

An Attention-Based Theory of Mental Accounting*

Botond Kőszegi[†] and Filip Matějka[‡]

August 12, 2018

Abstract

We analyze how an agent with costly attention optimally attends and responds to shocks in basic consumption decisions, explaining several types of mental accounting, and making additional predictions. When allocating consumption among goods with different degrees of substitutability, the agent may create budgets for the more substitutable products. But if the goods are complements, the agent — consistent with naive diversification — may choose a fixed, unconsidered mix of products. When managing her consumption from and transfers between an investment account and a checking account, the latter of which she has an incentive to balance, the agent’s marginal propensity to consume (MPC) out of shocks to the checking account is greater than her MPC out of shocks to the investment account. Furthermore, because paying attention to her budgeting is costly, the agent prefers to reduce spending risk, and she is more averse to risk in the checking account than in the investment account. As a result, she may optimally switch to a cheaper substitute product when a random price increase occurs.

Keywords: mental accounting, information acquisition, rational inattention, naive diversification.

* We thank Yuriy Gorodnichenko, Paul Heidhues, Rupal Kamdar, Marc Kaufmann, David Laibson, Andrei Shleifer, Philipp Strack, and seminar participants at Carnegie Mellon University, Cornell University, D-TEA, Harvard University, Princeton University, UC Berkeley, UC Davis, Washington University, and SITE for excellent comments. This project has received funding from the European Research Council under the European Union’s Horizon 2020 research and innovation programme (grant agreement No 678081) and the UNCE project (UNCE/HUM/035).

[†]Central European University. botondkoszegi@gmail.com

[‡]CERGE-EI, a joint workplace of Charles University in Prague and the Economics Institute of the Czech Academy of Sciences. filip.matejka@cerge-ei.cz

1 Introduction

Mental accounting — “the set of cognitive operations used by individuals and households to organize, evaluate, and keep track of financial activities” (Thaler, 1999) — is one of the oldest and most commonly invoked ideas in behavioral economics. The central component of mental accounting is the use of virtual (e.g., food vs. education) or real (e.g., checking vs. investment) accounts to guide consumption and savings choices without reoptimizing after every change in circumstances. A large empirical and experimental literature has identified types of reactions to shocks that violate the fungibility of money inherent in standard life-cycle models, but that can be interpreted in terms of mental accounting. Yet there is no theory that explains how a person can create separate mental accounts from fungible finances, and how this process interacts with her reactions to shocks.

In this paper, we develop a theory of mental accounting based on the idea that individuals find it too costly to think through all information relevant for consumption decisions, and therefore they pick and choose what information to pay attention to. When deciding how to allocate consumption between substitute products, the agent may find it optimal to think only about which of the more substitutable goods to consume, and not to think about tradeoffs between dissimilar goods. As a result, she may behave as if she had a fixed budget for the more substitutable goods. But if the goods are complements, the agent may find it optimal not to think about their relative consumption at all, only about how much to consume in total. Consistent with naive diversification, therefore, she may choose a fixed, unconsidered mix of products. When managing her consumption from and transfers between an investment account and a checking account, the latter of which she has an incentive to balance, the agent finds it more useful to pay attention to the checking account. As a result, her consumption is more responsive to the checking-account balance than to the investment-account balance. And because taking on expenditure risk entails having to pay extra attention to the checking-account balance, the agent could be more averse to risk in expenditures than in investments, and she may respond to random price increases by switching to cheaper substitutes.

After illustrating the logic of our budgeting and naive-diversification results in a simple example in Section 2, we turn to our general framework. Recognizing that the main questions of mental accounting — how a person allocates spending across multiple products or from multiple sources of income — are inherently multidimensional in nature, in Section 3 we develop a general theoretical methodology for solving rational-inattention problems with multidimensional decisions and

information. These methods are also of independent interest as they are likely to apply to many economic situations. To facilitate full analytical solutions, we make a few functional-form assumptions. We posit that the decisionmaker’s utility is quadratic in her action $\mathbf{y} \in \mathbb{R}^N$ and a shock $\mathbf{x} \in \mathbb{R}^J$, where the x_j are independent normal variables with the same variance. The decisionmaker can acquire any signals of her choice about \mathbf{x} , with the cost of information being equal to the reduction in the entropy of her beliefs. Based on the water-filling algorithm used in information theory (Telatar, 1999; Cover and Thomas, 2006), we show that the optimal information-acquisition strategy is to obtain independent normal signals of N or fewer orthogonal linear combinations of the x_j , with more precise signals about linear combinations that are “more important” in a specific way we identify (Proposition 1). Optimal actions, in turn, vary only along the same number of “important” directions, and while inattention to shocks implies that overall the agent underresponds to shocks, she overresponds to shocks in more important directions relative to those in less important directions (Proposition 2).

With our general results in hand, we analyze the implications of costly attention for consumption decisions. In Section 4, we consider how a person allocates expenditure to consumption goods when she faces uncertainty about her preferences or environment (e.g., consumption opportunities, prices) that she can partially resolve by devoting costly attention to it. The agent has utility function $-\sum_m y_m^2 - \sum_{m \neq n} \Theta_{mn} y_m y_n + \sum_m x_m y_m$, where the y_m are her consumption levels for N goods, the known Θ_{mn} are substitutability parameters, and the unknown x_m capture the uncertainty in her preferences or environment. To study how the substitutability of products affects decisionmaking, as well as for tractability, we assume that the goods can be grouped into different nested consumption categories, with two goods being more substitutable if they are in a smaller common category. For instance, food and theater could both be in the category “entertainment,” which itself could be part of the category “discretionary spending.” Then, the decisionmaker often behaves as if she had mental accounts: her total consumption in a category is fixed (Proposition 3). This means that (i) like with separable utility — which however we do not assume — consumption in a category is independent of shocks to other categories; and (ii) unlike with any natural utility function, total consumption is unresponsive, but individual consumption levels are smoothly responsive, to shocks within the category. Intuitively, the agent’s most important consideration is which of multiple highly substitutable products is worth buying. If she has sufficiently costly attention, therefore,

she thinks only about this consideration. As a result, she does not think about how much to consume, and hence her budget is fixed. Even when she does not have hard budgets, she pays more attention to the above consideration than to everything else, so she trades off within smaller categories relative to larger categories to a greater extent than a person with costless attention (Proposition 4).

The above results are consistent with evidence that many — especially lower-income — individuals and households separate expenditures into budgetary categories (Rainwater et al., 1959; Kahneman and Tversky, 1984; Lave, 1995; Ameriks et al., 2003; Antonides et al., 2011). Beyond explaining such findings, our model provides a novel prediction as to how products are grouped into mental accounts: according to their substitutability. Furthermore, we derive comparative statics on the extent of budgeting. If the agent faces a greater abundance of choice from closely substitutable products, she uses more of her attention to decide which of these products is most worth buying, so she is more likely to engage in budgeting. And if her tastes or the economic environment are more volatile (the variance of x_m is greater), then she must pay more attention, so she is less likely to engage in budgeting.

We also consider what happens when products are complements, and (similarly to the case of substitutes) they are the more so when they are in a smaller common category. In contrast to the above, the agent may now not think about tradeoffs within a category of products at all. Intuitively, because the optimal consumption levels for complementary products tend to move together, the agent does not think about their optimal relative consumption, only about how much she should consume in total. As a result, she chooses a fixed, unconsidered mix of products (Proposition 5). We argue that this prediction is consistent with the phenomenon of naive diversification in financial (Benartzi and Thaler, 2001, 2007) and consumption (Simonson, 1990) decisions. Hence, mental budgeting and naive diversification can be viewed as solutions to the same type of decisionmaking problem that apply in different circumstances.

In Section 5, we apply our framework to analyzing how a person treats income and expenditure shocks. We assume that the agent has two accounts, a checking account and an investment account, with the account balances subject to shocks due to the accrual of income and expenditure. The agent decides how much attention to allocate to the balances in her accounts, and given how much she knows, she chooses her consumption level and a costless transfer to or from her checking account.

Because overdrawing one’s checking account carries a financial penalty while leaving money in the account leads to foregone investment income, we posit that the agent has an incentive to balance the checking account. We show that the agent has a higher MPC out of shocks to the checking account than out of shocks to the investment account (Proposition 6). Intuitively, the agent’s incentive to balance the checking account leads her to pay more attention to it. Hence, she is more likely to notice shocks to the checking account, so her MPC out of these shocks is greater. A difference in MPC’s also arises if the agent first makes the transfer decision, and then allows herself to consume everything in the checking account (Proposition 7). These predictions explain evidence that individuals have lower MPC’s from income accruing to investment accounts than from cash receipts (Shefrin and Thaler, 1988; Hatsopoulos et al., 1989; Baker et al., 2007).

Finally, we observe that the agent’s attention costs generate a novel type of risk aversion. Because risky spending makes it necessary to pay more attention to the checking account or risk failing to balance the account — both of which are costly — the agent dislikes spending risk. But because she needs to pay less attention to her investment account, she is more willing to take risks in her investments, so that — consistent with evidence on the context-dependence of risk attitudes — her risk aversion is smaller for investments than for expenditure (Proposition 8). One way to reduce spending risk is to have nominal budgets, or — consistent with evidence by Hastings and Shapiro (2013) on gasoline purchases — to switch to cheaper substitutes when prices rise.

In Section 6, we briefly discuss theoretical research related to our paper. While some work (such as Shefrin and Thaler, 1988; Prelec and Loewenstein, 1998) explores other aspects of mental accounting, there is no previous micro-founded theory of the main phenomena — how the agent creates budgets from fungible finances, and how this affects her reactions to shocks — that we consider in this paper.

We conclude in Section 7 by discussing some issues missing from our paper. Most importantly, compelling evidence indicates that when individuals receive a subsidy that can only be spent on a subset of products, consumption of those products increases by more than if they receive cash, even among those who would have spent more than the subsidy on the allowed items anyway (Abeler and Marklein, 2016; Hastings and Shapiro, 2017). This phenomenon is usually interpreted in terms of mental accounting, but being about individuals’ reactions to ex-ante transfers, it cannot be modeled using our methodology based on ex-post shocks. Nevertheless, we argue that there is a

natural attention-based account for the phenomenon. We also add, however, that mental accounting is not solely about costly attention, and to the extent it is, individuals do not necessarily allocate attention in a fully rational way. But our mechanisms might be relevant (potentially more relevant) even if we allow for non-rational attention.

2 Budgeting and Naive Diversification: Example

In this section, we illustrate the logic of our budgeting result, and its relationship with naive diversification, using a simple example. We substantially generalize this example, and derive other predictions, in Section 4. Consider two goods, 1 and 2, for which the agent's tastes, x_1 and x_2 , are uncertain. The agent chooses the consumption levels y_1 and y_2 to maximize the expectation of

$$U(y_1, y_2, x_1, x_2) = x_1 y_1 + x_2 y_2 - \frac{y_1^2}{2} - \frac{y_2^2}{2} - \theta y_1 y_2, \quad (1)$$

where $\theta \in (-1, 1)$ captures the substitutability between the two goods. If $\theta > 0$, the goods are substitutes, and if $\theta < 0$, they are complements. The a-priori unknown x_1 and x_2 are independent and drawn from $N(0, 1)$; note that the disutility of spending money on the products can be subsumed into x_1 and x_2 .¹ Before choosing y_1 and y_2 , the agent can observe exactly one of x_1 , x_2 , $x_1 + x_2$, and $x_1 - x_2$: she can think about her taste for one of the goods, her total appetite for the two goods, or her relative taste for the two goods. We ask: what does the agent optimally choose to think about, and how does this affect her consumption patterns?

To facilitate an answer, we put the problem in a different form. Let $x_m = x_1 - x_2$, $x_p = x_1 + x_2$, $y_m = y_1 - y_2$, and $y_p = y_1 + y_2$. Up to a function of x_1 and x_2 — which the agent cannot influence — the objective (1) can then be written as

$$U(x_m, x_p, y_m, y_p) = -\frac{(x_m - (1 - \theta)y_m)^2}{2(1 - \theta)} - \frac{(x_p - (1 + \theta)y_p)^2}{2(1 + \theta)}.$$

¹ To see this, suppose that both products cost p , and the agent's utility function is $x'_1 y_1 + x'_2 y_2 - \frac{y_1^2}{2} - \frac{y_2^2}{2} - \theta y_1 y_2 - p(x_1 + x_2)$, where x'_1 and x'_2 are shocks to the (gross) marginal utilities of the two products. Setting $x_i = x'_i - p$, we get the same utility function as above.

To maximize the expectation of U conditional on her information, the agent therefore chooses

$$y_m = \frac{E[x_m|\text{info}]}{1 - \theta} \quad \text{and} \quad y_p = \frac{E[x_p|\text{info}]}{1 + \theta}.$$

The objective thus becomes

$$E[U] = -\frac{\text{var}(x_m|\text{info})}{2(1 - \theta)} - \frac{\text{var}(x_p|\text{info})}{2(1 + \theta)}. \quad (2)$$

The agent's optimal information-acquisition choice is now obvious. If the products are substitutes (i.e., $\theta > 0$, and therefore $1/(1 - \theta) > 1/(1 + \theta)$), then reducing the variance of x_m is the most valuable, so x_m is what the agent chooses to observe. Since x_m and x_p are independent, observing x_m provides no information about x_p , so $y_p = E[x_p]/(1 + \theta) = 0$. This means that $y_1 + y_2$ is constant, i.e., the agent has a fixed budget. Conversely, if the products are complements (i.e., $\theta < 0$, and therefore $1/(1 - \theta) < 1/(1 + \theta)$), then reducing the variance of x_p is the most valuable, so the agent chooses to observe x_p . As a result, she learns nothing about x_m , so $y_m = E[x_m]/(1 - \theta) = 0$. This means that $y_1 = y_2$, i.e., the agent naively diversifies, always choosing the goods in equal proportion.

It is worth considering briefly how correlation between x_1 and x_2 affects these results. We do so for the case of substitutes ($\theta > 0$). If x_1 and x_2 are positively correlated, then $\text{var}[x_p] > \text{var}[x_m]$, which by Equation (2) weakens the preference for observing x_m . If the correlation is sufficiently strong, then the agent prefers to observe x_p over x_m . Intuitively, if the tastes for two products are highly positively correlated, then the consumer is unlikely to learn much from thinking about which one she likes, so she does not think about this. Conversely, a negative correlation between x_1 and x_2 strengthens the preference for observing x_m .

Unfortunately, it is difficult to know what correlation between x_1 and x_2 is most realistic. When x_1 and x_2 represent taste shocks for similar products, it is plausible that they are positively correlated: one may like an Italian dinner because one feels like going out, in which case one would also like a movie. But a negative correlation is equally plausible: when one is in the mindset of a movie, one is less likely to be in the mindset of a play. And when x_1 and x_2 represent prices or consumption opportunities rather than tastes, it is unclear whether there should be any systematic correlation between them.

We should, however, emphasize an important implicit assumption of our framework: that the agent knows, or no longer wishes to think much about, the value of saving money. Since the value of saving affects the disutility of spending on both products equally, uncertainty about the value of saving is equivalent to a positive correlation between x_1 and x_2 . If the value of money is highly uncertain, therefore, the agent (preferring to learn about x_p) first wants to think about the value of money. In this sense, thinking about the importance of saving is a precursor to budgeting.

3 Theoretical Tools

In this section, we develop theoretical tools for analyzing rational-inattention models in which — as with mental accounting — the agent’s information and action are multi-dimensional.² Since these tools are potentially applicable to many economic settings, we present them in a general form. We lay out our results on mental accounting in a self-contained way, so readers not interested in the general tools can skip to Section 4.

3.1 Multi-Dimensional Rational Inattention

We model an agent who processes information about a random state of the world $\mathbf{x} \in \mathbb{R}^J$, and then chooses her action $\mathbf{y} \in \mathbb{R}^N$ to maximize the expectation of the utility function $U(\mathbf{y}, \mathbf{x})$ less the cost of information processing. For analytical convenience, we rewrite the agent’s strategy in an equivalent way: instead of thinking of her as choosing an information-acquisition strategy and an action strategy, we think of her as choosing a joint distribution $f(\mathbf{y}, \mathbf{x})$ between her action \mathbf{y} and the state \mathbf{x} .³ Given these considerations, the decisionmaker’s problem is

$$\max_{f \in \mathbb{D}(\mathbb{R}^N \times \mathbb{R}^J)} \int U(\mathbf{y}, \mathbf{x}) f(\mathbf{y}, \mathbf{x}) d\mathbf{x} d\mathbf{y} - \text{cost}(f), \quad (3)$$

$$\text{subject to } \int f(\mathbf{y}, \mathbf{x}) d\mathbf{y} = g(\mathbf{x}) \text{ for all } \mathbf{x}, \quad (4)$$

²For some previous applications of rational inattention, see Veldkamp (2006); Mackowiak and Wiederholt (2009); Woodford (2009); Luo and Young (2014); Matějka and McKay (2015); Caplin and Dean (2015); Matějka (2016).

³Following Kamenica and Gentzkow (2011) and Matějka and McKay (2015), any information-acquisition plus action strategy generates such a distribution f . Conversely, for any distribution f consistent with Bayesian rationality, there is a unique information-cost-minimizing strategy that implements f . This strategy assigns exactly one signal to each action. Intuitively, as long as the cost of information is a convex function of posteriors (as ours is), it is never optimal for different signal realizations to generate the same action.

where $cost(f)$ is the cost of information, $\mathbb{D}(\mathbb{R}^{N \times J})$ is the space of probability distributions over $\mathbb{R}^{N \times J}$, and g is the probability density function of the agent's prior. The constraint (4) captures Bayesian rationality, requiring consistency between prior and posterior belief.

To provide a full analytical solution, we restrict the general problem above to multivariate quadratic problems with Gaussian noise and an entropy-based cost of information. We assume that $U(\mathbf{y}, \mathbf{x}) = -\mathbf{y}'C\mathbf{y} + \mathbf{x}'B\mathbf{y}$, where $B \in \mathbb{R}^{N \times J}$, $C \in \mathbb{R}^{N \times N}$ and C is symmetric and positive definite. The matrix C summarizes complementarities within the vector of actions, while B summarizes interactions between states and actions.⁴ Furthermore, we let the prior uncertainty about \mathbf{x} be multivariate Gaussian with the variance-covariance matrix ψ . For simplicity and tractability, as well as to focus on the choice of attention driven by preferences only, we let $\psi = \sigma_0^2 I$. Extending our results to independent states with different variances or even to correlated states is straightforward (one can just rescale the variables for them to have the same variance, or apply a transformation of coordinates to the independent components), but the resulting expressions would be more complicated.

Finally, we posit a functional form for $cost(f)$ that is commonly used in the literature on rational inattention. We let $cost(f) = \lambda I(\mathbf{y}; \mathbf{x})$, where $\lambda \in \mathbb{R}_+$ and $I(\mathbf{y}; \mathbf{x})$ is the Shannon mutual information between \mathbf{y} and \mathbf{x} .⁵ Mutual information is the expected reduction of entropy H about \mathbf{x} upon taking \mathbf{y} ; or, equivalently, the expected reduction in the entropy of the agent's beliefs upon receiving information. As in the previous literature, there are two main reasons for using this specification. First, it has the basic intuitive property that basing actions on more precise beliefs is more costly. Second, it is highly tractable. In our setting, in particular, it is known that the agent chooses to collect Gaussian signals about \mathbf{x} (Sims, 2003). The resulting posterior beliefs are also Gaussian, with the variance-covariance matrix Σ subject to the agent's choice. This allows us to simplify the cost function to $cost(f) = (\lambda/2) \cdot (\log |\psi| - \log |\Sigma|)$.⁶

Of course, evidence on the precise form of attention costs is (to our knowledge) non-existent, so we cannot be sure that our cost function is empirically accurate. Nevertheless, because our

⁴ This form is a local approximation to any smooth and concave utility function, and it describes the first-order interactions between various elements of both the action and state vectors \mathbf{y} and \mathbf{x} . The utility function could also include linear terms of \mathbf{y} and \mathbf{x} , but the linear term of \mathbf{y} has the same effect as a shift of the coordinates of \mathbf{x} , and terms with \mathbf{x} only amount to an additive constant.

⁵ Formally, $I(\mathbf{y}; \mathbf{x}) = H[\mathbf{x}] - E_{\mathbf{y}}[H[\mathbf{x}|\mathbf{y}]]$, where $H[\mathbf{x}] = -\int_{\mathbf{x}} f(\mathbf{x}) \log(f(\mathbf{x})) d\mathbf{x}$ and $H[\mathbf{x}|\mathbf{y}] = -\int_{\mathbf{x}} f(\mathbf{x}|\mathbf{y}) \log(f(\mathbf{x}|\mathbf{y})) d\mathbf{x}$.

⁶The entropy of a multivariate normal $N(\mu, \Phi)$ of dimension J is $\frac{J}{2}(\log(2\pi) + 1) + \frac{1}{2} \log |\Phi|$.

main qualitative results are driven by patterns in the benefits of information rather than the costs of information, we do not view this issue as particularly problematic. Indeed, our cost function can be viewed as ideal for studying what information agents choose to get based on endogenous considerations about the benefits of information, and not based on exogenous assumptions about the costs of information. This is because all information has the same cost: what matters is the amount of uncertainty reduction, not what the uncertainty is about.

3.2 Optimal Lower-Dimensional Signals

We now explore the link between posterior beliefs Σ and expected utility. In the proof of Proposition 1 in Appendix A we describe the following steps in detail. We first show that the original problem (3)-(4) can be rewritten as:

$$\max_{\Sigma \preceq \psi} -E\left[(\tilde{\mathbf{x}} - \mathbf{x})'\Omega(\tilde{\mathbf{x}} - \mathbf{x})\right] + \frac{\lambda}{2} \log |\Sigma|, \quad (5)$$

where $\Omega = BC^{-1}B'/4$ and $\tilde{\mathbf{x}}$ is the random mean of posterior beliefs about \mathbf{x} , which depends on the realization of noise in signals. The first term in (5) is the expected loss from misperceptions ($\tilde{\mathbf{x}} - \mathbf{x}$) — the distribution of which is given by Σ . The matrix Ω translates misperceptions into utility losses. The second term in (5) is the cost of information, which includes here only the entropy of the posterior, since the entropy of the prior amounts to an additive constant. Finally, the condition $\Sigma \preceq \psi$ means that $(\psi - \Sigma)$ is required to be positive semi-definite: the prior cannot be more precise than the posterior.

Decomposition into 1D Problems. Let v^1, \dots, v^J be an orthonormal basis of eigenvectors of the loss matrix Ω (which is symmetric), with the eigenvalue corresponding to v^i denoted by Λ_i . The utility term in (5) can be conveniently expressed using the transformation of coordinates to this basis. Letting $(\tilde{\mathbf{x}} - \mathbf{x}) = \sum_i \tilde{\eta}_i v^i$, we have

$$(\tilde{\mathbf{x}} - \mathbf{x})'\Omega(\tilde{\mathbf{x}} - \mathbf{x}) = \left(\sum_i \tilde{\eta}_i v^i\right)'\Omega\left(\sum_i \tilde{\eta}_i v^i\right) = \sum_i \Lambda_i \tilde{\eta}_i^2.$$

The eigenvalue Λ_i is thus a scaling parameter for how uncertainty about the linear combination $(v^i \cdot \mathbf{x})$ translates into losses. Now the expectation of $\tilde{\eta}_i^2$ is by definition the posterior variance of $v^i \cdot (\tilde{\mathbf{x}} - \mathbf{x})$. Since the x_i are i.i.d. with prior variance σ_0^2 , the random variables $(v^i \cdot \mathbf{x})$ are also i.i.d.

with prior variance σ_0^2 . Let us denote the posterior variance of $(v^i \cdot \mathbf{x})$ by $\sigma_i^2 \leq \sigma_0^2$. In the proof we show that Σ must be diagonal in the basis of the eigenvectors, and thus $\log |\Sigma| = \sum_i \log \sigma_i^2$. The agent's problem therefore reduces to

$$\max_{\sigma_i^2 \leq \sigma_0^2} \left(- \sum_i \Lambda_i \sigma_i^2 + \frac{\lambda \log \sigma_i^2}{2} \right).$$

This can now be solved separately for each i , reducing the initial multi-dimensional problem to a system of one-dimensional problems. As a result, the optimal information-acquisition strategy has a strikingly simple form:

Proposition 1 (Information Acquisition). *The optimal information-acquisition strategy is to acquire independent signals of $v^i \cdot \mathbf{x}$ such that the posterior variance of $v^i \cdot \mathbf{x}$ is $\min\{\sigma_0^2, \lambda/(2\Lambda_i)\}$.*

Intuitively, the agent processes more information about vectors in the space of \mathbf{x} that are more costly to misestimate. Specifically, if $\sigma_0^2 \leq \lambda/(2\Lambda_i)$, then the agent acquires no information about $v^i \cdot \mathbf{x}$; and if $\sigma_0^2 > \lambda/(2\Lambda_i)$, then she observes a signal about $v^i \cdot \mathbf{x}$ with precision chosen to bring the posterior variance of $v^i \cdot \mathbf{x}$ down to $\lambda/(2\Lambda_i)$. Hence, when the cost of information λ is high ($\lambda/(2\Lambda_i) > \sigma_0^2$ for all i), then the agent does not process any information. If the cost is somewhat lower, then the agent processes information about $v^i \cdot \mathbf{x}$ with the highest Λ_i , but she processes no other information. At even lower costs, the agent processes information about more $v^i \cdot \mathbf{x}$, etc.

Responsiveness of Actions. Next, we discuss implications for actions. We first recall Bayesian updating in one dimension with independent Gaussian signals s of the form $s = x + \epsilon$, where $x \sim N(0, \sigma_0^2)$ and $\epsilon \sim N(0, \sigma_\epsilon^2)$. If $\xi = 1 - (\sigma^2/\sigma_0^2)$, where σ^2 is posterior variance, then the posterior mean equals:

$$\tilde{x} = \xi s = \xi x + \xi(1 - \xi)\epsilon^1, \tag{6}$$

where $\epsilon^1 \sim N(0, 1)$ is unit noise. The variable ξ is the weight on the signal in Bayesian updating, and is increasing with the signal's precision.⁷ It equals the resulting responsiveness of posterior mean \tilde{x} to a state x and also represents the level of attention.

Now, coming back to the multidimensional problem, given that the vectors v^i constitute an orthonormal basis of \mathbb{R}^J , we can decompose any shock to \mathbf{x} into independent shocks along the

⁷Posterior variance is $\sigma^2 = (\sigma_0^2 \sigma_\epsilon^2)/(\sigma_0^2 + \sigma_\epsilon^2)$, and noise in signals and posterior can be rewritten as $\sigma^2 = (1 - \xi)\sigma_0^2$ and $\sigma_\epsilon^2 = \frac{(1-\xi)}{\xi}\sigma_0^2$.

$\{v^i\}_{i=1}^J$. We show in the proof of Proposition 2 that actions move linearly with posterior means, and then we apply the orthogonal decomposition and the insight leading to (6). We get:

$$\mathbf{y} = H\tilde{\mathbf{x}} = H \sum_i (v^i \cdot \tilde{\mathbf{x}})v^i = \sum_i \left(\xi_i v^i \cdot \mathbf{x} + \xi_i(1 - \xi_i)\epsilon^1 \right) h^i, \quad (7)$$

where $H = C^{-1}B'/2$ and $h^i = Hv^i$.

Expected actions conditional on realized \mathbf{x} equal:

$$E[\mathbf{y}|\mathbf{x}] = \sum_i \xi_i (v^i \cdot \tilde{\mathbf{x}}) h^i. \quad (8)$$

We define ε_i^λ as the average change in the action \mathbf{y} when \mathbf{x} changes in direction v^i by 1. We can think of it as the average responsiveness of the agent's behavior to shocks along v^i . Using this notation, the responsiveness under perfect information — when the agent has no attention costs — is ε_i^0 .

Proposition 2 (Optimal Actions).

1. *The space of actions is spanned by $\{Hv^i | \lambda/(2\Lambda_i) < \sigma_0^2\}$.*
2. *The agent underresponds to shocks relative to the perfect-information case:*

$$\varepsilon_i^\lambda = \xi_i(\Lambda_i)\varepsilon_i^0, \text{ where } \xi_i(\Lambda_i) = 1 - \frac{\min(\sigma_0^2, \lambda/(2\Lambda_i))}{\sigma_0^2}. \quad (9)$$

3. *In the range $\Lambda_i > \Lambda_j > \lambda/(2\sigma_0^2)$, the relative responsiveness $\varepsilon_i^\lambda/\varepsilon_j^\lambda$ is strictly increasing in λ .*

Part 1 says that the agent's action moves only along directions that are sufficiently important to pay attention to — that is, along the directions in which losses are highest. Part 2 states that the agent underresponds to shocks. This is a simple implication of costly attention: because the agent pays only partial attention to information, on average she does not notice the extent of shocks, so she does not respond as much as an agent with zero attention costs. More interestingly, Part 3 says that with costly attention, optimal behavior calls for concentrating reactions to shocks in directions that are the most important. As a result, the responsiveness to shocks along v^i relative to v^j is higher than under perfect information if and only if $\Lambda_i > \Lambda_j$.

4 Consumption Patterns

We now begin to apply the tools we have developed in Section 3 to basic consumption problems. In this section, we analyze how a person attends and responds to taste, price or consumption-opportunity shocks when choosing her consumption basket from many products with different degrees of substitutability or complementarity.

The agent faces a menu of N imperfect substitutes or complements, and chooses how much of each to consume. Her utility function over consumption is

$$-\underbrace{\sum_m y_m^2 - \sum_{m \neq n} \Theta_{mn} y_m y_n}_{\text{baseline utility}} + \underbrace{\sum_m x_m y_m}_{\text{shocks}}, \quad (10)$$

where y_m is the level of consumption of good m , x_m is a shock to the marginal utility of consuming good m , and $\Theta \in \mathbb{R}^N \times \mathbb{R}^N$ with $\Theta_{mm} = 1$ is a symmetric matrix that captures the substitutability patterns between the goods. We can interpret the uncertainty in x_m as capturing uncertainty about the agent's taste: she does not know what combination of restaurant dinners, laptops, housing amenities, kids' education, etc. maximizes her well-being. Alternatively, shocks to the marginal utility of consumption could come from shocks to prices or consumption opportunities. As in Section 3, we assume that the x_m are independent normal random variables with the same variance, and in order to know more about them, the agent can flexibly think about them.⁸ Her cost of attention is also the same as in Section 3. In our setting, this implies that the agent's posterior is jointly normal, and she maximizes the sum of her expected utility given her posterior beliefs plus $\lambda \log |\Sigma|/2$, where $\lambda \geq 0$ is a scalar and Σ is the variance-covariance matrix of her posterior.

4.1 Substitutes: Mental Budgeting

In order to study how the agent's consumption behavior and budgeting depend on the substitutability properties of the goods, we posit a simple structure for Θ . First, we assume that the goods

⁸ When uncertainty regarding the x_m corresponds to price or consumption-opportunity shocks, the question arises how the agent can think flexibly about them. For instance, she may be able to find out a product's price, but she cannot look up a noisy version of the price. Even so, there is clearly flexibility in what she thinks about. She may, for instance, find out which of two products is cheaper without thinking about her marginal utility from the two products much.

are substitutes ($\Theta_{mn} > 0$), and they can be grouped into $L \geq 1$ levels of categories. The level L captures the largest category (e.g., discretionary spending), which includes all N goods. The level $L-1$ is the set of second-largest categories (e.g., entertainment), and so on, with the smallest ($l = 1$) categories being individual consumption goods. We denote by $R^{k,l} \subset \{1, \dots, N\}$ a consumption category at level l . We assume that all categories at level l are of the same size ($|R^{k,l}| = |R^{k',l}|$ for all k, k', l), and that each category at level $l < L$ is a subset of a higher category (for each $l < L, k$, there is a k' such that $R^{k,l} \subset R^{k',l+1}$). For instance, dinners and movies could be subsets of the larger entertainment category. The substitutability of two goods is determined by the smallest category to which they both belong. Precisely, let γ^l be constants satisfying $0 \leq \gamma^L < \dots < \gamma^2 < 1$. For two goods m and n , let l be the smallest l' such that there is a k with $m, n \in R^{k,l'}$. Then, $\Theta_{mn} = \gamma^l$. This formulation captures the idea that a good is a better substitute for other goods in its category than for goods in a different category. For instance, a French dinner is a closer substitute to a Chinese dinner than to a movie.

Our main result in this subsection is:

Proposition 3 (Hard Budgeting of Substitute Products). *There are $\lambda_1, \dots, \lambda_L$ satisfying $\lambda_L < \dots < \lambda_1$ such that*

$$\lambda \geq \lambda_l \Leftrightarrow \sum_{m \in R^{k,l}} y_m = \text{constant for all } k. \quad (11)$$

Proposition 3 says that if her attention cost is sufficiently high, then the agent behaves as if she had a fixed mental budget for each l -category of products: her total consumption in each l -category is constant and hence independent of her shocks. In particular, for high attention costs ($\lambda \geq \lambda_1$) the agent does not process any information, so y_m is determined by her prior and hence constant for each m . More interestingly, for somewhat lower attention costs ($\lambda \in [\lambda_2, \lambda_1)$) the agent processes information about shocks to her relative tastes within the 2-categories, and substitutes consumption within these categories to keep total consumption in each 2-category fixed. For even lower attention costs ($\lambda \in [\lambda_3, \lambda_2)$), she keeps consumption within each 3-category constant, and substitutes within these categories. And so on. Hence, the higher is the cost of attention, the narrower are the agent's budgets.

To appreciate ways in which the agent's behavior differs from that of a classical decisionmaker, suppose that $\lambda_l < \lambda < \lambda_{l-1}$. Then, the agent budgets categories at level l , but not at lower levels,

giving rise to two related phenomena. First, the agent’s consumption decisions across l -categories are independent of each other. In a classical consumption problem, this is optimal only if the utility function is separable across l -categories. We do not impose such separability; indeed, with full information a change in any x_m in general affects the consumption of all products. Second, the agent’s total consumption in an l -category is independent of her shocks, but her consumption within the category responds smoothly to within-category shocks. This is in general not the case in a classical model even if the utility function is separable across categories.

The intuition is based on where the agent optimally directs her costly attention. That (say) goods 1 and 2 are close substitutes means that the agent suffers a greater utility loss from choosing the relative amounts of goods 1 and 2 in a suboptimal way than from choosing the relative amounts of (say) goods 1 and 3 in a suboptimal way, or from choosing the total consumption of goods 1 and 2 in a suboptimal way. With her attention being constrained, she thinks only about the most important thing, the relative consumption of goods 1 and 2. As a result, she fixes the total consumption of goods 1 and 2 at the ex-ante optimal level.

Budgeting leads to specific types of mistakes. Suppose, for instance, that the decisionmaker has a fixed budget for entertainment consumption. Then, even if unusually fun entertainment opportunities present themselves in a particular month (i.e., the x_m in the entertainment category all increase), she does not increase her consumption of entertainment. If she had unlimited attention, in contrast, she would respond to such positive shocks by increasing entertainment consumption.

Proposition 3 is roughly consistent with evidence that many individuals allocate consumption to categories governed by budgets. A stark manifestation of this phenomenon is that many households used to place budgets allocated for different purposes into different envelopes or tin cans (Rainwater et al., 1959; Lave, 1995). More recently, Ameriks et al. (2003) and Antonides et al. (2011) document that the mental budgeting (if not physical separation) of expenses is still common. Indeed, most of the many online financial management tools seem to presume that users of these tools want to set budgets for separate categories.⁹ Proposition 3 not only explains these findings in a natural way, it makes a specific prediction on how a person groups goods into budgets: according to their substitutability.¹⁰

⁹ As another manifestation, Kahneman and Tversky (1984) find in hypothetical choices that subjects are less willing to buy a ticket to a play after they have lost their original ticket than after they have lost an equivalent amount of money — presumably because their entertainment budget is more depleted in the former case.

¹⁰ An alternative explanation for why similar products are covered by the same budget is that they are less costly

A budget in Proposition 3 fixes total consumption, not total nominal spending. The two are equivalent if differences between products arise only due to taste or consumption-opportunity shocks. While they are not equivalent under price uncertainty, for two reasons outside the current model the agent may prefer a nominal budget even in the face of price uncertainty. First, setting a nominal budget rather than a consumption budget reduces spending risk, and we show in Section 5.2 that the agent benefits from this. Second, if uncertainty regarding the overall level of prices in a category is relatively small — a reasonable assumption in low-inflation countries — a nominal budget approximates a consumption budget, and might be much easier to implement in practice (Shafir et al., 1997).¹¹

As the flip side of the possibility of budgeting, Proposition 3 also says that if the agent’s attention cost (λ) is sufficiently low, then she does not have fixed budgets for l -categories. Our next proposition implies that even then, the agent can be viewed as having “soft” budgets. To state the result, let $Y^{k,l} = \sum_{m \in R^{k,l}} y_m$ be total consumption in a category $R^{k,l}$, and let $X^{k,l} = \sum_{m \in R^{k,l}} x_m$ be the total shock to the marginal utility of consuming in $R^{k,l}$. The responsiveness of category consumption to category marginal utility is then $\epsilon_B^l = \partial E[Y^{k,l} | X^{k,l} = x] / \partial x$. In this notation, Proposition 3 implies that if $\lambda \geq \lambda_l$, then $\epsilon_B^l = 0$. Equivalently, if $\lambda \geq \lambda_l$, then consumption of a good in an l -category is perfectly negatively correlated with total consumption of the other goods in the same l -category. Beyond these extreme results:

Proposition 4 (Soft Budgeting of Substitute Products). *In the range $\lambda < \lambda_l$:*

1. ϵ_B^l is strictly decreasing in λ .
2. $\epsilon_B^{l-1} / \epsilon_B^l$ is strictly increasing in λ .
3. For any k and different $m, n \in R^{k,l}$, the correlation between y_m and y_n is strictly decreasing in λ .

The first part formalizes a kind of soft budgeting: as the agent’s attention cost increases, her spending on a category becomes more and more rigid, even if not completely fixed. The intuition is simple: reacting to shocks presumes paying attention to shocks, so the costlier is attention, the

to compare: it is easier to decide whether one feels like Chinese or Italian than to decide whether one should eat out or buy a TV. To the extent that this is the case, it strengthens our result that similar products are grouped together, but by itself it does not explain mental budgeting.

¹¹ As a case in point, suppose that the units of consumption are not symmetric across consumption goods — e.g., one unit of y_1 is one theater play, but one unit of y_2 is three and a half movies. Then, a consumption budget specifies that a weighted sum of different consumption events is constant, while a nominal budget simply specifies total spending.




Figure 1: Joint distributions of consumption of goods 1 and 2 for $\sigma_0^2 = 1$ and $\gamma^2 = 1/2, \gamma^3 = 1/4$ under different costs of attention. Iso-density curves are shown.

smaller is the optimal response to shocks. The second part says that as the agent's attention cost increases, she reacts relatively more to shocks within smaller categories. Once again, this is a reflection of her directing her scarce attention to tradeoffs that matter more. And as the third part says, this results in more negative correlation between the consumption levels of two products in the same l -category.

Figure 1 illustrates Propositions 3 and 4 in a specific example. We consider four goods grouped into categories $\{1, 2\}$ and $\{3, 4\}$, and draw the joint distribution of y_1 and y_2 for different levels of λ . For costless attention ($\lambda = 0$), the agent might choose a range of consumption pairs in response to shocks. At the other extreme, if the decisionmaker's attention cost is very high ($\lambda = 1$), then consumption amounts are fixed. For lower, but relatively high attention costs ($\lambda = 0.75, 0.5$), the agent sets a budget for the two products, so her consumption is always on the same budget line. These situations correspond to Proposition 3. For even lower positive attention costs ($\lambda = 0.48, 0.45$), the agent starts substituting goods 1 and 2 with goods 3 and 4, but not as much as with costless attention, so the distribution of y_1 and y_2 is closer to a budget line than with costless attention. These situations correspond to Proposition 4.

Our theory has a number of economically relevant comparative-statics implications regarding how changes in preferences or the economic environment affect budgeting. We begin with changes in the substitutability of the products:

Corollary 1. *An increase in γ^l leads to an increase in λ_{l-1} and a decrease in λ_{l+s} for all $s \in \{0, \dots, L-l\}$.*

If products within an l -category become better substitutes, then the agent responds more to shocks within the l -category, becoming more willing to substitute between $(l-1)$ -categories. Hence, she is less likely to have a hard budget for $(l-1)$ -categories (i.e., λ_{l-1} increases). At the same time, because increasing consumption of one product quickly decreases the marginal utility of other products in the same l -category, increasing consumption in the entire category is subject to quickly decreasing marginal utility. This implies that it is less valuable to adjust the budget for an l -category to circumstances, so that the agent is more likely to have hard budgets for l -categories (i.e., λ_l decreases). The same is true for higher categories.

While it is difficult to test directly, Corollary 1 may help explain why lower-income households are more likely to create category-specific budgets (Rainwater et al., 1959; Antonides et al., 2011). The simple fact that a low-income household consumes closer to the bare minimum in some essential categories implies greater substitutability of products within such categories; e.g., if the household does not consume a food item, it can fall short of getting enough nutritious food, so its marginal utility of other food items rises sharply. A high-income household, in contrast, is never on the verge of getting too little food, so the substitutability between its relevant food choices is lower. Corollary 1 then implies that the low-income household is more likely to have a budget for food.¹²

Second, we consider how ex-ante known changes affect consumption. For instance, the agent's average taste may evolve over time, or the mean price of some products may change. Let \bar{x}_m denote the prior mean of x_m , and let $\bar{\mathbf{x}} = (\bar{x}_1, \dots, \bar{x}_N)$.

Corollary 2. *The agent responds to ex-ante changes as with perfect information:*

$$\frac{dE[\mathbf{y}]}{d\bar{\mathbf{x}}} = \frac{\Theta^{-1}}{2}.$$

¹² An alternative potential explanation for why lower-income consumers are more likely to budget is that they have higher attention costs (higher λ). This possibility is, however, difficult to evaluate based on existing evidence or arguments.

While the agent’s consumption budget does not respond to ex-post changes in circumstances, Corollary 2 implies that in general it does respond to ex-ante known changes. Suppose, for instance, that $\lambda \geq \lambda_l$, so that it is optimal for the agent to set budgets for each l -category. If the agent’s expected taste for all goods in category $R^{k,l}$ increases by the same amount while her mean tastes for other goods remain unchanged, then she sets a higher — but still fixed — budget for the category.

Third, we formalize a sense in which the decision environment can become more complicated for the agent due to an increasing abundance of choice. Suppose we copy each product multiple times, with the copies being closer substitutes to each other than any previous products were: they have substitutability parameter $\gamma^{2'}$ satisfying $\gamma^2 < \gamma^{2'} < 1$. The idea behind this formulation is that innovation leads to new, substitute products being developed, increasing the options available to the decisionmaker. Using primes for other variables in the new decision problem as well, this means that there are now $L + 1$ layers of consumption categories, with $\gamma^{l'} = \gamma^{l-1}$ for all $l > 2$. We have:

Corollary 3. *For all $l > 2$, $\lambda'_l < \lambda_{l-1}$.*

That is, living in a world with more options leads the agent to budget more. With products having close substitutes available, the agent has less of a need to trade off a product with distant substitutes, leading her to budget more.

Finally, we consider comparative statics with respect to volatility in the environment:

Corollary 4. *An increase in σ_0^2 increases ϵ_B^l and λ_l for all l .*

If the agent faces more uncertainty in what she should consume — for instance due to an increase in price volatility — then paying attention has greater return, so she pays more attention to everything. As a result, she tends to respond to shocks more, and she tends to use fixed budgets less.

4.2 Complements: Naive Diversification

We now study complementary products. Our analysis leads to a completely different prediction than in the substitutes case, thereby connecting the logic of mental accounting to the logic of naive diversification.

Formally, we modify our model by assuming that $\gamma^2 < \dots < \gamma^L < 0$, keeping all other assumptions unchanged. This means that products are arranged in a nested fashion into categories,

with products belonging to smaller categories being stronger complements in consumption. For instance, different features of a car (e.g., driving experience, seats, sound system) might be highly complementary to each other, but not to items in one’s clothing. To simplify our statement as well as to capture situations in which the products are ex ante equally desirable, we also assume that the x_m have the same mean.

In this situation, the agent simplifies her decision problem in the following way:

Proposition 5 (Naive Diversification). *There are $\lambda_2, \dots, \lambda_L$ satisfying $\lambda_2 > \dots > \lambda_L$ such that*

$$\lambda \geq \lambda_l \Leftrightarrow \text{for any } k \text{ and any } m, n \in R^{k,l}, y_m = y_n. \quad (12)$$

Proposition 5 says that if the agent’s attention cost is sufficiently high, then she chooses a fixed mix of products in category l that is not responsive to circumstances. This contrasts with the case of substitute products, where it was not the *mix*, but the *budget* that was fixed. Intuitively, because the optimal consumption levels for complementary products tend to move together, the agent does not think about their optimal relative consumption at all, only how much she should consume overall. Continuing with the example of cars, the agent does not think separately about the quality of the engine, seats, sound system, etc. she wants — she only thinks about whether she wants an economy or luxury car.

An important application of the above result is naive diversification in financial decisions, whereby a person chooses a simple mix of investments that is unlikely to be fully optimal. For instance, Benartzi and Thaler (2001) document that many employees in employer-based retirement savings plans divide their investments equally across available funds, and accordingly, employees invest more in stocks if there are more stock funds available. To see how our model can account for this phenomenon in a simple example, suppose that an investor with mean-variance preferences decides the amounts y_1 and y_2 to invest into two assets. There are two equally likely states, with asset 1’s net return being $x_1 + 1$ in state 1 and $x_1 - 1$ in state 2, and asset 2’s net return being $x_2 - 1$ in state 1 and $x_2 + 1$ in state 2. It is easy to check that the mean of the investor’s wealth is $y_1x_1 + y_2x_2$, and the variance is $(y_1 - y_2)^2$, so that this problem maps precisely to our formal framework. Hence, Proposition 5 predicts that an investor with sufficiently costly attention splits her investment equally between the two assets. More generally, because diversification is desirable, different investments are often complements, so Proposition 5 predicts that investors may diversify

naively.¹³

Investigating a similar phenomenon in a completely different domain, Simonson (1990) finds that when choosing items to consume at different future dates, individuals tend to seek variety more than their future preferences seem to justify.¹⁴ Because subjects are choosing only what to consume and not how much to consume, this setting does not exactly fit our framework. Nevertheless, a logic akin to that with complementary products can capture the flavor of why subjects naively diversify: because thinking about their future preferences is unlikely to uncover information that dominates their taste for variety, subjects do not think about this, and therefore choose a mix of products.¹⁵ Consistent with our perspective, Simonson argues that naive diversification is due to the combination of taste uncertainty and the desire to simplify the decision.

Our model predicts a type of naive diversification also for substitute products when the agent's attention cost is so high that she does not obtain any information. In this case, her consumption of all products is fixed at the ex-ante optimal level and is therefore completely unresponsive to circumstances. Furthermore, if the products are ex-ante identical, then their consumption levels are equal. But the more interesting type of naive diversification above, whereby the agent does pay some attention to her decision problem and still naively diversifies, is only possible for complementary products.

5 Reactions to Income and Expenditure Uncertainty

In this section, we develop a model of how the agent handles multiple accounts. Our model can be used to tie together evidence on three important issues: differential marginal propensities to consume out of different accounts, context-dependent attitudes toward risk, and the non-fungible treatment of price shocks.

¹³ In the illustrative example above, the complementarity of the two investments relies on the asset returns being negatively correlated. Even for uncorrelated or somewhat positively correlated asset returns, investments are complements if the investor's disutility from variance is strictly concave. Furthermore, with a precautionary savings motive, risky and safe investments are often complements.

¹⁴ In one study, for instance, students chose snacks to be received at the end of three different classes. When choosing the snacks one at a time at the beginning of these classes, 9% of subjects chose three different snacks. But when simultaneously choosing three snacks ahead of time, 64% of subjects chose three different snacks.

¹⁵ To the extent that in the sequential-choice conditions students know more about their momentary tastes, they are less likely to choose a mix of snacks.

5.1 Marginal Propensities to Consume

We begin with how a person treats income and expenditure accruing to different accounts. Suppose that the agent is consuming her wealth over $T + 1$ periods, 1 through $T + 1$, where $T > 1$. She has two accounts, a standard checking account and an investment account. We think of the investment account quite generally; it could, for instance, include human capital. The initial balances in the checking and investment accounts, x_1 and sx_2 , respectively, are uncertain with normal and independent distributions. For our formulation to be consistent with the theoretical tools developed in Section 3, we assume that x_1 and x_2 have the same variance; but we include the scaling factor $s > 0$ to allow for the possibility that the variance of the investment-account balance is different from that of the checking-account balance. In period 1, the decisionmaker processes information about (x_1, x_2) ; for instance, she could think about expenses she will have to pay from the checking account that month, look up how her investments are doing, or ponder what wages she might expect to earn in the future. Then, she chooses a costless transfer from the investment to the checking account, t , and a consumption level for that period, c_1 , that comes out of the checking account. In periods 2 through $T + 1$, the agent consumes her leftover wealth, and we assume for simplicity that she splits her wealth equally across these periods. Her utility from consumption c_τ in period τ is $-c_\tau^2$. Crucially, we assume that the agent faces a quadratic penalty if she does not balance her checking account in period 1. The penalty captures the idea that a consumer has an incentive to balance her checking account: overdrawing the account typically has direct financial costs in the form of penalties or punitive interest payments, while leaving money in the account has indirect financial costs in the form of foregone investment income. We make the unrealistic assumption that the penalty is symmetric only for tractability, so that the formulation fits into the quadratic setup of Section 3.¹⁶ Combining the above components, the agent's total utility is

$$-c_1^2 - T \cdot \underbrace{[(sx_2 + x_1 - c_1)/T]^2}_{=c_2=\dots=c_{T+1}} - \underbrace{b(x_1 - c_1 + t)^2}_{\text{penalty}}. \quad (13)$$

We find:

¹⁶ We have also considered a more realistic setting in which there is an interest-rate differential between the two accounts, there is a penalty for overdrawing the checking account, but there is no direct penalty for leaving money in the checking account. While we cannot solve such a model analytically, in numerical simulations our results continue to hold.

Proposition 6 (Differential MPC's). *The agent's MPC out of shocks to the checking account is greater than her MPC out of shocks to the investment account ($\partial E[c_1]/\partial x_1 > \partial E[c_1]/\partial(sx_2)$).*

To understand Proposition 6, suppose first that attention is costless and there is no penalty for failing to balance the checking account ($\lambda = b = 0$). The agent then simply sets $c_1 = (x_1 + sx_2)/(T + 1)$, equalizing consumption across periods. Confirming the standard logic, in this case the MPC out of the two accounts is the same. Next, suppose that attention is costly ($\lambda > 0$). To get as close as possible to her optimal consumption given the attention cost she pays, the agent acquires a signal of $x_1 + sx_2$ — she thinks about her total wealth, not separately about her two accounts. As a result, her MPC out of the two accounts is still the same. But when there is also an incentive to balance the checking account ($b > 0$), there is an extra reason to pay attention to this account. Once the agent does so, she is more likely to notice funds arriving in the checking account, and she is therefore more likely to consume out of these funds.¹⁷

Proposition 6 is consistent with Hatsopoulos et al. (1989), who find that stockholders spend some of the cash income they receive from takeovers, but they spend a negligible amount of the capital gains in the stock market. Similarly, Baker et al. (2007) document that individuals' marginal propensity to consume from dividends is higher than from other sources of stock returns.¹⁸

In our model above, the agent attempts to minimize the losses from failing to balance the checking account by paying more attention to the checking account. She may, however, adapt other modes of behavior for the same purpose. As a plausible possibility, she might develop a habit to consume exactly what is available in the checking account after the transfer and net of income and expenditure shocks. Formally, this modifies our model only in that once the agent chooses

¹⁷ One might conjecture that a model without costly attention but with costs of transferring money between the accounts has implications similar to ours. While we are not aware of such a model and therefore cannot fully evaluate the differences, it is clear that the predictions are not the same. As a simple example, consider a modification of our model in which the agent knows the balances in the two accounts, and take first the extreme case in which she cannot transfer money between the accounts in period 1. Still, if x_1 is sufficiently high to cover optimal consumption, the MPC out of the two accounts is the same. To go further, a model with positive transfer costs can generate a negative MPC out of the checking account. In particular, suppose that the transfer cost is a fixed cost (i.e., independent of the amount transferred), and consider a checking-account balance that is high relative to optimal consumption. An increase in this balance can induce the agent to transfer money into the investment account and at the same time decrease current consumption.

¹⁸ See also Shefrin and Thaler (1988) and the citations therein for additional evidence. There are some related findings, however, that the model does not account for. Without additional, arguably ad hoc assumptions, for instance, our model cannot explain contradictory evidence from public economics on whether one-time tax payments generate higher or lower additional consumption than many smaller payments. On the one hand, Feldman (2010) documents that people with the same total tax obligation spend more if they have lower tax withholding than if they receive a larger refund. On the other hand, Sahn et al. (2012) document that the impact of a fiscal stimulus is greater if received in the form of a single check than if received in the form of lower withholding.

t , consumption in period 1 is set to $c_1 = x_1 + t$. For an agent with unlimited attention — who therefore knows x_1 — such a change leaves the choice problem and hence consumption behavior unchanged. For an agent with costly attention, however, the logic of the problem is quite different. Nevertheless, we show that our result regarding the differential MPC's out of the two accounts continues to hold:

Proposition 7 (Differential MPC's 2). *The agent acquires more information about the checking-account than about the investment-account balance, and her MPC out of shocks to the checking account is greater than her MPC out of shocks to the investment account ($\partial E[c_1]/\partial x_1 > \partial E[c_1]/\partial(sx_2)$).*

Note first that the agent always avoids the penalty for failing to balance the checking account, suggesting that — because the motive to pay extra attention to the checking account arose from the desire to balance the account — she does not pay more attention to the checking account. But she now does so for a different reason. Since she will consume everything in the checking account and her optimal consumption in period 1 is less than half of her wealth, her optimal transfer is more sensitive to a \$1 change in the checking-account balance than to a \$1 change in the investment-account balance. Hence, she pays more attention to the former. Furthermore, any part of a shock to the checking account that she does not think about when making the transfer translates into period-1 consumption one-to-one, generating a higher MPC out of the checking account.

Some evidence can be brought to bear to partially discriminate between our two specifications. Empirical research we have cited documents not only differences in the marginal propensity to consume, but also that the marginal propensity out of the checking account can be high (see also Souleles et al., 2006; Kueng, forthcoming). This is consistent with our second model, but only consistent with our first model for a short horizon (low T). To see this, note that for large T even an agent with full information spends little of any increase in wealth in period 1. In the first model, costly attention only exacerbates this tendency, as the agent tends to notice only part of any increase in wealth. In contrast, because in the second model the agent depletes her checking account, her marginal propensity to consume can be high even for large T . Indeed, if her cost of attention λ is so high that she does not obtain information, then her MPC out of the checking account is 1.

5.2 Handling Risk

Our framework has potentially important implications for how a person treats risk. We isolate these implications using a simple variant of our two-account model. Suppose that the agent does not make a consumption decision in period 1, only a transfer decision. Furthermore, her consumption utility in period 2 is linear. We consider a non-degenerate shock Δx that has mean zero and is independent of x_1 and x_2 . An agent with unlimited attention would obviously be neutral to the risk, no matter which account it applies to. In contrast:

Proposition 8. *For any Δx , the agent's expected utility is (i) strictly higher with balances x_1, x_2 than with $x_1 + \Delta x, x_2$; and (ii) equal with balances x_1, x_2 and with $x_1, x_2 + \Delta x$.*

The proposition says that costly attention leads to differential risk aversion across accounts: the agent is more averse to risk in her checking account than to risk in her investment account. The intuition is simple: a shock to the checking account makes it more difficult to successfully balance the account, so that the agent must either pay more attention, or pay more in penalties. To avoid these costly consequences, she would rather avoid the shock to her checking account in the first place. This result identifies a novel source of aversion to risk: the agent dislikes risk not because it exposes her to a decrease in consumption that is painful due to a high marginal utility of consumption or sensation of loss, but because it forces her to pay more attention to her finances. Because a similar consideration does not arise for the investment account, the agent is less averse (in our case, neutral) to risk accruing to that account.

Proposition 8 may help explain various documented forms of aversion to risk in one's spending combined with less risk aversion in other decisions. In particular, rental cars, wired and mobile phones, and other services with positive marginal cost are often offered to individuals at a flat rate, presumably because consumers prefer it.¹⁹ Similarly, individuals pay a high premium for low deductibles in housing (Sydnor, 2010) and health (Bhargava et al., 2017) insurance. Yet individuals are willing to bear more risk in other decisions. For instance, while for deductibles Sydnor (2010) estimates a coefficient of relative risk aversion in the triple digits, based on hypothetical choices between large gambles on lifetime wealth, Barsky et al. (1997) measure an average coefficient of relative risk aversion of around 5.

¹⁹ See Herweg and Mierendorff (2013) for a detailed discussion. For evidence on such "flat-rate bias," see for instance Train (1991, p. 211) (as well as Train et al., 1987, 1989) in the context of telephone calling plans, and Lambrecht and Skiera (2006) for a general review.

Our result that the agent is averse to spending risk also implies a motive to set money budgets — simply because they can help reduce spending risk. As an explicit example, consider a variant of our model in which the agent makes two sequential decisions in period 1. The first decision is motivated by the setting of Hastings and Shapiro’s (2013) study of gasoline purchases. The decisionmaker needs to purchase one unit of gasoline, but she can choose between high-grade and low-grade gasoline. High-grade gasoline always carries a price premium of Δp , and the agent perceives that it has a value premium of Δv . The price of low-grade gasoline can be either p_L or $p_H = p_L + \Delta p$. The second decision is the transfer decision as above. Our key assumption is that when the agent chooses the transfer, she does not remember how much she spent on gasoline: it is just lumped in with the other shocks to x_1 . To remember her spending — or, equivalently, her checking-account balance — more precisely, she has to pay the attention cost as above. We assume that the agent carries out the gasoline-purchase strategy she chooses at the beginning.²⁰

Proposition 8 implies that if $\Delta v \approx \Delta p$ (i.e., the two grades are quite good substitutes), then the agent’s optimal strategy involves switching to low-grade gasoline when gasoline prices rise. This strategy, which is in line with the empirical findings of Hastings and Shapiro (2013), eliminates spending risk in gasoline and therefore lowers the risk in the checking-account balance.²¹ More generally, our framework suggests that to reduce spending risk without having to keep track of or think about too much information at each decision point, the agent wants to create budgets whose implementation relies on a limited amount of non-salient information.²²

²⁰ This assumption implicitly imposes two conditions. First, the agent must be able to remember her strategy — even though she cannot remember her gasoline spending when she makes her transfer decision. It is presumably easier to remember the general strategy of what to do for low and high gasoline prices than to keep track of idiosyncratic spending, especially if there are many other purchases to keep track of. Second, the agent does not reconsider her spending strategy. This seems reasonable for an agent with costly attention, as rethinking optimal behavior might be costly. But we discuss the possibility that she reevaluates her strategy in Footnote 22 below.

²¹ At the same time, Hirshman et al. (2018) document that for many products, the opposite of Hastings and Shapiro’s finding holds: as prices rise, individuals report more willingness to upgrade to a premium version of the product. This is consistent with diminishing sensitivity or relative thinking in prices. Our model cannot account for the conflicting findings in a precise manner, and more research on when budgeting versus relative thinking dominates seems necessary. But it is worth noting that unlike gas, the products Hirshman et al. consider (e.g., pens or milk) are likely to be in the same budget category with many other products, so one would not expect budgeting at the individual-product level.

²² If the agent does reevaluate her strategy, then an interesting time-inconsistency issue arises: is it credible to stick with her purchase strategy of switching down if and only if the prices are high? Suppose $\Delta v > \Delta p$, so the agent is tempted to keep buying the expensive gasoline when the price is high. If she does, she knows she might not pay attention later, increasing the probability that she overdraws the account. Given such a concern, for Δv sufficiently close to Δp it is credible to stick with her plan. In contrast, the strategy is not credible if $\Delta v < \Delta p$: then, she wants to buy cheap even if prices are low, and this does not increase the risk of overdrawing the account. So the strategy is credible only if the agent prefers high-grade gasoline.

The primary alternative explanation for many of the above basic patterns is loss aversion. As explained for instance by Kőszegi and Rabin (2007), loss aversion implies aversion to small-scale risk, and because for larger risk a person’s consumption utility comes to dominate loss aversion, aversion to larger-scale risk is lower. We view our model as complementing, not replacing loss aversion as an explanation for risk attitudes. In addition to strengthening the patterns implied by loss aversion, our framework makes two distinct predictions. First, our model says that a person is more averse to expenditure risk than investment risk; i.e., not only the scale, but also the location of a risk affects attitudes toward it. We are not aware of a precise test of this prediction, but find it plausible given existing evidence. For instance, Sydnor (2010) documents that homeowners exhibit extremely large risk aversion when choosing between annual deductibles of \$1,000 and \$500 for their home insurance. Yet the same individuals might exhibit significantly lower risk aversion when investing in individual stocks, which carries similar-scale risk. Second, while in models of loss aversion budgets are often exogenously given, our framework explains how the agent can separate fungible funds into different budgets, and makes predictions on the budgets she forms. Of course, these differences are in addition to our predictions regarding the agent’s MPC’s, which to our knowledge existing models of loss aversion also do not make.

6 Related Literature

In this section, we discuss theoretical work most closely related to our paper. While we point out other differences below, previous work provides no theory of the main phenomena — how the agent creates separate accounts out of fungible finances, and how this affects her responses to shocks — that we consider in this paper.

The main existing explanation for mental accounts is self-control problems — attempting to use budgets or accounts to mitigate overconsumption in the future. Our theory provides a different, complementary, reason for mental accounts, with a number of distinct features. Most importantly, in theories on mental accounting and self-control, money is exogenously assumed to be non-fungible in the sense that spending from different accounts is subject to different constraints or preferences. In our model, mental accounts emerge despite money being fully fungible. Due to the different foundation, research on mental accounting and self-control also does not generate many of our other predictions, such as the connection we find between mental budgeting and naive diversification.

In a classic paper, Shefrin and Thaler (1988) develop a life-cycle consumption-savings model in which the individual’s “planner” self would like to control the “doer” self’s tendency to consume too much. Shefrin and Thaler assume that the individual can separate money into different mental accounts, current spendable income, current assets, and future income. They exogenously assume that the MPC out of these accounts is different.

In the context of goal setting under self-control problems, Koch and Nafziger (2016) assume that an individual can decide between broad and narrow goals, and that falling short of one’s chosen goals creates sensations of loss. The motive to avoid such losses creates an incentive that mitigates self-control problems. A broad goal diversifies the risk of failure due to shocks, but it also lowers incentives because underperformance in one task can be offset by good performance in another task. As a result, narrow bracketing can be optimal. Similarly to our Corollary 4 (but for a completely different reason), the agent is more likely to set narrow goals when uncertainty is low. Hsiaw (2018) qualifies this insight for multi-stage projects when uncertainty is resolved over time, showing that an increase in late uncertainty makes incremental goals more appealing.

Galperti (forthcoming) compares good-specific and total-expenditure budgets for a person who is subject to self-control problems as well as both intratemporal and intertemporal taste shocks. Good-specific budgets can be useful for an agent with a mild self-control problem, as they help curb overconsumption of all goods. But for an agent with severe self-control problems, effective good-specific budgets would distort intratemporal consumption too much, so a total-expenditure budget is superior.

Prelec and Loewenstein (1998) propose a model of mental accounting in which paying for goods and experiences is painful but buffered by thoughts of future consumption, and the pleasure of consumption is lowered by thoughts of future payments. This model predicts a strong aversion to paying for consumption with debt, and has a number of other implications for how an individual might want to time payments relative to consumption. These results capture a completely different aspect of mental accounting than do we.

Gorman (1959) identifies circumstances under which it is optimal for a standard utility maximizer to make consumption decisions using a two-step procedure similar to that in Section 4, whereby she first allocates fixed budgets to different consumption categories, and then optimizes within each category given the allocated budget. Unlike in our model, the budgeting in the first

stage requires the agent to know with certainty all the relevant price indices for the categories, and there is no taste uncertainty. Even so, the conditions under which two-stage budgeting is optimal are extremely strict.²³

Our result that the agent may completely ignore some aspects of her decision environment is similar in spirit to the sparsity-based model of bounded rationality by Gabaix (2011), where the agent also ignores variables that are not important. In Gabaix’s setting, the variables that the agent may choose to look at are exogenously given, whereas in ours the agent can choose any combination of variables. We also apply the model to different questions than does Gabaix.

As we have mentioned in Section 5.2, models of loss aversion make some of the same predictions regarding risk attitudes as our model, although there are important differences as well. Loss aversion also has implications for consumption-savings behavior (e.g., Pagel, 2017) that are orthogonal to our results.

7 Conclusion

While our models explain a number of findings, there are phenomena that are usually interpreted in terms of mental accounting that we have not covered. The most important of these is the effect of special-purpose transfers — transfers that can only be used on a subset of products — on consumption. A basic implication of the rational consumer model with full information is that if such a transfer is inframarginal — i.e., if the consumer would have spent more than the transfer on the products in question — then it is equivalent to cash. Yet experimental work by Abeler and Marklein (2016) and empirical work by Hastings and Shapiro (2017) document that inframarginal transfers have large effects on the consumption of targeted products. Even when a transfer is not inframarginal, it can have a surprisingly large effect: for instance, incentives for health-improving behaviors that are minute relative to the health benefits can significantly influence behavior (Volpp et al., 2008; Dupas, 2014). It being about reactions to ex-ante known transfers rather than ex-post shocks, this phenomenon does not fit in the methodology of the current paper. There is, however, a plausible attention-based account for the phenomenon. Namely, there are many things that a person could consider doing, but that she considers not worthwhile to think about due to costly

²³ Interestingly, Gorman, as well as Strotz (1957) preceding him, take it as self-evident that individuals engage in the above type of mental accounting.

attention, and that she therefore does not do. Receiving a transfer or subsidy can induce the person to think about the potential benefits, increasing the effect of the transfer. In ongoing work, we formalize this mechanism, and also consider what it implies for the optimal design of transfers.

Of course, we do not believe that mental accounting is solely about costly attention. As Shefrin and Thaler (1988) and others have argued, a likely motive for creating mental accounts is self-control problems. It would be interesting to combine the attention-based account of mental accounting with the self-control-based account to identify possible interactions. For example, a person may use the costly nature of her attention to improve her self-control by creating plans that she is unwilling to reconsider later.

Finally, even to the extent that costly attention is a central component of mental accounting, it is not necessarily of the type we have modeled, where any piece of information has the same cost, and attention allocation is fully rational. Without assuming unrealistically large attention costs, for instance, our model is unlikely to explain why most individuals spend less than an hour on thinking about their retirement-account contribution rates (Benartzi and Thaler, 2007). Nevertheless, the mechanisms we have identified can be just as important with non-rational attention. Recognizing that the checking account is more important to know about, for instance, a consumer may (irrationally) decide not to look at the investment account at all, increasing the difference in MPC's we have found in Section 5.

References

- Abeler, Johannes and Felix Marklein**, “Fungibility, Labels, and Consumption,” *Journal of the European Economic Association*, 2016, 15 (1), 99–127.
- Ameriks, John, Andrew Caplin, and John Leahy**, “Wealth Accumulation and the Propensity to Plan,” *Quarterly Journal of Economics*, 2003, 118 (3), 1007–1047.
- Antonides, Gerrit, I. Manon de Groot, and W. Fred van Raaij**, “Mental Budgeting and the Management of Household Finance,” *Journal of Economic Psychology*, 2011, 32 (4), 546–555.
- Baker, Malcolm, Stefan Nagel, and Jeffrey Wurgler**, “The Effect of Dividends on Consumption,” *Brookings Papers on Economic Activity*, 2007, 38 (1), 231–292.
- Barsky, Robert, F. Thomas Juster, Miles Kimball, and Matthew Shapiro**, “Preference Parameters and Behavioral Heterogeneity: An Experimental Approach in the Health and Retirement Study,” *Quarterly Journal of Economics*, 1997, 112 (2), 537–579.
- Benartzi, Shlomo and Richard H. Thaler**, “Naive Diversification Strategies in Defined Contribution Saving Plans,” *American Economic Review*, 2001, 91 (1), 79–98.

- and —, “Heuristics and Biases in Retirement Savings Behavior,” *Journal of Economic Perspectives*, 2007, 21 (3), 81–104.
- Bhargava, Saurabh, George Loewenstein, and Justin Sydnor**, “Choose to Lose: Health Plan Choices from a Menu with Dominated Option,” *Quarterly Journal of Economics*, 2017, 132 (3), 1319–1372.
- Caplin, Andrew and Mark Dean**, “Revealed preference, rational inattention, and costly information acquisition,” *The American Economic Review*, 2015, 105 (7), 2183–2203.
- Cover, Thomas M. and Joy A. Thomas**, *Elements of Information Theory*, Hoboken, NJ: Wiley, 2006.
- Dupas, Pascaline**, “Getting Essential Health Products to their End Users: Subsidize, but How Much?,” *Science*, 2014, 345 (6202), 1279–1281.
- Feldman, Naomi E.**, “Mental Accounting Effects of Income Tax Shifting,” *The Review of Economics and Statistics*, 2010, 92 (1), 70–86.
- Gabaix, Xavier**, “A Sparsity-Based Model of Bounded Rationality,” 2011. Working paper, NYU Stern. URL: <http://pages.stern.nyu.edu/~xgabaix/papers/sparsebr.pdf>.
- Galperti, Simone**, “A Theory of Personal Budgeting,” *Theoretical Economics*, forthcoming.
- Gorman, W. M.**, “Separable Utility and Aggregation,” *Econometrica*, 1959, 27 (3), 469–481.
- Hastings, Justine S. and Jesse M. Shapiro**, “Fungibility and Consumer Choice: Evidence from Commodity Price Shocks,” *Quarterly Journal of Economics*, 2013, 128 (4), 1449–1498.
- and —, “How Are SNAP Benefits Spent? Evidence from a Retail Panel,” 2017. Working Paper.
- Hatsopoulos, George N., Paul R. Krugman, and James M. Poterba**, “Overconsumption: The Challenge to U.S. Economic Policy,” 1989. American Business Conference.
- Herweg, Fabian and Konrad Mierendorff**, “Uncertain Demand, Consumer Loss Aversion, and Flat-Rate Tariffs,” *Journal of the European Economic Association*, 2013, 11 (2), 399–432.
- Hirshman, Samuel, Devin Pope, and Jihong Song**, “Mental Budgeting versus Relative Thinking,” *American Economic Review Papers and Proceedings*, 2018, 108, 148–152.
- Hsiaw, Alice**, “Goal Bracketing and Self-Control,” *Games and Economic Behavior*, 2018, 111, 100–121.
- Kahneman, Daniel and Amos Tversky**, “Choices, Values, and Frames,” *American Psychologist*, 1984, 39 (4), 341–350.
- Kamenica, Emir and Matthew Gentzkow**, “Bayesian persuasion,” *The American Economic Review*, 2011, 101 (6), 2590–2615.
- Kőszegi, Botond and Matthew Rabin**, “Reference-Dependent Risk Attitudes,” *American Economic Review*, 2007, 97 (4), 1047–1073.

- Koch, Alexander K. and Julia Nafziger**, “Goals and Bracketing under Mental Accounting,” *Journal of Economic Theory*, 2016, 162 (C), 305–351.
- Kueng, Lorenz**, “Excess Sensitivity of High-Income Consumers,” *Quarterly Journal of Economics*, forthcoming.
- Lambrecht, Anja and Bernd Skiera**, “Paying Too Much and Being Happy About It: Existence, Causes, and Consequences of Tariff-Choice Biases,” *Journal of Marketing Research*, 2006, 43 (2), 212–223.
- Lave, Jean**, *Cognition in Practice: Mind, Mathematics, and Culture in Everyday Life*, Cambridge University Press, 1995.
- Luo, Yulei and Eric R Young**, “Signal extraction and rational inattention,” *Economic Inquiry*, 2014, 52 (2), 811–829.
- Mackowiak, Bartosz; and Mirko Wiederholt**, “Optimal Sticky Prices under Rational Inattention,” *The American Economic Review*, June 2009, 99.
- Matějka, Filip**, “Rationally Inattentive Seller: Sales and Discrete Pricing,” *Review of Economic Studies*, 2016, 83 (3), 1125–1155.
- Matějka, Filip and Alisdair McKay**, “Rational inattention to discrete choices: A new foundation for the multinomial logit model,” *American Economic Review*, 2015, 105 (1), 272–98.
- Pagel, Michaela**, “Expectations-Based Reference-Dependent Life-Cycle Consumption,” *Review of Economic Studies*, 2017, 84 (2), 885–934.
- Prelec, Drazen and George Loewenstein**, “The Red and the Black: Mental Accounting of Savings and Debt,” *Marketing Science*, 1998, 17 (1), 4–28.
- Rainwater, Lee, Richard P. Coleman, Gerald Handel, and W. Lloyd Warner**, *Workingman’s Wife: Her Personality, World and Life Style*, Oceana Publications, 1959.
- Sahn, Claudia R., Matthew D. Shapiro, and Joel Slemrod**, “Check in the Mail or More in the Paycheck: Does the Effectiveness of Fiscal Stimulus Depend on How It Is Delivered?,” *American Economic Journal: Economic Policy*, 2012, 4 (3), 216–250.
- Shafir, Eldar, Peter Diamond, and Amos Tversky**, “Money Illusion,” *Quarterly Journal of Economics*, 1997, 112 (2), 341–374.
- Shefrin, Hersh M. and Richard H. Thaler**, “The Behavioral Life-Cycle Hypothesis,” *Economic Inquiry*, 1988, 26 (4), 609–643.
- Simonson, Itamar**, “The Effect of Purchase Quantity and Timing on Variety-Seeking Behavior,” *Journal of Marketing Research*, 1990, 27 (2), 150–162.
- Sims, Christopher A.**, “Implications of Rational Inattention,” *Journal of Monetary Economics*, 2003, 50 (3), 665–690.
- Souleles, Nicholas S., Jonathan A. Parker, and David S. Johnson**, “Household Expenditure and the Income Tax Rebates of 2001,” *American Economic Review*, 2006, 96 (5), 1589–1610.

- Strotz, Robert H.**, “The Empirical Implications of a Utility Tree,” *Econometrica*, 1957, 25 (2), 269–280.
- Sydnor, Justin**, “(Over)insuring Modest Risks,” *American Economic Journal: Applied Economics*, 2010, 2 (4), 177–99.
- Telatar, Emre**, “Capacity of Multi-antenna Gaussian Channels,” *European transactions on telecommunications*, 1999, 10 (6), 585–595.
- Thaler, Richard H.**, “Mental Accounting Matters,” *Journal of Behavioral Decision Making*, 1999, 12 (3), 183–206.
- Train, Kenneth E.**, *Optimal Regulation*, Cambridge, MA: MIT Press, 1991.
- , **Daniel L. McFadden, and Moshe Ben-Akiva**, “The Demand for Local Telephone Service: A Fully Discrete Model of Residential Calling Patterns and Service Choices,” *RAND Journal of Economics*, 1987, 18 (1), 109–123.
- , **Moshe Ben-Akiva, and Terry Atherton**, “Consumption Patterns and Self-selecting Tariffs,” *The Review of Economics and Statistics*, 1989, 71 (1), 62–73.
- Veldkamp, Laura L.**, “Information markets and the comovement of asset prices,” *The Review of Economic Studies*, 2006, 73 (3), 823–845.
- Volpp, Kevin G., Leslie K. John, Andrea B. Troxel, Laurie Norton, Jennifer Fassbender, and George Loewenstein**, “Financial Incentive-Based Approaches for Weight Loss,” *Journal of the American Medical Association*, 2008, 300 (22), 2631–2637.
- Woodford, Michael**, “Information-constrained state-dependent pricing,” *Journal of Monetary Economics*, 2009, 56 (Supplement 1).

A LQ multivariate setup

Proof of Proposition 1. The utility quadratic utility function can be rewritten as

$$U(\mathbf{y}, \mathbf{x}) = -\left(\mathbf{y} - \frac{C^{-1}B'}{2}\mathbf{x}\right)' C \left(\mathbf{y} - \frac{C^{-1}B'}{2}\mathbf{x}\right) + \frac{\mathbf{x}'BC^{-1}B'\mathbf{x}}{4}. \quad (14)$$

If the posterior mean is $\tilde{\mathbf{x}}$, then the agent chooses an action (maximizing expected utility):

$$\mathbf{y} = \frac{C^{-1}B'}{2}\tilde{\mathbf{x}}. \quad (15)$$

This is because certainty equivalence applies in a quadratic setup. Plugging (15) into (14), the realized utility \tilde{U} for a state \mathbf{x} , but a posterior mean $\tilde{\mathbf{x}}$ is:

$$\tilde{U}(\tilde{\mathbf{x}}, \mathbf{x}) = -(\tilde{\mathbf{x}} - \mathbf{x})' \Omega (\tilde{\mathbf{x}} - \mathbf{x}) + \mathbf{x}' \Omega \mathbf{x}, \quad (16)$$

where $\Omega = BC^{-1}B'/4$. The first term is the loss from imperfect posterior beliefs, $(\tilde{\mathbf{x}} - \mathbf{x})$ is the misperception. Given the variance-covariance matrix Σ for the distribution of $(\tilde{\mathbf{x}} - \mathbf{x})$, the expectation of the first term equals the trace of $\Omega\Sigma$. Since the second term in (16) depends on the realized state \mathbf{x} only, i.e., it is independent of the agent's strategy, then the original (3)-(4) takes the form:

$$\max_{\psi \succeq \Sigma} -Tr(\Omega\Sigma) + \frac{\lambda}{2} \log |\Sigma|. \quad (17)$$

The second term in (17) is the cost of information, it is a log of the determinant of Σ .²⁴ The larger the posterior uncertainty is, the lower the cost. The cost term here includes entropy of the posterior only, since entropy of a fixed prior amounts to an additive constant only. The condition $\psi \succeq \Sigma$ requires that $(\psi - \Sigma)$ is positive semi-definite, which means that acquisition of Gaussian signals cannot make beliefs less precise, i.e., signals must have non-negative precision.

To explore what signals the agent collects, let us decompose the loss matrix Ω , which is symmetric and thus has an orthonormal basis of eigenvectors. Let $\Omega = U\Lambda U'$, where U is a unitary matrix (the columns of which are eigenvectors of Ω), and Λ is a diagonal matrix with its elements Λ_{ii} equal to the eigenvalues Λ_i of Ω .

$$\begin{aligned} -Tr(\Omega\Sigma) + \frac{\lambda}{2} \log |\Sigma| &= -Tr(U\Lambda U'\Sigma) + \frac{\lambda}{2} \log |\Sigma| = \\ &= -Tr(\Lambda U'\Sigma U) + \frac{\lambda}{2} \log |U'\Sigma U| = \\ &= -Tr(\Lambda S) + \frac{\lambda}{2} \log |S|, \end{aligned} \quad (18)$$

where $S = U'\Sigma U$ is the posterior variance-covariance matrix in the basis of eigenvectors of Ω . The condition $(\psi \succeq \Sigma)$ takes the form of $(U'\psi U \succeq S)$; note that $\psi = \sigma_0^2 I$.

Now we show by contradiction that S is diagonal. Let the optimal S were not diagonal, and let S^D be the matrix constructed from its diagonal, i.e., $S_{ii}^D = S_{ii}$ for all i and $S_{ij}^D = 0$ for all $i \neq j$.

²⁴Entropy of a multivariate $N(\mu, \Sigma)$ of dimension n is $\frac{n}{2}(\log(2\pi) + 1) + \frac{1}{2} \log |\Sigma|$.

First, since $\sigma_0^2 I - S$ is positive semi-definite, then $\sigma_0^2 I - S^D$ is also positive semi-definite. This is because for a diagonal S^D it suffices to check that $S_{ii}^D \leq \sigma_0^2$, which is implied by the fact that $\sigma_0^2 I - S$ is also positive semi-definite. Second, Hadamard's inequality implies:

$$\frac{\lambda}{2} \log |S| \leq \frac{\lambda}{2} \sum_i \log S_{ii} = \frac{\lambda}{2} \log |S^D|, \quad (19)$$

where the equality holds if and only if S is diagonal. Third, $Tr(\Lambda S) = Tr(\Lambda S^D)$, since Λ is diagonal. Therefore, putting this together implies that S not cannot be the optimum, since S^D delivers a higher objective due to the lower information cost, (19), and is feasible.

Therefore, S is diagonal. Using (18), the original problem takes the form:

$$\max_{S_{ii} \leq \sigma_0^2} \left(- \sum_{i=1}^N S_{ii} \Lambda_i + \frac{\lambda \log(S_{ii})}{2} \right). \quad (20)$$

The first order condition with respect to S_{ii} implies:

$$-\Lambda_i + \frac{\lambda}{2S_{ii}} = 0,$$

and the solution is

$$S_{ii} = \min\left(\sigma_0^2, \frac{\lambda}{2\Lambda_i}\right).$$

□.

Proof of Proposition 2. Part 1: Proposition 1 implies that the space of posterior means $\tilde{\mathbf{x}}$ is spanned by all eigenvectors v^i for which $\lambda/(2\Lambda_i) < \sigma_0^2$, the statement is then a trivial implication.

Part 2: Let $\xi_i = 1 - \frac{S_{ii}}{\sigma_0^2}$ be the relative reduction of uncertainty about the component $v^i \cdot \mathbf{x}$. ξ is also the linear weight on a signal (as opposed to on the prior) in Bayesian updating with Gaussian signals. This means that in 1D Bayesian updating, if the random variable $v^i \cdot \mathbf{x}$ moves by Δx , then the posterior mean about this variable moves in expectation by $\xi_i \Delta x$.

Since the agent chooses independent signals on $v^i \cdot \mathbf{x}$, Bayesian updating does in fact take the 1D form. Responsiveness then is:

$$\varepsilon_i^\lambda = \frac{|H \xi_i v^i|}{|v_i|} = \frac{|\xi_i H v^i|}{|v_i|} = \xi_i \varepsilon_i^0.$$

This equation together with Proposition 1 implies the expression (9).

Part 3: Differentiating $\varepsilon_i^\lambda/\varepsilon_j^\lambda$ with respect to λ then implies the statement. \square .

B Consumption and spending budgets

Let Θ have the structure described in the main text, i.e., given by symmetrically nested categories, and let $R^{k,l}$ denote a category number k on level l , size of which is r_l .

Lemma 1. Θ has a base of eigenvectors $\{v^{k,l,r'}\}_{l \in \{2..L\}, k \in \{1..N/r_l\}, r' \in \{1..(r_l-1)\}}$, and $(1, \dots, 1)$. $v^{k,l,r'}$ is associated with a category k of a level l , and it has the following properties

$$v_m^{k,l,r'} = 0 \quad \forall m \notin R^{k,l} \quad (21)$$

$$\sum_{m \in R^{k,l}} v_m^{k,l,r'} = 0 \quad (22)$$

$$v_m^{k,l,r'} = v_n^{k,l,r'} \quad \forall m, n; \exists k' : m, n \in R^{k',l-1}. \quad (23)$$

Moreover, eigenvalues μ^l of $v^{k,l,r'}$ are given by:

$$\mu^l = \sum_{n \in R_m^{l-1}} (\Theta_{m,n} - \gamma^l) = \mu^{l-1} + (\gamma^{l-1} - \gamma^l)r_{l-1},$$

and $\mu^2 = \gamma^1 - \gamma^2$.

Proof of Lemma 1. Let us fix m , and apply Θ to an eigenvector associated with $R^{k,l}$; we drop

the index m of the vector. Let R_m^{l-1} be the category on level $(l-1)$ that the good m belongs to.

$$\begin{aligned}
\sum_n \Theta_{m,n} v_n^{k,l} &= \left(\sum_{n \notin R^{k,l}} \Theta_{m,n} v_n^{k,l} \right) + \left(\sum_{n \in R^{k,l}/R_m^{l-1}} \Theta_{m,n} v_n^{k,l} \right) + \left(\sum_{n \in R_m^{l-1}} \Theta_{m,n} v_n^{k,l} \right) = \\
&= 0 + \gamma^l \left(\sum_{n \in R^{k,l}/R_m^{l-1}} v_n^{k,l} \right) + \left(\sum_{n \in R_m^{l-1}} \Theta_{m,n} v_n^{k,l} \right) = \\
&= \gamma^l \left(\sum_{n \in R^{k,l}} v_n^{k,l} - \sum_{n \in R_m^{l-1}} v_n^{k,l} \right) + \left(\sum_{n \in R_m^{l-1}} \Theta_{m,n} v_n^{k,l} \right) = \\
&= \gamma^l \left(0 - \sum_{n \in R_m^{l-1}} v_m^{k,l} \right) + \left(\sum_{n \in R_m^{l-1}} \Theta_{m,n} v_m^{k,l} \right) = \\
&= v_m^{k,l} \sum_{n \in R_m^{l-1}} (\Theta_{m,n} - \gamma^l)
\end{aligned}$$

The first equality is a simple decomposition into terms with elements within different categories. In the second, we used (21). Third is based on a decomposition of elements of $R^{k,l}$ into a sub-category with m and the other elements. The fourth equality uses (22) for the first term, and (23) is applied for the other two terms to substitute elements $v_j^{k,l}$ indexed by j by a constant $v_i^{k,l}$, since $v_j^{k,l}$ is constant in R_i^{l-1} .

Eigenvalue μ^l is therefore $\sum_{n \in R_m^{l-1}} (\Theta_{m,n} - \gamma^l)$. For $l=2$ the only sub-category including m is m itself, $\mu^2 = \gamma^1 - \gamma^2$. And for $l > 2$:

$$\begin{aligned}
\mu^l &= \mu^{l-1} + \gamma^{l-1} r_{l-2} - \gamma^l r_{l-1} - (r_{l-2} - r_{l-1}) \gamma^{l-1} = \\
&= \mu^{l-1} + (\gamma^{l-1} - \gamma^l) r_{l-1}.
\end{aligned} \tag{24}$$

Therefore, each $v^{k,l,r'}$ is an eigenvector, and they form a basis. This is because they are all mutually orthogonal. Vectors associated with distinct categories due to (21), and vectors associated with a nested categories due to (23) and (22). And for vectors of the same category, the dimensionality is due to (23) equal the number of sub-categories minus one lower dimensionality due to (22), $r_l/r_{l-1} - 1$. The total number of vectors associated with level $l > 1$ is $N/r_{l-1} - N/r_l$, and the total number of these orthogonal eigenvectors on all levels is $N - 1$, which together with the eigenvector $(1, \dots, 1)$ delivers N orthogonal eigenvectors, and thus a basis. \square

Proof of Proposition 3. This proposition follows from the propositions 1 and 2, and Lemma 1.

First, since for (10), $\Omega = \Theta^{-1}/4$, then Ω has the same eigenvectors as Θ , and the corresponding eigenvalues Λ_i are proportional to the inverse of the eigenvalue of Θ that is associated with the same vector. Specifically, for v^i associated with a level l , the eigenvalue is:

$$\Lambda_i = \frac{1}{4}(\mu^l)^{-1} \quad (25)$$

The expression below is derived from (25) and (9) with $\sigma_0^2 = \lambda/2\Lambda_i$ at which ξ_i hits zero.

$$\lambda_{l-1} = 2\sigma_0^2/\Lambda_i = \frac{\sigma_0^2}{2 \sum_{n \in R_m^{l-1}} (\Theta_{m,n} - \gamma^l)}, \quad (26)$$

where m is a good in $R^{k,l-1}$. We denote this threshold cost for attention to vectors associated with level l by λ_{l-1} rather than by λ_l , because for λ lower than this quantity the total consumption in each category on levels $(l-1)$ and lower is constant.

The eigenvectors satisfy (22), which means that the agent pays attention to vectors that keep "budgets" of the random elements of \mathbf{x} across the corresponding category fixed, but not across categories on lower levels. Moreover, since the action matrix H equals $\Theta^{-1}/2$, the eigenvectors of Ω are its eigenvectors as well. Thus, the fixed budgets of \mathbf{x} translate to the fixed budgets of actions \mathbf{y} across the same categories.

Notice that for nested substitutes, since $\gamma^l > \gamma^{l-1}$, eigenvalues μ^l of Θ are increasing in l , and thus eigenvalues of Ω are decreasing in l . □

Proof of Proposition 4. Equation (8) implies:

$$E[Y^{k,l}|\mathbf{x}] = \sum_{j \in R^{k,l}} \sum_i \xi_i 2\Lambda_i (v^i \cdot \mathbf{x}) v_j^i, \quad (27)$$

because $H = 2\Omega$ has the same eigenvectors as Ω , with eigenvalues equal to $2\Lambda_i$. Now, we use Lemma 1. First, (22) implies that contributions to the RHS of (27) of all vectors v^i associated with levels l and lower are zero, because $\sum_{j \in R^{k,l}} v_j^i = 0$. Next, the contribution of v^i on levels $l+1$ and higher is according to (23) proportional to $(v^i \cdot \mathbf{x}) v_j^i = \left(\text{const} \cdot X^{k,l} + \sum_{s \notin R^{k,l}} v_s^i \mathbf{x}_s \right) v_j^i$. Note also that v_j^i is constant for all $j \in R^{k,l}$ for such i associated with levels $l+1$ and higher.

The derivative of $E[Y^{k,l}|\mathbf{x}]$ with respect to $X^{k,l}$ thus is:

$$\epsilon_B^l = \sum_{m=l+1}^L \xi_m \Lambda_m \omega_m, \quad (28)$$

where ω_m 's are constants that depend on the number of categories on the level m level. The first part of the statement in Proposition 4 is now implied by (28) and by $\partial \xi_l / \partial \lambda = -1/\Lambda_l$ for $\lambda < 1/\Lambda_l$, see the expression for ξ in (9). We abuse the notation here by denoting Λ_l eigenvalues Λ_k such that v^k is a vector on the level l (note that all such Λ_k associated with the level l are equal). Let $\hat{L} \geq (l+1)$ be the largest s such that $\lambda < 1/\Lambda_s$.

$$\begin{aligned} \frac{\partial \epsilon_B^{l-1} / \epsilon_B^l}{\partial \lambda} &= \frac{-\epsilon_B^l \sum_{m=l}^{\hat{L}} \omega_m + \epsilon_B^{l-1} \sum_{m=l+1}^{\hat{L}} \omega_m}{(\epsilon_B^l)^2} = \\ &= \frac{\omega_l \left(\xi_l \Lambda_l \sum_{m=l}^{\hat{L}} \omega_m - \sum_{m=l+1}^{\hat{L}} \xi_m \Lambda_m \omega_m \right)}{(\epsilon_B^l)^2} > 0. \end{aligned}$$

The second equality is obtained by using (28), and the last inequality is implied by the fact that Λ_l is decreasing in l , see the proof of Proposition 3, and thus so is $\xi_l \Lambda_l$. This concludes the proof of the second statement.

Finally, we proceed with the statement regarding correlations. The variance-covariance matrix of posterior means (describing correlations of beliefs about x_i and x_j) is $P = (\psi - \Sigma)$. This matrix is diagonal in the basis of eigenvectors v^k , i.e, $P = UQU^{-1}$, where the columns of U are v^i . The diagonal elements of $Q_{kk} \equiv Q_k$ equal $\sigma_0^2 - \sigma_k^2$, which is the reduction of uncertainty about $v^k \cdot \mathbf{x}$. The reduction $Q_k = \max(0, \sigma_0^2 - \lambda/(2\Lambda_k))$ is weakly increasing in Λ_k and weakly decreasing in λ , see Proposition (1).

The resulting variance-covariance matrix of actions is $A = HPH'$, where $P_{ij} = \sum_k Q_k v_i^k v_j^k$, and v^k are eigenvectors of H , too, $2\Lambda_k$ is the eigenvalue. The matrix A thus is:

$$A_{ij} = \sum_k Q_k (2\Lambda_k)^2 v_i^k v_j^k. \quad (29)$$

And finally, for the correlation $\rho_{ij} = A_{ij} / \sqrt{A_{ii}A_{jj}} = A_{ij}/A_{ii}$, where the last equality is implied by the fact that the variances of uncertainty about i and j are the same. Equation (29) implies that

actions y_i and y_j are more positively correlated if more uncertainty is reduced in a direction of v^k , for which the signs of entries v_i^k and v_j^k are the same.

Next, we express derivative of the correlation:

$$\frac{\partial \rho_{ij}}{\partial \lambda} = \frac{\left(-\sum_k 2\Lambda_k v_i^k v_j^k\right) \left(\sum_k Q_k (2\Lambda_k)^2 (v_i^k)^2\right) - \left(\sum_k Q_k (2\Lambda_k)^2 v_i^k v_j^k\right) \left(-\sum_k 2\Lambda_k (v_i^k)^2\right)}{\sum_k Q_k (2\Lambda_k)^2 (v_i^k)^2},$$

where the sums are over all k such that $\lambda < 1/\Lambda_k$. Due to Lemma 1, the eigenvectors can be selected such that $v_i^k v_j^k = -1$ for some vectors v^k that are associated with the smallest level on which goods i and j are in the same category, let the level be l^* and the number of such vectors be ψ_{l^*} . Similarly, $v_i^k v_j^k = 1$ for some v^k that are associated with levels higher than l^* , and let ψ_s be the number of such vectors on the level s . For all other vectors v^k : $v_i^k v_j^k = 0$.

Let $\hat{L} \geq (l+1)$ be the largest s such that $\lambda < 1/\Lambda_s$. The numerator of the RHS of equation above then equals eight times the following quantity:

$$\begin{aligned} & \left(\psi_{l^*} \Lambda_{l^*} - \sum_{s=l^*+1}^{\hat{L}} \psi_s \Lambda_s \right) \left(\sum_{s=l^*}^{\hat{L}} \psi_s Q_s (\Lambda_s)^2 \right) - \\ & - \left(-\psi_{l^*} Q_{l^*} (\Lambda_{l^*})^2 + \sum_{s=l^*+1}^{\hat{L}} \psi_s Q_s (\Lambda_s)^2 \right) \left(-\sum_{s=l^*}^{\hat{L}} \psi_s \Lambda_s \right) = \\ & = \psi_{l^*} \Lambda_{l^*} \left(\sum_{s=l^*+1}^{\hat{L}} \psi_s Q_s (\Lambda_s)^2 \right) - \psi_{l^*} Q_{l^*} (\Lambda_{l^*})^2 \left(\sum_{s=l^*+1}^{\hat{L}} \psi_s \Lambda_s \right) = \\ & = \psi_{l^*} \Lambda_{l^*} \sum_{s=l^*+1}^{\hat{L}} \left(\psi_s Q_s (\Lambda_s)^2 - \psi_s Q_{l^*} \Lambda_s \Lambda_{l^*} \right) < 0. \end{aligned}$$

In the last step we used the fact that both $Q_s \Lambda_s$ are decreasing in the level s , see the proof of Proposition 3. This together with the positivity of the denominator of the RHS of (30) concludes the proof. \square

Proof of Corollary 1. This is an immediate implication of (26). In the denominator of the RHS of (26), γ^k enter in two ways. First, γ^l is subtracted, and also each parameter $\Theta_{m,n}$ equals γ^k for some $k \leq (l-1)$, since the sum is over n such that m and n are included in the same $(l-1)$ -category (or smaller ones).

Therefore, due to (26), if γ^l increases, then λ_{l-1} must increase. On the other hand, if γ^{l-s-1}

for $s \geq 0$ increases, then λ_{l-1} must decrease. The second part of the statement then follows from the substitution: $l' = (l - s - 1)$. \square

Proof of Corollary 2. This is implied by the equation (15) for $C = \Theta$ and $B = I$, and certainty equivalence. \square

Proof of Corollary 3. Using equation (26) we get

$$\lambda_{l'} = \frac{\sigma_0^2}{2 \sum_{n \in R'_m} (\Theta'_{m,n} - \gamma^l)}, \quad (30)$$

where Θ' corresponds to the new menu of products. In comparison to the expression for λ_{l-1} the difference is that there are more terms summed up in the denominator. All of the terms in the denominator for λ_{l-1} are included in the denominator for $\lambda_{l'}$ but some are added, since there is more copied products now. Finally, since all these terms are positive, then we get the desired statement. \square

Proof of Corollary 4. The equation (9) implies that an increase in σ_0^2 increases ξ_i for all i , which due to equation (28) increases ϵ_B^l (since $\Lambda_m \omega_m$ are positive and independent of it). The first equality in (26) then implies the second part of the statement. \square

Proof of Proposition 5. Analogous to the proof of Proposition 3. Eigenvectors take the same form, but ranking of magnitudes of eigenvalues, given by (26) is the opposite because for complements $\gamma^2 < \dots < \gamma^L < 0$. \square

C Income

Proof of Proposition 6. Equation (8) implies for the two derivatives:

$$\partial E[c_1] / \partial x_1 = \xi_1 v_1^1 (Hv^1)_1 + \xi_2 v_1^2 (Hv^2)_1, \quad (31)$$

$$\partial E[c_1] / \partial (sx_2) = \xi_1 \frac{v_2^1}{s} (Hv^1)_1 + \xi_2 \frac{v_2^2}{s} (Hv^2)_1. \quad (32)$$

We also know that under perfect information the derivatives are equal to a positive constant \hat{c} , since then both sources of income have the same effect. Let $\alpha_1 = v_1^1 (Hv^1)_1$ and $\alpha_2 = \frac{v_2^1}{s} (Hv^1)_1$,

and thus $v_1^2(Hv^2)_1 = \hat{c} - \alpha_1$ and $\frac{v_2^2}{s}(Hv^2)_1 = \hat{c} - \alpha_2$. The desired inequality, $\partial E[c_1]/\partial x_1 > \partial E[c_1]/\partial(sx_2)$, takes the form:

$$\xi_1\alpha_1 + \xi_2(\hat{c} - \alpha_1) > \xi_1\alpha_2 + \xi_2(\hat{c} - \alpha_2). \quad (33)$$

Let $a = 1/T$. The agent's utility function is $-c_1^2 - a(sx_2 + x_1 - c_1)^2 - b(x_1 - c_1 + t)^2$. Ignoring terms that only contain x_1 and x_2 and that the agent can therefore not influence, this utility function expands to

$$\begin{aligned} & -c_1^2 - ac_1^2 + 2asx_2c_1 + 2ax_1c_1 - bc_1^2 - bt^2 + 2bc_1t + 2bx_1c_1 - 2bx_1t \\ & = -(1+a+b)c_1^2 - bt^2 + 2bc_1t + (2a+2b)x_1c_1 - 2bx_1t + 2asx_2c_1 \\ & = - \begin{bmatrix} y_1 & t \end{bmatrix} C \begin{bmatrix} y_1 \\ t \end{bmatrix} + \begin{bmatrix} x_1 & x_2 \end{bmatrix} B \begin{bmatrix} y_1 \\ t \end{bmatrix}, \end{aligned}$$

where

$$C = \begin{bmatrix} 1+a+b & -b \\ -b & b \end{bmatrix}, \quad B = \begin{bmatrix} 2a+2b & -2b \\ 2as & 0 \end{bmatrix}.$$

Therefore, in the notation of Section 3,

$$\Omega = \begin{bmatrix} a^2 + ab + b & a^2s \\ a^2s & a^2s^2 \end{bmatrix}.$$

Without loss of generality, let v^1 be the principal eigenvector of Ω , and thus $\xi_1 > \xi_2$. Some algebra implies:

$$v^1 = \left\{ \frac{a^2 + b + ab - a^2s^2 + \sqrt{-4a^2(1+a)bs^2 + (b + ab + a^2(1+s^2))^2}}{2a^2s}, 1 \right\}.$$

Next, we show that $s \cdot v_1^1 > 1$ for all $s > 0$ and thus $v_1^1 > \frac{v_2^1}{s}$, since $v_2^1 = 1$. After multiplying both sides of the desired inequality $s \cdot v_1^1 > 1$ by $2a^2$, and subtracting $2a^2$, we get

$$-a^2 + b + ab - a^2s^2 + \sqrt{-4a^2(1+a)bs^2 + (b + ab + a^2(1+s^2))^2} > 0. \quad (34)$$

This always holds because

$$-(-a^2 + b + ab - a^2s^2)^2 + \left(-4a^2(1+a)bs^2 + (b + ab + a^2(1+s^2))^2\right) = 4a^2(a+1)b > 0.$$

The square root in (34), which is non-negative, is thus always larger than the absolute value of the remaining terms. Note that this also implies that $v_1^1 > 0$. Now, we can reverse the sequence of steps of equivalence and infer that $v_1^1 > \frac{v_2^1}{s}$ for all $s > 0$.

Similarly, in the notation of Section 3,

$$\left(Hv^1\right)_1 = \frac{a^2 + b + ab + a^2s^2 + \sqrt{-4a^2(1+a)bs^2 + (b + ab + a^2(1+s^2))^2}}{2a(1+a)s} > 0,$$

which implies $\alpha_1 = v_1^1 \left(Hv^1\right)_1 > 0$, and thus also $\alpha_1 > \alpha_2$, because $\alpha_1 = \alpha_2 \frac{v_1^1}{v_2^1}$.

Multiplying the inequality ($\alpha_1 > \alpha_2$) by $(\xi_1 - \xi_2)$, which we know is positive, we get an equivalent inequality:

$$\xi_1\alpha_1 - \xi_2\alpha_1 > \xi_1\alpha_2 - \xi_2\alpha_2.$$

Next, we subtract \hat{c} from both sides to get:

$$\xi_1\alpha_1 + (\hat{c} - \xi_2\alpha_1) > \xi_1\alpha_2 + (\hat{c} - \xi_2\alpha_2).$$

which is (33). □

Proof of Proposition 7. In this case, the utility function is

$$-c_1^2 - \frac{(sx_2 + x_1 - c_1)^2}{T}.$$

Plugging in $c_1 = x_1 + t$ gives

$$-(x_1 + t)^2 - \frac{(sx_2 - t)^2}{T},$$

which implies that

$$B = \begin{bmatrix} -2 \\ 2s/T \end{bmatrix}; \quad C = [(T+1)/T].$$

Given that Ω is a constant multiple of BB' and B is a column matrix, Ω has a single eigenvector

with positive eigenvalue, which is equal to B . Hence, the agent acquires information about only one variable, $(sx_2/T - x_1)$. Given that $T > 1$, this immediately implies that the agent acquires more information about x_1 than about sx_2 .

Denoting posterior mean by E_{post} , the optimal transfer is given by

$$t = \frac{T}{T+1} \cdot E_{post}[sx_2/T - x_1]$$

and

$$\begin{aligned} c_1 &= x_1 + \frac{T}{T+1} \cdot E_{post}[sx_2/T - x_1] = x_1 - E_{post}[x_1] + E_{post}[x_1] + \frac{T}{T+1} \cdot E_{post}[sx_2/T - x_1] \\ &= x_1 - E_{post}[x_1] + \frac{E_{post}[sx_2 + x_1]}{T+1}. \end{aligned}$$

Clearly, $\partial E[E_{post}[sx_2 + x_1]]/\partial x_1 = \partial E[E_{post}[sx_2 + x_1]]/\partial (sx_2)$, so the second term is on average equally responsive to shocks in the two balances. Furthermore, $\partial E[x_1 - E_{post}[x_1]]/\partial x_1 > 0$, completing the proof. \square

Proof of Proposition 8. (i) For any distribution of posterior beliefs the agent chooses with initial balance $x_1 + \Delta x$, there is an information-acquisition strategy with initial balance x_1 that has the same cost and generates strictly more precise posterior beliefs (due to the entropy-reduction based cost). Choosing this information-acquisition strategy leads to a strictly lower penalty, yielding strictly higher expected utility. (ii) Since x_2 is a constant in the agent's utility function that she cannot influence and it also does not interact with actions c or t in the utility, she does not acquire information about x_2 , and her optimal strategy and expected penalty are independent of her beliefs about the distribution of x_2 . Since she is also risk-neutral, she is indifferent to a mean-zero shock to the investment account. \square