

Behavioral Macroeconomics Via Sparse Dynamic Programming

Xavier Gabaix*

March 16, 2017

Abstract

This paper proposes a tractable way to model boundedly rational dynamic programming. The agent uses an endogenously simplified, or “sparse,” model of the world and the consequences of his actions and acts according to a behavioral Bellman equation. The framework yields a behavioral version of some of the canonical models in macroeconomics and finance. In the life-cycle model, the agent initially does not pay much attention to retirement and undersaves; late in life, he progressively saves more, generating realistic dynamics. In the consumption-savings model, the consumer decides to pay little or no attention to the interest rate and more attention to his income. Ricardian equivalence and the Lucas critique partially fail because the consumer may not pay full attention to taxes and policy changes. In a Merton-style dynamic portfolio choice problem, the agent endogenously pays limited or no attention to the varying equity premium and hedging demand terms. Finally, in the neoclassical growth model, agents act on a simplified model of the macroeconomy; in equilibrium, fluctuations are larger and more persistent.

*xgabaix@fas.harvard.edu. I thank David Laibson for a great many enlightening conversations about behavioral economics. For useful comments I also thank Adrien Auclert, Robert Barro, John Campbell, John Cochrane, Emmanuel Farhi, Peter Ganong, Mark Gertler, David Hirshleifer, Bentley MacLeod, Matteo Maggiori, Jonathan Parker, Michael Rockinger, Thomas Sargent, Josh Schwarzstein, Andrei Shleifer, Bruno Strulovici, Richard Thaler, Jessica Wachter, Michael Woodford, seminar participants at various seminars and conferences, and my discussants: Andrew Abel, Nick Barberis, Harrison Hong, Jennifer La’O, Ali Lazrak, Ricardo Reis, Alp Simsek, Alexis Toda. I thank Wu Di, Chenxi Wang and Jerome Williams for very good research assistance. I am grateful to the Dauphine-CAM foundation, the Institute for New Economic Thinking, the NSF (SES-1325181) and the Sloan Foundation for financial support.

1 Introduction

In economics, we build a simplified model of the world – we select “important” dimensions of the world, and know that the model is not literally the true world. However, we assume that agents in our model are not like us: that they comprehend the full complexity of their world. This asymmetry is of course a bit odd. In contrast, in the present paper, agents behave like us economists: they build a simplified model of the (model) world they live in, and use it to act.

This paper shows a tractable way to model such behavioral agents, in a fairly general class of dynamic problems. Those agents tend to be more realistic than rational agents, and their bounded rationality has important policy consequences.

I show how the framework applies to some of the canonical models in macro-finance, in partial and general equilibrium: consumption-saving problems, the baseline neoclassical growth model (the Cass-Koopmans model), general linear-quadratic problems, and dynamic investment in risky assets (Merton’s problem). The upshot is that we have a portable, fairly general structure that applies to some core machines of macroeconomics and allows to see where bounded rationality (BR) is important in those situations.

One of the persistent criticisms of traditional economics is the unrealism of the infinitely forward-looking agent who computes the whole equilibrium in her own head. This lack of realism has long been suspected to be the cause of some counterfactual predictions that we will review below. Behavioral economics aims to provide an alternative. The greatest successes of behavioral economics in the literature so far change the agents’ tastes (e.g. prospect theory or hyperbolic discounting) or their beliefs (e.g. overconfidence), while keeping the assumption of rationality. When tackling the rationality assumption, there is much less agreement, and the modelling of bounded rationality is much more piecemeal, different from one situation to the next.

This paper proposes a compromise that keeps much of the generality of the rational approach and injects some of the wisdom of the behavioral approach, mostly inattention and simplification. It does so by proposing a way to insert some bounded rationality into a large class of problems, the “recursive” contexts, i.e. with dynamic programming around a dynamic steady state.

To illustrate these ideas, let us consider a canonical consumption-savings problem. The agent maximizes utility from consumption, subject to a budget constraint, with a stochastic interest rate and stochastic income. In the rational model, the agent solves a complex DP problem with three state variables (wealth, income and the interest rate). This is a complex problem that requires a computer to solve.

How will a boundedly rational agent behave? I assume that the agent starts with a much simpler model, where the interest rate and income are constant – this is the agent’s “default” model. Only one state variable remains, his wealth. He knows what to do then (consume a certain fraction of his wealth, and permanent income), but what will he do in a more complex environment, with stochastic interest rate and stochastic income? In the sparse version, he considers parsimonious

enrichments to the value function, as in a Taylor expansion. He asks, for each component, whether it will matter enough for his decision. If a given feature (say, the interest rate) is small enough compared to some threshold (taken to be a fraction of standard deviation of consumption), then he drops the feature or partially attenuates it. The result is a consumption policy that pays partial attention to income and possibly no attention at all to the interest rate. This seems realistic.¹

The result is a sparse version of the traditional permanent-income model. We see that it is often *simpler* than the traditional model. Indeed, the agent typically ends up using a rule which is simpler (e.g., not paying attention to the interest rate).

I also present a behavioral version of a large class of models and work out, in detail, a BR version of most canonical of them, the neoclassical growth model of Cass-Koopmans. In this version, agents pay of lot of attention to their own variables, less to aggregate variables. One upshot is that with BR, macroeconomic fluctuations are larger and more persistent. I illustrate this proposition, and qualify it, as it appears to hold for most reasonable values of the parameters, but can be overturned for extreme values. To understand the simple idea, imagine first an economy with only one state variable, capital. It starts with a steady state amount of capital. Then, there is a positive shock to the endowment of capital. In a rational economy, agents would consume a certain fraction of it, say 6%, every period. That will lead the capital stock to revert quickly to its mean. However, in an economy with sparse agents, investors will not pay full attention to the additional capital. They will consume less of it than a rational agent would. Hence, capital will be depleted more slowly and will mean-revert more slowly. The shock has more persistent effects. Given that shocks are more persistent, past shocks accumulate more. Mechanically, this leads to larger average deviations of capital from its trend. As a consequence, the interest rate and GDP also have larger, and more persistent, deviations from trend.

The model allows us to express those ideas in simple, quantitative ways. It also allows us to explore them in rich environments.

Here are some conclusions for *individual decision-making in macroeconomic contexts*:²

1. Agents react more to near, rather than future, shocks. For instance, the (finitely lived) consumer has a higher marginal propensity to consume for current and near payments, rather than distant payments.
2. The Euler equation fails (under the objective model of the world). However, it holds under the agent's subjective, sparse model.
3. Agents start saving "too late" for retirement, and expenditure falls in the years before retirement (as the agent scrambles to save more for retirement).

¹In the language of Kahneman (2011, Chapter 8), the agent "substitutes" a complex problem by a simpler one – decision-making in a simplified world.

²The rest of the paper will give references about those stylized facts.

4. Agents accumulate a too small buffer of savings, as they are (partially) inattentive to the risk of income fluctuations.
5. Agents do not react much or at all to the interest rate, at least when small purchases are concerned. However, they tolerate a non-smooth consumption profiles. Note that in the rational model, the first fact would require a low IES (intertemporal elasticity of substitution), while the second fact would require a high IES.
6. The agent doesn't need to the whole actual macro equilibrium in his head before acting; he just uses a simplified model of the world, leading to simple policies, e.g. partially forward-looking consumption functions.
7. When choosing their portfolio, agents pay less attention to the hedging demand motive.

Here are some conclusions for *aggregate macroeconomics*:

8. Fiscal policy is more powerful because Ricardian equivalence partly fails. If the government gives a dollar today and takes it back later (plus interest), consumption today increases – though it should not react in the simplest rational model.³
9. The Lucas critique has less, or zero, bite. When policy changes are small and temporary, sparse agents' policy functions do not change, or change little.
10. GDP fluctuations are amplified and more persistent in the most basic DSGE model (compared to the rational benchmark).
11. This perspective revives the “old Keynesian” agent. The agent here looks like a hybrid between a neoclassical agent and an old Keynesian agent⁴, in the sense that he's myopic and adopts simple decision rules. However, unlike the truly old Keynesian agent, those decision rules are microfounded—as policies that are optimal under a simplified model of the world. Unlike the neoclassical and New Keynesian agent, this sparse agent is partially myopic and does not react to all things.

Those substantive conclusions are, I hope, of some interest. Still, the main contribution of this paper is its methodology. It develops a procedure that allows an economist to continuously transform a rational agent into a boundedly rational one. This leads to further substantive conclusions about the impact of bounded rationality on economic life, such as the ones above.

This paper is indebted to a series of ideas in behavioral modelling in micro and macro contexts. This literature will be discussed in Section 8, after the model is clear to the reader. The rest of the

³In Gabaix (2016), I also find that monetary policy, especially forward guidance, is less powerful. I do so in writing a behavioral version of the New Keynesian model, applying the techniques of the present paper.

⁴The canonical New Keynesian agent is basically the neoclassical agent, in a world of sticky prices. He's fully rational.

paper is as follows. Section 2 gives an elementary, but substantive, example – a behavioral version of the life cycle model. Section 3 presents the general procedure. Then, we apply it to a variety of canonical examples. Section 4 presents basic partial-equilibrium building blocks with behavioral agents: the basic consumption-savings problem, including variants such as the failure of Ricardian equivalence. Section 5 works out the baseline neoclassical growth model. Section 6 develops other models. Section 7 proposes discusses variants, and frequently asked questions. Section 8 contains the literature review. Section 9 concludes. The appendix and online appendix contain further proofs and extensions.

Notations. I will use superscript r for the traditional or rational model, d for the default model, sometimes s for the sparse or behavioral model. I will call $a \odot b := (a_i b_i)_{i=1\dots n}$ the component-wise product.

2 Behavioral Life-Cycle Model: A Gentle Introduction

As a simple introduction, I give a behavioral version of the most basic life-cycle model (Modigliani and Brumberg (1954)).⁵ The agent works for the first L periods of his life, then retires, and dies at period T . His utility is

$$\sum_{t=0}^{T-1} u(c_t).$$

The interest rate and the subjective discount rate are both 0. He receives income $y_t = \bar{y}$ when working ($t \in [0, L)$), and $y_t = \bar{y} + \hat{y}$ when retired ($t \in [L, T)$). Here $\hat{y} < 0$ captures the income loss during retirement. Income corresponds to the dotted line in Figure 1. Financial wealth w_t evolves as $w_{t+1} = w_t + y_t - c_t$, and the terminal condition is $w_T = 0$.

Let us first analyze a rational agent. At time 0, his resources are $\Omega_0 := w_0 + \sum_{\tau=0}^{T-1} y_\tau = w_0 + T\bar{y} - x$, where $x := -(T - L)\hat{y} > 0$ is the total income loss due to retirement. His consumption problem is $\max_{(c_t)_{0 \leq t < T}} \sum_{t=0}^{T-1} u(c_t)$ s.t. $\sum_{t=0}^{T-1} c_t = \Omega_0$, with $u' > 0, u'' < 0$. So he consumes a constant amount at all periods: $c_t = \frac{\Omega_0}{T} = \frac{w_0 - x}{T} + \bar{y}$. In particular:

$$c_0 = \frac{w_0 - x}{T} + \bar{y}.$$

The same reasoning holds starting at a date $t \leq L$. Then, there are only $T - t$ periods remaining, so the policy becomes:

$$c_t = \frac{w_t - x}{T - t} + \bar{y} \tag{1}$$

That policy guarantees a constant consumption $c_t = c_0$ over his lifetime.

There is a dynamic programming formulation that will be useful in the behavioral model. At

⁵It is a bit unusual to start a framework with an example, but this example will be useful to illustrate modelling choices.

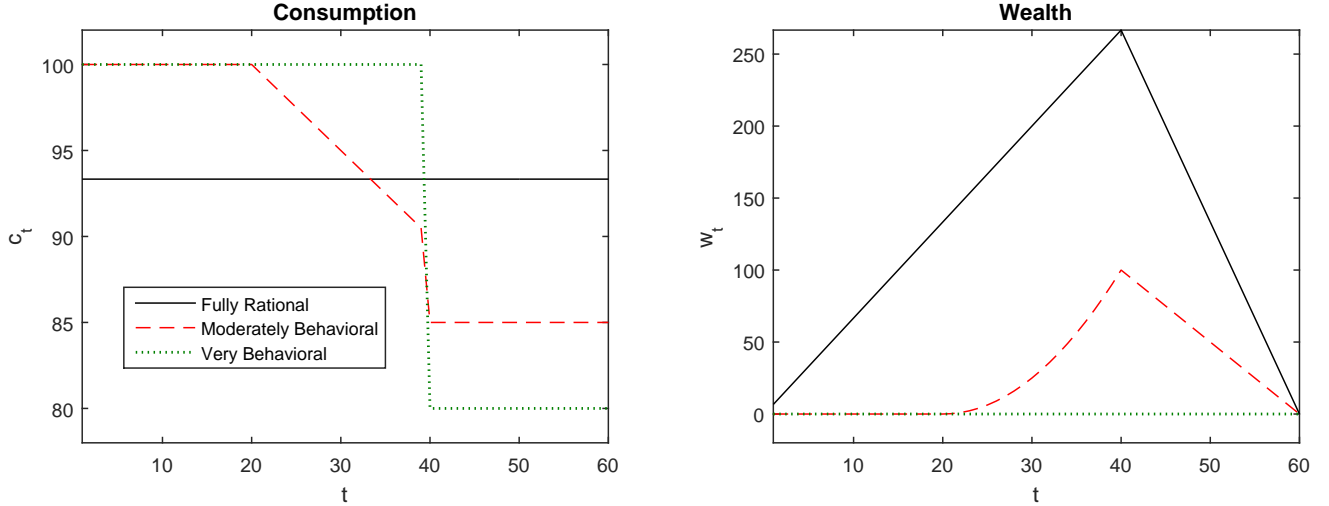


Figure 1: Consumption and wealth of behavioral life-cycle agents. *Notes.* Income is plotted in the dotted line – it is also the very behavioral agent’s consumption. The agent starts life at time 0, receives income $\bar{y} = 100$ while working (until period 40), and receives $\bar{y} + \hat{y} = 80$ in retirement (period 40 to 60). The solid line represents a fully rational agent (i.e. $\bar{\kappa} = 0$), the dashed line a moderately behavioral agent ($0 < \bar{\kappa} < \bar{\kappa}^*$, for a finite $\bar{\kappa}^* = |\hat{y}|$), and the dotted line a very behavioral agent ($\bar{\kappa} \geq \bar{\kappa}^*$), who just consumes current income. The moderately behavioral agent does not save for retirement at first, but starts saving before retirement. The right panel plots the wealth accumulated by the agent.

time t , the remaining lifetime utility is $(T - t) u(c_t)$. So, the value function is (for $t \leq L$)

$$V^r(w_t, x, t) = (T - t) u\left(\frac{w_t - x}{T - t} + \bar{y}\right) \quad (2)$$

where superscript r denotes the *rational* agent. This rational agent (1) satisfies the Bellman equation:⁶ $c_t = \arg \max_c v(c, w_t, x, t)$, where we define:

$$v(c, w_t, x, t) := u(c) + V^r(w_t + \bar{y} - c, x, t + 1) \quad (3)$$

Let us now consider the behavioral agent. I want to capture the idea that he does not fully perceive income loss associated with retirement. The systematic procedure will be justified and explained in the body of this paper, but here I just show how it applies in this case. The behavioral agent will consume:

$$c_t = \arg \max_c v(c, w_t, m_t x, t) \quad (4)$$

where $m_t \in [0, 1]$ denotes an (endogenous) attention to future retirement. If $m_t = 1$, the agent is

⁶Here I use the basic sparse max of Definition 3.4, where the agent uses a simplified model of his own future actions. If the agent has a “sophisticated” understanding of this future actions (as in the iterated procedure developed in Definition 3.8, the decision would be the same to the first order, i.e. up to $O(x^2)$ terms (see Section 3.5).

fully rational. However, if $m_t = 0$, the agent doesn't see retirement at all: he behaves as if there was no income loss from retirement in the future. The general model allows for partial attention m_t . Given m_t , the solution of (4) is:⁷

$$c_t = \frac{w_t - m_t x}{T - t} + \bar{y} \quad (5)$$

In the full model developed soon, attention m_t will come from the costs and benefits of attention, and will be expressed as follows. Calling $c_t(m_t) = \frac{w_t - m_t x}{T - t} + \bar{y}$, we define $c_t^d = c_t(0) = \frac{w_t}{T - t} + \bar{y}$, the “default” policy that corresponds to no attention to retirement; $c_t'(0) = \frac{-x}{T - t}$ the marginal impact of attention; and $v_{cc}^t = v_{cc}(c_t^d, w_t, 0, t) = (1 + \frac{1}{T - t - 1}) u''(c_t^d)$ the curvature of the objective function. The general procedure will give the attention at time t to the income loss during retirement:

$$m_t = \mathcal{A} \left(\frac{-v_{cc}^t c_t'(0)^2}{\kappa_t} \right) \quad (6)$$

for an attention function \mathcal{A} with values in $[0, 1]$ and a cost of cognition κ_t discussed soon. A little more calculation gives the value of attention as follows.

Proposition 2.1 (Lifecycle model, behavioral version) *In the behavioral life-cycle model, the optimal consumption policy is, before retirement ($t < L$),*

$$c_t = \frac{w_t - m_t x}{T - t} + \bar{y}.$$

with an attention $m_t = \mathcal{A} \left(\frac{-u''(c_t^d) x^2}{\kappa_t (T - t - 1)(T - t)} \right)$, and after retirement ($t \geq L$), $c_t = \frac{w_t}{T - t} + \bar{y} + \hat{y}$. Hence, when $\kappa_t > 0$, consumption weakly falls over time, and discretely falls at retirement. After retirement, consumption is constant.

I next present a numerical illustration. I use the attention function $\mathcal{A}(v) = \max \left(1 - \frac{1}{|v|}, 0 \right)$ from (15). I assume the following scaling of the cost: $\kappa_t = \bar{\kappa}^2 |u''(c_t^d)|$ with $c_t^d = \frac{w_t}{T - t} + \bar{y}$. This is largely for convenience, and it corresponds to a constant cost when utility is linear-quadratic. This simplifies the expressions without changing the economics much. Then (using continuous time to make expressions neat), the agent thinks about retirement as soon as $\frac{x}{T - t} \geq \bar{\kappa}$, i.e. at a time $s = \max \left(0, \min \left(L, T - \frac{x}{\bar{\kappa}} \right) \right)$. Solving for wealth, the value of consumption is (when $s \in (0, L)$;

⁷Proof: the first order condition of (4) is: $u'(c_t) = V_w^r = u' \left(\frac{w_t + \bar{y} - c_t - m_t x}{T - t - 1} + \bar{y} \right)$, i.e. $c_t = \frac{w_t + \bar{y} - c_t - m_t x}{T - t - 1} + \bar{y}$, i.e. $c_t = \frac{w_t - m_t x}{T - t} + \bar{y}$.

Section 14 in the appendix describes the whole solution, including in discrete time):⁸

$$c_t = \begin{cases} \frac{w_0}{T} + \bar{y} & \text{for } t < s \\ \frac{w_0}{T} + \bar{y} - \frac{2\bar{\kappa}^2}{x} (t - s) & \text{for } s \leq t < L \\ \frac{w_L}{T-L} + \bar{y} + \hat{y} & \text{for } t \geq L \end{cases}$$

where $w_L = (1 - \frac{L}{T}) w_0 + \frac{\bar{\kappa}^2}{x} (L - s)^2$.

Figure 1 plots the resulting consumption and wealth ($w_0 = 0$, $\bar{y} = 100$) for different levels of the cost of rationality κ . The solid line represents a fully rational agent (i.e. $\bar{\kappa} = 0$): he fully smooths consumption. The dotted line shows a very behavioral agent, who simply consumes current income ($\bar{\kappa} \geq \bar{\kappa}^* := \frac{x}{T-L} = |\hat{y}|$). The dashed line shows a moderately boundedly rational agent ($0 < \bar{\kappa} < \bar{\kappa}^*$).

At first, he does not save for retirement. He’s thinking “let me not think about that future loss of income at retirement, it’s so remote than it’s not worth it”. However, he does start saving at some point before retirement – at period 20 in this calibration. This is because in (5), $\left| \frac{\partial c_t}{\partial m_t} \right| = \frac{m_t x}{T-t}$, so that thinking more about retirement is more important as time goes by. So, at period 21 he thinks “OK, I should start thinking a bit about that retirement, and saves a bit”: Over time he thinks more about retirement, and saves more. At retirement, his consumption drops as he fully realizes that his income has fallen. This illustrates the smooth, partial myopia of this agent.

This paper is mostly theoretical, but it is worth asking the following question.

Is it indeed the case that people tend to save “too late” for retirement? The issue is still controversial, but let us consider three salient facts.

1. Expenditure declines after the age of 45 (Aguiar and Hurst 2013, Figure 1, for expenditures without housing services) – much like the behavioral agent of Figure 1.
2. There is a fall in expenditure when income predictably falls, again like in Figure 1:
 - (a) at retirement (Bernheim, Skinner and Weinberg 2001).⁹
 - (b) at the (predictable) expiration of unemployment benefits (Ganong and Noel 2015).
3. People say that they plan for retirement late in their working life, or not at all. For instance, 23% of the 18-29 year old say that have “figured out how much they need to save for retirement”, while 51% of the 45-59 year old say they have done so (Lusardi 2011).¹⁰

⁸With this scaling for κ_t , the expression for consumption is independent of the concave utility function u .

⁹In addition, Bernheim, Skinner and Weinberg (2001, Fig. 4) find that the drop is more pronounced for individuals with low wealth. This is what the present model predicts: controlling for lifetime incomes, agents with higher bounded rationality (higher κ) accumulate fewer assets and have a bigger drop in consumption at retirement. See also Kueng (2015).

¹⁰More indirectly, a whole literature on retirement plans supports the notion of inert decision-making for retirement. For instance, Beshears et al. (2015) find that the potency of framing and defaults is reduced as people approach retirement, consistent with the notion they think more carefully about retirement as they are near it.

Facts 1-3 arise naturally from the behavior as described in Figure 1. For Fact 3, if agents indeed have a κ too large, they don't plan for retirement at all, even right before it.

Facts 1-3 are each inconsistent with the plainest rational model.

Facts 1-2a can be made consistent with an *enriched* rational model. Fact 1 can be explained by introducing credit constraints and income risk (Gourinchas and Parker 2002). Facts 1 and 2a can be explained by observing that retired consumers buy more efficiently or buy fewer non-work goods (Aguiar and Hurst 2013). Fact 2b is harder to reconcile with fully forward-looking models, though qualitatively consistent with Figure 1. Fact 3 is not consistent with a rational model, but one could dismiss it by stating that agents' reports of what they think about are meaningless – a point of view I do not share.¹¹

Though these facts are not dispositive, they form, I submit, reasonable presumptive evidence for the idea that people are not fully far-sighted, even for retirement savings. The present framework could help write enriched empirical models allowing for both traditional factors and myopia, where the above-mentioned evidence could be systematically assessed.

This example illustrates the behavior and mechanics of the sparse agent. I now move on to a systematic formulation of that agent, which applies to much more general problems.

3 General Framework

3.1 The Sparse Max for Static Problems: Quick Review

To think about bounded rationality, the tractable dynamic framework laid out here is possible because it rests on a tractable static framework I laid out in previous work (Gabaix (2014)). I review it in this subsection. There, the core is a sparse max or smax operator, which is a generalized, behavioral version of the traditional max operator of maximization under constraints.

Let us review the sparse max when there is no budget constraint. The agent faces a maximization problem which is, in its rational version, $\max_a v(a, x)$, where a is an action and x a state variable.¹² There is an attention vector, m , and an attention-dependent extension of the utility function, $v(a, x, m)$. For instance, we will typically take

$$v(a, x, m) := v(a, m_1 x_1, \dots, m_n x_n) \tag{7}$$

to be the perceived utility function when the consumer is partially inattentive to x_i . When $m_i = 1$, the agent fully perceives dimension i ; when $m_i = 0$, the agent is fully inattentive to it. Attention

¹¹Note that a simple hyperbolic agent (Laibson 1997, O'Donoghue and Rabin 2001) would not behave this way. That agent has full foresight of future income flows, and sees equally perfectly pre- and post-retirement income. For instance, the policy of the log agent is $c_t = \frac{w_t + x + (T-t-1)\bar{y}}{1 + \beta(T-t-1)}$. However, with uncertainty and credit constraints, the hyperbolic agent can look closer to the one with the above behavior (Harris and Laibson 2001), though with much complexity and not because of limited foresight about future income.

¹²This utility may be a value function, as in (3).

generates an action, $a(x, m) := \arg \max_a v(a, x, m)$. There is a default attention vector m^d , taken to be 0 in most applications, and a default action $a^d := \arg \max_a v(a, x, m^d)$. I call $a_{m_i} = \frac{\partial a}{\partial m_i}$, evaluated at $(a, m) = (a^d, m^d)$, the normative impact on the action of a change in attention. Hence, $a_{m_i} = -v_{aa}^{-1} v_{am_i}$. When (7) holds, $a_{m_i} = a_{x_i}(x, m)|_{m=(1, \dots, 1), x=0} x_i$.

There is a nonnegative parameter κ , which is a cognition cost – formally, a taste for sparsity. When $\kappa = 0$, the agent is the traditional agent.¹³ The x_i are viewed by the agent as being drawn from a distribution with standard deviation σ_i .

Definition 3.1 (Sparse max operator, without a budget constraint) *The sparse max, $\text{smax}_{a, m|m^d} v(a, x, m)$, is defined by the following procedure.*

Step 1: Choose the attention vector m^ :*

$$m^* = \arg \min_{m \in [0, 1]^n} \sum_i \left[\frac{1}{2} \Lambda_{ii} (1 - m_i)^2 + \kappa g(m_i - m_i^d) \right] \quad (8)$$

with the cost-of-inattention factors $\Lambda_{ii} := -\mathbb{E}[a_{m_i} v_{aa} a_{m_i}]$, $g' > 0$.

Step 2: Choose the action

$$a^s = \arg \max_a v(a, x, m^*) \quad (9)$$

and set the resulting utility to be $v^s = v(a^s, x)$. In the expressions above, derivatives are evaluated at $m = m^d$ and $a^d = \arg \max_a v(a, x, m^d)$.

In other terms, the agent solves for the optimal m^* that trades off a proxy for the utility losses (the first term in the right-hand side of equation (8)) and a psychological penalty for deviations from a sparse model (the second term on the right-hand side of equation (8)).¹⁴ Then, the agent maximizes over the action a , as if m^* were the true model. The problem is solved by backward induction.

This leads to define the attention function:¹⁵

$$\mathcal{A}(v) := \arg \min_{m \in [0, 1]} \left\{ \frac{1}{2} |v| (1 - m)^2 + g(m) \right\}.$$

This represents the optimal attention to a variable with variance $|v|$, normalizing other factors to 1. Figure 2 plots typical shapes.

Then, the value of attention to dimension i is given by the key relation:

$$m_i^* = \mathcal{A}(-\mathbb{E}[a_{m_i} v_{aa} a_{m_i}] / \kappa). \quad (10)$$

¹³This is true unless the matrix Λ of Definition 3.1 is singular – in that case, the iterated sparse max of Definition 15.1 helps.

¹⁴In Gabaix (2014), I sum over all terms $\sum_{i, j} (1 - m_i) \Lambda_{ij} (1 - m_j)$, with $\Lambda_{ij} := -\mathbb{E}[a_{m_i} v_{aa} a_{m_j}]$. Here, I streamline the procedure, and let the agent consider only the diagonal terms. This is inessential, but is simpler.

¹⁵Also, if there are more than one optimum m 's, we take the largest one.

This formula gives a simple “plug and play” solution for the (potentially very complex) attention problem: to allocate attention to dimension i , just use (10). There is no need to come back (except for generalizations) to the background problem (8). The following Lemma derives a typical case.

Lemma 3.2 (Basic simple case for static smax) *In the case $v(a, x, m) = v(a, m_1x_1, \dots, m_nx_n)$, the smax operator yields:*

$$m_i^* = \mathcal{A}\left(-\sigma_i^2 a_{x_i} v_{aa} a_{x_i} / \kappa\right) \quad (11)$$

and

$$a^s = \arg \max_a v(a, m_1^*x_1, \dots, m_n^*x_n)$$

with $a_{x_i} = \frac{\partial a}{\partial x_i} = -v_{aa}^{-1} \cdot v_{a,x_i}$. In the expressions above, derivatives are evaluated at $x = 0$ and $a^d = \arg \max_a v(a, 0)$.

The intuition is that the x_i 's are truncated. If $|a_{x_i}|$ is small enough, so that x_i shouldn't matter much anyway, then $m_i^* = 0$, and the agent doesn't pay attention to x_i (if $m_i^d = 0$).

This leads to the defining the truncation function, with b the coefficient on the state variable x in a linear policy function:

$$\tau(b, k) := b\mathcal{A}\left(\frac{b^2}{k^2}\right) \quad (12)$$

It is the coefficient b , times the attention to the coefficient, divided by the scaled cognition cost k .

The following lemma gives a more explicit version of the action.

Lemma 3.3 *If the rational action is:*

$$a^r(x) = a^d + \sum_i b_i x_i + O(\|x\|^2)$$

then the sparse action is

$$a^s(x) = a^d + \sum_i \tau\left(b_i, \frac{\kappa_a}{\sigma_i}\right) x_i + O(\|x\|^2) \quad (13)$$

with $\kappa_a := (\kappa / |v_{aa}|)^{1/2}$.

When attention is chosen after seeing x (“ex post”), we use the same expressions, with $\sigma_i := |x_i|$. For instance, the ex-post action becomes:

$$a^s(x) = a^d + \sum_i \tau(b_i x_i, \kappa_a) + O(\|x\|^2) \quad (14)$$

In the “ex ante” procedure, the slope is chosen before seeing x_i . Hence, the policy is still linear in x_i , which makes that procedure useful in macro. In the “ex post” procedure, the truncation is




Figure 2: Three attention functions $\mathcal{A}_0, \mathcal{A}_1, \mathcal{A}_2$, corresponding to fixed cost, linear cost and quadratic cost respectively. We see that \mathcal{A}_0 and \mathcal{A}_1 induce sparsity – i.e. a range where attention is exactly 0. \mathcal{A}_1 and \mathcal{A}_2 induce a continuous reaction function. \mathcal{A}_1 alone induces sparsity and continuity.

chosen after seeing the x_i , and the policy is non-linear in x_i .

Attention and Truncation Functions. Here are some good truncation functions. In Gabaix (2014), I study attention functions $\mathcal{A}_\alpha(\sigma^2)$ corresponding to $g(m) = m^\alpha 1_{m>0}$. For instance, for the values $\alpha = 0, 1, 2$, we have:

$$\mathcal{A}_0(\sigma^2) = 1_{\sigma^2 \geq 2}, \quad \mathcal{A}_1(\sigma^2) = \max\left(1 - \frac{1}{\sigma^2}, 0\right), \quad \mathcal{A}_2(\sigma^2) = \frac{\sigma^2}{2 + \sigma^2} \quad (15)$$

The truncation functions $\tau_\alpha(b, k)$ is then (using (12)):¹⁶

$$\tau_0(b, k) = b \cdot 1_{b^2 \geq 2k^2}, \quad \tau_1(b, k) = b \max\left(1 - \frac{k^2}{b^2}, 0\right), \quad \tau_2(b, k) = \frac{b^3}{b^2 + k^2} \quad (16)$$

Figure 2 plots the attention functions, and Figure 3 the corresponding truncation functions.

3.2 Dynamic Programming: A Motivating Example

To motivate the general structure, let us start with a basic example, the consumption-savings problem. The agent has utility $\mathbb{E} \sum_{t=0}^{\infty} \beta^t c_t^{1-\gamma} / (1-\gamma)$. Wealth w_t , and the state variables evolve

¹⁶Another useful cost function is $g_{L_1}(m) = -\ln(1-m)$, which generates $\mathcal{A}_{L_1}(\sigma^2) = \max\left(1 - \frac{1}{|\sigma|}, 0\right)$, and $\tau_{L_1}(b, k) = \text{sign}(b) \max(|b| - |k|, 0)$. The subscript L_1 denotes that it often arises when doing an L_1 regularization, as in the sparsity literature in statistics (Tibshirani 1996, Candès and Tao 2006).




Figure 3: Three truncation functions. Because it gives sparsity and continuity, the τ_1 function is recommended.

as:

$$w_{t+1} = (1 + \bar{r} + \hat{r}_t)(w_t - c_t) + \bar{y} + \hat{y}_t. \quad (17)$$

$$\hat{r}_{t+1} = \rho_r \hat{r}_t + \varepsilon_{t+1}^r \quad (18)$$

$$\hat{y}_{t+1} = \rho_y \hat{y}_t + \varepsilon_{t+1}^y \quad (19)$$

That is, wealth at $t + 1$ is savings at t , $w_t - c_t$, invested at rate $r_t = \bar{r} + \hat{r}_t$, plus current income, $y_t = \bar{y} + \hat{y}_t$. Here, \hat{r}_t and \hat{y}_t are deviations of the interest rate and income from their means, respectively, and follow AR(1) processes, where ε_{t+1}^r and ε_{t+1}^y are disturbances with mean zero and no correlation across periods. For simplicity, assume here that $\beta R = 1$, where $R := 1 + \bar{r}$.

This is a complex problem, with 3 state variables

$$z_t := (w_t, \hat{r}_t, \hat{y}_t)$$

This is also a metaphor for a more complex model, which could have 30 or 300 state variables. What will the agent do at time 0?¹⁷

One thing I wish to capture is “the agent may not want to think about the interest rate”. The reader can introspect: when most people plan their vacation, do they think “now interest rates are high, so it’s a makes good sense to spend little this summer, and more next summer; this way we’ll respect our Euler equation, and will make sure that our consumption growth is high, in congruence with the current high interest rate?” Most people, and the reader, I imagine, do not do that (by the way, econometric evidence confirms that they don’t, see e.g. Hall 1988). Non-economists would find that depiction of them ludicrous. Accordingly, the sparse agent will be allowed not to think

¹⁷This paper could for instance help model simplified decision-making in economies with a large but finite number of firms linked in networks, where the state space has very high dimension (Acemoglu et al. 2012, Bigio and La’O 2016, Carvahlo and Grassi 2016, Caballero and Simsek 2013, Gabaix 2011).

about the interest rate — though he will think about it, say when buying a house, or when interest rates are very volatile.

Second, “the agent may wish to imagine simplified dynamics for the process”, e.g. he may replace the dynamics of this income, for instance, by a simpler process, e.g. imagine it will be roughly constant – without paying attention to the detailed stochasticity of the income process.

Before showing the equations, I propose an intuitive picture of what the agent’s world view is. I posit that the agent knows what to do in a simpler, default model. That is, he assumes that future interest rate and income will be constant. Then, the optimal consumption is

$$c^d(w_t) = \frac{\bar{r}w_t + \bar{y}}{R} \quad (20)$$

and the value function is $V^d(w_t) = \frac{R}{\bar{r}(1-\gamma)} \left(\frac{\bar{r}w_t + \bar{y}}{R}\right)^{1-\gamma}$. Then, the agent decides whether to enrich his very model where everything is constant. He asks “is it worth thinking about the interest rate?” To do so, he contemplates a one-variable enrichment, with the interest rate, and sees whether it’s worth it. If it isn’t, he settles for a model where the interest rate is constant. If it is, he enriches his model, with a non-constant interest rates.

Here is how I propose to capture those ideas. I posit that the agent contemplates a “simplifiable meta-model”. In our example, this is a transition function $F(c_t, z_t, m)$, whose components are (with $F = (F^w, F^r, F^y)$):¹⁸

$$w_{t+1} = F^w(c_t, z_t, m) := (1 + \bar{r} + m_r \hat{r}_t)(w_t - c_t) + \bar{y} + m_y \hat{y}_t \quad (21)$$

$$\hat{r}_{t+1} = F^r(c_t, z_t, m) := \rho_r(m) \hat{r}_t + m_{\sigma_r} \varepsilon_{t+1}^r, \quad \rho_r(m) := m_{\rho_r} \rho_r + (1 - m_{\rho_r}) \rho_r^d \quad (22)$$

$$\hat{y}_{t+1} = F^y(c_t, z_t, m) := \rho_y(m) \hat{y}_t + m_{\sigma_y} \varepsilon_{t+1}^y, \quad \rho_y(m) := m_{\rho_y} \rho_y + (1 - m_{\rho_y}) \rho_y^d \quad (23)$$

For instance, when $m_r = 0$, the agent doesn’t pay attention to the interest rate, while if $m_r = 1$, he will pay attention to 1. Likewise, m_y represents the agent’s attention to future income shocks: if $m_y < 1$ the agent pays little attention to future income streams – he is myopic towards them. Parameter m_{σ_y} in (23) represents the attention to noise: when $m_{\sigma_y} = 0$, the agent doesn’t pay attention to the stochasticity of the income—which will lead the agent to accumulate a too small buffer of savings and be vulnerable to shocks.¹⁹ Parameter m_ρ represents the attention to the fine structure of the income process: when $m_{\rho_y} = 0$, agent replaces the true autocorrelation ρ_y by another one, ρ_y^d , e.g. if $\rho_y^d > \rho_y$ and $m_{\rho_y} < 1$, the agent will think that income shocks are more persistent than they are.

¹⁸This is related to the “perceived law of motion” of the literature on learning (Evans and Honkapohja 2001). There are important differences: the concept of a simplified model also holds in static contexts, without learning; the procedure is not about learning, but rather about the simplification of a known (but complex) already-learned model; and in dynamic programming the sparse agent needs to have a model of his *own* future actions.

¹⁹Section 12.3 of the online appendix develops this. See Lusardi, Schneider and Tufano (2011) for evidence of an insufficient buffer stock of saving. If agents underperceive the risks to income, that motivates policy of extra assistance, e.g. extra unemployment benefits.

Here

$$m := (m_y, m_{\rho_y}, m_{\sigma_y}, m_r, m_{\rho_r}, m_{\sigma_r})$$

is the subjective parametrization of the world, and $\mu = (1, 1, 1, 1, 1, 1)$ is the objective parametrization. Instead of thinking that the transition function is $z_{t+1} = F(c_t, z_t, \mu)$, the agent will decide to imagine that it is $z_{t+1} = F(c_t, z_t, m)$.

So far, we have described a structure. The next section defines the sparse max. In particular, what does it mean to “use a simplified model”, what’s the agent’s anticipation of his future actions, so we have a well-defined notion of dynamic programming? Next, how will attention be allocated? Then, we will study consequences of this behavior for classic macro questions.

3.3 Sparse Dynamic Programming: Basic Definition

The state z_t , the action a_t and the i.i.d. innovation ε_{t+1} are vectors. There are T periods, where T could be infinite.

The rational problem. The agent’s rational problem is:

$$\max_{(a_t)_{0 \leq t < T}} \sum_{t=0}^{T-1} \beta^t u(a_t, z_t) \quad \text{s.t.} \quad z_{t+1} = F^z(a_t, z_t, \varepsilon_{t+1}) \quad (24)$$

and a terminal condition $z_T \in \mathcal{F}^T$ for a given set \mathcal{F}^T .²⁰

The rational version of the dynamic programming (DP) problem is a series of value functions $V^{r,t}$ satisfying the Bellman equation:

$$V^{r,t}(z) = \max_a \{u(a, z) + \beta \mathbb{E} [V^{r,t+1}(F^z(a, z, \varepsilon_{t+1}))]\} \quad (25)$$

for $t = 0, \dots, T - 1$, and with $V^{r,T}(z) = 0$. A policy is then a function $a(z)$. Actually, one can drop the index t in the explicit formulation of (25), if state vector includes the calendar date t , e.g. if we can write $z_t = (w, x, t)$ and V depends on that calendar date component. This way, the traditional Bellman equation can be simply written without explicit t superscript:

$$V^r(z) = \max_a \{u(a, z) + \beta \mathbb{E} [V^r(F^z(a, z, \varepsilon_{t+1}))]\} \quad (26)$$

The sparse max version. In the smax version, we are given attention-augmented utility and transition functions $u(a, z, m)$ and $F^z(a, z, \varepsilon_{t+1}, m)$ – in a way that will be illustrated later.²¹

²⁰With infinite horizon, the transversality condition is typically $\limsup_{T \rightarrow \infty} Z_T \in \mathcal{F}^T$.

²¹One can take $u(a, z; m) = u(a, m \odot z, t)$ and for the k -th component of vector F^z

$$F^{z_k}(a, z, \varepsilon, m) = F^{z_k}(a, m^k \odot z, \varepsilon)$$

where $m^k \in \mathbb{R}_+^{\dim z}$ denotes the attention to factors; generally $(m_k^k)^d = 1$: when predicting the future values of variables z_k , full attention is paid to its initial value.

We are also given a “default proxy value function”, $V^p(z)$. Typically, it is just the rational value function ($V^p = V^r$), i.e. the function assuming that the agent will behave rationally afterwards – in the simplified model contemplated by the agent (this will be very clear in the 3-period example of Proposition 4.1).²² We could also have V^p to be the objective value function (where the agent has rational expectations about his own boundedly rational behavior – something I derive in Section 3.4). But, the two differ only by second order terms (see Proposition 3.11), so this assumption makes materially little difference. Since it is simpler, and (I will argue) more realistic, I recommend taking $V^p = V^r$.

The agent’s action is as follows.²³

Definition 3.4 (Action in sparse dynamic programming) *The action of the basic sparse max behavioral agent is:*

$$a(z, V^p) = \arg \operatorname{smax}_{a; m | m^d} \{u(a, z, m) + \beta \mathbb{E}[V^p(F^z(a, z, \varepsilon_{t+1}, m))]\}. \quad (27)$$

Hence, the agent maximizes his perceived flow utility function, and his perceived continuity value function. The degree of sophistication in those perceptions is controlled by vector m . When the cost of rationality κ is 0 (and $m = (1, \dots, 1)$), and $V^p = V^r$, the agent is just the rational agent.

3.4 More Advanced Notion: Iterated Dynamic Sparse Max

The rest of this section examines advanced notions, so the reader is encouraged to skip it in the first reading.

Given the decision function $a(z) := a(z, V^p)$ from Definition 3.4, the agent obtains an objective value function $V^o(z)$, which satisfies:

$$V^o(z) = u(a(z), z) + \beta \mathbb{E}[V^o(F^z(a(z), z, \varepsilon_{t+1}))] \quad (28)$$

Does this equation admit a well-defined solution V^o ? This is not problematic with a finite-horizon (it is calculated by backwards induction). With an infinite horizon, we adapt the machinery of Bellman operators.

Definition 3.5 (Bellman operator with sparse action) *Given a policy $a(z, V^p)$ from Definition 3.4, the Bellman operator $\mathcal{T}(V, V^p)$ is defined by:*

$$\mathcal{T}(V, V^p)(z) := u(a(z, V^p), z, \mu) + \beta \mathbb{E}[V(F(a(z, V^p), z, \mu))] \quad (29)$$

²²It could also be some approximation of the objective value function—much like in chess, people and computers use something akin to proxy value functions, that encode distant behavior. See Gabaix, Laibson, Moloche and Weinberg (2006) for an algorithm using an early version of this idea of proxy value function.

²³If there are domain conditions for $a(Z, V^p)$, we just add them to the definition of $a(Z, V^p)$, as in the static smax with budget constraints (Gabaix 2014).

Given this, (28) can be written as:

$$V^o = \mathcal{T}(V^o, V^p) \quad (30)$$

This operator $\mathcal{T}(V, V^p)$ has the usual good properties, recorded in the following lemma.²⁴

Lemma 3.6 (Monotone contraction) *The operator $\mathcal{T}(V, V^p)$ is a monotone β -contraction as a function of V . More explicitly: for any functions V, \tilde{V} (i) $\|\mathcal{T}(V, V^p) - \mathcal{T}(\tilde{V}, V^p)\|_\infty \leq \beta \|V - \tilde{V}\|_\infty$; (ii) if $V(z) \leq \tilde{V}(z)$ for all z , then $\mathcal{T}(V, V^p)(z) \leq \mathcal{T}(\tilde{V}, V^p)(z)$ for all z .*

This implies that a solution to (30) exists.

Lemma 3.7 (Existence of a value function) *Given Lemma 3.6, there is a unique solution V^o of the fixed point relation:*

$$V^o = \mathcal{T}(V^o, V^p).$$

The objective value function is the a unique fixed point V^o of the equation $V^o = \mathcal{T}(V^o, V^p)$.

Hence so far, given an original proxy value V^p , we defined a policy $a(z, V^p)$, which generates a new value function V^o . We can iterate the process. This generates the “ q -iterated” smax action.

Definition 3.8 (Action in sparse dynamic programming, q -iterated) *The basic (0-iterated) dynamic smax action is the action $a(z, V^p)$ from Definition 3.4. The q -iterated dynamic smax action is $a(z, V^{(q)})$, where $V^{(0)} = V^p$, and for $q > 1$, $V^{(q)}$ is characterized by $V^{(q)} = \mathcal{T}(V^{(q)}, V^{(q-1)})$.*

In some sense, the basic sparse max agent is naive about his future actions, while the q -iterated agent is more sophisticated. In practice, we will just take the basic action of Definition 3.4 in most problems of interest: this captures the essence of the economics, while keeping the model quite easy to use. Working through examples below will suggest that the higher-iterations are quite demanding in rationality.

With a finite horizon when $q \geq T - 1$, the q -iterated value $V^{(q)}$ satisfies:

$$V(z) = \operatorname{smax}_{a; m | m^d} \{u(a, z, m) + \beta \mathbb{E}V(F^z(a, z, \varepsilon_{t+1}, m))\}. \quad (31)$$

That is, we obtain the agent’s full value function. This is the same formulation as in the rational version, but with a smax rather than a max operator. In that formulation, the BR agent is very sophisticated (perhaps too much so) about his own future behavior: he sees how much he will see how much he will see (etc. – iterated q times) future inattention.

²⁴It is tempting to define $V = \mathcal{T}(V, V)$, but the operator $V \mapsto \mathcal{T}(V, V)$ does not (prima facie at least) satisfy the good properties of Lemma 3.6. Only $V \mapsto \mathcal{T}(V, V^p)$ does.

With finite horizon, the definition gives a construction of the value function by backward induction: starting from $V^T = 0$, we successively calculate V^{T-1}, \dots, V^0 (note that the time t is inside vector z).²⁵

3.5 Some Tools for Sparse Dynamic Programming

This subsection presents tools to compute BR dynamic programming. The reader is invited to skim it, read the main examples shown later, and then come back to it with those examples in mind. The results are proven for finite-horizon problems. I conjecture that they hold for infinite-horizon problems under some reasonable assumptions.²⁶

Taylor expansion of policy and value functions. We decompose the vector of state variables into: $z = (w, x)$, where w is a vector of variables that are fully taken into account in the default model (including possibly calendar time), while x is a vector of variables not taken into account in the default model. To capture this, I assume throughout the paper that $u(a, w, x, m)$ and $F(a, w, x, m)$ are independent of m when $x = 0$. I also assume throughout the paper that $F^x(a, w, x, m) = 0$ when $x = 0$, so small x 's at t generate small x 's at $t + 1$.

I will frequently assume the following “local autonomy of the disturbance” condition:

$$F_a^x(a, w, x) = 0 \text{ at } (a, x) = (a^d(w), 0) \tag{32}$$

where F_a^x is derivative with respect to a of the law of motion of x . This says that when $x = 0$, a small change in action a doesn't affect it directly, i.e. x is locally independent of the agent's actions. This is for instance the case for macroeconomic disturbances, e.g. if $x_{t+1} = \rho_x x_t + \varepsilon_{t+1}^x$, as we postulated for some variables above (e.g. (18)).

For simplicity, I assume that the functions are infinitely differentiable, and so is the attention function (the online appendix weakens those conditions in Section 16.1; the exposition is then heavier).

The main tools are the following.

Proposition 3.9 (Obtaining the sparse policy from the rational policy, to the first order). *Assume the local autonomy condition (32), and that $u(a, w, x, m) = u(a, w, m \odot x)$ and $F(a, w, x, m) = F(a, w, m \odot x)$. Consider the first order expansion of the optimal rational policy for small x ,*

$$a^r(w, x) = a^d(w) + \sum_i b_i(w) x_i + O(\|x\|^2)$$

²⁵I contemplated taking (31) as a starting point for dynamic programming, but it proved psychologically too complicated, and mathematically hard to handle, due to the unusual fixed point in V .

²⁶See Harris and Laibson (2001, 2013) for a related analysis showing the difficulty of the mathematics of infinite horizon dynamic programming when the agents misoptimize. Fortunately, those difficulties dissolve with a finite horizon. Still, extending the analysis here to an infinite horizon seems like an interesting, difficult problem.

Then, the sparse policy is, with ex-ante attention allocation:

$$a^s(w, x) = a^d(w) + \sum_i \tau \left(b_i(w), \frac{\kappa_a}{\sigma_i} \right) x_i + O(\|x\|^2) \quad (33)$$

and with ex-post attention allocation:

$$a^s(w, x) = a^d(w) + \sum_i \tau(b_i(w) x_i, \kappa_a) + O(\|x\|^2) \quad (34)$$

This proposition will be quite useful. To derive policies, first we can simply do a Taylor expansion of the rational policy around the default model, and then truncate term by term.

The following proposition indicates that to calculate the leading terms of the sparse policy, one can do a simple Taylor expansion around the default model, $V^d(w)$. One does not need to calculate explicitly the full rational policy $V^r(w, x)$.

Proposition 3.10 (Simple procedure to calculate the rational or sparse policy, to the first order) *Assume the local autonomy condition (32). To calculate the policies $a^s(w, x)$ and $a^r(w, x)$ up to second order terms, one can simply do a Taylor expansion around the default model with value function $V^d(w)$ (which gives V_x and V_{wx}), following the procedure outlined in Section 10.1. In particular, one does not need to fully solve the rational model.*

I now present some basic facts that are helpful in thinking about those policy functions.

Proposition 3.11 *For small x , we have: $V(w, x) = V^r(w, x) + R(w, x)$, where the residual $R(w, x) = O(\|x\|^2)$ for x close to 0. In other words, the sparse value function and the rational value functions differ only by second order terms in x .*

This basically generalizes the envelope theorem.²⁷ It implies that, at $x = 0$:

$$V_w = V_w^r, \quad V_{ww} = V_{ww}^r, \quad V_x = V_x^r \quad V_{wx} = V_{wx}^r \quad (35)$$

Typically $V_{xx} \neq V_{xx}^r$, however.

We have the following Lemma.²⁸

Lemma 3.12 (Close policies give very close value functions) *Suppose a policy $a(w, x)$ such that $a(w, x) = a^d(w) + O(\|x\|)$. Then, $V^o(w, x|a(\cdot)) = V^r(w, x) + O(\|x\|^2)$.*

This means that if the policy is approximately correct, up to first order terms, then the value function is approximated by the correct optimal policy, up to second order terms. This is again the envelope theorem.

²⁷The intuition is that as x is small, and the action is close to the optimum, we get only second order losses $O(\|x\|^2)$ from misoptimization.

²⁸Here $V^o(w, x|a(\cdot))$ is the policy induced by a , closely along the lines of Lemma 3.7.

Lemma 3.13 (Close proxy value functions give close policies) *Consider two proxy value functions $V^p, V^{p'}$ such that, $V^p(w, x) = V^{p'}(w, x) + O(\|x\|^2)$, and assume (32). Then, $a(w, x, V^p) = a(w, x, V^{p'}) + O(\|x\|^2)$.*

This lemma means that to know the optimal policy up to second order terms, we just need to have a proxy value function that is accurate up to second order terms. This is intuitive, but the proof reveals that condition (32) is required.

The above two lemmas imply that sophistication vs naiveté lead to the same policies, up to second order terms. Given also the naive / basic action ($q = 0$) is simpler, this is another reason for modeling agents with the basic policy.²⁹

Proposition 3.14 (Sophistication vs Naiveté makes only a second order difference in actions) *Consider the q -iterated agent of Definition 3.8. Suppose that the level 0 agent used a proxy value function V^p such that $V^p(w, x) = V^r(w, x) + O(\|x\|^2)$. Then, the actions at higher levels of sophistication $q \geq 0$ differ only by second order terms from the naive action: $a(Z, V^{(q)}) = a(Z, V^{(0)}) + O(\|x\|^2)$.*

For instance, in the life-cycle model of Section 2, I assumed basic ($q = 0$) agents, who project them selves are rational. If I had assumed partially (low $q > 0$) or fully sophisticated agents ($q \geq T - 1$), who keenly understand their future bounded rationality, Proposition 3.14 says that the consumption would have been the same, up to $O(x^2)$ terms.

4 Intertemporal Consumption: Behavioral Version

I now work out a few explicit examples, starting from very simple ones to build the intuition.

4.1 The life-cycle problem revisited

4.1.1 Detailed analysis with three periods

Here I revisit the life-cycle model of Section 2, in more detail. I take a 3-period version ($L = 2$, $T = 3$).³⁰ Utility is $\sum_{t=0}^2 u(c_t)$, and there is no discounting ($R = \beta = 1$). The agent starts with an endowment w_0 , has no regular income ($\bar{y} = 0$) and receives x at time 2; x is known at time 0, but the agent “may not think about it”. For instance, x could represent a negative income shock, such as a tax to pay, or a decrease in income as retirement. Calling w_t the wealth at the beginning of period t , the budget constraints at times $t = 0, 1, 2$ are: $w_1 = w_0 - c_0$, $w_2 = w_1 - c_1$, $0 = w_2 + x - c_2$.

²⁹Hyperbolic discounting naïve and sophisticated agents often act quite differently (Laibson 1997, O’Donoghue and Rabin 2001). However, in simple consumption-savings problems, the discrepancy between the two types’ consumption is still $O((1 - \beta)^2)$, where β is the hyperbolicity parameter, so quite modest if β is close to 1.

³⁰Section 10.3 gives more complements for the T period model.

How much attention will the agent pay to time-2 payment x ? First, let us observe that a rational agent smooths consumption: lifetime resources are $w_0 + x$ (initial wealth w_0 and time-2 payment x), and they should be consumed equally in the three periods:

$$c_t^r = \frac{w_0 + x}{3} \text{ for } t = 0, 1, 2$$

The corresponding dynamic policy (expressed with w_t rather than w_0) is:

$$c_0^r = \frac{w_0 + x}{3}, \quad c_1^r = \frac{w_1 + x}{2}, \quad c_2^r = w_2 + x.$$

e.g. at time 1, the life-time remaining resources ($w_1 + x$) should be divided equally among the two remaining periods.

I first state the behavioral policy, then derive it. The derivation is instructive.³¹

Proposition 4.1 *Take the 3-period life-cycle problem. The BR policy is*

$$c_0 = \frac{w_0 + m_0 x}{3}, \quad c_1 = \frac{w_1 + m_1 x}{2}, \quad c_2 = w_2 + x.$$

where the attention values are $m_0 = \mathcal{A}\left(\frac{1}{6\kappa}u''\left(\frac{w_0}{3}\right)\sigma_x^2\right)$ and $m_1 = \mathcal{A}\left(\frac{1}{2\kappa}u''\left(\frac{w_1}{2}\right)\sigma_x^2\right)$. If $|x|$ is not too large, they satisfy $m_0 \leq m_1 \leq 1$, i.e. the agent reacts more to near variables than distant variables.

Derivation of Proposition 4.1 We apply the smax procedure of Definition 3.4, using backward induction.

At time 2, the agent consumes all his disposable wealth:

$$V^2(w_2, x) = u(w_2 + x)$$

At time 1, the agent's problem is:

$$\text{smax}_{c_1; m_1} v^2(c_1, x, m_1) \text{ with } v^1(c_1, x, m_1) := u(c_1) + V^2(w_1 - c_1, m_1 x).$$

The FOC $v_{c_1}^1 = 0$ reads: $u'(c_1) = V_w^2(w_1 - c_1, m_1 x) = u'(w_1 - c_1 + m_1 x)$, so $c_1 = w_1 - c_1 + m_1 x$, and

$$c_1 = \frac{w_1 + m_1 x}{2} \tag{36}$$

The agent pays partial attention m_1 to the time-2 income x .

To calculate attention m_1 , we apply (11). Noting that $v_{cc}^2(c, x, m_1)|_{m_1=0} = 2u''(c_1^d)$, where

³¹I use the notation $\sigma_x^2 = \mathbb{E}[x^2]$ to indicate the prospective magnitude of x ; with ex post attention, we simplify have $\sigma_x^2 = x^2$.

$c^d = \frac{w_1}{2}$ is the optimal consumption with $m_1 = 0$, we have $m_1 = \mathcal{A} \left(\frac{1}{\kappa} 2u''(c_1^d) \text{var} \left(\frac{x}{2} \right) \right)$, so

$$m_1 = \mathcal{A} \left(\frac{1}{2\kappa} u'' \left(\frac{w_1}{2} \right) \sigma_x^2 \right). \quad (37)$$

At time 0, the agent does $\text{smax}_{c_0, m} v^0(c_0, x, m_0)$, with

$$v^0(c_0, w_0, x, m_0) := u(c_0) + V^{1,p}(w_0 - c_0, m_0 x) \quad (38)$$

where I use Definition 3.4, with $V^p = V^r$, the value function where the agent projects that he will be rational in the future:³²

$$V^{1,p}(w, x) = 2u \left(\frac{w+x}{2} \right). \quad (40)$$

The FOC is $v_{c_0}^0 = 0$ with, i.e. $u'(c_0) = V_w^{1,p}(w_0 - c_0, m_0 x) = u' \left(\frac{w_0 - c_0 + m_0 x}{2} \right)$, i.e. $c_0 = \frac{w_0 - c_0 + m_0 x}{2}$ and

$$c_0 = \frac{w_0 + m_0 x}{3} \quad (41)$$

To determine attention m_0 , we again use (11); we calculate:

$$v_{cc}^0 = u''(c^d) + V_{ww|m=0}^{1,p} = u''(c^d) + \frac{1}{2} u''(c^d) = \frac{3}{2} u''(c^d)$$

so that $m_0 = \mathcal{A} \left(\frac{1}{\kappa} v_{cc}^0 \text{var}(\partial_{m_0} c_0) \right) = \mathcal{A} \left(\frac{1}{\kappa} \frac{3}{2} u''(c^d) \text{var} \left(\frac{x}{3} \right) \right)$, i.e.

$$m_0 = \mathcal{A} \left(\frac{1}{6\kappa} u'' \left(\frac{w_0}{3} \right) \sigma_x^2 \right). \quad (42)$$

The consumptions are:

$$c_0(w_0, x) = \frac{w_0}{3} + \frac{m_0}{3} x, \quad c_1(w_0, x) = \frac{w_0}{3} + \left(\frac{m_1}{2} - \frac{m_0}{6} \right) x, \quad c_2(w_0, x) = \frac{w_0}{3} + \left(1 - \frac{m_0 + m_1}{2} \right) x. \quad (43)$$

Comparing (37) and (42), we see that $m_0 \leq m_1$ iff $\frac{1}{6} |u''(\frac{w_0}{3})| \leq \frac{1}{2} |u''(\frac{w_1}{2})|$. When $x = 0$, this is automatically verified, as $\frac{w_0}{3} = \frac{w_1}{2}$. Hence, we have $m_0 \leq m_1$ iff x is not too large.³³

□

³²In terms of the iterated sparse max of Definition 3.8, this is the value function under procedure with $q = 0$ iteration, t. If we used the procedure with $q = 1$ -iteration, the agent uses the objective value function, which is:

$$V^1(w_1, x) = u \left(\frac{w_1 + m_1(w_1)x}{2} \right) + u \left(\frac{w_1 + (2 - m_1(w_1))x}{2} \right) \quad (39)$$

As per the envelope theorem, the two value functions differ only by second order terms: $V^1(w_1, x) = V^{1,p}(w, x) + O(x^2)$.

³³If x was very large and positive, we could have the following effect: the agent realizes at time 1 that he's actually quite wealthy, so pays less attention. This effect needs a very large x , so is not operative in most situations.

4.1.2 A few features of a sparse agent

This example and the life-cycle model of Section 2 illustrates a few general features.

Sparse agents are globally patient like rational agents, but still myopic to a variety of small future shocks. Indeed, agents here invest their wealth w_t very patiently, exactly like rational agents: they fully smooth it over the horizon. At the same time, they tend to be myopic about future small shocks (the time-2 shock x). In other terms, in the present model, agents are only partially myopic (e.g. don't react to a scheduled increase in taxes). This behavior cannot be captured with a model simply assuming a low discount factor β .

The Euler equation fails. The Euler equation holds under the BR-perceived consumption, but not under the actual consumption. For instance, at time 1, if $m_1 = 0$, then the agent expects to consume $(c_1, c_2) = (\frac{w_1}{2}, \frac{w_1}{2})$, but actually consumes $(c_1, c_2) = (\frac{w_1}{2}, \frac{w_1}{2} + x)$. The traditional Euler equation basically only holds if agents are exactly rational (Hall 1978), so it is a fragile way to model agents. In contrast, sparse consumption functions are a more robust way of modelling them.

The first and second welfare theorems fail. This 3-period example features a simple economy, “production” being the linear storage technology allowing to transfer the good 1-for-1 across periods. Here, we do not have a Pareto-optimum, as the agent fails to maximize.³⁴ Likewise, if the first welfare theorem fails, typically the second welfare theorem fails. However, some optimum tax policy can generally restore efficiency.

Agents react more to “near” shocks than to “distant” shocks (in math, $m_0 \leq m_1$). The main reason is that, normatively, the shock x should impact c_0 as $\frac{x}{3}$ ($\frac{\partial c_0(w_0, x)}{\partial x} = \frac{1}{3}$), while it should impact c_1 as $\frac{x}{2}$ ($\frac{\partial c_1(w_1, x)}{\partial x} = \frac{1}{2}$). Hence, attention to the last period shock x is lower at earlier dates ($t = 0$) than at late dates ($t = 1$).

This example suggests a few interesting variants. I discuss some of them in Section 7.1.

4.2 Consumption-savings with infinite horizon

Here I propose a version of the permanent income consumption problem of Section 3.2.

Formalism. We have x_t the state vector of disturbances, which follows a process: $x_{t+1} = F^x(x_t, \varepsilon_{t+1}^x)$. The interest rate and income deviations from their means are a linear function of that state vector: $\hat{r}_t = k^r \cdot x_t$, $\hat{y}_t = k^y \cdot x_t$ for some vectors k^r, k^y . The perceived law of motion for wealth is kept throughout as (21). However, the perceived law of motion for x_t can be very general, e.g. $x_{t+1} = A(m)x_t + \sigma(m)\varepsilon_{t+1}$ for some parametrizations of the matrix and stochasticity.

As a concrete example, let us review the basic example with the full formalism (which the reader

³⁴In Gabaix (2014, Proposition 8) the first welfare theorem (in a static economy) does hold if agents have the same misperceptions: the reason is that in that setup, the agent is assumed to fully perceive consumptions (a defensible assumption in a static context) – whereas in the present model he doesn't. For instance, at time 1 the fully naïve agent projects that $c_1 = c_2 = \frac{w_1}{2}$, whereas in fact $c_2 = \frac{w_1}{2} + x$.

is encouraged to skip at first). The state vector is:³⁵

$$z_t = (w_t, \hat{r}_t, \hat{y}_t, \rho_r, \sigma_r, \rho_y, \sigma_y) \quad (44)$$

Here σ^y parametrizes the stochasticity of income, which may be set to zero in a simplified model. The simplifiable meta-model is: $\hat{z}_{t+1} = F^z(c_t, z_t, \varepsilon_{t+1}, m)$, where the components are: for wealth, (21):

$$w_{t+1} = F^w(c_t, z_t, \varepsilon_{t+1}, m) := (1 + \bar{r} + m_r \hat{r}_t)(w_t - c_t) + \bar{y} + m_y \hat{y}_t,$$

for the interest rate and income, using the notation X a shorthand for r or y ,

$$\hat{X}_{t+1} = F^X(c_t, z_t, \varepsilon_{t+1}, m) := (m_{\rho_X} \rho_X + (1 - m_{\rho_X}) \rho_X^d) \hat{X}_t + (m_{\sigma_X} \sigma_X) \varepsilon_{t+1}^X$$

for parameters $z_i = \rho_r, \sigma_r, \rho_y, \sigma_y$,

$$z_{i,t+1} = F^{z_i}(c_t, z_t, \varepsilon_{t+1}, m) := m_{z_i} z_{it} + (1 - m_{z_i}) z_i^d.$$

For instance, if $m_{\sigma_y} = 0$, the agent projects income to have 0 stochasticity in the future.

Results. I start with a simple lemma describing the rational policy, using Taylor expansions.³⁶ To signify “up to second order term”, I use the notation $O(\|x\|^2)$, where $\|x\|^2 := \mathbb{E}[\hat{y}_\tau^2] / \bar{y}^2 + \mathbb{E}[\hat{r}_\tau^2] / \bar{r}^2$ (the constants \bar{y}, \bar{r} are just here to keep valid units). Recall that $R := 1 + \bar{r}$.

Lemma 4.2 (Traditional rational consumption function) *In the rational policy, the optimal consumption is: $c_t = c_t^d + \hat{c}_t$, with $c_t^d = \frac{rw_t + \bar{y}}{R}$ and*

$$\hat{c}_t = \mathbb{E}_t \left[\sum_{\tau \geq t} \frac{b_r(w_t) \hat{r}_\tau + b_y \hat{y}_\tau}{R^{\tau-t+1}} \right] + O(\|x\|^2), \quad (45)$$

$$b_r(w_t) := \frac{\bar{r}}{R} (w_t - \bar{y}) - \psi c^d, \quad b_y := \frac{\bar{r}}{R}. \quad (46)$$

The associated value function is

$$V^r(w_t, x_t) = \frac{1}{\mu_t} u(\mu_t w_t + \nu_t) + O(\|x\|^2) \quad (47)$$

$$\mu_t := \frac{\bar{r}}{R} + (1 - \psi) \frac{\bar{r}}{R} \mathbb{E}_t \left[\sum_{\tau \geq t} \frac{\hat{r}_\tau}{R^{\tau-t+2}} \right], \quad \nu_t := \bar{y} + \mathbb{E}_t \left[\sum_{\tau \geq t} \frac{\frac{-\psi \bar{y}}{R} \hat{r}_\tau + b_y \hat{y}_\tau}{R^{\tau-t+1}} \right]$$

³⁵If $\rho_r, \sigma_r, \rho_y, \sigma_y$ are correctly perceived, then we can just take $z_t = (w_t, \hat{r}_t, \hat{y}_t)$.

³⁶In independent work, Auclert (2015) contains a similar Taylor expansion for the rational case, with more general assets.

Consumption reacts to future interest rates and income changes, according to the usual income and substitutions effects (multiplied by ψ).³⁷

The behavioral policy is then as follows.

Proposition 4.3 (Sparse consumption function) *In the behavioral model, $c_t = c_t^d + \hat{c}_t$, with $c_t^d = \frac{rw_t + \bar{y}}{R}$ and*

$$\hat{c}_t^s = \mathbb{E}_t^s \left[\sum_{\tau \geq t} \frac{b_r(w_t) m_r \hat{r}_t + b_y m_y \hat{y}_\tau}{R^{\tau-t+1}} \right] + O(\|x\|^2) \quad (48)$$

where \mathbb{E}_t^s is the expectation taken with respect to the agent's beliefs under the subjective model.

The policy of the behavioral agent is the policy of a rational agent, under a subjective model, and with partial inattention to income and interest rate.

Now, I record the important case of AR(1) processes.

Lemma 4.4 *In the AR(1) model (18)-(19), the rational policy is:*

$$c^r(w_t, \hat{y}_t, \hat{r}_t) = c^d(w_t) + B_r(w_t) \hat{r}_t + B_y \hat{y}_t + O(\|x\|^2) \quad (49)$$

$$B_r(w_t) := \frac{b_r(w_t)}{R - \rho_r}, \quad B_y := \frac{b_y}{R - \rho_y}. \quad (50)$$

and the value function as in (47) with

$$\mu_t = \frac{\bar{r}}{R} + (1 - \psi) \frac{\bar{r}}{R^2} \frac{\hat{r}_t}{R - \rho_r}, \quad \nu_t = \bar{y} + \frac{-\psi \bar{y}}{R} \frac{\hat{r}_t}{R - \rho_r} + b_y \frac{\hat{y}_t}{R - \rho_y}$$

Proposition 4.5 *Under the AR(1) model, the behavioral policy is:*

$$c^s(w_t, \hat{y}_t, \hat{r}_t) = c^d(w_t) + B_r^s(w_t) \hat{r}_t + B_y^s \hat{y}_t + O(\|x\|^2) \quad (51)$$

with $B_r^s(w_t) = \frac{b_r(w_t) m_r}{R - \rho_r(m)}$ and $B_y^s = \frac{b_y m_y}{R - \rho_y(m)}$.

Endogenizing attention, we have $m_X = \mathcal{A} \left(\frac{\sigma_X^2 B_X(w_t)^2 |v_{cc}|}{\kappa} \right)$, for $X = r$ and $X = y$ as in (11).³⁸ This shows a “feature-by-feature” truncation. It is useful because it embodies in a compact way the policy of a sparse agent in quite a complicated world. Note that the agent can solve this problem without solving the 3-dimensional (and potentially 21-dimensional, say, if there are 20

³⁷The term μ_t is the MPC, and contains substitutions and income effects from wealth before receiving the innovation in income, $w_t - \bar{y}$. The term ν_t contains the baseline income \bar{y} , expected value of future income \hat{y}_t , and substitution effects interest rates associated with the income stream of \bar{y} per period. There is no income effect for \bar{y} , as no interest rate was going to accrue to \bar{y} .

³⁸Call $v(c) = u(c) + V^r((1 + \bar{r} + \hat{r}_t)(w_t - c) + \bar{y} + \hat{y}_t)$. I use the well-known fact that in the limit of small time intervals, $v_{cc} = u_{cc}$. See e.g. Section 13.1.




Figure 4: This figure shows the attention to the interest rate (m_r) and to income (m_y) as a function of the cost of thinking, $\bar{\kappa}$.

state variables besides wealth) problem. Only local expansions and truncations are necessary.³⁹ Attention to the interest will become higher if the interest rate is very volatile (σ_r high, e.g. as in hyperinflations).⁴⁰

Numerical illustration *Numerical illustration.* To get a feel for the effects, consider a calibration with (using annual units): $\gamma = 1$, $\bar{r} = 5\%$, $w_t = 2c_t^d$, $c_t^d = 1$, $\sigma_r = 0.8\%$, $\sigma_y = 0.2$, $\rho_y = 0.95$, $\rho_r = 0.7$: as income shocks (roughly corresponding to “carrier risk”) are persistent, they are important to the consumer’s welfare.⁴¹ I use the τ_1 truncation function. I parametrize $\kappa = \bar{\kappa}^2 |u''(c_t^d)|$, as in Section 2 (see also Section 10.2.2).

Figure 4 shows the attention to the interest rate and income. At $\kappa = 0$, the agent is fully rational, attention is 1. For higher $\bar{\kappa}$, attention drops, more so for the interest rate. For $\bar{\kappa}$ above 0.02, the agent stops thinking about the interest rate. Because the interest rate makes him change his consumption by less than 2% on average, it’s optimal for him not to think about the interest rate. Still, he thinks a lot about his income. Attention to current income falls when $\bar{\kappa} \simeq 0.1$, i.e. then the agent truncates features that make him change his consumption by less than 10%. Note that this “source-specific” selective attention could not be rationalized by just a fixed adjustment cost to consumption, which affect all causes of consumption change equally.

The same reasoning holds in every period. The above describes a practical way to do sparse dynamic programming. In some cases, this is simpler than the rational way (as the agent does not

³⁹Note here that the active decision is that of consumption, not savings. For most variables (except current income), it does not matter: the impact of interest rates, future taxes, future income shocks etc. are the same whether a sparse agent uses the consumption frame or saving frame. See Section 12.7 for a discussion: the consumption frame is arguably better, as it yields great utility when $\phi_y > r$, income shocks are less than extremely persistent.

⁴⁰Likewise, in the model, a wealthy rentier, with much financial wealth but little labor or state income, will endogenously pay more attention to the interest rate, as it is more important for her consumption (and will bemoan at low interest rates, which lead her to cut on consumption).

⁴¹This numerical example is a rough AR(1) rendition of the “permanent-transitory models” in the literature.

need to solve for the equilibrium), and it may also be more sensible.

This section has developed the sparsely behavioral version of a basic machine of macroeconomics: the consumption function of an agent in a world of dynamic interest rates and income shocks. This is useful, because from this machine can be used in a host of situations. I turn to two such situations before turning its use in general equilibrium (Section 5).

4.3 Application: Ricardian Equivalence

Intuitively, a sparse agent will violate Ricardian equivalence as he is partially myopic. To see more precisely, I keep interest rates constant, and call \hat{g}_t the transfers from the government to the agent at time t . To see this, call w_{t-} the financial wealth at the beginning of the period (before government transfers) and specialize Proposition 4.3, setting $\hat{y}_\tau := \hat{g}_{\tau+1}$. We obtain the consumption:

$$\hat{c}_t^s = b_y (w_{t-} + \hat{g}_t) + \sum_{\tau>t} b_y m_g \frac{\hat{g}_\tau}{R^{\tau-t}} \quad (52)$$

Suppose that the government gives \hat{g} to the agent at t , and $-\hat{g}R^T$ in T periods. A rational consumer would not change consumption, as the present value is unchanged (Barro 1974). However, a BR consumer increases consumption by $\frac{\bar{r}}{R} (1 - m_g) \hat{g}$: the positive shock increases it by $\frac{\bar{r}}{R} \hat{g}$, and the negative shock decreases it by $\frac{\bar{r}}{R} m_g \hat{g}$.

Proposition 4.6 (Failure of Ricardian equivalence) *A behavioral consumer increases consumption at t by $\frac{\bar{r}}{R} (1 - m_g) \hat{g}$. This way, Ricardian equivalence does not hold, unless attention is full. The further away the increase in taxes (keeping their present value constant), the lower the reaction.*

Section 12.6 develops in details the dynamics of attention to future taxes: in particular, how attention becomes higher close to the tax hike.

4.4 Application: Lucas Critique

With sparse agents, the Lucas (1976) critique stops applying – or at least applies less. Lucas’ critique is that if the parameters of the world (e.g. tax) change, people’s reactions functions will change. However, with sparse agents, this is not true: when changes are small, agents choose keep their default “policy functions”. For instance, suppose that there is a small, temporary consumption tax \mathcal{T}_t , in the permanent income model above, so that the perceived law of motion for wealth becomes:

$$w_{t+1} = F^w (c_t, z_t, m) := (1 + \bar{r} + m_r \hat{r}_t) (w_t - (1 + m_\mathcal{T} \mathcal{T}_t) c_t) + \bar{y} + m_y \hat{y}_t$$

If the tax is small, agent will pay 0 attention to it ($m_\mathcal{T} = 0$), and their policy function will not change. The aggregate outcome will change, because agents will be poorer – but the policy function will not change.

5 Neoclassical Growth Model: A Behavioral Version

I now move on to general equilibrium, and study a behavioral version of the baseline neoclassical growth model, the Cass-Koopmans model.⁴²

5.1 Setup

The utility function is still $\mathbb{E} \sum_{t=0}^{\infty} \beta^t \frac{C_t^{1-\gamma}}{1-\gamma}$, and we again call $\psi = \frac{1}{\gamma}$. The aggregate capital stock follows:

$$K_{t+1} = f(K_t, L) + (1 - \delta) K_t - C_t + \varepsilon_{t+1} \quad (53)$$

where ε_{t+1} are mean-zero shocks. This way, there is just one state variable in the economy, the capital stock. In the most basic neoclassical model, ε_{t+1} is always 0, and L is fixed.⁴³ I define $f(K) := f(K, L) - \delta K$, which is output net of the capital depreciation at the fixed labor supply, so that $K_{t+1} = f(K) + K_t - C_t + \varepsilon_{t+1}$.

This is a textbook example, which introduces generations of students to macroeconomics (Blanchard-Fischer (1989, Chapter 2), Acemoglu (2009, Chapter 8), Romer (2012, Chapter 2)). In that tradition, we have infinitely-rational forward looking agents that calculate the whole macroeconomic equilibrium in their heads. They are supposed to use the dynamics:

$$\frac{\dot{c}_t}{c_t} = \frac{1}{\psi} (f'(K_t) - \rho), \quad \dot{K}_t = f(K_t) - c_t$$

and given K_0 , find the unique c_0 that leads to a non-explosive path (the saddle path, see Figure 5). The psychology of that is arguably quite alien to any human's intuition. Still, it's on that strange core model that much of dynamic macroeconomics rests. Hence, it is useful to develop an alternative to that model, something I now do.

Let us first review the mechanics of convergence. If there were no shocks, the economy would be at the steady state, with capital stock K^* . I use the hat notation for deviations from the mean, e.g. $\widehat{K}_t = K_t - K^*$. The law of motion for capital (53) is, in linearized form:

$$\widehat{K}_{t+1} = (1 + r) \widehat{K}_t - \widehat{C}_t + \varepsilon_{t+1} \quad (54)$$

where r is the steady state interest rate, $1 + r = f'(K_*)$.

As there is one state variable, the linear policy function of the agent (rational or not) is:

$$\widehat{C}_t = b \widehat{K}_t \quad (55)$$

⁴²For simplicity, I remove here the "growth" part, but it is straightforward to add.

⁴³I keep the noise ε_{t+1} because it makes the model a true DSGE model, and because it yields a scale for typical business-cycle fluctuations that it useful to determine attention.




Figure: Transitional dynamics in the baseline neoclassical growth model

Figure 5: This Figure shows the traditional approach to the neoclassical growth model. The agent is supposed to find the unique saddle path leading to non-explosive dynamics. Arguably, this is psychologically quite absurd. The present paper proposes a more behavioral approach. Source: Acemoglu (2009).

for some constant b to be determined.

Plugging this into (54) we obtain: $\widehat{K}_{t+1} = (1 + r - b) \widehat{K}_t + \varepsilon_{t+1}$, i.e.

$$\widehat{K}_{t+1} = (1 - \phi) \widehat{K}_t + \varepsilon_{t+1} \quad (56)$$

where ϕ is the speed of mean-reversion:

$$\phi = b - r. \quad (57)$$

When agents are more reactive to shocks (when b is higher), the economy mean-reverts faster to the steady state (ϕ is higher).

Squaring equation (56), we obtain: $var \widehat{K}_{t+1} = (1 - \phi)^2 var \widehat{K}_t + \sigma_\varepsilon^2$. As in the steady state, $var \widehat{K}_{t+1} = var \widehat{K}_t$, $var \widehat{K}_t = \frac{\sigma_\varepsilon^2}{1 - (1 - \phi)^2}$, i.e., in the limit of small time intervals:

$$\sigma_K = \frac{\sigma_\varepsilon}{\sqrt{2\phi}} \quad (58)$$

When shocks mean-revert more slowly (lower ϕ), the average deviation of the stock price from trend is higher (shocks “pile up” more).

The steady state. *Rational agent.* The rational agent has a value function $V(K_t)$, which satisfies:

$$V(K) = \max_c \{u(c) + \beta \mathbb{E}[V(K + f(K) - c + \tilde{\varepsilon})]\} \quad (59)$$

The steady state is at $K = K_*$, $C_t = C_*$ with:⁴⁴

$$\beta(1 + \bar{r}) = 1 \tag{60}$$

which determines K_* , the gross interest rate $1 + \bar{r} = f'(K_*)$ and consumption is $C_* = F(K_*)$.

Behavioral agent.

I consider here the decentralized approach. In a population with mass 1 and aggregate capital stock K_t , I consider an individual represent agent, who has his own wealth k_t , under control. In equilibrium $k_t = K_t$. But when deciding on his own consumption, and agent is an infinitesimal price taker, and takes the macro variables (K_t) as unaffected by his own actions.

Suppose that the agent has a consumption rule that's consistent with a steady state growth, i.e. $c_t = \bar{y} + \mu k_t$, where the MPC is $\mu = 1 - \beta$. By the analytics above (equation (20)), this is the optimal policy in the steady state.⁴⁵ Then, wealth evolves as $k_{t+1} = (1 + \bar{r})(k_t + \bar{y} - c_t) = (1 + \bar{r})(1 - \mu)k_t$. So, at a steady state, \bar{r} adjusts so that $1 = (1 + \bar{r})(1 - \mu) = (1 + \bar{r})\beta$. This gives the same condition (60), hence the same steady state as in the rational case.

The behavioral agent here inherits the same steady state as the neoclassical agent, with $\beta(1 + \bar{r}) = 1$ at the steady state. Only the dynamics around the steady state are different. I now turn to them. For simplicity, I use the notation r for the steady state interest rate \bar{r} .

5.2 Convergence: Boundedly Rational Version

The agent's wealth evolves as:

$$k_{t+1} = (1 + r_t)(k_t + y_t - c_t)$$

where $y_t = f(K_t) - K_t f'(K_t)$ is labor income, and $r_t = f'(K_t)$ is the interest rate. Taylor expansions around the steady state yield:

$$\hat{r}_t = f''(K^*) \hat{K}_t, \quad \hat{y}_t = -K^* f''(K^*) \hat{K}_t \tag{61}$$

Rather than positing that the agent correctly sees a whole saddle path, I posit that in the agent's model, capital evolves with a subjective speed of mean-reversion ϕ^s : $\mathbb{E}_t [\hat{K}_{t+1}] = (1 - \phi^s) \hat{K}_t$. I again parametrize it as:

$$\phi^s = (1 - m_\phi) \phi^d + m_\phi \phi$$

where ϕ is the true speed of mean-reversion (which will be determined as an endogenous outcome) and ϕ^d is a default value – perhaps coming from some empirical experience, saying for instance that “business cycles” have a half-life of a few years.

For simplicity, I posit that the agent pays the same attention m_K to all macro variables.⁴⁶

⁴⁴The traditional proof is by taking the derivative w.r.t. K in (59), so $V'(K_*) = \beta(1 + F'(K_*))V'(K_*)$.

⁴⁵Here wealth is taken at the beginning of the period, before receiving labor income.

⁴⁶It is easy to generalize to the case where people's attention to income and the interest rate differ.

Lemma 4.4 reads here (using the continuous time notation):

$$\widehat{c}_t = r\widehat{k}_t + \frac{r}{r + \phi^s} m_K \widehat{y}_t + \frac{rk^* - c^*\psi}{r + \phi^s} m_K \widehat{r}_t$$

and using (61) and (63) we obtain:

$$\widehat{c}_t = r\widehat{k}_t + \frac{-rK^* f''(K^*) + (rk^* - c^*\psi) f''(K^*)}{r + \phi^s} m_K \widehat{K}_t$$

i.e., as $k^* = K^*$,

$$\widehat{c}_t = r\widehat{k}_t + \frac{\xi}{r + \phi^s} m_K \widehat{K}_t \quad (62)$$

with

$$\xi := -\psi C_* f''(K_*) \geq 0 \quad (63)$$

Hence, in the aggregate $\widehat{C}_t = b\widehat{K}_t$ with $b = r + \frac{\xi}{r + \phi^s} m_K$, and using (57) we obtain the actual speed of mean-reversion of the economy:

$$\phi = m_K \frac{\xi}{r + \phi^s}. \quad (64)$$

Rational agent. The rational agent has a correct model ($\phi^s = \phi$, $m_K = 1$). So the speed of mean-reversion is $\phi = \phi^r$; (64) gives

$$\phi^r = \frac{\xi}{r + \phi^r} \quad (65)$$

whose solution is $\phi^r = \frac{-r + \sqrt{r^2 + 4\xi}}{2}$.

Behavioral agent. Let us now endogenize m_ϕ and m_K . Given (62), $\frac{\partial \widehat{c}_t}{\partial \phi} |_{m_K=0} = 0$: at the default model, he doesn't react to ϕ , the speed of mean reversion of aggregate variables, given he's not even thinking about aggregate variables (here, just capital). This means (using equation (11)) that $m_\phi = 0$: the optimal refinement in thinking about the speed of mean reversion is 0. Hence, $\phi^s = \phi^d$, the projected speed of mean-reversion (ϕ^s) is the default one (ϕ^d).⁴⁷ Given (64), this implies that the actual speed of mean-reversion is:

$$\phi = m_K \phi_0, \quad \phi_0 := \frac{\xi}{r + \phi^d}. \quad (66)$$

Note that in the rational expectations case $\phi^d = \phi$, which remains a useful benchmark. The next Proposition studies the more general case, in which $\phi^d \geq \min(\phi, \phi^r)$.

⁴⁷With the iterated sparse max of Definition 15.1 of the online appendix, we can have $m_\phi > 0$.

Proposition 5.1 (Fluctuations are larger and more persistent with sparse agents) *Suppose that $\phi^d \geq \min(\phi, \phi^r)$ and $\kappa > 0$. Then $\phi < \phi^r$: in the sparse economy, GDP fluctuations have slower mean-reversion and higher variance than in the rational economy.*

Consumption is less volatile in the sparse than in the rational economy, if $\phi \geq r$, which seems empirically valid.

To get a quantitative expression for ϕ , I use the attention function $\mathcal{A}_1(v) = \max(1 - \frac{1}{v}, 0)$ from (15).

Proposition 5.2 *The speed of mean-reversion of the economy is: $\phi = \frac{\phi_0}{1 + \frac{B}{\phi_0}}$ with $\phi_0 := \frac{\xi}{r + \phi^d}$, $B := \frac{2\kappa}{|u_{cc}|\sigma_\varepsilon^2}$. It is lower (and aggregate fluctuations are larger) when agents' bounded rationality κ is stronger.*

There are two conclusions from this analysis.

First, bounded rationality generates more volatile and persistent fluctuations.

Second, and more importantly, we have a structure that gives an alternative to the rational general equilibrium model. In addition, the worldview of the agent is arguably more sensible. In the behavioral model, the agent pays only partial attention to aggregate state variables and has a simplified understanding of their dynamics: this is, arguably, how most people orient themselves, without a full, structural model of the causes of aggregate dynamics (which no one knows anyway, including economists). For instance, they say “healthy business cycles typical mean-revert in about 4 years” (which corresponds here to taking $\phi^d = \frac{1}{4}$), and they forecast based on that benchmark. People can act with that simplified understanding.

Of course, this does not say that the model does make better predictions – though the generation of higher volatility of GDP is intriguing. It does mean that we have a modelling structure that could be promising in future quantitative work.

5.3 Generalization: Recursive Competitive Equilibrium with Behavioral Agents

Here I define the behavioral extension of the Prescott-Mehra (1980) recursive competitive equilibrium – which is a formalization of the intuitive concept of dynamic equilibrium. It formalizes what we did in Section 5.2. It is more of methodological interest. It shows the “template” of dynamic equilibrium with behavioral in other contexts. The key issue is that the agents' behavior affect aggregate future outcomes, but agents may have only a simplified understanding of that.

We are given a steady state capital \bar{K} , a default law of motion $K' = G^d(K)$, with $G^d(K) = \bar{K} + \phi^d(K - \bar{K})$. A given m gives perceived laws of motions for private wealth $k' = F^k(c, k, K, m)$ and for the aggregate variables $K' = G(K, m)$.⁴⁸ In turn they generate a default value function

⁴⁸The present paper offers a systematic way in which agents simplify the moments they use to think about general

$V^r(k, K | F^k(\cdot, m), G(\cdot, m))$, which is the rational value function under those laws of motion. The proxy value function chosen is then: $V^p = V^r$ or, depending on the context, a second-order approximation of it.⁴⁹

In the context of the Cass-Koopmans above the definition is as follows. This template can be generalized quite easily to more state variables, several types, flexible labor supply etc. This mostly just takes more space (see Ljungqvist and Sargent, Chapter 12).

Definition 5.3 (Behavioral recursive competitive equilibrium) *A behavioral recursive competitive equilibrium is*

1. A set of prices $r(K), \omega(K)$, and perceived prices $r(K, m), \omega(K, m)$ with $m := (m_K, m_K^2, m_{K'}, m_\varepsilon)$, with:

$$r(K, m) = \bar{r} + m_{K'} r'(\bar{K})(K - \bar{K}) + m_{K^2} [r(K) - \bar{r} - r'(\bar{K})(K - \bar{K})] \quad (67)$$

and the same for $\omega(K, m)$.

2. A set of policies: consumption $c(k, K)$ and firms's factor demand: $k^D(K), L^D(K)$
3. A perceived law of motion for private capital: $k' = F^k(c, k, K, m)$

$$F^k(c, k, K, m) := (1 + r(K, m))(k - c + \omega(K, m)\bar{L}) + m_\varepsilon \varepsilon_{t+1}$$

4. An actual law of motion for the aggregate capital: $K' = G(K)$
5. A perceived law of motion for aggregate capital: $K' = G(K, m) = m_{K'} G(K) + (1 - m_{K'}) G^d(K)$ such that
6. The agent chooses consumption sparsely and optimally:

$$c(k, K, m) = \arg \operatorname{smax}_{c, m} \{u(c) + \beta V^r(F^k(c, k, K, m), G(K, m) | F^k(\cdot, m), G(\cdot, m))\}$$

where $V^r(k, K | F^k(\cdot, m), G(\cdot, m))$ is the rational value function associated with the perceived law of motions F^k and G .

7. Firms maximize profit: $k^D, L^D \in \arg \max_{k^D, L^D} \{f(k^D, L^D) - r(K)k^D - \omega(K)L^D\}$
8. Markets clear: $k^D(K) = K, L^D(K) = \bar{L}$
9. The actual law of motion is consistent with individual choice: $G(K) = f(K, \bar{L}) - c(K, K) + \varepsilon_K$.

equilibrium. This could be useful to systematize procedures such as the one in Krusell and Smith (1998). See also Barberis, Shleifer and Vishny (1998) and Hong, Stein and Yu (2007) for an early analyses of financial equilibrium when agents use a simplified model.

⁴⁹I.e., $V^p(k, K | F(\cdot, m), G(\cdot, m)) = V^r(k, K | F(\cdot, m), G(\cdot, m)) + O(|K - \bar{K}|^2)$.

To interpret (67), consider the case where the agent doesn't pay attention to nonlinear terms ($m_{K^2} = 0$),⁵⁰ doesn't pay attention to stochastic noise in the future ($m_\varepsilon = 0$) and uses a simplified, AR(1) view of aggregate dynamics ($m_{K'} = 0$). Then, his model of the world is exactly as in Section 5.2. His proxy value function $V^r(k, K | F^k(\cdot, m), G(\cdot, m))$ is the proxy value function under that simplified model, where everything is an AR(1). The value function was already calculated in Lemma 4.4, using $\rho_r = \rho_y = 1 - \phi$.

6 Behavioral Version of a Few Other Models

To make sure that the model is widely applicable, I developed a behavioral version of a few other important machines of dynamic economics. I summarize them here, while the online appendix has the specifics. I start with the Merton problem, which illustrates a case where risk matters for the agent's decision.

6.1 Merton Portfolio Choice Model

A core model in dynamic finance is the Merton (1971) portfolio choice problem.⁵¹ The agent's utility is: $\mathbb{E} \left[\frac{1}{1-\gamma} \int_0^\infty e^{-\rho s} c_s^{1-\gamma} ds \right]$, and his wealth w_t evolves according to:

$$dw_t = (-c_t + rw_t) dt + w_t \theta_t (\pi_t dt + \sigma dZ_t)$$

where θ_t is the allocation to equities. The equity premium $\pi_t = \bar{\pi} + \hat{\pi}_t$ has a variable part $\hat{\pi}_t$, which follows

$$d\hat{\pi}_t = -\Phi \hat{\pi}_t dt - \chi_t \sigma dZ_t^1 + \sigma'_\pi dZ_t^2$$

where the return is $d\tilde{r}_t = (r_t + \pi_t) dt + \sigma dZ_t^1$. The parameter $\chi_t \geq 0$ indicates that equity returns mean-revert: good returns today lead to lower returns tomorrow. That will create a hedging demand term. I again call $\psi = \frac{1}{\gamma}$ the intertemporal elasticity of substitution.

So, in theory, the agent should take into account both the variable equity premium $\hat{\pi}_t$ and also the "hedging demand," the demand due to the fact that equity returns mean-revert.

Proposition 6.1 (Behavioral dynamic portfolio choice) *The fraction of wealth allocated to equities is, with $\theta_* := \frac{\bar{\pi}}{\gamma\sigma^2}$:*

$$\theta_t^s = \theta_* + \tau \left(\frac{\hat{\pi}_t}{\gamma\sigma^2}, \kappa \right) + \tau \left(\frac{H_t}{\gamma\sigma^2}, \kappa_\theta \right)$$

⁵⁰Here m_{K^2} is the attention to non-linear terms. If $m_{K^2} = 0$, the agent simply linearizes the impact of the state variable on the interest rate. When using the plain sparse max (as opposed to the iterated sparse max of Definition 15.1), the algorithm will select $m_{K^2} = 0$, i.e. remove the nonlinear terms.

⁵¹Section 12.1 of the online appendix studies a more general Merton's problem with many assets and factors.

while consumption is: $c_t^s = \mu_t^s w_t$ with

$$\mu_t^s = \mu_* + \tau \left((1 - \psi) \frac{\mu_*}{\mu_* + \Phi} \theta_* \widehat{\pi}_t, \kappa_c/w \right) + \tau \left((1 - \psi) \theta_* H_t, \kappa_c/w \right)$$

where H_t is the hedging demand term

$$H_t = - (1 - \gamma) \theta_* \frac{1}{\mu_* + \Phi} \sigma^2 \chi_t$$

Proposition 6.1 describes the action of a sparse agent. When $\kappa = 0$ we recover the fully rational agent, as in Barberis (2000), Campbell and Viceira (2002) and Wachter (2010), with the notoriously sophisticated hedging demand term H_t . When $\kappa > 0$ increases, portfolio choice becomes insensitive to the change in the equity premium, $\widehat{\pi}_t$, and the agent thinks less about the mean-reversion of asset, the H_t terms. In addition, the agent is here allowed to be “schizophrenic:” he pays attention to different things as a consumer than as an investor.⁵² For instance, suppose that the price-earnings ratio of equities is very high and expected equity returns are very low; then, the agent will decrease her exposure to equities ($\theta < \theta_*$), but not reduce her consumption ($\mu_t^s = \mu_*$). The online appendix (Section 12.1) develops the Merton portfolio choice model in much greater generality, with many assets and factors.

6.2 Other models

Linear quadratic models Many economic problems can be conveniently expressed as linear-quadratic (LQ) models (Ljungqvist and Sargent 2012). Section 12.2 derives how to systematically derive a BR version of those models. The models are very tractable (there is no matrix Riccati equation to solve). Hence, the model offers a way to quite generally have behavioral version of linear-quadratic problems.

Precautionary savings The consumer may save more when the future is uncertain, a phenomenon known as “precautionary savings” (e.g. Carroll 1997). However, a BR consumer may not think much or at all about the stochasticity of his income and under-accumulate a buffer of savings. This is developed in Section 12.3.

Real Investment Model In macro-style investment models, the firm optimizes on investment, with adjustment costs. A behavioral version of this problem is in Section 12.4. The firm uses a simplified model of future profits, which can lead it astray (see Greenwood and Hanson (2015)).

⁵²To do that, formally, I use the variant of sparse max specified in Gabaix (2014, Section XV.D of the online appendix), which specifies a different attention for each type of action (formally, for each coordinate of the action vector).

The Becker-Murphy model of rational addiction The Becker-Murphy (1988) model of rational addiction is a peak of the use of rationality in economics. Section 12.5 gives a behavioral version of it. Qualitative evidence in favor of the model (the fact that future increase in prices lower consumption today) are also consistent with this BR version - it shows that agent are at least *partially* rational (as in the present model), not that they are *fully* rational (as assumed by Becker-Murphy). This distinction is important: if people are BR enough, they'd be better off under a high tax, or a ban, of the addictive substance – while the optimal tax is 0 in the Becker-Murphy model.⁵³

7 Discussion

7.1 Model Enrichments

Here I discuss different potential enrichments of the model; the online appendix (Section 15) discusses more. I discuss them at first in the context of the 3-period model of Section 4.1.

What happens if the agent takes into account the future costs of thinking? One could imagine that the agent will take into consideration the future costs of thinking in this decision, and in the value function. Of course, this becomes a very sophisticated agent then – far from a BR agent. However, for completeness, I explore this direction in the online appendix, Section 15.3. The consequences are very minor.⁵⁴ This is why I don't incorporate this feature in the basic model.

More sophisticated behavioral agents We can illustrate the degrees of sophistication of Definition 3.8 in the 3-period model of Section 4.1. In the basic smax ($q = 0$), the time-1 value function is $V^{t=1,(q=0)}(w_1, x) = 2u\left(\frac{w_1+x}{2}\right)$. With one iteration, the time-1 value function is:

$$V^{t=1,(q=1)}(w_1, x) = u\left(\frac{w_1 + m_1^*(w_1, \sigma_x^2)x}{2}\right) + u\left(\frac{w_1 + (2 - m_1^*(w_1, \sigma_x^2))x}{2}\right) = V^{t=1,(q=0)}(w_1, x) + O(x^2).$$

which gives the agent a more sophisticated understanding of his future actions.

The value functions using $q = 1$ or $q = 0$ are the same, up to second order terms, and so are all consumptions. In addition, m_0^* is the same under both models. Hence, as in Proposition 3.14, the naive / sophisticated difference changes only second order terms. This is why I have selected the naive case ($q = 0$ iterations) as the baseline case. It is simpler, more plausible cognitively, and yields the same result as the sophisticated case, up to second order terms.⁵⁵

⁵³This analysis is in the spirit of Gruber and Köszegi (2001), who study a hyperbolic discounting addict, rather than a boundedly rational one in the sense of this paper.

⁵⁴First, the basic smax agent does not change his action at all. What changes is on the the q -iterated agent with $q \geq 1$. Second, in accordance with Section 3.5, this change makes only a second order change in the action.

⁵⁵Why not model iterated expectations, such as the agent's perception at time 0 of his perception at time 2 of

Enriched models of attention. Suppose that x is very negative, and the agent doesn't see it at first (still $w_0 + x > 0$, so the agent can avoid starvation). Will the agent detect the starvation that menaces him? Section 15.1 of the online appendix gives a simple fix, in which the attention at time 1 becomes:

$$m_1^* = \mathcal{A} \left(\frac{\sigma_x^2}{2\kappa} \max \left\{ \left| u'' \left(\frac{w_1}{2} \right) \right|, \frac{1}{K} \left| u'' \left(\frac{w_1}{2} + x \right) \right| \right\} \right)$$

for some $K > 1$. In the second term of the max, the agent takes into account the extra income x . One could imagine other variants as well, which are discussed in Section 15 of the online appendix.

7.2 Frequently Asked Questions

How do markets clear? They clear as usual: prices adjust so that supply equals demand.⁵⁶ We need at least one agent, in each market, to pay some attention to the price. For instance, in the Cass-Koopmans model of Section 5, firms pay attention to the interest rate (in their loan demand), even though consumers might potentially not pay attention to the interest rate (in their loan supply). So, $r_t = f'(K_t)$ survives, but the Euler equation doesn't.

Does the model only generate underreaction? The model most directly gives underreaction, but can indirectly generate overreaction.⁵⁷ Suppose that an analyst sees a company C with a surge in profits. She'll increase her estimate of its value. However, suppose that she forgets that the competitors of company C will try to imitate it, this way eroding C 's future profit (see Hanson and Greenwood 2015). The analyst overreacts to the surge in profits because she neglected the competition's response. More generally neglecting mitigating factors (i.e., negatively correlated additional effects) leads to overreaction. Likewise, a consumer overreacts to an income shock if she underpays attention to the fact that that shock is very transitory.

What about other types of dynamic problems, with discrete actions? To have a viable path to bounded rationality, or artificial intelligence, we need environments with enough structure.⁵⁸ Here, I rely a particular structure: "smooth" environments, where actions and variables are continuous, functions are smooth, and the world lives close to a well-understood steady state. This means that this theory – in its present form – does not have much to say about discrete environments with essentially no structure, like the traveling-salesman problem.⁵⁹ Fortunately,

his perception at time 5? The short answer is that this leads to a combinatorial explosion of the complexity of the model. See Section 15.7.1.

⁵⁶See Gabaix (2014, Section IV) for a systematic analysis in static contexts.

⁵⁷See also Hirshleifer, Lim and Teoh (2009).

⁵⁸There is no known good general-purpose algorithm for worlds with no structure and high dimensionality. There may not be any, if $P \neq NP$. Also, one will never be able to calculate the cognitive cost of the "rational limit" while problem $P \neq NP$ is open; most experts think it is likely to remain open for decades.

⁵⁹One could imagine that some variant could be helpful, e.g. by smoothing or doing a convex relation of discrete problems.

worlds with smooth structure are very relevant in economics, especially macro-finance, and the present framework allows us to make progress on those.

What happens with developed financial markets? For instance, in the life-cycle problem of Section 2, what would happen if the x was securitized (so that not just bonds, but also claims on x can be traded)? Here are some short remarks – I plan to offer a serious treatment of sparse finance in a future paper. First, all basic arbitrages are excluded: as x is deterministic, if x is not priced at its fair value, there’s an infinite payoff from trading it: so, agent will pay attention to that arbitrage, and x will be traded at its fair value, x . Second, even though the agent as an investor may fully take x into account, as a consumer he might not fully take it into account.⁶⁰

Can’t one explain anything with bounded rationality? Let us observe that traditional economics models, when free, can “explain almost anything”. For instance, the Arrow-Debreu model has a great many degrees of freedom. However, there are two disciplines (besides intuition, a soft form of empirical knowledge): hard empirical constraints on the parameters (e.g. on the production functions) and constraints on the number of free parameters that can be tuned around an agreed-upon baseline model.

I adopt the same attitude here. Take the life-cycle model of Figure 1. There is one free parameter ($\bar{\kappa}$) to fit a whole lifetime: the model is parsimonious. Next, the model is tightly constrained in its prediction: it can only give shapes similar to the ones in the Figure 1: first consumption is flat, then declines before retirement, drops at retirement, and is flat thereafter.

This said, there are many extensions that one can think about to enrich the sparse agent (e.g. varying defaults, persistence in attention, heterogeneity in rationality etc.). It means that it’s a framework with intriguing features worth exploring. But to keep the parsimony, I think it is useful to coalesce around a very simple default structure. Here I reported a minimalist one with I found useful across a range of problems.

Isn’t there much arbitrariness in terms of the “default”, and what can non-attended to? The Bayesian counterpart of the “default” is the “prior.”⁶¹ There is no good theory of the “prior,” but researchers still move on with “reasonable” priors. Likewise, in information economics researchers posit what is known to the agent and what information requires a mental cost.

The situation with sparsity and its default is very similar. I use “reasonable” default and try to make “reasonable” assumptions about what is perfectly perceived (e.g. time t) and is not (income after retirement).

An open question for much of economics is a way to systematize such “intuitively reasonable” choices. In the meantime, despite the fact that the model needs inputs, the model produces a good

⁶⁰See Proposition 6.1 and its discussion.

⁶¹The default is a deterministic point estimate (which is useful to model bounded rationality), the prior is a probabilistic distribution.

amount of value added, e.g. in Figure 1.

As there is one way to be rational, and many ways to be boundedly in rational theory, how do we know how to proceed? It is true that there are many ways to deviate from pure rationality, but some are much more important than others (perhaps they are on a power law, like so many things). Here, I focus of one feature: people use a simplify their picture of the world. There are others (e.g. loss aversion, hyperbolic discounting) and together they constitute a useful account of major deviations from pure rationality.

This agent is quite sophisticated – almost Bayesian. How do we account for big mistakes then? We can try to account for big mistakes and miscalibrations in at least two ways. First, agents don't have good priors (“hyperparameters” in hierarchical Bayes models) on which information is useful. For instance, they (ex post faultily) don't think that thinking about future retirement will help them much, that “it's just too hard”.⁶² Second, they could have a high cost of cognition in some domains: in the “proportional thinking” specification of Section 10.2.2, agent eliminate considerations that matter for less than 5% of the decision ($\bar{\kappa} = 5\%$ then). A medium-run research goal is to measure more systematically attention and costs of cognition, informed by models such as the present one.⁶³

Furthermore, sparse agents display only intermediate sophistication. This paper differs from models that demand Bayesian updating, with optimal use of information (and agents pay optimal costs to acquire it, e.g. as in Veldkamp 2010). Behavioral agents like my sparse agents do not do Bayesian updating.⁶⁴ For instance, if they see their income go up, Bayesian agents should infer “probably, the interest rate is high” (as income and interest rates tend to have positive correlation), but sparse agents don't do that inference.

Are those agents easier or more complex to use than rational agents? Historically, behaviorally agents have often been more difficult to use (for the economist deriving their actions) than rational agents. However, sparse agents are often easier to use. For instance, they don't need to solve the full equilibrium in their head – they use a sensible consumption rule (e.g. (51)), which may not be exactly optimum in their context, but does guide them through life. Also, they may use a reduced number of variables (e.g. their decision depends on two variables rather than 20). Hence, they are easier to use.

⁶²Farhi-Gabaix (2015, Section 6.1) discusses models with suboptimal allocation of attention.

⁶³See Chetty, Looney and Kroft (2009), Brown, Hossain and Morgan (2010), Allcott and Taubinsky (forth.), Taubinsky and Rees-Jones (2015) for recent progress.

⁶⁴It would be interesting to formalize systematically a form of “lazy pseudo-Bayesian updating”, but I differ that to future research.

8 Literature Review

This paper builds on two strands of the literature: on a relatively recent strand of basic behavioral modelling and on an older set of ideas tailored to specific macroeconomic contexts.

One key difference here is that the model is quite systematic. It applies, in a unified manner, to a wide class of models. This generality is possible because I build on earlier work (Gabaix 2014) that developed a quite systematic way to think about boundedly rational optimization in static contexts. There, I proposed a “sparse max”, a behavioral, less than fully rational and attentive version of the traditional max operator. That paper was concerned with static cases, while here I explore dynamic ones. That paper allows to give a behavioral formulation of some basic chapters of the microeconomic textbooks (consumer theory, equilibrium theory, Arrow-Debreu). The present paper gives a behavioral version of basic chapters of the macroeconomics textbook. A related paper, Farhi and Gabaix (2015) proposes a behavioral version of textbook taxation theory. Hence, we have a more unified view of bounded rationality in micro and macro, whereas other papers, while insightful, are more piecemeal.

The “micro” sparse max relies on a strand of paper thinking about basic behavioral modelling, in particular inattention, salience. See, for instance, Bordalo, Gennaioli and Shleifer (2013), Chetty, Looney and Kroft (2009), Gabaix and Laibson (2006), Kőszegi and Szeidl (2013), Schwartzstein (2014) and Woodford (2012).

The “macro” dimension of the sparse max relates to papers which offer a variety of devices to model some form of bounded rationality and are tailored to macro contexts (see Veldkamp (2011) and Woodford (2013) for reviews). To examine the pros and cons of models, I use the following criteria in Table 1:

1. Do the models also apply to bounded rationality in static contexts? If the goal (which is mine, at least) is to have a quite unified way of talking about bounded rationality in both static (e.g., doing static consumer theory with n goods) and dynamic contexts, this is a desirable property.
2. Can the model keep a representative agent? This is very desirable, as most clean theory keeps the representative agent (or at most a small number of types of agents).
3. Do we keep deterministic actions, or do actions have to be random?
4. Do we have source-dependent inattention?
5. Is the model Bayesian? Most economists would think it’s a good feature. As the goal is to explore bounded rationality, I submit that it is worthwhile to entertain non-Bayesian models.
6. Can the model handle a tractable behavioral version of a general textbook static problems, like $\max_{c_1, \dots, c_n} u(c_1, \dots, c_n)$ s.t. $\sum_i p_i c_i = w$? A tractable dynamic model handling general

Table 1: Different Models of Myopia

	BR in static contexts	Rep. agent allowed	Determ. actions	Source dependent attention	Bayesian	Textbook micro
Rule-of-thumb	No	No	Yes	Yes	No	No
Hyperbolics	No	Yes	Yes	No	Yes	No
Credit constraints	No	No	Yes	No	Yes	No
Imperfectly updated info.	No	No	Yes	Yes	Yes	No
Noisy info.	Yes	No	No	Yes	Yes	No
Entropy penalty	Yes	No	No	No	Yes	No
BGS salience	Yes	Yes	Yes	Yes	No	No
Sparsity	Yes	Yes	Yes	Yes	No	Yes

problems such as (24) should a fortiori be able to work in static context. Only the sparse max is – so far – able to do that. Other approaches might do that, but now they appear too complicated to do so.

Rule-of-thumb consumers. Campbell and Mankiw (1989) assume that a fraction λ are “rule-of-thumb” consumers who just consume their current income, $c_t = \bar{y} + \hat{y}_t$. Those agents could be interpreted as “super-sparse,” and pay attention to no other variables, while perceiving income shocks as being fully persistent.⁶⁵ In a sense, this research vastly generalizes the Campbell-Mankiw agents, which have proved a useful benchmark. Hence, aggregate consumption is $\hat{c}_t = \lambda \hat{y}_t + (1 - \lambda) \frac{\bar{r}}{R} \sum_{\tau \geq t+1} \frac{1}{R^{\tau-t}} \hat{y}_\tau$, and we do obtain the higher MPC for current income than for discounted future income.

Hyperbolic consumers. Laibson (1997) models hyperbolic consumers (fully rational or fully naive), and O’Donoghue and Rabin (2001) model partially naive consumers. There are some key differences. One is that a hyperbolic consumer is still a wealth-maximizer, so his consumption function depends only on the present value of future income, without a higher marginal propensity to consume out of current rather than future income (Barro 1999) – unless they are credit constrained (then hyperbolicity amplifies the effect of credit constraints). Another is that sparse agents have source-dependent myopia, unlike (the most common) hyperbolic consumers.

Credit constrained consumers. Those consumers are rational, but credit constrained now or in the near future. As a result, their MPC out of current income is higher than out of future income. A big difference with sparse agents is that again sparse agents have source-dependent myopia. Hsieh (2003) provides evidence for source-dependent myopia: his consumers consume “regularly arriving” government transfers (from the Alaskan oil fund) rationally (see however Kueng 2015), but react in a myopic way (high MPC) to other transitory income shocks.

⁶⁵Actually, the “fully” persistent part is not necessary. It is enough that the agent’s consumption is an affine function of their current income (and perhaps wealth), to the exemption of forward-looking variables.

Imperfectly updated information. Agents update their information set every now and then, e.g. when hitting a threshold, at fixed or random time intervals (Caballero 1995, Gabaix and Laibson 2001, Mankiw and Reis 2002, Reis 2006, Abel, Eberly and Panageas 2013) or in other imperfect ways (Sargent 1993, Furster, Hébert and Laibson 2012, Schwarzstein 2014). This captures a real phenomenon, but does not apply to the most basic setup, e.g. static ones. Those models are also very complex in rich environments.

Other behavioral models of salience. Bordalo, Gennaioli and Shleifer (BGS, 2013) are pursuing an influential research program modeling the impact of “salience” in a vast number of settings (see also Kőszegi and Szeidl (2011)). In a rough sense, their salience transforms initial variables x_i into salience-adjusted variables $m_i x_i$, with a salience function different determining the m_i (which can be greater than 1) in a way different from the sparsity setup, but related to it (also, sparsity concentrates on continuous actions, BGS on discrete choices). Modulo this change, and the current proposal should apply to BGS’ setup, making it dynamic.

Entropy penalty. In the Sims (2003) approach to inattention, the penalty is an entropy penalty, which results in parsimonious (but very complex to solve, see Matejka and Sims (2010)) models, which do not allow for source-dependent inattention. Other papers (Maćkowiak and Wiederholt (forth.), Woodford (2012)) deviate from this doctrine by accepting source-dependent inattention. Even though there are many differences of details, there is a common spirit. In noisy-signal theory actions are stochastic: the same problems lead to stochastic actions by agents. As a result, we lose the representative agent in the noisy-signals models. It can be recovered by taking linear approximations. In contrast, in sparsity theory actions are deterministic, which is one factor making this model easier to analyze. In addition, and relatedly, the sparsity model applies well to basic micro (Gabaix 2014), which is so far too difficult for noisy-signal models.

Near rational approach. The “near rational” approach says that agents will lose at most ε utils: it is often useful (Akerlof and Yellen 1985, Chetty 2012, Kueng 2015, Parker 2015), but it does not offer a precise model of which actions people will take. The statistics used by the sparsity model are those of the near rational approach (utility losses as in (10)), but sparsity makes a definite prediction about agents’ actions.

The Hansen-Sargent (2007) robustness framework generates distorted subjective beliefs; they exhibit a sophisticated form of risk aversion and pessimism, but not myopia and simplification per se.

9 Conclusion

I presented a practical way to do boundedly rational dynamic programming. It is portable and tractable and sheds light on a number of puzzles reviewed in the introduction.

Given that it is easy to use and sensible, this model may be useful to study other issues in macroeconomics and finance.

10 Appendix: Complements

10.1 Tools to Expand a Simple Model Into a More Complex one

Calculating the marginal impact of a new variable: $V_x, V_{w,x}$ Here I develop the method to derive the Taylor expansion of a richer model, when starting from a simpler one. The methods are entirely paper and pencil.⁶⁶

The state variables evolve according to:

$$w' = F^w(w, x, a), \quad x' = F^x(w, x, a)$$

where the ' denote next-period variables. Also, x is stable around 0: $F^x(w, 0, a) = 0$.

Consider the fully rational model:

$$V^r(w, x) = \max_a u(w, x, a) + \beta \mathbb{E} V^r(F^w(w, x, a), F^x(w, x, a)) \quad (68)$$

We start with a simpler model, where x is always 0: $x \equiv 0, F^{x'} = 0$, i.e.

$$V^d(w) = \max_a u(w, 0, a) + \beta \mathbb{E} V^d(F^{w'}(w, 0, a))$$

We use the notation, for a function $f(w, a)$,

$$D_w f = \partial_w f + (\partial_a f) \partial_w a^d(w) \quad (69)$$

which is the total derivative with respect to w , e.g. the full impact of a change in w , including the impact it has on a change in the action a .⁶⁷

Bellman method

Lemma 10.1 *The impact of a change x on the value function and its derivative is, respectively:*

$$V_x = \frac{u_x + \beta V_{w'}' F_x^w(w, x, a)}{1 - \beta F_x^x} \quad (70)$$

$$V_{w,x} = \frac{D_w u_x + \beta D_w [F_x^w(w, 0, a) V_{w'}'(w', 0)] + \beta V_{x'}' F_{xw}^x}{1 - \beta F_x^x D_w w'} \quad (71)$$

where all derivatives are evaluated at $(a, x) = (a^d(w), 0)$.

⁶⁶They draw from the techniques surveyed by Judd (1998, Chapter 14), who has a more computer-based perspective. Section 15.9 extends this to continuous time.

⁶⁷The notations are a bit awkward when variables are multidimensional, and can be made clearer by using tensors. To keep things elementary, I don't explicitly use tensors here.

Proof The proof yields the general method of calculation. We shut down uncertainty and differentiate the rational Bellman equation (68), first with respect to the new variable x (using the envelope theorem):

$$V_x(w, x) = u_x + \beta V'_{w'} F_x^w + \beta V'_{x'} F_x^x$$

which yields the announced expression for V_x . Then we take the total derivative w.r.t. w :

$$V_{w,x}(w, x) = D_w u_x + \beta D_w [V'_{w'} F_x^w(w, x, a)] + \beta F_x^x V'_{w',x'} D_w w' + \beta V'_{x'} D_w F_x^x$$

so

$$V_{w,x}(w, 0) = \frac{D_w u_x + \beta D_w [F_x^w(w, 0, a) V'_{w'}(w', 0)] + \beta V'_{x'} D_w F_x^x}{1 - \beta F_x^x D_w w'}$$

□

Lemma 10.2 *Assume the local autonomy condition (32), $F_a^x = 0$, and consider some value function $V(w, x)$. Then, the impact of a change x on the optimal action is (around $x = 0$):*

$$a_x = -\Psi_a^{-1} \Psi_x \quad (72)$$

with $\Psi(a, x) = u_a + \beta V'_{w'} F_a^w$, and

$$\Psi_a = u_{aa} + \beta F_a^w V'_{w'w'} F_a^w + \beta V'_{w'} F_{aa}^w, \quad \Psi_x = u_{ax} + \beta V'_{w'} F_{ax}^w + \beta F_x^w V'_{w'w'} F_a^w + \beta F_x^x V'_{w'x'} F_a^w. \quad (73)$$

evaluated at $(a, x) = (a^d(w), 0)$. They depend only on the transition functions and the derivatives of the simpler baseline value function $V^d(w')$.

Proof The first order condition for a is $\Psi(a, x) = 0$ with

$$\Psi(a, x) = u_a + \beta V'_{w'} F_a^{w'} + \beta V'_{x'} F_a^{x'}$$

The rest follows by the implicit function theorem: terms in V_{xx} drop out because $F_a^{x'} = 0$ at the default action. □

When the local autonomy condition (32) doesn't hold, a term V_{xx} appears. Then, the situation is more complex, and requires solving for a fixed point, in the form of a matrix Ricatti equation.

Feynman-Kac method In some cases, it is useful to do the same via a Feynman-Kac type of approach.⁶⁸ Here we view x_t as exogenous, i.e. assume $x_{t+1} = F^x(x_t)$. Calling w the initial

⁶⁸I call it “Feynman-Kac” because this approach deals particularly well with stochastic problems.

condition for wealth, the Lagrangian is:

$$L = \sum_{t=0}^{\infty} \beta^t u(w_t, x_t, a_t) + \sum_{t=1}^{\infty} \beta^t q_t^w (-w_t + F^{w_t}(w_{t-1}, x_{t-1}, a_{t-1})) + q_0^w (-w_0 + w)$$

where $q_t^w = V_w(w_t)$ are the Lagrange multipliers. At the optimum, the agent solves $V(w) = \max_{(a_t, w_t)_{t \geq 0}} L$. This implies that $L_{a_t} = L_{w_t} = 0$. The envelope theorem gives:

$$V_{x_t} = L_{x_t} = \beta^t [u_x(w_t, x_t, a_t) + \beta q_{t+1}^w F_x^{w_{t+1}}(w_t, x_t, a_t)]$$

so that, using the total derivative notation (69),

$$V_{x_t, w_0} = \beta^t \frac{Dw_t}{Dw_0} \Big|_{x=0} \left[u_x(w_t, x_t, a_t^d(w_t)) \Big|_{x=0} + \beta V_w(w_{t+1}) F_x^{w_{t+1}}(w_t, x_t, a_t^d(w_t)) \Big|_{x=0} \right]$$

or in short

$$V_{x_t, w_0} = \beta^t \frac{Dw_t}{Dw_0} D_{w_t} [u_x^t + \beta V_{w_{t+1}}^{t+1} F_x^{w_{t+1}}] \quad (74)$$

Application. In the consumption problem with $\beta R = 1$, let us derive again at the impact of a one-time change of interest rate \hat{r}_t from Lemma 4.2. Under the default model $c_t = c_0$ and $w_t = w_0$. So $\frac{Dw_t}{Dw_0} = 1$, $u_{x_t} = 0$, $V_w^d(w_t) = u'(c_t) = \left(\frac{\bar{r}w_t + \bar{y}}{R}\right)^{-\gamma}$, and given $F^{w_{t+1}} = (1 + r_t)(w_t - c_t) + \bar{y}$,

$$F_{r_t}^{w_{t+1}} = w_t - c_t = w_t - \frac{rw_t + \bar{y}}{R} = \frac{w_t - \bar{y}}{R} = \frac{w_t^-}{R}$$

using the notation $w_t^- := w_t - \bar{y}$ for the beginning of period wealth. Using (74) gives:

$$\begin{aligned} V_{x_t, w_0} &= \beta^t D_{w_t} \left[\beta \left(\frac{\bar{r}w_{t+1} + \bar{y}}{R} \right)^{-\gamma} \frac{w_t^-}{R} \right] = \frac{1}{R^{t+2}} D_{w_t} \left[\left(\frac{\bar{r}w_{t+1} + \bar{y}}{R} \right)^{-\gamma} w_t^- \right] \\ &= \frac{1}{R^{t+2}} \left(-\gamma \frac{\bar{r}}{R} c_{t+1}^{-\gamma-1} \frac{Dw_{t+1}}{Dw_t} w_t^- + c_t^{-\gamma} \right) = \frac{1}{R^{t+2}} c_0^{-\gamma-1} \left(-\gamma \frac{\bar{r}}{R} w_0^- + c_0 \right) \end{aligned}$$

as under the default model $c_t = c_0$. As time-0 consumption satisfies $u_{c_0} = V_{w_0}$, we have $u_{cc} \partial_{\hat{r}_t} c_0 = \partial_{\hat{r}_t} V_{w_0}$, and

$$\partial_{\hat{r}_t} c_0 = \frac{\partial_{\hat{r}_t} V_{w_0}}{u_{cc}} = \frac{\frac{1}{R^{t+2}} c_0^{-\gamma-1} \left(-\gamma \frac{\bar{r}}{R} w_0^- + c_0 \right)}{-\gamma c_0^{-\gamma-1}} = \frac{1}{R^{t+2}} \left(\frac{\bar{r}}{R} w_0^- - \psi c_0 \right) = \frac{1}{R^{t+2}} \left(\frac{r}{R} (w_0 - \bar{y}) - \psi c_0 \right).$$

which gives again Lemma 4.2 (the income part being easy as always).

10.2 Extension: Proportional Thinking and Scale-Free κ

Here are some notes for a practical use of the model.

10.2.1 Using the model without κ

First, one could imagine directly measuring m . For instance, take the basic life-cycle model of Section 2. Suppose we had data on retirement saving, one could measure the m_t of say high-educated people vs less educated people, at different points in time t . That would be interesting. This way, we could back out the correct attention function, or see its determinants over time – a useful research enterprise.

Second, suppose that we keep κ as a more basic parameter, and measure κ it – for one given type of decision. For instance, in the life-cycle model, one could study just one variable, consumption, fitting a κ for it, say over a lifetime.

One important difficulty is that κ is not “scale-free” – it depends on the units of utility. To remedy that problem, I discussed earlier a “scale-free κ ” (Gabaix 2014), which I state and generalize here.

10.2.2 Scale-free κ and Proportional Thinking

Scale-free parametrization with a one-dimensional action Suppose first that there is just a one-dimensional action. The following formulation is easy to microfoundation, and proves sensible and convenient:⁶⁹

$$\kappa = (\bar{\kappa} a^{\natural})^2 |v_{aa}(a^d(x, m^d))| \quad (75)$$

where $\bar{\kappa}$ is a unitless parameter, which might portable from one situation to the next (Rabin (2013) advocates such portability of parameters). Here a^{\natural} is the “natural” or “typical scale” of the action. In music, \natural means the “natural” pitch of a note, so perhaps it’s a mnemonic symbol for the “natural scale” of variable).

What is a natural scale? In many cases, this is reasonably clear. If the action is a purchase c_k in dollar of good k , then $a^{\natural} = c^d$. If the action is a transform $f(c_k)$ of that purchase, then $a_k^{\natural} = |f'(c_k^d)| c_k^{\natural}$. As a result, if we consider log consumption, the natural scale is $a_k^{\natural} = 1$.⁷⁰

Formulation (75) implies that attention to dimension i is:⁷¹

$$m_i = \mathcal{A} \left(\frac{\mathbb{E}[a_{m_i}^2]}{(\bar{\kappa} a^{\natural})^2} \right)$$

In turns, we obtain a rule (with $\mathcal{A} = \mathcal{A}_1$): “Keep only dimensions of reality that matter for more

⁶⁹Section 15.5 provides a microfoundation, which I outline here. Mental effort creates trembling in the action, with a standard deviation $\sigma_a = \sqrt{2\bar{\kappa} a^{\natural}} (\sum_i g(m_i))^{1/2}$ proportional to the typical scale of the action (as in much of psychophysics, e.g. in the Weber-Fechner law), and increasing in mental activity m . The resulting utility losses are then $\frac{1}{2} v_{aa} \sigma_a^2 = \kappa \sum_i g(m_i)$ with the value of κ in (75).

⁷⁰In addition, if a is presented as belong to a natural range $a \in [\underline{a}, \bar{a}]$ (as in many lab experiments): then we can take $a^{\natural} = \frac{1}{2} (|\bar{a} + \underline{a}| + |\bar{a} - \underline{a}|)$, which reflects the absolute value of the midpoint, and the width of the range. In other contexts, we can take $a^{\natural} = \text{var}(a^r(x))^{1/2}$.

⁷¹I use $m_i^* = \mathcal{A}(-\mathbb{E}[a_{m_i} u_{aa} a_{m_i}]/\kappa)$, from (10).

than $\bar{\kappa}$ of the typical scale a^\natural of the decision”.

For instance, a reasonable parameter might be $\bar{\kappa} = 5\%$. People then eliminate all considerations that matter for less than 5% of the decision, whether it is for a small decision (e.g., buying at the supermarket) or a big decision (e.g., buying a car).

Some evidence consistent with that is presented by Samuelson and Zeckhauser (1988), who show similar percentage price dispersion between cheap and expensive goods, and by Tversky and Kahneman (1981): people consider a \$5 discount more worthy of an extra shopping trip if it is for a \$15 calculator than for a \$125 jacket. Bordalo, Gennaioli and Shleifer (2013) and Bushong, Rabin and Schwartzstein (2015) have emphasized the importance of this proportional thinking.

I conjecture that this feature is a good benchmark which would be interesting to evaluate empirically (I do not claim it will work perfectly, but I conjecture that it will hold more likely than the polar opposite prediction that people would be 100 times more precise for a good that costs 100 times more).

Scale-free parametrization with many actions. We can generalize the idea above to many actions. To do that, formally, I use a variant of sparse max, which specifies a different attention for each type of action (formally, for each coordinate of the action vector).⁷²

Definition 10.3 (Sparse max: scale-free with action-specific attention). *When it is scale-free, with action-specific attention, the sparse max $\text{smax}_{a; m|a^\natural}^{SF} v(a, x, m)$ is defined as follows.*

Step 1: For each dimension k of the action, the attention vector $m^k = (m_i^k)_{i=1\dots n}$ is given by:

$$m^* = \arg \min_{m \in [0,1]^n} \sum_i \left[\frac{1}{2} (a_{m_i}^k)^2 (1 - m_i)^2 + (\bar{\kappa} a^{k, \natural})^2 g(m_i) \right] \quad (76)$$

Step 2: The k -th dimension of the action (a_k^) is the optimal under the model m^k , given the other dimensions of actions (a_{-k}^*).*

$$a_k^* \in \arg \max_{a_k} v((a_k, a_{-k}^*), x, m^k). \quad (77)$$

Practically, Definition 10.3 yields an explicit formula for attention:

$$m_i^k = \mathcal{A} \left(\frac{\mathbb{E} \left[(a_{m_i}^k)^2 \right]}{(\bar{\kappa} a^{k, \natural})^2} \right) \quad (78)$$

⁷²See Gabaix (2014, Section XV.D of the online appendix)

and a first order condition for action is, for each k :⁷³

$$v_{a^k}(a^*, x, m^k) = 0 \quad (79)$$

In conclusion: based on the psychological evidence on proportional thinking, it seems plausible that the above formulation will calibrate reasonable well, with one scale-free parameter $\bar{\kappa}$, which I conjecture to be around $\bar{\kappa} = 5\%$ – if people omit factors that matter by less than 5%.

For concreteness, here is an example. Suppose that there are two actions, a^1 , a^2 , and:

$$v(a, x) = -\frac{1}{2} \sum_{k=1}^2 w^k \left(a^k - \sum_i b_i^k x_i \right)^2$$

where $w^k > 0$ is a scale, b^k is a vector. Then, the attention-augmented utility function is $v(a, x, m) = -\frac{1}{2} \sum_{k=1}^2 w^k (a^k - \sum_i b_i^k m_i x_i)^2$. The basic sparse max yields:

$$m_i = \mathcal{A} \left(-\mathbb{E} [a_{m_i} v_{aa} a_{m_i}] \frac{1}{\kappa} \right) = \mathcal{A} \left(\frac{1}{\kappa} \sum_{k=1}^2 w_k (b_i^k)^2 \sigma_{x_i}^2 \right)$$

and $a^k = \sum_i b_i^k m_i x_i$. In contrast, the scaled sparse max of Definition 10.3 gives:

$$m_i^k = \mathcal{A} \left(-\mathbb{E} [a_{m_i}^k v_{a^k a^k} a_{m_i}^k] \frac{1}{\kappa} \right) = \mathcal{A} \left(\frac{(b_i^k)^2 \sigma_{x_i}^2}{(a^{k, \dagger})^2 \bar{\kappa}^2} \right)$$

and $a^k = \sum_i b_i^k m_i^k x_i$ for $k = 1, 2$. Here, the attention is independent of w^k : for action k , the agent makes sure to include all considerations i that makes him change his decision by more than $\bar{\kappa}$ times the typical scale $a^{k, \dagger}$.

10.3 Complements: Life-cycle model with T periods

Here are some more precisions to see how the intuitive procedure of Section 2 is an outcome of the general formalism. The action is consumption c_t and the state vector is

$$z_t = (w_t, x, t)$$

Note that $x^* := (T - L) \hat{y}$ is the objective value, so $z_0 = (w_0, x^*, 0)$. But it is useful to consider x a variable.

The objective laws of motion are expressed in $F^z(c, z_t)$ ⁷⁴. So, the rational agent at time t solves

⁷³If several actions are possible (which is not a generic situation) then for instance we take one closest to the default action by some metric.

⁷⁴Formally,

$$F^w(c, z_t) = w_t + \bar{y} + \frac{x}{T - L} \mathbf{1}_{t \geq L} - c_t, \quad F^x(c, z_t) = x, \quad F^t(c, z_t) = t + 1.$$

$c_t^r = \arg \max_t v(c, z_t)$, with v defined in (3): $v(c, z_t) := u(c) + V^r(w_t + \bar{y} - c, x, t + 1)$.

For the behavioral agent, the attention-augmented law of motion is $F^z(z_t, m) := F^z(w_t, mx, t)$ before retirement (for $t < L$): only x is misperceived. After retirement, I assume correct perception, $F^z(z_t, m) := F^z(z_t)$.⁷⁵ The proxy value function is $V^p = V^r$ (equation (2)). Then, Definition 3.4 gives:

$$c_t = \arg \operatorname{smax} v(c, z_t, m) \text{ with } v(c, z_t, m) := v(c, w_t, mx, t) \quad (80)$$

When we apply the sparse max of Definition 3.1. Step 1 gives the value of m_t in (6). Step 2 gives the value of c_t in (6).

11 Appendix: Proofs

The online appendix contains most proofs.

Proof of Lemma 3.3 The rational reaction function satisfies: $a^r(x) = a^d + \sum_i b_i x_i + \lambda(x)$ for a function $\lambda(x) = O(\|x\|^2)$. So, $\partial a / \partial x_i = b_i$ and:

$$m_i^* = \tau \left(1, \frac{\kappa_a}{\sigma_i \cdot \partial a / \partial x_i} \right) = \tau \left(1, \frac{\kappa_a}{\sigma_i \cdot b_i} \right)$$

We shall use the notation $\bar{\lambda}(x) := \lambda((m_i^* x_i)_{i=1 \dots n})$, which also satisfies $\bar{\lambda}(x) = O(\|x\|^2)$. The sparse reaction function is:

$$\begin{aligned} a^s(x) &= \arg \max_a u(a, m_1^* x_1, \dots, m_n^* x_n) = a^s(m_1^* x_1, \dots, m_n^* x_n) \\ &= a^d + \sum_i b_i m_i^* x_i + \lambda((m_i^* x_i)_{i=1 \dots n}) = a^d + \sum_i b_i \tau \left(1, \frac{\kappa_a}{\sigma_i \cdot b_i} \right) x_i + \bar{\lambda}(x) \\ &= a^d + \sum_i \tau \left(b_i, \frac{\kappa_a}{\sigma_i} \right) x_i + \bar{\lambda}(x) = a^d + \sum_i \tau \left(b_i, \frac{\kappa_a}{\sigma_i} \right) x_i + O(\|x\|^2). \end{aligned}$$

Proof of Proposition 4.3 Suppose first that the problem is deterministic. Given attention m_r, m_y , the agent's problem is identical to the rational agent's problem, in a world in which the true interest rate is $\bar{r} + m_r \hat{r}_t$, and true income is $\bar{y} + m_y \hat{y}_t$, and the transition function is the one perceived by the agent. Hence, calling $c_0^r(w_0, (\hat{r}_\tau, \hat{y}_\tau)_{\tau \geq 0})$ the rational consumption given flows $c_0^s(w_0, (\hat{r}_\tau, \hat{y}_\tau)_{\tau \geq 0}) = c_0^r(w_0, (m_r \hat{r}_\tau, m_y \hat{y}_\tau)_{\tau \geq 0})$. Using the Taylor expansion in Lemma 4.2 gives the result. When the problem is stochastic, as in the proof of Lemma 4.2, we take the certain equivalent at each period.

⁷⁵Formally, there is an element of arbitrariness: one could conceive a partial misperception of time t , or misperception income after retirement. A medium-run open question is a way to endogenize such "intuitively reasonable" choices. Still, the model produces a lot of reasonable behavior "on its own", e.g. in Figure 1.

Proof of Proposition 5.1 If $\phi^d \geq \phi^r$, we have $\phi_0 = \frac{\xi}{r+\phi^d} \leq \frac{\xi}{r+\phi^r} = \phi^r$, hence $\phi = m_K \phi_0 < \phi_0 \leq \phi^r$. If $\phi^d \geq \phi$ then $\phi^0 = \frac{\xi}{r+\phi^d} \leq \frac{\xi}{r+\phi}$, so that $\phi = m_K \phi_0 < \phi_0 \leq \frac{\xi}{r+\phi}$, and $\phi(r+\phi) < \xi$. Given that $\xi = \phi^r(r+\phi^r)$, that implies $\phi < \phi^r$.

By (58) $\phi < \phi^r$ implies that fluctuations in capital $\text{var}(\hat{K}_t)$ are more volatile in the sparse economy. Given $Y_t = f(K_t) + \delta K_t$, we have $\hat{Y}_t = (f'(K_*) + \delta) \hat{K}_t$. So the same properties holds for GDP fluctuations \hat{Y}_t as capital \hat{K}_t , as the two are proportional.

For consumption have $\hat{C}_t = (r+\phi) \hat{K}_t$, so $\sigma_C = (r+\phi) \sigma_K = (r+\phi) \frac{\sigma_\varepsilon}{\sqrt{2\phi}}$ from (58). The latter function is increasing in ϕ for $\phi \geq r$. Hence, σ_C is higher in the sparse economy ($r \leq \phi \leq \phi^r$). Next, $\hat{I}_t = \hat{Y}_t - \hat{C}_t = (f'(K_*) + \delta - (r+\phi)) \hat{K}_t = (\delta - \phi) \hat{K}_t$, so $\sigma_{\hat{I}_t} = |\delta - \phi| \sigma_K = |\delta - \phi| \frac{\sigma_\varepsilon}{\sqrt{2\phi}}$ a U-shaped function of ϕ with minimum at $\phi = \delta$. So, if $\phi^r \leq \delta$, investment is more volatile in the sparse economy.

Proof of Proposition 5.2 In continuous-time limit, the second derivative of the value normalized function $v(c) := [u(c) dt + \beta V(k_t + y_t - (1+r_t)dt) c_t dt, K_{t+1}) - V(k_t, K_t)] / dt$ is $v_{cc} = u_{cc}$, this is a quite general phenomenon. Applying (10) with $c_K = \phi_0$, we obtain:

$$m_K = \mathcal{A} \left(\frac{|u_{cc}| \phi_0^2 \sigma_K^2}{\kappa} \right) = \mathcal{A} \left(\frac{|u_{cc}| \phi_0^2 \sigma_\varepsilon^2}{2\phi\kappa} \right) = \mathcal{A} \left(\frac{\phi_0^2}{B\phi} \right), \quad B := \frac{2\kappa}{|u_{cc}| \sigma_\varepsilon^2}$$

using (58), $\sigma_K^2 = \frac{\sigma_\varepsilon^2}{2\phi}$. Using $\phi = \phi_0 m_K = \phi_0 \max \left(1 - \frac{B\phi}{\phi_0^2}, 0 \right)$, so $\phi = \frac{\phi_0}{1 + \frac{B}{\phi_0}}$.

References

Abel, Andrew, Janice C. Eberly, and Stavros Panageas, “Optimal Inattention to the Stock Market with Information Costs and Transactions Costs,” *Econometrica*, 81 (2013), 1455–1481.

Acemoglu, Daron, *Introduction to Modern Economic Growth* (Princeton, NJ: Princeton University Press, 2008).

Acemoglu, Daron, Vasco M. Carvalho, Asuman Ozdaglar, and Alireza Tahbaz-Salehi. 2012. “The network origins of aggregate fluctuations.” *Econometrica*, 80(5): 1977–2016.

Akerlof, George, and Janet Yellen, “Can Small Deviations from Rationality Make Significant Differences in Economic Equilibria?” *American Economic Review*, 75 (1985), 708–720.

Allcott, Hunt, and Dmitry Taubinsky, “Evaluating Behaviorally-Motivated Policy: Experimental Evidence from the Lightbulb Market,” *American Economic Review* (forthcoming).

Alvarez, Fernando, Luigi Guiso and Francesco Lippi, “Durable Consumption and Asset Management with Transaction and Observation Cost,” *American Economic Review*, 102 (2012), 2859–96.

Auclert, Adrien, “Monetary Policy and the Redistribution Channel,” Working Paper, 2015.

Baker, Malcolm, Stefan Nagel and Jeffrey Wurgler, “The Effect of Dividends on Consumption,” *Brookings Papers on Economic Activity*, 38 (2007), 231–292.

- Barberis, Nicholas, “Investing for the Long Run when Returns are Predictable,” *Journal of Finance* 55 (2000), 225-264.
- Barberis, Nicholas, Andrei Shleifer, and Robert Vishny, “A model of Investor Sentiment” *Journal of Financial Economics* 49 (1998), 307-343.
- Barro, Robert, “Are Government Bonds Net Wealth?” *Journal of Political Economy*, 82 (1974), 1095-1117.
- Becker, Gary S., and Kevin M. Murphy, “A Theory of Rational Addiction,” *Journal of Political Economy*, 94 (1988), 675-700.
- Beshears, John , James Choi, David Laibson, Brigitte C. Madrian, and Sean Wang, “Who is Easier to Nudge?” Working Paper, 2015.
- Bernheim, B. Douglas, Jonathan Skinner, and Steven Weinberg, “What Accounts for the Variation in Retirement Wealth among US Households?” *American Economic Review*, 91 (2001), 832-857.
- Bigio, Saki and Jennifer La’O, “Financial Frictions in Production Networks,” Working Paper, Columbia Business School, 2016.
- Blanchard, Olivier and Stanley Fischer, *Lectures on Macroeconomics* (Cambridge, MA: MIT Press, 1989).
- Bordalo, Pedro, Nicola Gennaioli and Andrei Shleifer, “Salience and Consumer Choice,” *Journal of Political Economy*, 121 (2013), 803-843.
- Brown, Jennifer, Tanjim Hossain, and John Morgan, “Shrouded Attributes and Information Suppression: Evidence from the Field.” *Quarterly Journal of Economics* 125 (2010), 859-876.
- Bushong, Benjamin, Matthew Rabin, Joshua Schwartzstein, “A Model of Relative Thinking,” Working Paper, 2015.
- Caballero, Ricardo, “Near-Rationality, Heterogeneity, and Aggregate Consumption,” *Journal of Money, Credit and Banking* 27 (1995), 29–48.
- Caballero, Ricardo and Alp Simsek, “Fire Sales in a Model of Complexity,” *Journal of Finance*, 68.6 (2013), 2549-2587.
- Caballero, Ricardo, and Emmanuel Farhi, “The Safety Trap,” Working Paper, 2015.
- Campbell, John Y., and N. Gregory Mankiw, “Consumption, Income, and Interest Rates: Reinterpreting the Time Series Evidence,” *NBER Macroeconomics Annual* (1989), 185-216.
- Campbell, John and Luis Viceira, *Strategic Asset Allocation* (Oxford, UK: Oxford University Press, 2002).
- Candès, Emmanuel, and Terence Tao, “Near-Optimal Signal Recovery from Random Projections: Universal Encoding Strategies?” *IEEE Transactions on Information Theory*, 52 (2006), 5406-5425.
- Carroll, Christopher, “Buffer-Stock Saving and the Life Cycle/Permanent Income Hypothesis,” *Quarterly Journal of Economics*, 112 (1997), 1-55.
- Carvalho, Vasco and Basile Grassi. “Large Firm Dynamics and the Business Cycle,” Working

paper, 2016.

Chetty, Raj, “Bounds on Elasticities with Optimization Frictions: A Synthesis of Micro and Macro Evidence on Labor Supply,” *Econometrica*, 80 (2012), 969–1018.

Chetty, Raj, Adam Looney, and Kory Kroft, “Salience and Taxation: Theory and Evidence,” *American Economic Review* 99 (2009), 1145–1177.

Cochrane, John, “The Sensitivity of Tests of the Intertemporal Allocation of Consumption to Near-Rational Alternatives,” *American Economic Review*, 79 (1989), 319–337.

Croce, Mariano, Martin Lettau, and Sydney C. Ludvigson, “Investor Information, Long-run Risk, and the Term Structure of Equity,” *Review of Financial Studies*, 28 (2015), 706–742.

Evans, George, and Seppo Honkapohja, *Learning and Expectations in Macroeconomics* (Princeton, NJ: Princeton University Press, 2001).

Farhi, Emmanuel and Xavier Gabaix, “Optimal Taxation with Behavioral Agents,” NBER Working Paper 21524, 2015.

Fuster, Andreas, Benjamin Hébert and David Laibson, “Natural Expectations, Macroeconomic Dynamics, and Asset Pricing,” *NBER Macroeconomics Annual*, 26 (2012), 1–48.

Gabaix, Xavier. 2011. “The granular origins of aggregate fluctuations.” *Econometrica*, 79(3): 733–772.

Gabaix, Xavier, “A Sparsity-Based Model of Bounded Rationality,” *Quarterly Journal of Economics*, 129 (2014), 1661–1710.

Gabaix, Xavier, “A Behavioral New Keynesian Model” Working Paper, 2016.

Gabaix, Xavier, and David Laibson, “The 6D Bias and the Equity Premium Puzzle,” *NBER Macroeconomics Annual*, 16 (2002), 257–312.

Gabaix, Xavier, and David Laibson, “Shrouded Attributes, Consumer Myopia, and Information Suppression in Competitive Markets,” *Quarterly Journal of Economics*, 121 (2006), 505–40.

Gabaix, Xavier, David Laibson, Guillermo Moloche and Stephen Weinberg, “Costly Information Acquisition: Experimental Analysis of a Boundedly Rational Model,” *American Economic Review*, 96 (2006), 1043–1068.

Ganong Peter and Pascal Noel, “How Does Unemployment Affect Consumer Spending?” Working paper, Harvard, 2015

Gourinchas, Pierre-Olivier and Jonathan A. Parker, “Consumption Over The Life Cycle,” *Econometrica*, 70 (2002), 47–89.

Greenwood, Robin, and Samuel G. Hanson, “Waves in Ship Prices and Investment,” *Quarterly Journal of Economics* 130(2015), 55–109.

Hall, Robert, E., “Stochastic Implications of the Life Cycle—Permanent Income Hypothesis: Theory and Evidence,” *Journal of Political Economy*, 86 (1978), 971–987.

Harris, Christopher, and David Laibson, “Hyperbolic Discounting and Consumption,” *Proceedings of the 8th World Congress of the Econometric Society*, 2001.

Harris Christopher, and David Laibson, “Instantaneous Gratification,” *Quarterly Journal of*

Economics, 128 (2013), 205-248.

Hirshleifer, David, Sonya Lim and Siew Teoh, "Driven to Distraction: Extraneous Events and Underreaction to Earnings News," *Journal of Finance*, 64 (2009), 2289-2325.

Hong, Harrison, Jeremy C. Stein, and Jialin Yu, "Simple Forecasts and Paradigm Shifts," *The Journal of Finance*, 62.3 (2007), 1207-1242.

Hsieh, Chang-Tai, "Do Consumers React to Anticipated Income Changes? Evidence from the Alaska Permanent Fund," *American Economic Review*, 93 (2003), 397-405.

Johnson, David S., Jonathan A. Parker, and Nicholas S. Souleles, "Household Expenditure and the Income Tax Rebates of 2001," *American Economic Review*, 96 (2006), 1589-1610.

Judd, Kenneth, *Numerical Methods in Economics* (Cambridge, MA: MIT Press, 1998).

Kahneman, Daniel, *Thinking, Fast and Slow* (New York, NY: Farrar, Straus and Giroux, 2011).

Kőszegi, Botond, and Adam Szeidl, "A Model of Focusing in Economic Choice," *Quarterly Journal of Economics*, 128 (2013), 53-104.

Krusell, Per and Anthony Smith, "Rules of Thumb in Macroeconomic Equilibrium: A Quantitative Analysis," *Journal of Economic Dynamics and Control* 20 (1996), 527-558.

Krusell, Per and Anthony Smith, "Income and Wealth Heterogeneity in the Macroeconomy," *Journal of Political Economy*, 106 (1998), 867-896.

Kueng, Lorenz, "Explaining Consumption Excess Sensitivity with Near-Rationality: Evidence from Large Predetermined Payments," NBER Working Paper 21772, 2015.

Laibson, David, "Golden Eggs and Hyperbolic Discounting," *Quarterly Journal of Economics*, 112 (1997), 443-477.

Ljungqvist, Lars and Thomas Sargent, *Recursive Macroeconomic Theory* (Cambridge, MA: MIT Press, 2012).

Lucas, Robert E, "Econometric Policy Evaluation: A Critique," *Carnegie-Rochester Conference Series on Public Policy*, 1 (1976), 19-46.

Lusardi, Annamaria, "Americans' Financial Capability," NBER Working Paper 17103, 2011.

Lusardi, Annamaria, Daniel Schneider, and Peter Tufano, "Financially Fragile Households: Evidence and Implications," *Brookings Papers on Economic Activity*, (2011), 83-134.

Maćkowiak, Bartosz and Mirko Wiederholt, "Business Cycle Dynamics Under Rational Inattention," *Review of Economic Studies*, 82 (2015), 1502-1532.

Mankiw, N. Gregory and Ricardo Reis, "Sticky Information Versus Sticky Prices: A Proposal To Replace The New Keynesian Phillips Curve," *Quarterly Journal of Economics*, 117 (2002), 1295-1328.

Matejka, Filip and Christopher Sims, "Discrete Actions in Information-Constrained Tracking Problems," Working Paper, Princeton University, 2010.

Merton, Robert, "Optimum Consumption and Portfolio Rules in a Continuous-Time Model," *Journal of Economic Theory*, 3 (1971), 373-413.

Modigliani, Franco, and Richard H. Brumberg, "Utility Analysis and the Consumption Function:

an Interpretation of Cross-Section Data,” in Kenneth K. Kurihara, ed., *Post-Keynesian Economics* (New Brunswick, NJ.: Rutgers University Press, 1954), 388–436.

Mullainathan, Sendhil, Joshua Schwartzstein, and Andrei Shleifer, “Coarse Thinking and Persuasion,” *Quarterly Journal of Economics*, 123 (2008), 577–619.

O’Donoghue, Ted, and Matthew Rabin, “Choice and Procrastination,” *Quarterly Journal of Economics*, 116 (2001), 121-160.

Pagel, Michaela, “Expectations-Based Reference-Dependent Life-Cycle Consumption,” Working Paper, Columbia, 2014.

Parker, Jonathan, “Why Don’t Households Smooth Consumption? Evidence from a 25 Million Dollar Experiment,” Working Paper, MIT, 2015.

Prescott, Edward, and Rajnish Mehra, “Recursive Competitive Equilibrium: The Case of Homogeneous Households,” *Econometrica*, 48.6 (1980), 1365-79.

Rabin, Matthew, “Incorporating Limited Rationality into Economics,” *Journal of Economic Literature*, 51(2013), 528-43.

Reis, Ricardo, “Inattentive Consumers,” *Journal of Monetary Economics*, 53 (2006), 1761-1800.

Romer, David, *Advanced Macroeconomics*, (New York, NY: McGraw-Hill, 2012).

Samuelson, W. and Richard Zeckhauser, “Status Quo Bias in Decision Making,” *Journal of Risk and Uncertainty*, 1 (1998), 7–59.

Sargent, Thomas J, *Bounded Rationality in Macroeconomics*, (Oxford, UK: Oxford University Press, 1993).

Schwartzstein, Joshua, “Selective Attention and Learning,” *Journal of the European Economic Association*, 12 (2014), 1423-1452.

Sims, Christopher, “Implications of Rational Inattention,” *Journal of Monetary Economics*, 50 (2003), 665–690.

Taubinsky, Dmitry and Alex Rees-Jones, “Attention Variation and Welfare: Theory and Evidence from a Tax Salience Experiment”, Working Paper, 2015.

Thaler, Richard H, “Mental Accounting and Consumer Choice,” *Marketing Science*, 4 (1985), 199-214.

Tibshirani, Robert, “Regression Shrinkage and Selection via the Lasso,” *Journal of the Royal Statistical Society B*, 58 (1996), 267–288.

Tversky, Amos, and Daniel Kahneman, “The Framing of Decisions and the Psychology of Choice,” *Science*, 211.4481 (1981), 453-458.

Veldkamp, Laura, *Information Choice in Macroeconomics and Finance* (Princeton: Princeton University Press, 2011).

Wachter, Jessica, “Asset Allocation,” *Annual Review of Financial Economics* 2 (2010), 175-206.

Woodford, Michael, “Inattentive Valuation and Reference-Dependent Choice,” Working Paper, Columbia, 2012.

Woodford, Michael, “Macroeconomic Analysis Without the Rational Expectations Hypothesis,”
Annual Review of Economics, 5 (2013), 303–46.