FL, third edition, 2013.

Heikki Haario, Eero Saksman, and Johanna Tamminen. Adaptive proposal distribution for random walk metropolis algorithm. *Computational Statistics*, 1998.

Heikki Haario, Eero Saksman, and Johanna Tamminen. An adaptive metropolis algorithm. *Bernoulli*, 7(2), 2001.

Robert McCulloch and Ruey Tsay. Bayesian inference and prediction for mean and variance shifts in autoregressive time series. *Journal of the American Statistical Association*, 88(423):968–978, 1993.

Gareth Roberts, Andrew Gelman, and Walter Gilks. Weak convergence and optimal scaling of random walk metropolis algorithms. *The Annals of Applied Probability*, 7(1):110–120, 1997.

Richard Stone. Linear expenditure systems and demand analysis: An application to the pattern of british demand. *The Economic Journal*, 64(255):511–527, 1954.