

Behavioral Macroeconomics Via Sparse Dynamic Programming

Xavier Gabaix*

28 September 2023

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. This paper was a background paper for my Marshall Lecture at the EEA meetings in August 2023. First draft: March 2012. 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, Drew Fudenberg, Peter Ganong, Mark Gertler, David Hirshleifer, Bentley MacLeod, Matteo Maggiori, Jonathan Parker, Nicola Pavoni, Zhaonan Qu, Michael Rockinger, Thomas Sargent, Josh Schwarzstein, Andrei Shleifer, Joel Sobel, Bruno Strulovici, Richard Thaler, Jessica Wachter, Michael Woodford, seminar participants at various seminars and conferences, and my discussants: Andrew Abel, Nick Barberis, Thomas Chaney, Harrison Hong, Jennifer La’O, Ali Lazrak, Ricardo Reis, Alp Simsek, Alexis Toda. I thank Wu Di, Lilian Hartmann, Chenxi Wang and Jerome Williams for very good research assistance. I am grateful to the Dauphine-CAM foundation, the Ferrante Fund at Harvard, the Institute for New Economic Thinking, the NSF (SES-1325181), 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 modeling 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 a 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. In the language of Kahneman (2011)(chapter 8), the agent “substitutes” a complex problem by a simpler one – decision-making in a simplified world.
 2. The rest of the paper will give references about those stylized facts.

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).
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.

3. In Gabaix (2020), 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.

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

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 modeling 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 paper is as follows. Section 2 gives an elementary, but substantive, example – a behavioral version 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, 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

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

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:

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

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 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 consumes:

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

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

6. Here I use the basic sparse max of Definition 2, 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 4, the decision would be the same to the first order, i.e. up to $O(x^2)$ terms (see Section 3.5).

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




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.

$\frac{w_t - m_t x}{T-t} + \bar{y}$, 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 $v_{cc}^t = v_{cc}(c_t^d, w_t, 0, t) = \left(1 + \frac{1}{T-t-1}\right) u''(c_t^d)$ 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.⁸ A little more calculation gives the value of attention as follows.

PROPOSITION 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.

8. As we will see below, the numerator inside (6) represents the utility gains from optimizing (up to a factor of two), which are related to the curvature of the value function, v_{cc}^t , and how much the agent’s consumption changes when she is more attentive, $c_t'(0)$.

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 $s \in (0, L)$; Section E in the appendix describes the whole solution, including in discrete time):⁹

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

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

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:

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

- (a) at retirement (Bernheim et al. (2001)).¹⁰
 - (b) at the (predictable) expiration of unemployment benefits (Ganong and Noel (2019)).
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 and Mitchell (2011)).¹¹

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 (Aguilar 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.¹²

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.

10. In addition, Bernheim et al. (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 (2018).

11. More 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.

12. Note 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.

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, 2019)). 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.¹³ 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

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

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 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.¹⁴ The x_i are viewed by the agent as being drawn from a distribution with standard deviation σ_i .

DEFINITION 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

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

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

14. This is true unless the matrix Λ of Definition 1 is singular – in that case, the iterated sparse max of Definition F.1 helps.

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)).¹⁵ 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:¹⁶

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

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:

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

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 1. (Basic simple case for static smax) *In the case $v(a, x, m) = v(a, m_1 x_1, \dots, m_n x_n)$, the smax operator yields:*

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

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 $m_i^d = 0$).

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

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

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 2. *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 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)):¹⁷

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

17. Another 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), Candes and Tao (2006)).




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.




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

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 as:

$$w_{t+1} = (1 + \bar{r} + \hat{r}_t)(w_t - c_t) + \bar{y} + \hat{y}_t. \quad (16)$$

$$\hat{r}_{t+1} = \rho_r \hat{r}_t + \varepsilon_{t+1}^r \quad (17)$$

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

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?¹⁸

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 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

$$c^d(w_t) = \frac{\bar{r}w_t + \bar{y}}{R} \tag{19}$$

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

18. This 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 (2020), Carvalho and Grassi (2019), Caballero and Simsek (2013), Gabaix (2011)).

function $F(c_t, z_t, m)$, whose components are (with $F = (F^w, F^r, F^y)$):¹⁹

$$w_{t+1} = F^w(c_t, z_t, m) := (1 + \bar{r} + m_r \hat{r}_t)(w_t - c_t) + \bar{y} + m_y \hat{y}_t \quad (20)$$

$$\hat{r}_{t+1} = F^r(c_t, z_t, m) := \rho_r(m) \hat{r}_t + m_{\sigma_r} \varepsilon_{t+1}^r \quad (21)$$

$$\hat{y}_{t+1} = F^y(c_t, z_t, m) := \rho_y(m) \hat{y}_t + m_{\sigma_y} \varepsilon_{t+1}^y \quad (22)$$

where $\rho_r(m) := m_{\rho_r} \rho_r + (1 - m_{\rho_r}) \rho_r^d$ and $\rho_y(m) := m_{\rho_y} \rho_y + (1 - m_{\rho_y}) \rho_y^d$.

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 (22) 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.²⁰ 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.

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.

19. This is related to the “perceived law of motion” of the literature on learning (Evans and Honkapohja (2001) Evans and Honkapohja (2013)). 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.

20. Section C.3 of the online appendix develops this. See Lusardi et al. (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.

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

$$\max_{(a_t)_{0 \leq t < T}} \sum_{t=0}^{T-1} \beta^t u(a_t, z_t) \text{ s.t. } z_{t+1} = F^z(a_t, z_t, \varepsilon_{t+1}) \quad (23)$$

and a terminal condition $z_T \in \mathcal{F}^T$ for a given set \mathcal{F}^T .²¹

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

$$V^{r,t}(z) = \max_a \{u(a, z) + \beta \mathbb{E} [V^{r,t+1}(F^z(a, z, \varepsilon_{t+1}))]\} \quad (24)$$

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 (24), 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:

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

The sparse max version. In the smax version, we are given attention-augmented utility and transition functions in a way that will be illustrated later, e.g. in Proposition 2:²²

$$u(a, z, m), \quad F^z(a, z, \varepsilon_{t+1}, m) \quad (26)$$

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 6).²³ 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 4), so this assumption makes materially little difference. Since it is simpler, and (I will argue) more realistic, I recommend taking $V^p = V^r$.

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

22. One 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.

23. It 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 et al. (2006) for an algorithm using an early version of this idea of proxy value function.

The agent’s action is as follows.²⁴

DEFINITION 2. (*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 2, 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. It uses some technical assumptions, made throughout this section. Formally, we assume that in this section, and in particular in Definition 3, u , V and V^p are bounded functions, so that the operator \mathcal{T} maps pairs of bounded functions into bounded functions. We also make the very mild assumption that the expectation in (29) is well-defined—which is true under weak technical conditions.²⁵

DEFINITION 3. (*Bellman operator with sparse action*) Given a policy $a(z, V^p)$ from Definition 2, the Bellman operator $\mathcal{T}(V, V^p)$ is defined by:

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

Given this, (28) can be written as:

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

24. If 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).

25. The simplest would be to assume a finite number of potential values for the noise z , for instance.

This operator $\mathcal{T}(V, V^p)$ has the usual good properties, recorded in the following lemma.²⁶

LEMMA 3. (Monotone contraction) *The operator $\mathcal{T}(V, V^p)$ is a monotone β -contraction as a function of V . More explicitly: for any bounded 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 4. (Existence of a value function) *Given Lemma 3, there is a unique bounded 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 4. (Action in sparse dynamic programming, q -iterated) The basic (0-iterated) dynamic smax action is the action $a(z, V^p)$ from Definition 2. 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 2 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

26. It 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. Only $V \mapsto \mathcal{T}(V, V^p)$ does.

his own future behavior: he sees how much he will see how much he will see (etc. – iterated q times) future inattention.

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).²⁷

3.5. Some General Features of Sparse Dynamic Programming

This subsection presents some general features of sparse dynamic programming, including tools to compute it easily, and derives quite systematically the difference between the actions of a sparse agent and of a rational agent. It derives the general form of predictions (e.g. in Proposition 2) that appear in the rest of the paper. Still, as this section is a bit dry, the reader may wish 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.²⁸

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:

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

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. (17)). This decoupling of the change in the variable x_t from action a_t brings a lot of tractability to the system—otherwise things

27. I 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 .

28. See Harris and Laibson (2001, 2013) for a related analysis showing the ty 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. Zhaonan Qu and I have results along those lines, but they is quite involved, and too technical for the present paper.

would much more complicated (a matrix Ricatti equation would need to be solved for, see Section A.1).

For simplicity, I assume that the functions are infinitely differentiable, and so is the attention function (the online appendix weakens those conditions in Section G.1; the exposition is then heavier).

The main tools are the following.

PROPOSITION 2. (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 ,*

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

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. (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 A.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. The first three (Proposition 4 and Lemmas 5-6) are very simple—basically they rely on the envelope theorem—and would apply any other

“reasonable” models of misoptimization where the behavioral action differs from the rational action by some small term (order $O(\|x\|)$).²⁹

PROPOSITION 4. *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 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.³⁰

LEMMA 5. (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.

LEMMA 6. (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.³¹

29. The 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.

30. Here $V^o(w, x|a(\cdot))$ is the policy induced by a , closely along the lines of Lemma 4.

31. Hyperbolic discounting naive 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.

PROPOSITION 5. (Sophistication vs Naiveté makes only a second order difference in actions) *Consider the q -iterated agent of Definition 4. 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 5 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 ($L = 2, T = 3$).³² 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: $w_1 = w_0 - c_0$, $w_2 = w_1 - c_1$, $0 = w_2 + x - c_2$.

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:

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

32. Section A.3 gives more complements for the T period model.

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.³³

PROPOSITION 6. *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 6. We apply the smax procedure of Definition 2, 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 $c_1^d = \frac{w_1}{2}$ 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

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

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

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

where I use Definition 2, with $V^p = V^r$, the value function where the agent projects that he will be rational in the future:³⁴

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

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

34. In terms of the iterated sparse max of Definition 4, this is the value function under procedure with $q = 0$ iteration, t. If we used the procedure with $q = 1$ -iteration, the agent

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

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

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

$$c_0(w_0, x) = \frac{w_0}{3} + \frac{m_0}{3}x, \quad (43)$$

$$c_1(w_0, x) = \frac{w_0}{3} + \left(\frac{m_1}{2} - \frac{m_0}{6}\right)x, \quad (44)$$

$$c_2(w_0, x) = \frac{w_0}{3} + \left(1 - \frac{m_0 + m_1}{2}\right)x. \quad (45)$$

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.³⁵

□

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 .

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: $V^1(w_1, x) = V^{1,p}(w, x) + O(x^2)$.

35. If 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.

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 (1988)), so it is a fragile way to model agents. In contrast, sparse consumption functions are a more robust way of modeling 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.³⁶ 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 (20). 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 is encouraged to skip at first). The state vector is:³⁷

$$z_t = (w_t, \hat{r}_t, \hat{y}_t, \rho_r, \sigma_r, \rho_y, \sigma_y) \quad (46)$$

36. In 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 naive agent projects that $c_1 = c_2 = \frac{w_1}{2}$, whereas in fact $c_2 = \frac{w_1}{2} + x$.

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

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, (20):

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

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.³⁸ 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 7. (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*

$$\hat{c}_t = \mathbb{E}_t \left[\sum_{\tau \geq t} \frac{b_r(w_t) \hat{r}_\tau + b_y \hat{y}_\tau}{R^{\tau-t+1}} \right] + O(\|x\|^2), \quad (47)$$

$$b_r(w_t) := \frac{\bar{r}}{R} (w_t - \bar{y}) - \psi c_t^d, \quad b_y := \frac{\bar{r}}{R}. \quad (48)$$

The associated value function is

$$V^r(w_t, x_t) = \frac{1}{\mu_t} u(\mu_t w_t + \nu_t) + O(\|x\|^2) \quad (49)$$

where

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

and

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

.

38. In independent work, Auclert (2019) 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).³⁹

The behavioral policy is then as follows.

PROPOSITION 7. (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*

$$\hat{c}_t^s = \mathbb{E}_t^s \left[\sum_{\tau \geq t} \frac{b_r(w_t) m_r \hat{r}_t + b_y m_y \hat{y}_\tau}{R^{\tau-t+1}} \right] + O(\|x\|^2) \quad (50)$$

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 8. *In the AR(1) model (17)-(18), the rational policy is:*

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

$$B_r(w_t) := \frac{b_r(w_t)}{R - \rho_r}, \quad B_y := \frac{b_y}{R - \rho_y}. \quad (52)$$

and the value function as in (49) 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 8. *Under the AR(1) model, the behavioral policy is:*

$$c^s(w_t, \hat{y}_t, \hat{r}_t) = c^d(w_t) + B_r^s(w_t) \hat{r}_t + B_y^s \hat{y}_t + O(\|x\|^2) \quad (53)$$

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).⁴⁰ This shows a “feature-by-feature” truncation. It is useful

39. The 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}_τ , 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} .

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

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 state variables besides wealth) problem. Only local expansions and truncations are necessary.⁴¹ Attention to the interest will become higher if the interest rate is very volatile (σ_r high, e.g. as in hyperinflations).⁴²

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.⁴³ I use the τ_1 truncation function. I parametrize $\kappa = \bar{\kappa}^2 |u''(c_t^d)|$, as in Section 2 (see also Section A.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 $\bar{\kappa} \simeq 0.1$,

41. Note 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 C.7 for a discussion: the consumption frame is arguably better, as it yields great utility when $\varphi_y > r$, income shocks are less than extremely persistent.

42. Likewise, 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).

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

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 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 7, setting $\hat{y}_\tau := \hat{g}_{\tau+1}$. We obtain the consumption:

$$\hat{c}_t^s = b_y (w_{t-} + \hat{g}_t) + \sum_{\tau > t} b_y m_g \frac{\hat{g}_\tau}{R^{\tau-t}} \quad (54)$$

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 9. (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 C.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 Jr (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 it 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.⁴⁴

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:

$$K_{t+1} = f(K_t, L) + (1 - \delta) K_t - C_t + \varepsilon_{t+1} \quad (55)$$

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.⁴⁵ 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 and Fisher (1989) (chapter 2), Acemoglu (2008) (Chapter 8), Romer (2018)). 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.

44. For simplicity, I remove here the “growth” part, but it is straightforward to add.

45. I 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 8.1 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 (2008).

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. $\hat{K}_t = K_t - K^*$. The law of motion for capital (55) is, in linearized form:

$$\hat{K}_{t+1} = (1 + r) \hat{K}_t - \hat{C}_t + \varepsilon_{t+1} \quad (56)$$

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:

$$\hat{C}_t = b \hat{K}_t \quad (57)$$

for some constant b to be determined.

Plugging this into (56) we obtain: $\hat{K}_{t+1} = (1 + r - b) \hat{K}_t + \varepsilon_{t+1}$, i.e.

$$\hat{K}_{t+1} = (1 - \varphi) \hat{K}_t + \varepsilon_{t+1} \quad (58)$$

where φ is the speed of mean-reversion:

$$\varphi = b - r. \quad (59)$$

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 (58), we obtain: $\text{var} \hat{K}_{t+1} = (1 - \varphi)^2 \text{var} \hat{K}_t + \sigma_\varepsilon^2$. As in the steady state, $\text{var} \hat{K}_{t+1} = \text{var} \hat{K}_t$, $\text{var} \hat{K}_t = \frac{\sigma_\varepsilon^2}{1 - (1 - \varphi)^2}$, i.e., in the limit of

small time intervals:

$$\sigma_K = \frac{\sigma_\varepsilon}{\sqrt{2\varphi}} \quad (60)$$

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:

$$V(K) = \max_c \{u(c) + \beta \mathbb{E}[V(K + f(K) - c + \tilde{\varepsilon})]\} \quad (61)$$

The steady state is at $K = K_*$, $C_t = C_*$ with:⁴⁶

$$\beta(1 + \bar{r}) = 1 \quad (62)$$

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 (19)), this is the optimal policy in the steady state.⁴⁷ 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 (62), 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:

$$\hat{r}_t = f''(K^*) \hat{K}_t, \quad \hat{y}_t = -K^* f''(K^*) \hat{K}_t \quad (63)$$

46. The traditional proof is by taking the derivative w.r.t. K in (61), so $V'(K_*) = \beta(1 + F'(K_*))V'(K_*)$.

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

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 : $\mathbb{E}_t [\hat{K}_{t+1}] = (1 - \varphi^s) \hat{K}_t$. I again parametrize it as:

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

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.⁴⁸ Lemma 8 reads here (using the continuous time notation):

$$\hat{c}_t = r \hat{k}_t + \frac{r}{r + \varphi^s} m_K \hat{y}_t + \frac{r k^* - c^* \psi}{r + \varphi^s} m_K \hat{r}_t$$

and using (63) and (65) we obtain:

$$\hat{c}_t = r \hat{k}_t + \frac{-r K^* f''(K^*) + (r k^* - c^* \psi) f''(K^*)}{r + \varphi^s} m_K \hat{K}_t$$

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

$$\hat{c}_t = r \hat{k}_t + \frac{\xi}{r + \varphi^s} m_K \hat{K}_t \quad (64)$$

with

$$\xi := -\psi C_* f''(K_*) \geq 0 \quad (65)$$

Hence, in the aggregate $\hat{C}_t = b \hat{K}_t$ with $b = r + \frac{\xi}{r + \varphi^s} m_K$, and using (59) we obtain the actual speed of mean-reversion of the economy:

$$\varphi = m_K \frac{\xi}{r + \varphi^s}. \quad (66)$$

Rational agent. The rational agent has a correct model ($\varphi^s = \varphi$, $m_K = 1$). So the speed of mean-reversion is $\varphi = \varphi^r$; (66) gives

$$\varphi^r = \frac{\xi}{r + \varphi^r} \quad (67)$$

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

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

Behavioral agent. Let us now endogenize m_φ and m_K . Given (64), $\frac{\partial \hat{c}_t}{\partial \varphi}|_{m_K=0} = 0$: 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 $m_\varphi = 0$: the optimal refinement in thinking about the speed of mean reversion is 0. Hence, $\varphi^s = \varphi^d$, the projected speed of mean-reversion (φ^s) is the default one (φ^d).⁴⁹ Given (66), this implies that the actual speed of mean-reversion is:

$$\varphi = m_K \varphi_0, \quad \varphi_0 := \frac{\xi}{r + \varphi^d}. \quad (68)$$

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

PROPOSITION 10. (Fluctuations are larger and more persistent with sparse agents) *Suppose that $\varphi^d \geq \min(\varphi, \varphi^r)$ and $\kappa > 0$. Then $\varphi < \varphi^r$: 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 $\varphi \geq r$, 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 11. *The speed of mean-reversion of the economy is: $\varphi = \frac{\varphi_0}{1 + \frac{B}{\varphi_0}}$ with $\varphi_0 := \frac{\xi}{r + \varphi^d}$, $B := \frac{2\kappa}{|u_{cc}|\sigma_\varepsilon^2}$. 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

49. With the iterated sparse max of Definition F.1 of the online appendix, we can have $m_\varphi > 0$.

$\varphi^d = \frac{1}{4}$), 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 modeling structure that could be promising in future quantitative work. Acemoglu and Jensen (2023) propose a further analysis of the neoclassical growth model with behavioral features.

5.3. Generalization: Recursive Competitive Equilibrium with Behavioral Agents

Here I define the behavioral extension of the Prescott and 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 $G^d(K) = \bar{K} + \varphi^d(K - \bar{K})$. A given m gives perceived laws of motions for private wealth $k' = F^k(c, k, K, m)$ and for the aggregate variables $K' = G(K, m)$.⁵⁰ In turn they generate a default value function $V^r(k, K | F^k(\cdot, m), G(\cdot, m))$, 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.⁵¹

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 (2012), Chapter 12).

DEFINITION 5. (*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:

$$r(K, m) = \bar{r} + m_K r'(\bar{K})(K - \bar{K}) + m_{K^2} [r(K) - \bar{r} - r'(\bar{K})(K - \bar{K})] \quad (69)$$

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

50. The present paper offers a systematic way in which agents simplify the moments they use to think about general equilibrium. This could be useful to systematize procedures such as the one in Krusell and Smith (1998). See also Barberis et al. (1998) and Hong et al. (2007) for an early analyses of financial equilibrium when agents use a simplified model.

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

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 $c(k, K, m)$ sparsely and optimally as:

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

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

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

To interpret (69), consider the case where the agent doesn't pay attention to nonlinear terms ($m_{K^2} = 0$),⁵² 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 $V^r(k, K \mid F^k(\cdot, m), G(\cdot, m))$ is the proxy value function under that simplified model, where everything is an AR(1). The value function was already calculated in Lemma 8, using $\rho_r = \rho_y = 1 - \varphi$.

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.

52. Here 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 F.1), the algorithm will select $m_{K^2} = 0$, i.e. remove the nonlinear terms.

6.1. Merton Portfolio Choice Model

A core model in dynamic finance is the Merton (1971) portfolio choice problem.⁵³ 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:

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

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

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

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

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

where H_t is the hedging demand term

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

Proposition 12 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 and Yogo (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, $\hat{\pi}_t$, and the agent thinks less about the mean-reversion of asset, the H_t terms. The fact that people are relatively insensitive

53. Section C.1 of the online appendix studies a more general Merton's problem with many assets and factors.

to the changes in the risk premium is a empirical central finding in Giglio et al. (2021). It also means that people’s demand for stocks is inelastic, so that flows in the stock markets have a large impact on prices (see Gabaix and Koijen (2022) and the references therein). Hence a simple model of inattention is useful to model basic facts of economic life.

In addition, the agent is here allowed to be “schizophrenic:” he pays attention to different things as a consumer than as an investor.⁵⁴ 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 C.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 C.2 derives how to systematically derive a BR version of those models. The models are very tractable (there is no matrix Ricatti 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 and Samwick (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 C.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 C.4. The firm uses a simplified model of future profits, which can lead it astray (see Greenwood and Hanson (2015)).

The Becker-Murphy model of rational addiction. The Becker and Murphy (1988) model of rational addiction is a peak of the use of rationality in economics. Section C.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

54. To 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).

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.⁵⁵

7. Discussion

7.1. Model Enrichments

Here I discuss different potential enrichments of the model; the online appendix (Section F) 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 F.3. The consequences are very minor.⁵⁶ 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 4 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:

$$\begin{aligned} V^{t=1,(q=1)}(w_1, x) &= u\left(\frac{w_1 + m_1^*(w_1, \sigma_x^2)x}{2}\right) + u\left(\frac{w_1 + (2 - m_1^*(w_1, \sigma_x^2))x}{2}\right) \\ &= V^{t=1,(q=0)}(w_1, x) + O(x^2). \end{aligned} \tag{70}$$

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 5, 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.⁵⁷

55. This 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.

56. First, 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.

57. Why not model iterated expectations, such as the agent's perception at time 0 of his perception at time 2 of his perception at time 5? The short answer is that this leads to a combinatorial explosion of the complexity of the model. See Section F.7.1.

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 F.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 F of the online appendix.

7.2. Frequently Asked Questions

How do markets clear? They clear as usual: prices adjust so that supply equals demand.⁵⁸ 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.⁵⁹ 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 Greenwood and Hanson (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.⁶⁰ 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

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

59. See also Hirshleifer et al. (2009), and Section 2.3.5 of Gabaix (2019).

60. There 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.

problem.⁶¹ Fortunately, 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.⁶²

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.”⁶³ 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).

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

62. See Proposition 12 and its discussion.

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

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 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 on 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” .⁶⁴ Second, they could have a high cost of cognition in some domains: in the “proportional thinking” specification of Section A.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.⁶⁵

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 Van Nieuwerburgh and Veldkamp (2010)). Behavioral agents like my sparse agents do not do Bayesian updating.⁶⁶ 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. (53)), which may not be exactly optimum in their context, but does guide them through life. Also,

64. Farhi and Gabaix (2020), (Section 6.1) discusses models with suboptimal allocation of attention.

65. See Chetty et al. (2009), Brown et al. (2010), Allcott and Taubinsky (2015), Rees-Jones and Taubinsky (2020) for recent progress.

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

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.

8. Literature Review

This paper builds on two strands of the literature: on a relatively recent strand of basic behavioral modeling 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 companion piece, Gabaix (2020), uses ideas in this paper to propose a behavioral version of the New Keynesian model. A related paper, Farhi and Gabaix (2020), proposes a behavioral version of textbook taxation theory. Hence, we have a more unified view of bounded rationality in micro and macro. In Gabaix (2019), which is a survey / manifesto, I try to convey how this and related approach to behavioral inattention unify tractably many behavioral biases, and mesh well with normal economics.

The “micro” sparse max relies on a strand of paper thinking about basic behavioral modeling, in particular inattention, salience. See, for instance, Bordalo et al. (2018), Chetty et al. (2009), Gabaix and Laibson (2006), Khaw et al. (2021), Kőszegi and Szeidl (2013), Lian (2021), Schwartzstein (2014).

Since this paper circulated (in the Spring 2012), a number of “behavioral macro” papers have been written. One particularly fruitful theme has been the application to monetary policy, including the “forward guidance puzzle” (McKay et al. (2016), Del Negro et al. (2023)). In Gabaix (2020) I proposed a tractable approach, building on the present paper: the inattention is there on the dynamics of the state vector, i.e. shrinking the ρ_r and ρ_y in (16)-(16), and all similar variables, by a “cognitive discounting” multiplicative factor \bar{m} (so that $\rho_r^s = \bar{m}\rho_r$). This makes agents dampened a bit the immediate future, and a lot the distant future. This proves useful to offer a unified treatment of agent’s myopia for monetary and fiscal policy, as well as optimal reaction by the government. Other approaches are incomplete information Angeletos and Lian (2018), and level- k thinking (Farhi and Werning (2019)), limited planning horizons (Woodford (2019)), and credit constraints (Bilbiie (2021)). While those approaches differ in a variety of small ways, Meggiorini (2023)

notes that the sparsity approach is indeed particularly tractable, which allows her, using it (especially in the “cognitive discounting” version) to develop a behavioral enrichment of the Smets and Wouters (2007) model, which widely used for policy analysis. This brings us closer to behavioral macro for usable for policy analysis.

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)). 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 problems such as (23) 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.⁶⁷ In a sense, this research vastly generalizes the Campbell and Mankiw (1989), 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

67. Actually, 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.

TABLE 1. Different Models of Myopia

	BR in static contexts	Rep. agent allowed	Determ. actions	Source dependent attention	Bayesian	Textbook micro
Rule-of-thumb			✓	✓		
Hyperbolics		✓	✓		✓	
Credit constraints			✓		✓	
Imperfectly updated info.			✓	✓	✓	
Noisy info.	✓			✓	✓	
Entropy penalty	✓				✓	
BGS / salience	✓	✓	✓	✓		
Sparsity	✓	✓	✓	✓		✓

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 (e.g. Kaplan et al. (2018)). 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 (2018)), but react in a myopic way (high MPC) to other transitory income shocks.

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 (2002), Mankiw and Reis (2002), Reis (2006), Abel et al. (2013)) or in other imperfect ways (Sargent et al. (1993), Fuster et al. (2012), Schwartzstein (2014), Bushong et al. (2021)). 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, though some pockets of tractability are being elaborated (Angeletos and Huo (2021)).

Other behavioral models of salience. Bordalo et al. (2013, 2018) are pursuing an influential research program modeling the impact of “salience” in a vast number of settings (see also Kőszegi and Szeidl (2013)). 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 Jung et al. (2019)) models, which do not allow for source-dependent inattention. Other papers (Maćkowiak and Wiederholt (2015), Woodford (2020)) 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 (2018)), 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 and Sargent (2008) 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 is likely useful to study other issues in macroeconomics and finance. One limitation of the proposed approach is its reliance on linearization and small deviations from the steady state. Overcoming this limitation would be a fruitful research challenge.

References

Abel, Andrew B., Janice C. Eberly, and Stavros Panageas (2013). “Optimal inattention to the stock market with information costs and transactions costs.” *Econometrica*, 81(4), 1455–1481.

- Acemoglu, Daron (2008). *Introduction to modern economic growth*. Princeton university press.
- Acemoglu, Daron, Vasco M Carvalho, Asuman Ozdaglar, and Alireza Tahbaz-Salehi (2012). “The network origins of aggregate fluctuations.” *Econometrica*, 80(5), 1977–2016.
- Acemoglu, Daron and Martin Kaae Jensen (2023). “Equilibrium Analysis in Behavioural One-Sector Growth Models.” *Review of Economic Studies*.
- Achdou, Yves, Jiequn Han, Jean-Michel Lasry, Pierre-Louis Lions, and Benjamin Moll (2022). “Income and wealth distribution in macroeconomics: A continuous-time approach.” *The Review of Economic Studies*, 89(1), 45–86.
- Aguiar, Mark and Erik Hurst (2013). “Deconstructing Life Cycle Expenditure.” *Journal of Political Economy*, 121(3), 437–492.
- Akerlof, George A. and Janet L. Yellen (1985). “Can small deviations from rationality make significant differences to economic equilibria?” *The American Economic Review*, 75(4), 708–720.
- Allcott, Hunt and Dmitry Taubinsky (2015). “Evaluating behaviorally motivated policy: Experimental evidence from the lightbulb market.” *American Economic Review*, 105(8), 2501–2538.
- Angeletos, George-Marios and Zhen Huo (2021). “Myopia and anchoring.” *American Economic Review*, 111(4), 1166–1200.
- Angeletos, George-Marios and Chen Lian (2018). “Forward guidance without common knowledge.” *American Economic Review*, 108(9), 2477–2512.
- Auclert, Adrien (2019). “Monetary policy and the redistribution channel.” *American Economic Review*, 109(6), 2333–2367.
- Bansal, Ravi and Amir Yaron (2004). “Risks for the long run: A potential resolution of asset pricing puzzles.” *The journal of Finance*, 59(4), 1481–1509.
- Barberis, Nicholas (2000). “Investing for the long run when returns are predictable.” *The Journal of Finance*, 55(1), 225–264.
- Barberis, Nicholas, Andrei Shleifer, and Robert Vishny (1998). “A model of investor sentiment.” *Journal of financial economics*, 49(3), 307–343.
- Barro, Robert J (1974). “Are Government Bonds Net Wealth?” *Journal of Political Economy*, 82(6), 1095–1117.
- Barro, Robert J (1999). “Ramsey meets Laibson in the neoclassical growth model.” *The Quarterly Journal of Economics*, 114(4), 1125–1152.
- Barro, Robert J (2009). “Rare disasters, asset prices, and welfare costs.” *American Economic Review*, 99(1), 243–264.
- Basak, Suleyman and Georgy Chabakauri (2010). “Dynamic mean-variance asset allocation.” *The Review of Financial Studies*, 23(8), 2970–3016.
- Becker, Gary S and Kevin M Murphy (1988). “A theory of rational addiction.” *Journal of political Economy*, 96(4), 675–700.
- Bernheim, B Douglas, Jonathan Skinner, and Steven Weinberg (2001). “What accounts for the variation in retirement wealth among US households?”

- American Economic Review*, 91(4), 832–857.
- Beshears, John, James J Choi, Joshua Hurwitz, David Laibson, and Brigitte C Madrian (2015). “Liquidity in retirement savings systems: An international comparison.” *American Economic Review*, 105(5), 420–425.
- Bigio, Saki and Jennifer La’o (2020). “Distortions in production networks.” *The Quarterly Journal of Economics*, 135(4), 2187–2253.
- Bilbiie, Florin O (2021). “Monetary Policy and Heterogeneity: An Analytical Framework.” *CEPR Discussion Paper No. 12601*.
- Blanchard, Olivier and Stanley Fisher (1989). *Lectures on Macroeconomics*. The MIT Press, Cambridge, Mass.
- Bordalo, Pedro, Nicola Gennaioli, and Andrei Shleifer (2013). “Salience and Consumer Choice.” *Journal of Political Economy*, 121(5), 803–843.
- Bordalo, Pedro, Nicola Gennaioli, and Andrei Shleifer (2018). “Diagnostic expectations and credit cycles.” *The Journal of Finance*, 73(1), 199–227.
- Brown, Jennifer, Tanjim Hossain, and John Morgan (2010). “Shrouded attributes and information suppression: Evidence from the field.” *The Quarterly Journal of Economics*, 125(2), 859–876.
- Bushong, Benjamin, Matthew Rabin, and Joshua Schwartzstein (2021). “A model of relative thinking.” *The Review of Economic Studies*, 88(1), 162–191.
- Caballero, Ricardo J. (1995). “Near-Rationality, Heterogeneity, and Aggregate Consumption.” *Journal of Money, Credit and Banking*, 27(1), 29–48.
- Caballero, Ricardo J and Alp Simsek (2013). “Fire sales in a model of complexity.” *The Journal of Finance*, 68(6), 2549–2587.
- Campbell, John Y. and N. Gregory Mankiw (1989). “Consumption, Income, and Interest Rates: Reinterpreting the Time Series Evidence.” *NBER Macroeconomics Annual*, 4, 185–216.
- Campbell, John Y and Luis M Viceira (2002). *Strategic asset allocation: portfolio choice for long-term investors*. Clarendon Lectures in Economic.
- Candes, Emmanuel J. and Terence Tao (2006). “Near-optimal signal recovery from random projections: Universal encoding strategies?” *IEEE Transactions on Information Theory*, 52(12), 5406–5425.
- Carroll, Christopher D. and Andrew A. Samwick (1997). “The nature of precautionary wealth.” *Journal of Monetary Economics*, 40(1), 41–71.
- Carvalho, Vasco M. and Basile Grassi (2019). “Large Firm Dynamics and the Business Cycle.” *American Economic Review*, 109(4), 1375–1425.
- Chetty, Raj (2012). “Bounds on elasticities with optimization frictions: A synthesis of micro and macro evidence on labor supply.” *Econometrica*, 80(3), 969–1018.
- Chetty, Raj, Adam Looney, and Kory Kroft (2009). “Salience and taxation: Theory and evidence.” *American economic review*, 99(4), 1145–1177.
- Cochrane, John H. (1989). “The Sensitivity of Tests of the Intertemporal Allocation of Consumption to Near-Rational Alternatives.” *The American Economic Review*, 79(3), 319–337.

- Del Negro, Marco, Marc P. Giannoni, and Christina Patterson (2023). “The Forward Guidance Puzzle.” *Journal of Political Economy Macroeconomics*, 1(1), 43–79.
- Evans, George W and Seppo Honkapohja (2001). *Learning and Expectations in Macroeconomics*. Princeton University Press.
- Evans, George W. and Seppo Honkapohja (2013). *Rethinking Expectations : The Way Forward for Macroeconomics*, chap. 2. Learning as a Rational Foundation for Macroeconomics and Finance, pp. 68–111. Princeton University Press.
- Farhi, Emmanuel and Xavier Gabaix (2020). “Optimal Taxation with Behavioral Agents.” *American Economic Review*, 110(1), 298–336.
- Farhi, Emmanuel and Iván Werning (2019). “Monetary Policy, Bounded Rationality, and Incomplete Markets.” *American Economic Review*, 109(11), 3887–3928.
- Fuster, Andreas, Benjamin Hebert, and David Laibson (2012). “Natural expectations, macroeconomic dynamics, and asset pricing.” *NBER Macroeconomics Annual*, 26(1), 1–48.
- Fuster, Andreas, David Laibson, and Brock Mendel (2010). “Natural expectations and macroeconomic fluctuations.” *The Journal of Economic Perspectives*, 24(4), 67–84.
- Gabaix, Xavier (2009). “Linearity-Generating Processes: A Modelling Tool Yielding Closed Forms for Asset Prices.” *Working paper*.
- Gabaix, Xavier (2011). “The granular origins of aggregate fluctuations.” *Econometrica*, 79(3), 733–772.
- Gabaix, Xavier (2012). “Variable Rare Disasters: An Exactly Solved Framework for Ten Puzzles in Macro-Finance.” *Quarterly Journal of Economics*, 127(2), 645–700.
- Gabaix, Xavier (2014). “A sparsity-based model of bounded rationality.” *Quarterly Journal of Economics*, 129(4), 1661–1710.
- Gabaix, Xavier (2019). “Behavioral Inattention.” *Handbook of Behavioral Economics*, 2, 261–344.
- Gabaix, Xavier (2020). “A behavioral New Keynesian model.” *American Economic Review*, 110(8), 2271–2327.
- Gabaix, Xavier and Ralph SJ Koijen (2022). “In search of the origins of financial fluctuations: The inelastic markets hypothesis.” *NBER Working Paper No. 28967*.
- Gabaix, Xavier and David Laibson (2002). “The 6D Bias and the Equity-Premium Puzzle.” *NBER Macroeconomics Annual*, 16, 257–312.
- Gabaix, Xavier and David Laibson (2006). “Shrouded Attributes, Consumer Myopia, and Information Suppression in Competitive Markets.” *Quarterly Journal of Economics*, 121(2), 505–540.
- Gabaix, Xavier, David Laibson, Guillermo Moloche, and Stephen Weinberg (2006). “Costly Information Acquisition: Experimental Analysis of a Boundedly Rational Model.” *American Economic Review*, 96(4), 1043–1068.

- Ganong, Peter and Pascal Noel (2019). “Consumer spending during unemployment: Positive and normative implications.” *American economic review*, 109(7), 2383–2424.
- Giglio, Stefano, Matteo Maggiori, Johannes Stroebel, and Stephen Utkus (2021). “Five facts about beliefs and portfolios.” *American Economic Review*, 111(5), 1481–1522.
- Glimcher, Paul W. (2011). *Foundations of Neuroeconomic Analysis*. Oxford University Press.
- Gourinchas, Pierre-Olivier and Jonathan A Parker (2002). “Consumption over the life cycle.” *Econometrica*, 70(1), 47–89.
- Greenwood, Robin and Samuel G Hanson (2015). “Waves in ship prices and investment.” *The Quarterly Journal of Economics*, 130(1), 55–109.
- Gruber, Jonathan and Botond Köszegi (2001). “Is Addiction Rational? Theory and Evidence.” *Quarterly Journal of Economics*, 116(4), 1261–1303.
- Hall, Robert E (1988). “Intertemporal substitution in consumption.” *Journal of political economy*, 96(2), 339–357.
- Hansen, Lars Peter and Thomas J Sargent (2008). *Robustness*. Princeton university press.
- Harris, Christopher and David Laibson (2001). “Dynamic choices of hyperbolic consumers.” *Econometrica*, 69(4), 935–957.
- Harris, Christopher and David Laibson (2013). “Instantaneous gratification.” *The Quarterly Journal of Economics*, 128(1), 205–248.
- Hirshleifer, David, Sonya Seongyeon Lim, and Siew Hong Teoh (2009). “Driven to distraction: Extraneous events and underreaction to earnings news.” *The Journal of Finance*, 64(5), 2289–2325.
- Hong, Harrison, Jeremy C Stein, and Jialin Yu (2007). “Simple forecasts and paradigm shifts.” *The Journal of Finance*, 62(3), 1207–1242.
- Hsieh, Chang-Tai (2003). “Do consumers react to anticipated income changes? Evidence from the Alaska permanent fund.” *American Economic Review*, 93(1), 397–405.
- Judd, Kenneth L (1998). *Numerical methods in economics*. MIT press.
- Jung, Junehyuk, Jeong Ho Kim, Filip Matějka, and Christopher A Sims (2019). “Discrete actions in information-constrained decision problems.” *The Review of Economic Studies*, 86(6), 2643–2667.
- Kahneman, Daniel (2011). *Thinking, fast and slow*. Macmillan.
- Kaplan, Greg, Benjamin Moll, and Giovanni L Violante (2018). “Monetary policy according to HANK.” *American Economic Review*, 108(3), 697–743.
- Khaw, Mel Win, Ziang Li, and Michael Woodford (2021). “Cognitive imprecision and small-stakes risk aversion.” *The Review of Economic Studies*, 88(4), 1979–2013.
- Köszegi, Botond and Adam Szeidl (2013). “A Model of Focusing in Economic Choice.” *Quarterly Journal of Economics*, 128(1), 53–104.
- Krusell, Per and Anthony A Smith, Jr (1998). “Income and wealth heterogeneity in the macroeconomy.” *Journal of political Economy*, 106(5),

- 867–896.
- Kueng, Lorenz (2018). “Excess sensitivity of high-income consumers.” *The Quarterly Journal of Economics*, 133(4), 1693–1751.
- Laibson, David (1997). “Golden Eggs and Hyperbolic Discounting.” *Quarterly Journal of Economics*, 112(2), 443–478.
- Lian, Chen (2021). “A theory of narrow thinking.” *Review of Economic Studies*, 88(5), 2344–2374.
- Ljungqvist, Lars and Thomas J. Sargent (2012). *Recursive Macroeconomic Theory*. The MIT Press, URL <http://www.jstor.org/stable/j.ctt5vjq05>.
- Lucas Jr, Robert E (1976). “Econometric policy evaluation: A critique.” In *Carnegie-Rochester conference series on public policy*, vol. 1, pp. 19–46, North-Holland.
- Lusardi, Annamaria and Olivia S Mitchell (2011). “Financial literacy and retirement planning in the United States.” *Journal of pension economics & finance*, 10(4), 509–525.
- Lusardi, Annamaria, Daniel Schneider, Peter Tufano, Adair Morse, and Karn M. Pence (2011). “Financially Fragile Households: Evidence and Implications.” *Brookings Papers on Economic Activity*, pp. 83–150.
- Maćkowiak, Bartosz and Mirko Wiederholt (2015). “Business Cycle Dynamics Under Rational Inattention.” *Review of Economic Studies*, 82(4), 1502–1532.
- Malamud, Semyon and Grigory Vilkov (2018). “Non-myopic betas.” *Journal of Financial Economics*, 129(2), 357–381.
- Mankiw, N. Gregory and Ricardo Reis (2002). “Sticky Information versus Sticky Prices: A Proposal to Replace the New Keynesian Phillips curve.” *Quarterly Journal of Economics*, 117(4), 1295–1328.
- McKay, Alisdair, Emi Nakamura, and Jón Steinsson (2016). “The Power of Forward Guidance Revisited.” *American Economic Review*, 106(10), 3133–3158.
- Meggiorini, Greta (2023). “Behavioral New Keynesian Models: An empirical assessment.” *Journal of Macroeconomics*, 77, 103538.
- Merton, Robert C (1971). “Optimum consumption and portfolio rules in a continuous-time model.” *Journal of Economic Theory*, 3(4), 373–413.
- Modigliani, Franco and Richard H. Brumberg (1954). “Utility analysis and the consumption function: an interpretation of cross-section data.” In *Post-Keynesian Economics*, edited by Kenneth K. Kurihara, pp. 388–436. Rutgers University Press, New Brunswick, NJ.
- O’Donoghue, Ted and Matthew Rabin (2001). “Choice and procrastination.” *The Quarterly Journal of Economics*, 116(1), 121–160.
- Prescott, Edward C. and Rajnish Mehra (1980). “Recursive Competitive Equilibrium: The Case of Homogeneous Households.” *Econometrica*, 48(6), 1365–1379.
- Rabin, Matthew (2013). “An approach to incorporating psychology into economics.” *American Economic Review*, 103(3), 617–622.

- Rees-Jones, Alex and Dmitry Taubinsky (2020). “Measuring ‘schmeduling’.” *The Review of Economic Studies*, 87(5), 2399–2438.
- Reis, Ricardo (2006). “Inattentive consumers.” *Journal of Monetary Economics*, 53(8), 1761–1800.
- Romer, David (2018). *Advanced Macroeconomics*. New York: McGraw Hill Companies.
- Samuelson, William and Richard Zeckhauser (1988). “Status quo bias in decision making.” *Journal of risk and uncertainty*, 1, 7–59.
- Sargent, Thomas J et al. (1993). *Bounded rationality in macroeconomics*. Oxford University Press.
- Schwartzstein, Joshua (2014). “Selective attention and learning.” *Journal of the European Economic Association*, 12(6), 1423–1452.
- Sims, Christopher A (2003). “Implications of rational inattention.” *Journal of Monetary Economics*, 50(3), 665–690.
- Smets, Frank and Rafael Wouters (2007). “Shocks and Frictions in US Business Cycles: A Bayesian DSGE Approach.” *American Economic Review*, 97(3), 586–606.
- Thaler, Richard (1985). “Mental Accounting and Consumer Choice.” *Marketing Science*, 4(3), 199–214.
- Tibshirani, Robert (1996). “Regression Shrinkage and Selection via the Lasso.” *Journal of the Royal Statistical Society. Series B (Methodological)*, 58(1), 267–288.
- Tversky, Amos and Daniel Kahneman (1981). “The framing of decisions and the psychology of choice.” *science*, 211(4481), 453–458.
- Van Nieuwerburgh, Stijn and Laura Veldkamp (2010). “Information acquisition and under-diversification.” *The Review of Economic Studies*, 77(2), 779–805.
- Veldkamp, Laura L. (2011). *Information Choice in Macroeconomics and Finance*. Princeton University Press.
- Wachter, Jessica A and Motohiro Yogo (2010). “Why do household portfolio shares rise in wealth?” *The Review of Financial Studies*, 23(11), 3929–3965.
- Wang, Chong, Neng Wang, and Jinqiang Yang (2016). “Optimal consumption and savings with stochastic income and recursive utility.” *Journal of Economic Theory*, 165, 292–331.
- Woodford, Michael (2013). “Macroeconomic Analysis without the Rational Expectations Hypothesis.” *Annual Reviews of Economics*, 5, 303–346.
- Woodford, Michael (2019). “Monetary policy analysis when planning horizons are finite.” *NBER Macroeconomics Annual*, 33(1), 1–50.
- Woodford, Michael (2020). “Modeling imprecision in perception, valuation, and choice.” *Annual Review of Economics*, 12, 579–601.