

Thinking on Their Feet: Along Main Street*

Sergiy Verstyuk[†]

This draft: 4 November 2017

Abstract

This paper considers the problem of learning and decision-making in a dynamic stochastic economic environment by agents subject to information processing constraints. An agent endogenously chooses to operate in terms of a simplified model of the economy, which implies: a delayed, if at all, updating of the estimates of evolving states/random variables' conditioning parameters; as well as the entropy reduction, or even its complete “folding” that drops the less important variables from the agent's approximating model. Specifically, parameter learning is implemented relying on computational complexity theory, which produces a constrained version of the standard Kalman filter. The latter leads to a less than one-for-one reaction to the newly observed information, without the need to postulate e.g. habit formation; which is responsible for an underreaction to permanent parameter changes (“stickiness”), as well as for an overreaction to transitory shocks (“overshooting”). In a standard stochastic growth model with government transfers, such agents may fail to realize that a fiscal expansion now necessitates a compensatory fiscal contraction later, which implies the effectiveness, in certain sense, of the fiscal stimulus policy (albeit at the expense of efficiency losses) and a violation of the Ricardian equivalence. Numerical simulations suggest high fiscal multipliers, with the effects relatively stronger at times of economic recession. Being the outcomes of endogenous choices of rational agents, these results are immune to the Lucas critique.

*Preliminary version. Acknowledgements: to be added.

[†]Contact address: verstyuk@cmsa.fas.harvard.edu.

Contents

1	Introduction	1
1.1	Literature	2
2	Model	3
2.1	Setup and infeasible problem	3
2.2	Feasible problem (formulation)	4
2.3	Feasible problem (solution)	9
3	Discussion	12
4	Numerical exercise	14
5	Conclusion	18
A	Generalized information processing capacity	19
A.1	Formulation	19
A.2	Corresponding optimization problem	20
B	Some formulas for stochastic relationships	21
C	Optimality and steady state conditions	23
D	Parameter values used	25
E	Baseline scenario simulations (for comparison)	26
F	Proof of Proposition 1	26
G	Proof of Proposition 2	27
H	Proof of Proposition 3	29
I	Proof of Proposition 4	30

The chief difficulty Alice found at first was in managing her flamingo: she succeeded in getting its body tucked away, comfortably enough, under her arm, with its legs hanging down, but generally, just as she had got its neck nicely straightened out, and was going to give the hedgehog a blow with its head, it would twist itself round and look up in her face, with such a puzzled expression that she could not help bursting out laughing: and when she had got its head down, and was going to begin again, it was very provoking to find that the hedgehog had unrolled itself, and was in the act of crawling away: besides all this, there was generally a ridge or furrow in the way wherever she wanted to send the hedgehog to, and, as the doubled-up soldiers were always getting up and walking off to other parts of the ground, Alice soon came to the conclusion that it was a very difficult game indeed.

—Lewis Carroll, *Alice's Adventures in Wonderland*

1 Introduction

We are interested in the behavior of economic agents whose information-processing capacities are limited. Consequently, they resort to subjective simplification of the external stochastic environment when solving their optimization problem as well as making their economic and financial decisions. This study is an extension of the earlier work (Verstyuk, 2017); now we consider an economic environment with richer non-i.i.d. stochastic dynamics where the process of learning the evolving parameters/state variables takes center stage.

Specifically, we take a real business cycle/stochastic growth model as a workhorse, and extend it with costly information-processing (plus government transfers standing in for fiscal policy as well as a storage technology standing in for savings/investments in riskless bonds). The agent chooses a simplified, approximating model of the economy and ignores some dimensions of the stochastic environment he/she operates in; and even to those dimensions he does not ignore, adjusts only with a delay. As a result, the agent commits systematic, “forced” mistakes. In particular, he is unable to (i.e., endogenously chooses not to) take into account that the fiscal expansion today must be balanced by a fiscal contraction of comparable magnitude later. This triggers the breakdown of the “Ricardian equivalence” and engenders the effectiveness of “fiscal stimulus”.

Thus, the constraints on agents’ information-processing capacities imply that fiscal policy becomes potent: “stimulus works”. But at the same time these constraints prevent the agents from achieving a “first-best” optimum, as a result the optimal fiscal policy can only strive for “second-best” outcomes.

It is important to emphasize that the Lucas critique does not apply in our model: if the fiscal stimulus is large and/or long enough, the agents readjust (it would be both optimal and feasible for them to do so).

In our numerical simulations, the economic effects are potentially powerful: the “fiscal multiplier” for consumption is well above 1. Also, the effects are state-dependent: for instance, they are stronger at times of economic recession.

Turning to methodological aspects, while in a less involved case of i.i.d. stochastic dynamics of Verstyuk (2017) we relied on the methods based on information theory, in this work we adopt the approaches from computational complexity theory and theoretical computer science more broadly. (In terms of computational complexity, these two formalisms roughly correspond to the concepts of communication complexity and time complexity, respectively; moreover, in our framework they can be straightforwardly reconciled and formulated in a unified fashion, as we show in Appendix §A.)

Last but not least, our relatively more abstract theoretical results include the reformulation and refinement of the Kalman filter that accounts for computational constraints. Also, within our framework the information-processing capacity constraints in general induce endogenously such biases in agents’ perception and behavior as “stickiness” (i.e., “unreaction” to parameter/state variable changes) and “overshooting” (i.e., “overreaction” to noise). (Note that delayed response occurs here without the need to resort to adjustment costs or appeal to habit-formation arguments.)

1.1 Literature

Broadly speaking, we aim to improve our understanding of some puzzling aspects in the consumption (and savings/investment) responses to income shocks, bringing about such fundamental issues as the “propensity to consume” and “Ricardian equivalence”. This is a long-standing research problem, stretching to—we limit ourselves to modern formulations only—classical Friedman (1957) and Campbell and Deaton (1989) (also see Campbell and Cochrane, 1999; Carroll et al., 2000; Fuhrer, 2000; or, alternatively, Carroll, 2001) as well as Barro (1974). More recently, a careful empirical analysis by Parker (2017) examines the households’ propensity to consume the 2008 economic stimulus payments in the US, paying a particular attention to the decision-theoretic aspects involved.

It is the policy facet of the above consumption and income interrelationship that interests us most. Specifically, this work models the effect of fiscal policy on aggregate economic outcomes, for instance it deals with the notion of a “fiscal multiplier”. The relevant fresh theoretical research in this area includes Woodford (2011) as well as Farhi and Werning (2016). The empirical studies have been undertaken using both structural (e.g., Lucas, 2016; Faria-e-Castro, 2017) as well as non-structural (e.g., Auerbach and Gorodnichenko, 2012) approaches. A good concise overview is available from the Congressional Budget Office (2012).

The channel through which fiscal policy becomes effective here rests not on market structure (a combination of monopolistic competition and nominal rigidities) or financial (net worth/balance sheet constraints) considerations, but rather on decision-theoretic underpinnings. Such a mechanism is based on the optimizing behavior of agents, thus

departing from and developing further the approach based on postulated “rule-of-thumb” behavior (due to Campbell and Mankiw, 1989; also see Mankiw, 2000, and Galí et al., 2007), and generally sharing the thinking and ambition of Gabaix (2014a, 2016b).

Our mechanism behind the potency of fiscal policy is one specific example of the broader mechanisms that restrain the (extent of) decision-makers’—and thus endogenous economic variables’—responses to exogenous shocks. The potential importance of such “dampening” effects have been examined recently in Gabaix (2016a, 2016b), Cochrane (2016) Angeletos and Lian (2016, 2017) as well as Farhi and Werning (2017); but also see an earlier proposal on modeling this kind of “stickiness” in economic dynamics by Sims (1998).

The conceptual approach builds on Verstyuk (2017). In contrast to what has been done there, here to formalize the information processing capabilities of decision-makers in the context of learning tasks the concept of computational complexity that we use is the complexity of mathematical operations (i.e., time complexity rather than communication complexity). For a comprehensive textbook introduction, see Arora and Barak (2009) as well as Cormen et al. (2009).

2 Model

2.1 Setup and infeasible problem

The problem formulation is as follows (using the same notation as in Verstyuk, 2017). The agent’s **original/infeasible problem** is:

$$v^\sharp(\mathbf{K}_{t-1}, \mathbf{z}_t, \boldsymbol{\mu}_{t|t}) = \max_{C_t, \mathbf{K}_t, \mathbf{L}_t} \{u(C_t, \mathbf{L}_t) + \beta \mathbb{E}_t^g [v(\mathbf{K}_t, \mathbf{z}_{t+1}, \boldsymbol{\mu}_{t+1|t+1})]\} \quad \{\mathcal{P}_K\}$$

subject to

$$C_t + \mathbf{V}_t^\top \mathbf{q}_t + \mathbf{1}^\top \mathbf{K}_t = (\mathbf{V}_t + \mathbf{D}_t)^\top \mathbf{q}_{t-1} + (\mathbf{r}_t + (\mathbf{1} - \boldsymbol{\delta}))^\top \mathbf{K}_{t-1} + \mathbf{w}_t^\top \mathbf{L}_t =: W_t,$$

where C_t is a sum of perfectly substitutable consumption goods produced by K production technologies. Utility function is

$$u(C_t) = \varpi \frac{C_t^{1-\gamma}}{1-\gamma} - \sum_{k=1}^K \frac{L_{k,t}^{1+\eta_k}}{1+\eta_k}.$$

The production function is CRS Cobb-Douglas,

$$Y_{k,t} = (Z_{k,t} - \check{Z}_k) K_{k,t-1}^{\alpha_k} L_{k,t}^{1-\alpha_k}, \quad \forall k \in \{1, \dots, K\},$$

where $Z_{k,t}$, $K_{k,t-1}$ and $L_{k,t}$ are, respectively, technological productivity, capital and labor (with production function-specific normalizing constant \check{Z}_k). The firm’s problem gives

$$\begin{aligned} r_{k,t} K_{k,t-1} &= \alpha_k Y_{k,t}, & \forall k \in \{1, \dots, K\}, \\ w_{k,t} L_{k,t} &= (1 - \alpha_k) Y_{k,t}, & \forall k \in \{1, \dots, K\}, \end{aligned}$$

where the returns on capital \mathbf{r}_t and wage rates \mathbf{w}_t equal the factor marginal products. The firms rent capital from consumers-investors-workers, and are owned by the latter. Since firms are competitive, profits and dividends \mathbf{D}_t , and hence firm value \mathbf{V}_t , are all identically 0, and can be dropped from the agent's budget constraint.

The stochastic process for productivity \mathbf{Z}_t is captured by the following dynamic system, defined in terms of $z_{k,t} := \ln Z_{k,t}$. The state equation is

$$\boldsymbol{\mu}_t = \mathbf{A}_\mu \boldsymbol{\mu}_{t-1} + \boldsymbol{\varepsilon}_{\mu,t},$$

where

$$\boldsymbol{\mu}_t = \begin{bmatrix} \boldsymbol{\mu}_{0,t} \\ \boldsymbol{\mu}_{1,t} \end{bmatrix}, \quad \mathbf{A}_\mu = \begin{bmatrix} \mathbf{I}_K & \mathbf{0}_{K \times K} \\ \mathbf{I}_K & \boldsymbol{\rho} \end{bmatrix}, \quad \boldsymbol{\varepsilon}_{\mu,t} = \begin{bmatrix} \boldsymbol{\varepsilon}_{\mu 0,t} \\ \boldsymbol{\varepsilon}_{\mu 1,t} \end{bmatrix} \sim \mathcal{N} \left(\begin{bmatrix} \mathbf{0}_{K \times 1} \\ \mathbf{0}_{K \times 1} \end{bmatrix}, \begin{bmatrix} \boldsymbol{\Sigma}_{\varepsilon \mu 0} & \mathbf{0}_{K \times K} \\ \mathbf{0}_{K \times K} & \boldsymbol{\Sigma}_{\varepsilon \mu 1} \end{bmatrix} \right);$$

and the observation equation is

$$\mathbf{z}_t = \mathbf{A}_z \boldsymbol{\mu}_t + \boldsymbol{\varepsilon}_{z,t},$$

where

$$\mathbf{A}_z = \begin{bmatrix} \mathbf{0}_{K \times K} & \mathbf{I}_K \end{bmatrix}, \quad \boldsymbol{\varepsilon}_{z,t} \sim \mathcal{N}(\mathbf{0}_{K \times 1}, \boldsymbol{\Sigma}_{\varepsilon z}).$$

Henceforth, notation $\mathbf{0}_{r \times c}$ denotes a matrix of zeros with dimensionality $r \times c$; notation \mathbf{I}_r denotes an identity matrix of dimensionality r .

The current state $\boldsymbol{\mu}_t$ is unobserved and must be learned from observed variables \mathbf{z}_t using a (standard) Kalman filter, producing $\boldsymbol{\mu}_{t|t}$:

$$\begin{aligned} \boldsymbol{\mu}_{t|t-1} &= \mathbf{A}_\mu \boldsymbol{\mu}_{t-1|t-1}, \\ \boldsymbol{\Upsilon}_{\mu,t|t-1} &= \mathbf{A}_\mu \boldsymbol{\Upsilon}_{\mu,t-1|t-1} \mathbf{A}_\mu^\top + \boldsymbol{\Sigma}_{\varepsilon \mu}, \\ \boldsymbol{\mu}_{t|t} &= \boldsymbol{\mu}_{t|t-1} + \mathbf{B}_t (\mathbf{z}_t - \mathbf{A}_z \boldsymbol{\mu}_{t|t-1}), \\ \boldsymbol{\Upsilon}_{\mu,t|t} &= (\mathbf{I}_K - \mathbf{B}_t \mathbf{A}_z) \boldsymbol{\Upsilon}_{\mu,t|t-1}, \end{aligned}$$

where Kalman gain \mathbf{B}_t solves

$$\mathbf{B}_t := \arg \left\{ \min_{\mathbf{B}_t} \mathbb{E}^g \left[\|\boldsymbol{\mu}_t - \boldsymbol{\mu}_{t|t}\|_2^2 \right] \right\}$$

and equals

$$\mathbf{B}_t = \boldsymbol{\Upsilon}_{\mu,t|t-1} \mathbf{A}_z^\top (\mathbf{A}_z \boldsymbol{\Upsilon}_{\mu,t|t-1} \mathbf{A}_z^\top + \boldsymbol{\Sigma}_{\varepsilon z})^{-1}.$$

Intuitively, the motive for estimating $\boldsymbol{\mu}_{t|t}$ is to improve the forecast (i.e., reduce the conditional variance) of \mathbf{z}_{t+1} .

2.2 Feasible problem (formulation)

Solving the optimization problem is affected by the potentially binding information-processing capacity constraints: 1) the learning of the parameters of the relevant probability distributions is subject to the mathematical operations complexity constraint; 2)

the maximization of the value function and the evaluation of the integral on the RHS of the Bellman equation is subject to the mutual information constraint. The agent's **generalized/feasible problem** is therefore:

$$v(\mathbf{K}_{t-1}, \hat{\mathbf{z}}_t, \tilde{\boldsymbol{\mu}}_{t|t}) = \max_{C_t, \mathbf{K}_t, \mathbf{L}_t} \{u(C_t, \mathbf{L}_t) + \beta \mathbb{E}_t^h [v(\mathbf{K}_t, \hat{\mathbf{z}}_{t+1}, \tilde{\boldsymbol{\mu}}_{t+1|t+1})]\} \quad \{\mathcal{P}_{\mathcal{K}\mathcal{T}\mathcal{I}}\}$$

subject to

$$C_t + \mathbf{V}_t^\top \mathbf{q}_t + \mathbf{1}^\top \mathbf{K}_t = (\mathbf{V}_t + \mathbf{D}_t)^\top \mathbf{q}_{t-1} + (\hat{\mathbf{r}}_t + (\mathbf{1} - \boldsymbol{\delta}))^\top \mathbf{K}_{t-1} + \hat{\mathbf{w}}_t^\top \mathbf{L}_t =: \hat{W}_t,$$

with utility function and firm production relations as above. One module is the familiar consumption-investment-labor problem we have seen in part §2.1.

We are given the probability distribution $g(\mathbf{z}_{t+1}|\boldsymbol{\mu}_{t+1})$ (which can be more usefully reformulated as $g(\mathbf{z}_{t+1}|\tilde{\boldsymbol{\mu}}_{t+1|t})$, as is done later) that represents the dynamic system introduced in §2.1; and the distribution $h(\hat{\mathbf{z}}_{t+1}|\hat{\boldsymbol{\mu}}_{t+1})$ (also, $h(\hat{\mathbf{z}}_{t+1}|\hat{\boldsymbol{\mu}}_{t+1|t})$) is a marginal of $f(\mathbf{z}_{t+1}, \hat{\mathbf{z}}_{t+1}|\boldsymbol{\mu}_{t+1}, \hat{\boldsymbol{\mu}}_{t+1})$ (which is defined below).

Objects $\tilde{\boldsymbol{\mu}}_{t|t}$ and $h(\hat{\mathbf{z}}_t|\hat{\boldsymbol{\mu}}_{t|t})$ solve an auxiliary complexity/information processing problem (which in turn consists of learning/updating and informational/evaluation-optimization sub-parts)

$$\min_{\hat{\mathbf{B}}_t, f(\cdot, \cdot)} \mathbb{E}_t^f [d(v^\#(\mathbf{K}_t, \mathbf{z}_{t+1}, \boldsymbol{\mu}_{t+1|t+1}), v(\mathbf{K}_t, \hat{\mathbf{z}}_{t+1}, \tilde{\boldsymbol{\mu}}_{t+1|t+1}))] \quad \{\mathcal{P}_{\mathcal{T}\mathcal{I}}\}$$

subject to constraints

$$\begin{aligned} \mathcal{T}(\tilde{\boldsymbol{\mu}}_{t+1|t}, \hat{\boldsymbol{\mu}}_{t+1|t} \| \mathbf{z}_t, \tilde{\boldsymbol{\mu}}_{t|t-1}) &\leq \kappa_{\mathcal{T}}, \\ \mathcal{I}(g(\mathbf{z}_{t+1}|\tilde{\boldsymbol{\mu}}_{t+1|t}); h(\hat{\mathbf{z}}_{t+1}|\hat{\boldsymbol{\mu}}_{t+1|t})) &\leq \kappa_{\mathcal{I}} \end{aligned}$$

(see Appendix §A for the definitions and additional formal details).

Here, information constraint $\mathcal{I}(g(\cdot); h(\cdot)) \leq \kappa$ means that the chosen simplified distribution $h(\cdot)$, which approximates a given complex distribution $g(\cdot)$ and is used for finding optimal choices of the control variables $\{C_t^*, \mathbf{K}_t^*, \mathbf{L}_t^*\}$, does not exceed the stochastic evaluation capacity bound $\kappa_{\mathcal{I}}$ (this is based on the concept of mutual information and entropy more broadly, details can be found in, e.g., Cover and Thomas, 2006; most closely it is related to the concept of communication complexity, see, e.g., Arora and Barak, 2009). It was introduced and discussed in Verstyuk (2017). Intuitively, capacity $\kappa_{\mathcal{I}}$ may be thought of as the complexity of the (approximating) probability distribution that can be feasibly evaluated and used in stochastic optimization.

Additionally, learning/updating constraint $\mathcal{T}(\tilde{\boldsymbol{\mu}}_{t+1|t}, \hat{\boldsymbol{\mu}}_{t+1|t} \| \cdot) \leq \kappa_{\mathcal{T}}$ means that the execution of the Kalman filter updating and mean-adjustment procedures producing, respectively, $\tilde{\boldsymbol{\mu}}_{t+1|t}$ and $\hat{\boldsymbol{\mu}}_{t+1|t}$, does not exceed the mathematical operations capacity bound $\kappa_{\mathcal{T}}$ (this is close to the concept of time complexity, see Arora and Barak, 2009).¹ Intuitively, capacity $\kappa_{\mathcal{T}}$ may be thought of as the number of the (approximating) model's

¹Perhaps a more familiar way of formalizing the notion of time complexity of mathematical operations is specifying the complexity class a function/procedure belongs to, e.g., $\mathcal{O}(K^3)$ or $\Theta(K^3)$. But that approach provides instead the number of operations' bound defined only asymptotically up to a normalizing constant, as a function of algorithm's dimensionality, i.e., $2K$ in our case.

parameters that can feasibly be re-calculated and used for updating the stochastic optimization problem's setup.²

The complexity of updating conditional means is measured as

$$\mathcal{T}(\tilde{\boldsymbol{\mu}}_{t+1|t}, \hat{\boldsymbol{\mu}}_{t+1|t} \| \mathbf{z}_t, \tilde{\boldsymbol{\mu}}_{t|t-1}) := \text{time}(\{\mathcal{P}_{\rightsquigarrow}, \mathcal{P}_{\wedge}\});$$

where procedure $\mathcal{P}_{\rightsquigarrow}$, which will be defined soon, stands (with a slight misuse of notation) for the operations

$$\tilde{\boldsymbol{\mu}}_{t+1|t} := \mathbf{A}_{\mu} (\tilde{\boldsymbol{\mu}}_{t|t-1} + \tilde{\mathbf{B}}_t (\mathbf{z}_t - \mathbf{A}_z \tilde{\boldsymbol{\mu}}_{t|t-1})),$$

and \mathcal{P}_{\vee} , which will be derived later in Proposition 2, stands for

$$\hat{\boldsymbol{\mu}}_{t+1|t} := \tilde{\boldsymbol{\mu}}_{t+1|t} + \check{\boldsymbol{\mu}}_{t+1},$$

while function $\text{time}(\cdot)$ returns the number of primitive arithmetic operations to be applied to individual scalar elements to produce an output (in our case, $\{\tilde{\boldsymbol{\mu}}_{t+1|t}, \hat{\boldsymbol{\mu}}_{t+1|t}\}$) from the input ($\{\mathbf{z}_t, \tilde{\boldsymbol{\mu}}_{t|t-1}\}$) using the procedure(s) specified in the function's argument (as well as the parameters such as $\mathbf{A}_{\mu}, \mathbf{A}_z, \tilde{\mathbf{B}}_{t+1}, \check{\boldsymbol{\mu}}_{t+1}$ that are left implicit). Think of it not as a physical time, but as a processing time as, for example, provided by on-demand computing services.

To be specific, in this problem

$$\text{time}(\{\mathcal{P}_{\rightsquigarrow}, \mathcal{P}_{\vee}\}) = \left(((2k)^2)k + k + (k^2)2k + 2k + (2k)^3 \right) + (2k), \quad (1)$$

where non-negative integer k is the size of vectors that procedures $\{\mathcal{P}_{\rightsquigarrow}, \mathcal{P}_{\vee}\}$ operate on (see Cormen et al., 2009, for details on the complexity of matrix algebra operations). Here, the first term on the right-hand-side of the expression is the number of operations required to update $\tilde{\boldsymbol{\mu}}_{t+1|t}$ (following the equations in $\mathcal{P}_{\rightsquigarrow}$). While the second term is the number of operations required to update $\hat{\boldsymbol{\mu}}_{t+1|t}$ (using its constructive definition, \mathcal{P}_{\vee} , from Proposition 2). Altogether, the right-hand side of the expression specifies the total number of operations required to update k of the problem's dimensions (witness the ‘‘curse of dimensionality’’ due to the presence of cubic terms).

Above, $\tilde{\boldsymbol{\mu}}_{t|t}$ comes from the ancillary Kalman filter updating procedure $\mathcal{P}_{\rightsquigarrow}$ performed every period $t > 1$:

$$\begin{aligned} \tilde{\boldsymbol{\mu}}_{t|t-1} &= \mathbf{A}_{\mu} \tilde{\boldsymbol{\mu}}_{t-1|t-1}, & \{\mathcal{P}_{\rightsquigarrow}\} \\ \boldsymbol{\Upsilon}_{\mu,t|t-1} &= \mathbf{A}_{\mu} \boldsymbol{\Upsilon}_{\mu,t-1|t-1} \mathbf{A}_{\mu}^{\top} + \boldsymbol{\Sigma}_{\varepsilon\mu}, \\ \tilde{\boldsymbol{\mu}}_{t|t} &= \tilde{\boldsymbol{\mu}}_{t|t-1} + \tilde{\mathbf{B}}_t (\mathbf{z}_t - \mathbf{A}_z \tilde{\boldsymbol{\mu}}_{t|t-1}), \\ \boldsymbol{\Upsilon}_{\mu,t|t} &= (\mathbf{I}_K - \tilde{\mathbf{B}}_t \mathbf{A}_z) \boldsymbol{\Upsilon}_{\mu,t|t-1}, \end{aligned}$$

²In our specific application, the above concepts of computational complexity are essentially equivalent, both communication and time complexity can be mapped to each other, or to the concept of space complexity. This equivalence allows defining comparable measurement scales and, taking it further, a common information-processing capacity bound; as is done in Appendix §A. For ease of exposition, however, in the main text we disentangle these two notions.

where $\tilde{\mathbf{B}}_t$ is defined later. This procedure constitutes a (constrained) Kalman filter-based algorithm for parameter learning/updating.

The complexity problem is solved at the outset, in period $t = 1$, and its solution is used for optimizing the agent's consumption-investment-labor problem from then on, in periods $t > 1$.

As specified above, the parameters defining our stationary ergodic dynamic system $(\mathbf{A}_\mu, \mathbf{A}_z; \Sigma_{\varepsilon\mu}, \Sigma_{\varepsilon z})$ are fixed and known. However, in a realistic application, the parameters assumed so far as fixed may change. For example, a regime change (cyclical upswing/downturn, the start/end of hyperinflation, etc.) may severely perturb the first and second moments of the dynamic system considered. Such a disruption would prompt the agent to re-solve the informational and learning problem $\mathcal{P}_{\mathcal{I}}$ afresh. The costs of re-solution are high (due to the excess strain on available computational capacity and the need to fine-tune new "learning and informational policy" rules leading to temporarily larger decision errors, etc.³), so can not be incurred routinely; but they may nevertheless be outweighed by the benefits if the parameter change is large. The following Proposition offers a strategy for detecting a (potentially non-ergodic) perturbation to such a dynamic system.

Proposition 1 (Re-Solution of Complexity Problem: Threshold Strategy). *Assume that (i) the mean and variance parameters $(\mathbf{A}_\mu, \mathbf{A}_z; \Sigma_{\varepsilon\mu}, \Sigma_{\varepsilon z})$ may change within some parameter set (Θ) with a "small" but strictly positive probability; and that (ii) upon detection (with sufficient significance level defined below) of a parameter change, it is optimal to re-solve the complexity problem $\mathcal{P}_{\mathcal{I}}$ immediately. Then, the following sequential testing strategy gives an optimal trigger for re-solving the problem $\mathcal{P}_{\mathcal{I}}$ afresh (and for re-setting the value of current t to 0): a cumulative (for $s < t - 1$) absolute error between the expected and realized/experienced value functions*

$$\varsigma_{s:t-1} := \frac{1}{t-s} \sum_{\iota=s}^{t-1} \left| v(\mathbf{K}_\iota, \hat{\mathbf{z}}_{\iota+1}, \tilde{\boldsymbol{\mu}}_{\iota+1|\iota+1}) - \mathbb{E}_\iota^h \left[v(\mathbf{K}_\iota, \hat{\mathbf{z}}_{\iota+1}, \tilde{\boldsymbol{\mu}}_{\iota+1|\iota+1}) \right] \right|$$

that exceeds a chosen threshold, $\varsigma_{s:t-1} > \varsigma_a$. The threshold value is $\varsigma_a := \varsigma_{0:s-1}/a$, with $\varsigma_{0:s-1}$ for $0 < s \leq t - 1$ serving the role of the inference sample defined analogously to

³Our understanding is that there are some fixed costs of (re-) solving the complexity problem $\mathcal{P}_{\mathcal{I}}$. We do not explicitly quantify these costs and leave them abstract because we are only interested in the conceptual point here. One can think about them as, first, costs of information gathering (data collection, parameter estimation) in monetary and physical time terms. Second, computational costs $\text{time}(\mathcal{P}_{\mathcal{I}})$ of solving for new learning and informational policy rules. During the process, some default values of the control variables are used, for example, the previous period's values (i.e., 0 optimizing iterations are performed starting from the last period's values, see Verstyuk, 2017, Appendix §F), or a pre-defined minimal consumption level along with investment solely in a riskless capital/storage technology and employment solely in a firm with riskless productivity process. Third, economic, including opportunity, costs in terms of losses due to decision mistakes committed during the interval of fine-tuning the new policy rules (say, by reinforcement learning); as well as due to postponement of urgent decisions (e.g., investment into new technology or new employment/hiring) until the re-solution is over.

testing sample $\varsigma_{s:t-1}$, and the scalar α denoting a chosen Type-I error rate (significance level).

Proof. See Appendix §F. □

Assumption (ii) can be motivated by, say, the extreme aversion to ambiguity induced by unknown values of the mean and variance parameters.

Intuitively, the large deviations between observed and expected decision criterion (i.e., value function) suggest to the agents that their existing model/solution has broken down. In other words, the agents react not to a series of large “bad” (or “good”, as the treatment is symmetric) shocks and resulting losses (gains), but rather to their failure in understanding the causes of such deviations.⁴

Thus, a “regime change” (e.g., an increase in temporal/spatial dependence or a volatility rise) prompts agents to re-solve their model afresh. One particularly important implication is that the behavior of agents in our framework becomes immune to the Lucas critique. For instance, it ensures that a systematic policy to stimulate aggregate demand can not be effective in the long run.

To sum up, the architecture/hierarchy of decision-making in our framework comprises two modules that together solve problem $\mathcal{P}_{\mathcal{K}\mathcal{T}\mathcal{I}}$ for all periods:

- (i) in period $t = 0$, the auxiliary complexity sub-problem $\mathcal{P}_{\mathcal{T}\mathcal{I}}$ is (re-) solved subject to learning/updating constraint on $\mathcal{T}(\cdot)$ and information constraint on $\mathcal{I}(\cdot; \cdot)$ (as well as using the eigendecomposition procedure \mathcal{P}_{\boxtimes} if necessary, see Verstyuk, 2017);
- (ii) in each period $t > 0$, the consumption-investment-labor sub-problem of the main problem $\mathcal{P}_{\mathcal{K}\mathcal{T}\mathcal{I}}$ is solved subject to the budget constraint, utilizing solutions to $\mathcal{P}_{\mathcal{T}\mathcal{I}}$ from (i) above, as well as using the Kalman filter updating procedure \mathcal{P}_{∞} , the mean-adjustment procedure \mathcal{P}_{\vee} and the expectation operator $E_t^h[\cdot]$ evaluation procedure \mathcal{P}_f (once per each optimization iteration).

In this write-up, the optimization problems are signified by \mathcal{P} with letter subscripts, while the ancillary procedures are instead signified by \mathcal{P} with pictogram subscripts.⁵

⁴Such periods of paradigm shifts and the revision of established approaches in order to “face reality” are important in real life, and might be one of the ingredients behind the dynamics involving the phenomena usually referred to as “sentiment swings”.

⁵Note that in the interest of transparency and easier exposition, we only impose bounds on the learning of stochastic parameters and on the manipulation of stochastic objects in the optimization problem. However, we do not restrict other computations such as dealing with the budget constraint. One way to accept such an approach is by assuming that computations not accounted for explicitly here have been prioritized in the hierarchy of capacity demands, and now we are only allocating the remaining, marginal capacity.

2.3 Feasible problem (solution)

Assume a distance function of the L^2 norm (squared) form: $d(v^\#, v) := (v^\# - v)^2$. Since we do not know the analytical expression for the value function in this problem, for tractability we replace it with wealth (which is somewhat motivated by a special case from Verstyuk, 2017), effectively assuming one to be proportional to the other.⁶ This yields $(v^\# - v)^2 \approx (W - \hat{W})^2$. Then, the following Proposition can be stated.

Proposition 2 (Distortion Function). *The above distortion function, in the context of problem $\mathcal{P}_{\mathcal{KT}\mathcal{I}}$, given the distributional assumptions (as well as (i) replacement of the value function with the wealth variable, (ii) the dominance of the first-order effects of \mathbf{z}_{t+1} and $\hat{\mathbf{z}}_{t+1}$ on wealth, and (iii) imposing the motive of minimization of maximum loss for possible combinations of chosen control variables and realized state variables), can be reformulated as follows:*

$$\begin{aligned} d(v^\#(\mathbf{K}_t, \mathbf{z}_{t+1}, \boldsymbol{\mu}_{t+1|t+1}), v(\mathbf{K}_t, \hat{\mathbf{z}}_{t+1}, \tilde{\boldsymbol{\mu}}_{t+1|t+1})) &\approx (W_{t+1} - \hat{W}_{t+1})^2 \approx \\ &\approx \|\mathbf{z}_{t+1} - \hat{\mathbf{z}}_{t+1} + \mathbf{A}_z \check{\boldsymbol{\mu}}_{t+1}\|_2^2, \end{aligned}$$

where the conditional mean $\hat{\boldsymbol{\mu}}_{t+1|t}$ of the simplified random variable $\hat{\mathbf{z}}_{t+1}$ equals

$$\hat{\boldsymbol{\mu}}_{t+1|t} := \tilde{\boldsymbol{\mu}}_{t+1|t} + \check{\boldsymbol{\mu}}_{t+1}, \quad \{\mathcal{P}_V\}$$

with the bias term

$$\check{\boldsymbol{\mu}}_{t+1} := \begin{bmatrix} \mathbf{0}_{K \times K} \\ \frac{1}{2} \text{diag}^{-1}(\boldsymbol{\Sigma}_{z,t+1|t} - \hat{\boldsymbol{\Sigma}}_{z,t+1|t}) \end{bmatrix},$$

which uses $\boldsymbol{\Sigma}_{z,t+1|t}$ and $\hat{\boldsymbol{\Sigma}}_{z,t+1|t}$ to denote the conditional variance-covariance matrixes for \mathbf{z}_{t+1} and $\hat{\mathbf{z}}_{t+1}$, respectively.

Proof. See Appendix §G. □

The technical assumptions (ii) and (iii) mentioned in the statement of the Proposition are fairly mild, and are flashed out in more detail in the arguments of the proof. The Gaussian distribution for the simplified random variable $\hat{\mathbf{z}}_{t+1}$ implicitly postulated above is verified later. Appendix §B lists the definitions and decomposition formulas relevant to this part.

The main result of Proposition 2 is the sum-of-squares representation of the distance function, which will let us stay within the Gaussian environment facilitating subsequent derivations and arguments. Another result is the introduction of the bias term $\check{\boldsymbol{\mu}}_{t+1}$ that compensates for the difference between variance-covariance matrixes $\boldsymbol{\Sigma}_{z,t+1|t}$ and $\hat{\boldsymbol{\Sigma}}_{z,t+1|t}$, as well as the definition of the respective ancillary mean-adjustment procedure \mathcal{P}_V .⁷

⁶An alternative route would be to deal with the second-order Taylor expansion of the two value functions above.

⁷For instance, without the stochastic volatility optimal mean adjustment $\check{\boldsymbol{\mu}}_{t+1}$ is constant; if, additionally, $\hat{\boldsymbol{\Sigma}}_{z,t+1|t} = \mathbf{0}$, then $\check{\boldsymbol{\mu}}_{t+1}$ ensures that $\hat{\mathbf{Z}}_{t+1}$ under $h(\hat{\mathbf{z}}_{t+1}|\hat{\boldsymbol{\mu}}_{t+1|t})$ equals (almost surely) to the expected value of \mathbf{Z}_{t+1} under $g(\mathbf{z}_{t+1}|\check{\boldsymbol{\mu}}_{t+1|t})$.

Before solving the complexity problem $\mathcal{P}_{\mathcal{I}\mathcal{L}}$, we use Proposition 2 to split it into two.

Proposition 3 (Complexity Problem Split into Two Sub-problems). *For the distortion function above, the complexity problem $\mathcal{P}_{\mathcal{I}\mathcal{L}}$ splits into learning/updating and informational/evaluation-optimization sub-problems,*

$$\min_{\tilde{\mathbf{B}}_t} \mathbb{E}_t^f [\|\boldsymbol{\mu}_t - \tilde{\boldsymbol{\mu}}_{t|t}\|_2^2] = \text{tr}(\boldsymbol{\Upsilon}_{\boldsymbol{\mu},t|t}) \quad \{\mathcal{P}_{\mathcal{T}}\}$$

subject to

$$\mathcal{T}(\tilde{\boldsymbol{\mu}}_{t+1|t}, \hat{\boldsymbol{\mu}}_{t+1|t} \| \mathbf{z}_t, \tilde{\boldsymbol{\mu}}_{t|t-1}) \leq \kappa_{\mathcal{T}};$$

and

$$\min_{f(\cdot, \cdot)} \mathbb{E}_t^f [\|\mathbf{z}_{t+1} - \hat{\mathbf{z}}_{t+1} + \mathbf{A}_z \check{\boldsymbol{\mu}}_{t+1}\|_2^2] = \text{tr}(\boldsymbol{\Psi}_{\mathbf{z},t+1|t}) \quad \{\mathcal{P}_{\mathcal{I}}\}$$

subject to

$$\mathcal{I}(g(\mathbf{z}_{t+1} | \tilde{\boldsymbol{\mu}}_{t+1|t}); h(\hat{\mathbf{z}}_{t+1} | \hat{\boldsymbol{\mu}}_{t+1|t})) \leq \kappa_{\mathcal{I}}.$$

Proof. See Appendix §H. □

Solution to the informational sub-problem $\mathcal{P}_{\mathcal{I}}$ is analogous to Verstyuk (2017, see Propositions 4 and 5 there), with an exception that mean $\boldsymbol{\mu}_{t+1}$ is also stochastic now. The learning sub-problem $\mathcal{P}_{\mathcal{T}}$ is attended to next.

Proposition 4 (Solution to Learning Problem). *In the context of problem $\mathcal{P}_{\mathcal{K}\mathcal{T}\mathcal{I}}$, the solution to the learning sub-problem $\mathcal{P}_{\mathcal{T}}$ is given by the constrained Kalman gain matrix $\tilde{\mathbf{B}}_t$ defined as*

$$\tilde{\mathbf{B}}_t := \begin{cases} \mathbf{B}_t & \text{if } k^* = K, \\ \mathbf{B}_t^* & \text{if } 0 < k^* < K, \\ \mathbf{0}_{2K \times K} & \text{if } k^* = 0, \end{cases}$$

where \mathbf{B}_t is the unconstrained Kalman gain defined earlier; and in the leading case solution $\tilde{\mathbf{B}}_t$ simplifies to

$$\mathbf{B}_t^* := \begin{bmatrix} \mathbf{B}_{[1:k^* \times 1:K],t} \\ \mathbf{0}_{(K-k^*) \times K} \\ \mathbf{B}_{[(K+1):(K+k^*) \times 1:K],t} \\ \mathbf{0}_{(K-k^*) \times K} \end{bmatrix}.$$

Pre-sorting the (diagonal) elements that form the variance-covariance matrix of update errors $\boldsymbol{\Upsilon}_{\boldsymbol{\mu},t|t}$ in the order of variability, the cut-off parameter k^* is defined by:

$$\begin{aligned} \{\Upsilon_{kk}\}_1^K &:= \text{sortdescending}(\{\Upsilon_{kk}\}_1^K), \\ k^* &:= \arg \min_{k \in \{1, \dots, K\}} \{\Upsilon_{kk} \mid \mathcal{T}(\tilde{\boldsymbol{\mu}}_{t+1|t}, \hat{\boldsymbol{\mu}}_{t+1|t} \| \mathbf{z}_t, \tilde{\boldsymbol{\mu}}_{t|t-1}) \leq \kappa_{\mathcal{T}}\} = \\ &= \arg \min_{k \in \{1, \dots, K\}} \{\Upsilon_{kk} \mid ((2k)^2)k + k + (k^2)2k + 2k + (2k)^3\} + (2k) \leq \kappa_{\mathcal{T}}\}. \end{aligned}$$

Proof. See Appendix §I. □

In the Proposition's formulation, notation $\mathbf{M}_{[r1:r2 \times c1:c2]}$ denotes a partition of matrix \mathbf{M} starting at position $(r1, c1)$ and ending at position $(r2, c2)$. Basically, in the end certain rows in the Kalman gain matrix are wiped out.⁸

In the definition of k^* above, we measure the complexity of updating parameters as specified in part §2.2. The left-hand side of the conditioning inequality specifies the number of the problem's dimensions that can be updated given capacity bound $\kappa_{\mathcal{T}}$.

The basic idea of the mechanism is that due to capacity bounds, which are restricting the scope of updating of the stochastic environment's parameters (i.e., $\tilde{\boldsymbol{\mu}}_{t|t}$ and $\hat{\boldsymbol{\mu}}_{t|t}$), such an updating is performed only for a subset of the most important variables (in our illustrative application, the ones with the largest variability).

For clarity, the practical implications of Proposition 4 are distilled into two corollaries.

Corollary 1 (Solution to Learning Problem: Stickiness/Inertia). *The updating of the estimates of the random variables' conditional first moments by the (constrained) Kalman filter from Proposition 4 is determined by the sensitivity of $\tilde{\boldsymbol{\mu}}_{t|t}$ to period- t shocks to $\boldsymbol{\mu}_t$, which is captured by a bounded expression (for a non-explosive matrix \mathbf{A}_z)*

$$\frac{\partial \tilde{\boldsymbol{\mu}}_{t|t}}{\partial \boldsymbol{\varepsilon}_{\mu,t}} = \tilde{\mathbf{B}}_t \mathbf{A}_z \in [0, 1).$$

Thus, updating is either slow and staggered ($0 < \partial \tilde{\boldsymbol{\mu}}_{t|t,k} / \partial \boldsymbol{\varepsilon}_{\mu,t,k} < 1, \forall k \leq k^$) or completely absent and rigid ($\partial \tilde{\boldsymbol{\mu}}_{t|t,k} / \partial \boldsymbol{\varepsilon}_{\mu,t,k} = 0, \forall k > k^*$).*

Proof. The equation follows from \mathcal{P}_{∞} , with Proposition 4 giving the bounds on the row sums of $\tilde{\mathbf{B}}_t$. □

Corollary 2 (Solution to Learning Problem: Overshooting/Momentum). *The misadjustment of the estimates of the random variables' conditional first moments by the (constrained) Kalman filter from Proposition 4 is determined by the sensitivity of $\tilde{\boldsymbol{\mu}}_{t|t}$ to period- t shocks to \mathbf{z}_t , which is captured by a bounded expression*

$$\frac{\partial \tilde{\boldsymbol{\mu}}_{t|t}}{\partial \boldsymbol{\varepsilon}_{z,t}} = \tilde{\mathbf{B}}_t \in [0, 1).$$

Thus, the receptiveness to observation noise, and the resulting misadjustment, arises in some cases ($0 < \partial \tilde{\boldsymbol{\mu}}_{t|t,k} / \partial \boldsymbol{\varepsilon}_{\mu,t,k} < 1, \forall k \leq k^$).*

⁸The usage of filtering/signal extraction methods to deal with learning tasks is common in the literature (e.g., see Veronesi, 1999, 2000, for the case of discretely distributed signals; or Pástor and Veronesi, 2003, for continuously distributed random variables case). But, to the best of our knowledge, Teixeira et al. (2008) is the only relevant reference proposing a somewhat similar to ours modification of the standard Kalman filter.

Proof. Same as the proof of Corollary 1. □

Therefore, the upshot of implementing the learning process—that is, in the case when the capacity constraint actually allows learning—by means of the (constrained) Kalman filter is that even though the Kalman gain reacts to shocks, it does so less than one-for-one. On the one hand, this leads to only a partial reaction to changes in the underlying parameter of interest, so full “digestion” of news takes time (Corollary 1). On the other hand, this engenders a non-zero reaction to noise, which results in erroneous changes in parameter estimates that are unwarranted by the underlying fundamentals (Corollary 2).

3 Discussion

As a result of the solution to sub-problems $\mathcal{P}_{\mathcal{I}}$ and $\mathcal{P}_{\mathcal{T}}$, each element k of \mathbf{z}_{t+1} is transformed (“simplified”) in one of the three possible ways (listed below in the order of increasing demands for capacity, starting with $\kappa_{\mathcal{T}}$ and then following with $\kappa_{\mathcal{I}}$):⁹

- (i) mean $\hat{\mu}_{k,t+1|t}$ is not updated at every $t > 1$ and is fixed at its long-run level $\hat{\mu}_{k,t+1|t} := \hat{\mu}_k$, while variable $\hat{z}_{k,t+1}$ is non-stochastic and is fixed at its mean $\hat{\mu}_{k,t+1|t} = \hat{\mu}_k$;
- (ii) mean $\hat{\mu}_{k,t+1|t}$ is updated at every $t > 1$, while variable $\hat{z}_{k,t+1}$ is non-stochastic and is fixed at its mean $\hat{\mu}_{k,t+1|t}$;
- (iii) mean $\hat{\mu}_{k,t+1|t}$ is updated at every $t > 1$, while variable $\hat{z}_{k,t+1}$ is stochastic and distributed as $\mathcal{N}(\hat{\mu}_{k,t+1|t}, \hat{\sigma}_{k,t+1|t}^2)$.

It is the first two cases that interest us in this paper.¹⁰ Here, the conceptually crucial role of simplification is manifested in dropping some random variables from the agent’s approximating model (following the random variable’s “dispersion folding” Corollary 1 in Verstyuk, 2017). Due to the effects of entropy/variance reduction (resulting in overconfidence and, in turn, categorization), such random variables are replaced with non-stochastic objects, that is by their—sufficiently biased—means.¹¹

For example, in a bivariate case the capacity-constrained agent treats the “folded” random variable as non-stochastic and ends up working with a simple univariate approximate model. Say, when a bullet (“projectile”) is shot at a certain angle with a certain speed, classical Newtonian mechanics stipulates how to model and predict its motion towards the target based on the initial velocity/mass as well as the wind speed/direction. However, in

⁹Cf. spike-and-slab variable selection (Mitchell and Beauchamp, 1988; George and McCulloch, 1993; Madigan and Raftery, 1994), as well as least absolute shrinkage and selection operator, or Lasso (Tibshirani, 1996).

¹⁰One interesting variant of the third case is examined in Verstyuk (2017).

¹¹Note that this is a mechanism similar to the sparsity logic in Gabaix (2016b), albeit derived by different formal means.

the “first approximation”, we can ignore the wind drift; thus reducing two fundamental factors/forces (plus noise that represents the factors we have abstracted away from, e.g. force of gravity, Coriolis force, Magnus effect, etc.) to one factor (plus noise, but now with the other factor subsumed into it too). As a result, we obtain a simplified model, where the less important “causal branches” of the driving forces are (subjectively) “pruned”.

Another role of simplification is manifested in updating the estimates of the random variables’ first moments by the constrained Kalman filter, which reacts to shocks less than one-for-one, if at all. This is a mechanism for implementing the parameter learning in a way that is conceptually consistent with a broader framework, and which allows to deal with rich dynamics enabled by a non-i.i.d. environment.¹²

The resulting learning dynamics produces “underreaction” to a random variable’s mean changes, as the Kalman gain reacts to them only partially (i.e., “stickiness”, or “inertia”). It also produces “overreaction” to a random variable’s measurement/observation shocks (“overshooting”, or “momentum”). In this paper, we are interested in the “stickiness” phenomenon, particularly with respect to productivity process z_t ; and we ignore the issue of “overshooting” here.¹³

Such a nuanced assortment of effects implied by the learning processes has important macroeconomic implications. First, “folding” and ignoring the dynamics of positively autocorrelated random variables suggests an answer to the so-called Deaton’s paradox. While permanent income hypothesis (Friedman, 1957) provides an argument why agents’ consumption is empirically found to be less volatile than their current income, Campbell and Deaton (1989) uncovered the fact that consumption instead “underreacts” to permanent income shocks (“excess smoothness of consumption puzzle”).¹⁴ From the perspective of our approach, though, this can be explained by the agent’s “folding” of the positively correlated, i.e. long-lasting and in some sense permanent, income-shocks variable. (Which

¹²Note that the dynamics of slow, or even completely absent, updating of the mean could have been modeled differently. Here it is implemented as a parameter learning problem performed via the (constrained) Kalman filter.

Alternatively, it could have been obtained as an information acquisition/observation problem by applying the canonical “rational inattention” theory—which penalizes the costs of observing realized random variables—to signals about the realization of the mean. (Although canonically in the literature it is applied to realizations of random variables themselves.) Arguably, in the specific application considered here our formulation seems more plausible (and tractable beyond the linear Gaussian case).

¹³“Stickiness” produces a (subjective) “hysteresis” effect: agents perceive random variables’ moments as more persistent than they really are (or their distributions as “more identical” than they really are); for example, here this applies to productivity process z_t . “Overshooting” may produce another kind of (objective) “hysteresis” effect: an initial, exogenous shock may lead to further, endogenous shocks of the same sign; e.g., in an asset pricing application whenever asset demand depends positively on the above random shocks, thus generating a positive feedback loop from (subjective) fundamentals to (objective) prices.

¹⁴To be precise, they find, in some sense, “excess sensitivity of consumption” to income innovations; but accounting for the fact that income innovations are persistent, that sensitivity is not strong enough, ultimately resulting in “excess smoothness”.

is consistent with Campbell and Deaton’s, 1989, original call for rationalizing these findings with some mechanism that generates slow adjustment of consumption to innovations in income, such as liquidity constraints (see e.g. Carroll, 2001) or inertia/habit formation (Campbell and Cochrane, 1999; Carroll et al., 2000; Fuhrer, 2000). Note, however, that our mechanism for “underreaction” differs from that of canonical habit-formation, even though their theoretical effects (and, presumably, empirical manifestations) look quite similar.)

Second, “folding” and failing to update the dynamic parameters of negatively autocorrelated random variables suggests a reason for the effectiveness of fiscal policy. Ricardian equivalence (see Barro, 1974, for its first modern formulation) states that a rational agent anticipates the fact that expansionary fiscal policy necessarily implies compensatory fiscal contraction in the future, leaving the permanent income unchanged; hence such an agent saves rather than consumes the income sourcing from the fiscal stimulus. In our framework, though, a capacity-constrained agent may (choose to) ignore this temporal regularity and effectively “underreact” to a predictable future reversal; as a result he gets “fooled” by such oscillating transitory shocks, which in turn induces the potency of fiscal policy.

Lastly, it is worthwhile to emphasize that the effects of learning (“stickiness”/“inertia”) and variance-reduction (“folding”) each produce a kind of (subjective) “dampening” effect on the stochastic dynamics of the model’s exogenous variables, being the source of decision rules and decision outcomes that depart from “full rationality”. See Gabaix (2016a, 2016b), Cochrane (2016), Angeletos and Lian (2016, 2017), Farhi and Werning (2017) for potential explanatory power of such dampening.

4 Numerical exercise

The economic idea is that an agent turns a blind eye to, or “folds” and ignores the evolving dynamics of, one stochastic source of income; thus (rationally) misinterpreting its current source and future consequences, which ultimately results in effective/non-neutral fiscal policy (i.e., violation of Ricardian equivalence). Technically, the idea is that the agent treats as constant the mean of the above income source, and thus reacts only to a materialized shock (income/resources), but not to a change in conditioning information (new state). The latter would have allowed to anticipate the decline in the next period’s mean (due to negative autocorrelation), but otherwise the agent is influenced by the stimulus.

We restrict the number of production functions-technologies to $K = 3$. The first technology is a standard productive technology.

The second technology is key in this paper, it embodies a lump-sum government fiscal transfers process as a proxy for fiscal policy. Within the model introduced above, it is implemented by fixing $K_{2,t-1}$ and $L_{2,t}$ for all t at some constant values (we chose these to

be 1). Setting \check{Z}_1 to a positive value (we use 1) and the autoregressive coefficient ρ_2 in \mathbf{A}_μ to a negative value ensures the government budget balance over the long term, as positive transfers are eventually offset with negative ones. The constrained agent knows that the conditional distribution reflects the long-run budget balance restriction, but because of the information-processing costs chooses to ignore this regularity and fixes the mean of z_2 at its unconditional long-run level.

The third technology represents storage/savings stockpile (or investments in riskless bonds), where the agents can preserve goods without any risk, but where they can not borrow from. This is implemented by fixing $(Z_{3,t} - \check{Z}_3)$ and $L_{3,t}$ for all t at some constant values (here, 1); then, $K_{3,t}$ is a non-negative control variable. The third technology is introduced so as not to force investments in the first technology following a negative wealth shock.

An interesting situation that we focus on is based on the following qualitative structure (quantitative values of the parameters used are presented in Tables D.1 and D.2 of Appendix §D). First, the variance of the first technology is much larger than that of the second one.

Second, the mathematical operations capacity $\kappa_{\mathcal{T}}$ is binding to the extent that the constrained agent when making decisions does not update the mean of $z_{2,t}$ and fixes it at the constant long-run level (so, $k^* = 1$). The stochastic evaluation capacity $\kappa_{\mathcal{I}}$ is not binding though, hence the corresponding variance remains unchanged (so, $\hat{\Sigma}_{z,2} := \Sigma_{z,2}$, and thus $\hat{\mu}_{2,t} := \tilde{\mu}_{2,t}$).¹⁵

The model’s optimality and steady state conditions are available in Appendix §C. Model solutions that are presented below were obtained numerically, by a second-order perturbation (e.g., see Fernandez-Villaverde et al., 2016).¹⁶ (These numerical results are still preliminary.)

Now, we consider the effect of a “large” positive fiscal transfer shock (“stimulus”) at period $t = 2$ followed by a “large” negative fiscal shock (“payback”, i.e. tax) at some distant period $t = 7$. (Technically, these shocks arise as impulses to $\varepsilon_{\mu 1,2}$.) The above “large” shocks overlay a sequence of “small” mean-zero shocks, resulting in the empirical autoregressive coefficient of -0.5 over the full simulated time series of $z_{2,t}$ (i.e., in line with the theoretical value $\rho_2 := -0.5$ in Table D.2). (All other shocks are zero throughout, which completes the stochastic structure of z_t .) Figure 1 contrasts the reactions of

¹⁵Note that in the interest of higher transparency, we do not take the logic of capacity limits further to a case when both capacity constraints are binding, including the stochastic evaluation capacity $\kappa_{\mathcal{I}}$ to the extent that the agent “folds” $z_{2,t}$, simplifying it to a non-stochastic object and fixing it at the (adjusted) long-run mean (which would have implied $\hat{\Sigma}_{z,2} = 0$, i.e. $\hat{\Sigma}_{\varepsilon\mu 0,2} = \hat{\Sigma}_{\varepsilon\mu 1,2} = \hat{\Sigma}_{\varepsilon z,2} := 0$; and thus $\hat{\mu}_{2,t} := \tilde{\mu}_{2,t} - \tilde{\mu}_{2,t}$). For instance, this spares us of the need to find an optimal value of the mean of $\hat{z}_{2,t}$.

¹⁶An exact closed-form solution is available whenever the standard formulation admits one; but in the multivariate case with a non-i.i.d. log-Normal productivity and a capital depreciation rate less than 1, which is necessary for practically sensible results, even the standard problem does not admit a closed-form solution.

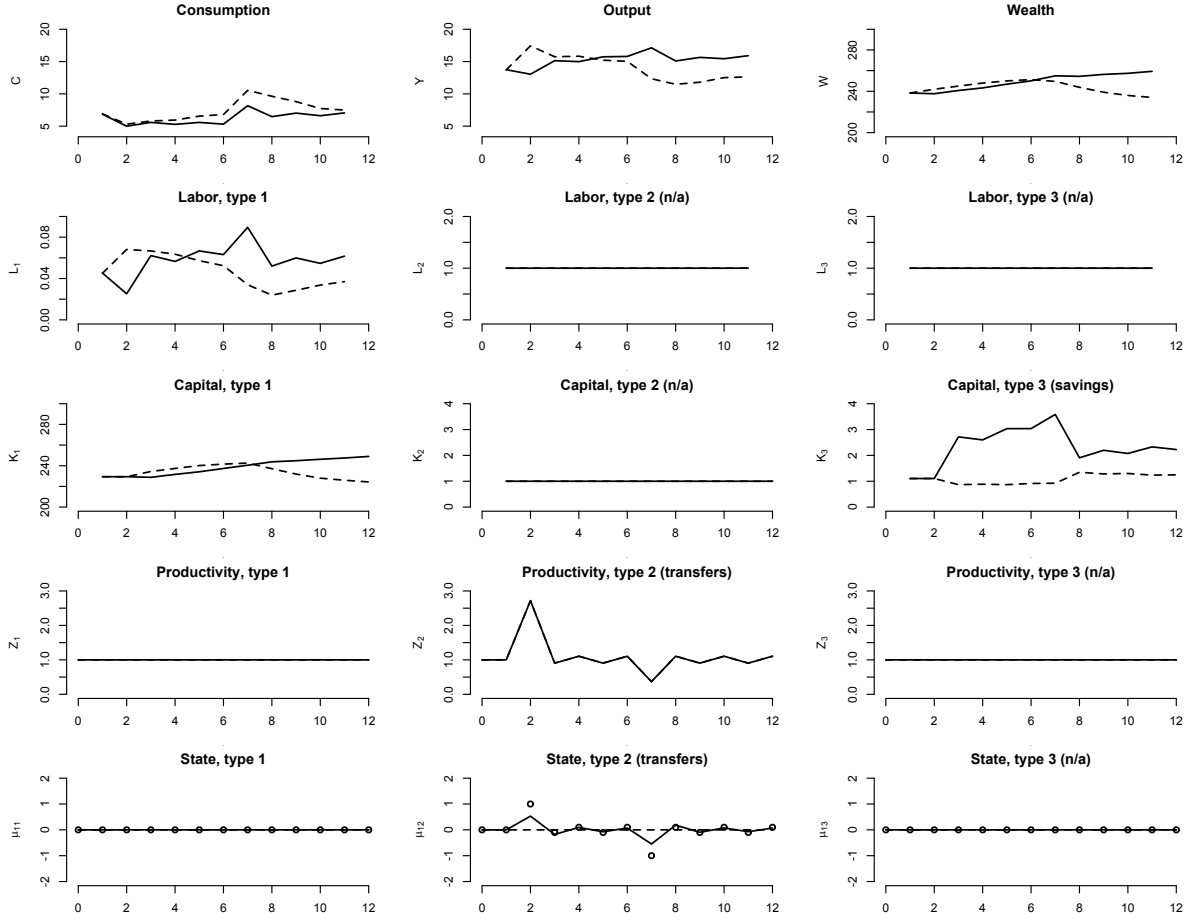


Figure 1: Simulation of a fiscal transfer scenario (an impulse response over a selected interval): decisions and outcomes of unconstrained (solid) vs. constrained (dashed) agents (in the top three rows), agents' corresponding observed variables (fourth row), unobserved state variables (circles) with agents' corresponding inferences (fifth row) (parameters used given in Tables D.1 and D.2; remaining endogenous variables C_t , $L_{1,t}$, $K_{1,t}$ and $K_{3,t}$ initialized at steady-state values; exogenous random variables μ_t and z_t are 0 except for the shocks described in the text).

the constrained (with $k^* = 1$) and the standard, unconstrained ($k^* = 2$) agents. (No-stimulus counterparts result in “boring” flat dynamics for both kind of agents, Figure E.1 in Appendix §E features such a baseline scenario.)

Comparing constrained and unconstrained agents' learning and decision, as well as the eventual economic outcomes, we find the following differences:

- (i) the constrained agent does not update his estimate of $z_{2,t}$'s conditional mean (middle panel in the last row);
- (ii) consumption of the constrained agent is higher than that of the unconstrained one following a positive fiscal transfer;
- (iii) consumption of the constrained agent is relatively less smooth;

- (iv) savings in the storage technology $K_{3,t}$ are relatively lower for the constrained agent (right panel in the third row);
- (v) investments into productive technology $K_{1,t}$ are relatively higher in the short-run but lower in the medium-run for the constrained agent (left panel in the third row);
- (vi) labor paths emulate those of $K_{1,t}$ (left panel in the second row);
- (vii) output and wealth of the constrained agent are relatively higher in the short-run but lower in the medium-run.

Summarizing the key results, we emphasize the following ones. First, when learning the current state, the constrained agent (who has $k^* = 1$ and $\hat{\Sigma}_{z,2} = \Sigma_{z,2}$) does not ascribe the increase in $z_{2,t}$ to a change in $\mu_{1,2,t}$ due to the $\varepsilon_{\mu_{1,2,t}}$ shock. Instead, the agent takes it as a positive wealth shock, and allocates the extra resources to higher consumption and more investments into productive capital. Second, due to these learning inaccuracies that result in “second-best” decisions, the constrained agent fares worse in the medium/long run (which is not the case in the no-shocks baseline, where no new information has to be processed).^{17,18}

These results reflect, in some sense, the violation of Ricardian equivalence: the constrained agent ignores the fact that the positive fiscal transfer will have to be repaid, so the increase in output does not reflect an increase in permanent income and should not be directed to consumption.

Such systematic, “forced” mistakes play a role in the short run; but if they become large enough to affect the long-run outcomes, eventually they do get corrected (to the extent permitted by available capacity). The constrained agent ignores the dynamics in the mean of $z_{2,t}$ only if the dynamic system he based his solution of the informational and learning problem $\mathcal{P}_{\mathcal{TI}}$ on remains unperturbed. Otherwise, the agent may choose to re-solve the informational and learning problems afresh. For example, the situation when the government resorts to a stimulating fiscal policy to the extent that public debt sustainability issues arise translates within our framework to a regime change with a shift in the second moment $\Sigma_{\varepsilon\mu}$ or a jump in the first moment parameter $\mu_{0,2,t}$; the resulting accumulation of decision mistakes would exceed the threshold of Proposition 1 and would prompt the agent to re-solve the problem $\mathcal{P}_{\mathcal{TI}}$.

Before finishing, a couple of remarks are in order. First, in our “proof-of-concept” demonstration we are mostly concerned about a qualitative picture. However, quantitative magnitudes of the effects are potentially substantial. In particular, with the parameterization considered in our numerical exercise, “fiscal multipliers” for the consumption

¹⁷There is no government consumption in our setting, so a classical “crowding-out” effect does not apply. But the “second-best” economic outcomes discussed above can be thought of as somewhat similar kind of efficiency losses.

¹⁸Lastly, when interpreting Figure 1, note that even though the income effect dominates the substitution effect in our application (as $\gamma > 1$), without shocks to $Z_{1,t}$ it does not play a role.

of a constrained agent are well above 1 (and easily reach 5 on some horizons).¹⁹ Second, the fiscal transfer effects considered in our numerical simulation exercise are state/regime-dependent: under parameterization considered here, they turn out to be stronger at times of economic bust (operationally implemented via negative shocks to $Z_{1,t}$, following the standard definition of a recession as two periods of negative growth; not shown on plots above). That said, the results mentioned in these finishing remarks are merely suggestive, and more research is warranted.

5 Conclusion

This study shows how the limitations in information processing of economic agents (consumers-investors-workers) affect their decisions and resulting economic outcomes. Specifically, we focus on the problem of learning the evolving parameters/current state, which in the context of a relatively standard stochastic growth model with government transfers leads to a violation of Ricardian equivalence and engenders the potency of fiscal policy.

Some scope for further macroeconomic applications and generalizations is given below. Firstly, here the government borrowing is assumed to be done from outside of the economy (say, foreign markets), but the mechanism survives and should have a qualitatively similar effect if the borrowing is done from the agents themselves.

Secondly, a similar mechanism can generate monetary policy non-neutrality. Furthermore, note that such a model would not suffer from the so-called forward guidance puzzle of DSGE models (Del Negro et al., 2012), which imply a counterfactual hypersensitivity of economic outcomes in the near term to announcements about monetary policy in the very distant future.

¹⁹So-called “fiscal multipliers” are not some structural parameters, but rather empirical or theoretical sensitivities/“elasticities” that are conditioned on economic specifications, parameter values, the relevant equilibrium, realized state, etc. Traditionally, a fiscal multiplier is defined as $\sum_{s=t}^{t+T} (Y_s - Y) / \sum_{s=t}^{t+T} (G_s - G)$, where the activity measure Y_s is usually output or consumption, fiscal instrument G_s is usually government expenditures or taxes/transfers, and variables without time subscripts are some benchmarks such as steady-state values. When horizon T is “low”, we are speaking of impact/instantaneous/period multipliers, when T is “high” we call them summary/integrated/cumulative multipliers.

A Generalized information processing capacity

As mentioned in the main text, two complexity concepts used in this work—time complexity and communication complexity/entropy—are in some cases equivalent and can be mapped into each other. Hence, it is natural to consider a setup with common information-processing capacity bound utilized both for updating the conditioning parameters by \mathcal{P}_{\sim} and \mathcal{P}_V , which appear in the learning problem $\mathcal{P}_{\mathcal{T}}$, as well as for evaluation of expectations by \mathcal{P}_f , which appears in the informational problem $\mathcal{P}_{\mathcal{I}}$, thus establishing a meaningful trade-off and introducing the issue of efficient capacity allocation between two auxiliary sub-problems.

Usually, these complexity concepts are defined initially in discrete terms: this is conceptually simpler, and also has more solid mathematical foundations, at least in the case of information-theoretic formulations (see Arora and Barak, 2009; Cover and Thomas, 2006). However, their continuous counterparts are often introduced due to the latter’s higher analytical tractability.

In the case of communication complexity, dealing with higher quantization (finesness) of probability space requires more communication operations, and taking quantization to the limit produces continuous objects that are convenient analytically. In the case of time complexity, dealing with richer parameterization of the economic system requires more mathematical operations, and taking parameterization to the limit would similarly produce asymptotical objects that are convenient to work with.

Throughout this Appendix, we adopt the formulations in discrete terms for easier exposition and higher transparency. Thus, we stay with the low-parameterization pre-limit model as far as the learning problem is concerned; as to the informational problem, we are also dealing with discrete random variables and with the corresponding definition of information-theoretic notions such as Shannon entropy (which, for instance, implies that the entropy measure stays on a $[0, \infty)$ range).

However, in the main text, where we disentangle two complexity concepts, we resort to continuous formulations in the case of communication complexity/entropy and the informational problem in the interest of analytical convenience, but we work with discrete formulations in the case of time complexity and the learning problem due to their, arguably, higher economic plausibility and interpretability.

A.1 Formulation

The information processing capacity can be utilized for performing (i) operations with deterministic objects, dealing with—possibly stochastic, but already realized—parameters that are taken as given (in our case used for learning/updating the conditioning parameters); and (ii) operations with stochastic objects, dealing with—not yet realized—random variables that are viewed as exhibiting non-degenerate stochasticity (in our case used for evaluation of expectations). Now we present a unified approach to quantifying the

complexity of these two information-processing tasks.

Fundamentally, this equivalence/reduction becomes possible due to an intrinsically algorithmic nature of time complexity, connecting it to Solomonoff-Kolmogorov-Chaitin complexity notion (for definition, see Rissanen, 2007), which is in turn intimately related to Shannon entropy (for details, see Leung-Yan-Cheong and Cover, 1978).

Capacity accounting: Extending the approach of Verstyuk (2017, Appendix §F) and keeping the same notation as there, available full computation channel capacity \mathcal{K}_C (assuming it is responsible for the binding constraint) can be allocated between the demands of two tasks above:

$$\begin{aligned} \mathcal{K}_C \geq & \left\{ \mathcal{T}(\tilde{\boldsymbol{\mu}}_{t+1|t}, \hat{\boldsymbol{\mu}}_{t+1|t} \| \mathbf{z}_t, \tilde{\boldsymbol{\mu}}_{t|t-1}) \times \log |\mathcal{A}| \right\} + \\ & + \left\{ n_\iota (\hat{n}_d \times \mathcal{I}(\ddot{g}(\mathbf{z}_{t+1} | \tilde{\boldsymbol{\mu}}_{t+1|t}); \ddot{h}(\hat{\mathbf{z}}_{t+1} | \hat{\boldsymbol{\mu}}_{t+1|t})) + \text{overhead}) \right\}, \end{aligned} \quad (\text{A.1})$$

where the second term on the right-hand side is unchanged from Verstyuk (2017) (except that \mathbf{z}_{t+1} and $\hat{\mathbf{z}}_{t+1}$ now have non-constant means), while the first term is the measure of complexity of updating the conditional means.

Specifically, object $\mathcal{I}(\ddot{g}(\cdot); \ddot{h}(\cdot))$ is a mutual information between the input distribution $\ddot{g}(\cdot)$ and the output distribution $\ddot{h}(\cdot)$, and it is measured in bits. Object $\mathcal{T}(\tilde{\boldsymbol{\mu}}_{t+1|t}, \hat{\boldsymbol{\mu}}_{t+1|t} \| \cdot)$ is the number of primitive arithmetic operations required to update the specified parameters, and after scaling by the codeword length $\log |\mathcal{A}|$ (which is determined by the “native” alphabet \mathcal{A} used for all operations), the output of such operations is also measured in bits. Thus, available full computation channel capacity \mathcal{K}_C is measured in bits (per period).

Rearranging, we can formulate the (generalized) information processing capacity constraint

$$\begin{aligned} & A_{\mathcal{T}} \times \mathcal{T}(\tilde{\boldsymbol{\mu}}_{t+1|t}, \hat{\boldsymbol{\mu}}_{t+1|t} \| \mathbf{z}_t, \tilde{\boldsymbol{\mu}}_{t|t-1}) + \\ & + \mathcal{I}(\ddot{g}(\mathbf{z}_{t+1} | \tilde{\boldsymbol{\mu}}_{t+1|t}); \ddot{h}(\hat{\mathbf{z}}_{t+1} | \hat{\boldsymbol{\mu}}_{t+1|t})) \leq \frac{\mathcal{K}_C}{\hat{n}_d \times n_\iota} - \frac{\text{overhead}}{\hat{n}_d} =: \kappa, \end{aligned} \quad (\text{A.2})$$

where κ denotes effective information processing capacity bound, and where the leading constant equals

$$A_{\mathcal{T}} := \frac{\log |\mathcal{A}|}{\hat{n}_d \times n_\iota}.$$

A.2 Corresponding optimization problem

A more general complexity/information processing problem is

$$\min_{\tilde{\mathbf{B}}_t, \tilde{f}(\cdot, \cdot)} \mathbb{E}_t^f \left[d(v^\sharp(\mathbf{K}_t, \mathbf{z}_{t+1}, \boldsymbol{\mu}_{t+1|t+1}), v(\mathbf{K}_t, \hat{\mathbf{z}}_{t+1}, \tilde{\boldsymbol{\mu}}_{t+1|t+1})) \right] \quad \{\mathcal{P}_{\mathcal{T}\mathcal{I}+}\}$$

subject to the information processing capacity constraint

$$A_{\mathcal{T}} \times \mathcal{T}(\tilde{\boldsymbol{\mu}}_{t+1|t}, \hat{\boldsymbol{\mu}}_{t+1|t} \| \mathbf{z}_t, \tilde{\boldsymbol{\mu}}_{t|t-1}) + \mathcal{I}(\ddot{g}(\mathbf{z}_{t+1} | \tilde{\boldsymbol{\mu}}_{t+1|t}); \ddot{h}(\hat{\mathbf{z}}_{t+1} | \hat{\boldsymbol{\mu}}_{t+1|t})) \leq \kappa$$

(as well as the necessary technical restrictions).

Forming the usual Lagrangian functional, minimization with respect to $\tilde{\mathbf{B}}_t$ requires

$$\frac{\partial \mathcal{L}}{\partial k} := 0$$

(by accounting for the number of non-zero elements of $\tilde{\mathbf{B}}_t$, together with those of $\check{\boldsymbol{\mu}}_{t+1}$, as given by expression for time $(\{\mathcal{P}_{\sim}, \mathcal{P}_{\vee}\})$ in §2.2; and with the ordering of $\tilde{\mathbf{B}}_t$'s elements as, e.g., in Proposition 4). Minimization with respect to $\check{f}(\cdot, \cdot)$ requires

$$\frac{\delta \mathcal{L}}{\delta \check{f}(\mathbf{x}, \hat{\mathbf{x}})} := 0.$$

These two first-order conditions define the trade-off between two uses of the capacity resource; and together with the capacity constraint they allow to solve for the interior optimum in $\tilde{\mathbf{B}}_t$ and $\check{f}(\cdot, \cdot)$, as well as to obtain the quantities

$$\begin{aligned} \mathcal{T}(\tilde{\boldsymbol{\mu}}_{t+1|t}, \hat{\boldsymbol{\mu}}_{t+1|t} \| \mathbf{z}_t, \tilde{\boldsymbol{\mu}}_{t|t-1}) \Big|_{\tilde{\mathbf{B}}_t(k^*), \check{f}^*(\cdot, \cdot)} &=: \kappa_{\mathcal{T}}, \\ \mathcal{I}(\check{g}(\mathbf{z}_{t+1} | \tilde{\boldsymbol{\mu}}_{t+1|t}); \check{h}(\hat{\mathbf{z}}_{t+1} | \hat{\boldsymbol{\mu}}_{t+1|t})) \Big|_{\tilde{\mathbf{B}}_t(k^*), \check{f}^*(\cdot, \cdot)} &=: \kappa_{\mathcal{I}}, \end{aligned}$$

where $\kappa_{\mathcal{T}}$ denotes what can be understood as an effective mathematical operations capacity bound, and $\kappa_{\mathcal{I}}$ an effective stochastic evaluation capacity bound.

The solution gives $0 \leq \kappa_{\mathcal{T}} \leq \kappa_{\mathcal{T}}^{\#} < \infty$ and $0 \leq \kappa_{\mathcal{I}} \leq \kappa_{\mathcal{I}}^{\#} < \infty$ for some no-longer-binding (sharp) upper limits $\kappa_{\mathcal{T}}^{\#}$ and $\kappa_{\mathcal{I}}^{\#}$ (but with $-\infty < \kappa_{\mathcal{I}} \leq \kappa_{\mathcal{I}}^{\#} < \infty$ in the main text due to scaling particularities of the differential entropy that is used there). In the interior solution case both of the above first-order conditions hold with equality.

The capacity bounds $\kappa_{\mathcal{T}}$ and $\kappa_{\mathcal{I}}$ used in two disentangled complexity constraints of the main body can be understood as if they were obtained by solving the problem above. (Note, however, that in the main body's §2 we focus on a relatively more interesting situation in which both of the capacity bounds are non-trivial, i.e. located away from the lower limits of the admissible range.)

B Some formulas for stochastic relationships

The following relationships between original/unconstrained and simplified/constrained stochastic objects hold.

Definitions of random variables:

$$\begin{aligned} \mathbf{z}_{t+1} &:= \mathbf{A}_z \boldsymbol{\mu}_{t+1} + \boldsymbol{\varepsilon}_{z,t+1} = \mathbf{A}_z \boldsymbol{\mu}_{t+1|t} + \mathbf{A}_z \check{\boldsymbol{\zeta}}_{\boldsymbol{\mu},t+1|t} + \boldsymbol{\varepsilon}_{z,t+1} = \\ &= \mathbf{A}_z \tilde{\boldsymbol{\mu}}_{t+1|t} + \mathbf{A}_z \check{\boldsymbol{\zeta}}_{\boldsymbol{\mu},t+1|t} + \boldsymbol{\varepsilon}_{z,t+1}, \\ \hat{\mathbf{z}}_{t+1} &:= \mathbf{A}_z \hat{\boldsymbol{\mu}}_{t+1} + \hat{\boldsymbol{\varepsilon}}_{z,t+1} = \mathbf{A}_z \tilde{\boldsymbol{\mu}}_{t+1|t} + \mathbf{A}_z \hat{\boldsymbol{\zeta}}_{\boldsymbol{\mu},t+1|t} + \mathbf{A}_z \check{\boldsymbol{\mu}}_{t+1} + \hat{\boldsymbol{\varepsilon}}_{z,t+1}. \end{aligned}$$

Definitions of conditional simplified/constrained means:

$$\begin{aligned}\hat{\boldsymbol{\mu}}_{t+1} &:= \tilde{\boldsymbol{\mu}}_{t+1|t} + \hat{\boldsymbol{\zeta}}_{\mu,t+1|t} + \check{\boldsymbol{\mu}}_{t+1}, \\ \hat{\boldsymbol{\mu}}_{t+1|t} &:= \mathbb{E}_t^h [\hat{\boldsymbol{\mu}}_{t+1}] = \tilde{\boldsymbol{\mu}}_{t+1|t} + \check{\boldsymbol{\mu}}_{t+1}.\end{aligned}$$

Definitions (with some notation abuse) of conditional means/states' misestimates (note that $\boldsymbol{\mu}_{t+1|t}$ and $\tilde{\boldsymbol{\mu}}_{t+1|t}$ are known at period t):

$$\begin{aligned}\boldsymbol{\zeta}_{\mu,t+1|t} &:= \boldsymbol{\mu}_{t+1} - \boldsymbol{\mu}_{t+1|t}, \\ \tilde{\boldsymbol{\zeta}}_{\mu,t+1|t} &:= \boldsymbol{\mu}_{t+1} - \tilde{\boldsymbol{\mu}}_{t+1|t}, \\ \hat{\boldsymbol{\zeta}}_{\mu,t+1|t} &:= \widehat{\boldsymbol{\mu}_{t+1} - \tilde{\boldsymbol{\mu}}_{t+1|t}}.\end{aligned}$$

Definitions of conditional variances:

$$\begin{aligned}\boldsymbol{\Upsilon}_{\mu,t+1|t} &:= V_t^g [\boldsymbol{\zeta}_{\mu,t+1|t}], \\ \tilde{\boldsymbol{\Upsilon}}_{\mu,t+1|t} &:= V_t^g [\tilde{\boldsymbol{\zeta}}_{\mu,t+1|t}], \\ \hat{\boldsymbol{\Upsilon}}_{\mu,t+1|t} &:= V_t^h [\hat{\boldsymbol{\zeta}}_{\mu,t+1|t}].\end{aligned}$$

Definitions of approximation errors (note that states' misestimates $\boldsymbol{\zeta}_{\mu,t+1|t}$ and $\hat{\boldsymbol{\zeta}}_{\mu,t+1|t}$ form an additional shock term and eventually enter the ‘‘broad’’ approximation error):

$$\begin{aligned}\boldsymbol{\epsilon}_{\mu,t+1|t} &:= \boldsymbol{\zeta}_{\mu,t+1|t} - \hat{\boldsymbol{\zeta}}_{\mu,t+1|t}, \\ \boldsymbol{\epsilon}_{z,t+1} &:= \boldsymbol{\epsilon}_{z,t+1} - \hat{\boldsymbol{\epsilon}}_{z,t+1}.\end{aligned}$$

Variance decompositions for random variables:

$$\begin{aligned}\boldsymbol{\Sigma}_{z,t+1|t} &= \hat{\boldsymbol{\Sigma}}_{z,t+1|t} + \boldsymbol{\Psi}_{z,t+1|t}; \\ \boldsymbol{\Sigma}_{z,t+1|t} &= \mathbf{A}_z \boldsymbol{\Upsilon}_{\mu,t+1|t} \mathbf{A}_z^\top + \boldsymbol{\Sigma}_{\epsilon z}, \\ \hat{\boldsymbol{\Sigma}}_{z,t+1|t} &= \mathbf{A}_z \hat{\boldsymbol{\Upsilon}}_{\mu,t+1|t} \mathbf{A}_z^\top + \hat{\boldsymbol{\Sigma}}_{\epsilon z}.\end{aligned}$$

Variance decomposition for conditional means/states:

$$\boldsymbol{\Upsilon}_{\mu,t+1|t} = \hat{\boldsymbol{\Upsilon}}_{\mu,t+1|t} + \boldsymbol{\Psi}_{\mu,t+1|t}.$$

Variance decomposition for shocks:

$$\boldsymbol{\Sigma}_{\epsilon z} = \hat{\boldsymbol{\Sigma}}_{\epsilon z} + \boldsymbol{\Psi}_{\epsilon z}.$$

Variance decomposition for approximation errors:

$$\boldsymbol{\Psi}_{z,t+1|t} = \mathbf{A}_z \boldsymbol{\Psi}_{\mu,t+1|t} \mathbf{A}_z^\top + \boldsymbol{\Psi}_{\epsilon z}.$$

C Optimality and steady state conditions

The Lagrangian functional for problem $\mathcal{P}_{\mathcal{K}\mathcal{T}\mathcal{I}}$ produces the following first-order necessary conditions:

$$\begin{aligned} 0 &=: \frac{\partial \mathcal{L}}{\partial K_{1,t}} = -u_C(C_t, \mathbf{L}_t) + \beta \mathbf{E}_t^h \left[v_{K1}(\mathbf{K}_t, \hat{\mathbf{z}}_{t+1}, \tilde{\boldsymbol{\mu}}_{t+1|t+1}) \right], \\ 0 &=: \frac{\partial \mathcal{L}}{\partial K_{3,t}} = -u_C(C_t, \mathbf{L}_t) + \beta \mathbf{E}_t^h \left[v_{K3}(\mathbf{K}_t, \hat{\mathbf{z}}_{t+1}, \tilde{\boldsymbol{\mu}}_{t+1|t+1}) \right], \\ 0 &=: \frac{\partial \mathcal{L}}{\partial L_{1,t}} = u_L(C_t, \mathbf{L}_t) + u_C(C_t, \mathbf{L}_t) \hat{w}_{1,t}; \end{aligned}$$

as well as the corresponding envelope conditions:

$$\begin{aligned} v_{K1}(\mathbf{K}_t, \hat{\mathbf{z}}_{t+1}, \tilde{\boldsymbol{\mu}}_{t+1|t+1}) &:= u_C(C_t, \mathbf{L}_t)(\hat{r}_{1,t} + (1 - \delta_1)), \\ v_{K3}(\mathbf{K}_t, \hat{\mathbf{z}}_{t+1}, \tilde{\boldsymbol{\mu}}_{t+1|t+1}) &:= u_C(C_t, \mathbf{L}_t)(\hat{r}_{3,t} + (1 - \delta_3)). \end{aligned}$$

The optimality conditions then include 2 Euler equations and 1 labor supply equation

$$\begin{aligned} \varpi C_t^{-\gamma} &= \beta \mathbf{E}_t^h \left[\varpi C_{t+1}^{-\gamma} (\hat{r}_{1,t+1} + (1 - \delta_1)) \right], \\ \varpi C_t^{-\gamma} &= \beta \mathbf{E}_t^h \left[\varpi C_{t+1}^{-\gamma} (\hat{r}_{3,t+1} + (1 - \delta_3)) \right], \\ L_{1,t}^{\eta_1} &= \varpi C_t^{-\gamma} \hat{w}_{1,t}; \end{aligned}$$

plus returns on capital and wage rate from the firm's problem

$$\begin{aligned} \hat{r}_{1,t+1} K_{1,t} &= \alpha_1 Y_{1,t+1}, \\ \hat{r}_{3,t+1} K_{3,t} &= \alpha_3 Y_{3,t+1}, \\ \hat{w}_{1,t} L_{1,t} &= (1 - \alpha_1) Y_{1,t}; \end{aligned}$$

as well as the budget constraint (with $k \in \{1, 2, 3\}$)

$$C_t + \mathbf{1}^\top \mathbf{K}_t = (\hat{\mathbf{r}}_t + (\mathbf{1} - \boldsymbol{\delta}))^\top \mathbf{K}_{t-1} + \hat{\mathbf{w}}_t^\top \mathbf{L}_t = \hat{W}_t;$$

where also $K_{2,t} = L_{2,t} = 1$ and $L_{3,t} = 1$ for all periods t .

Or, in a shorter form, for all t ,

$$\begin{aligned} C_t^{-\gamma} &= \beta \mathbf{E}_t^h \left[C_{t+1}^{-\gamma} \left(\alpha_1 \frac{Y_{1,t+1}}{K_{1,t}} + (1 - \delta_1) \right) \right], \\ C_t^{-\gamma} &= \beta \mathbf{E}_t^h \left[C_{t+1}^{-\gamma} \left(\alpha_3 \frac{Y_{3,t+1}}{K_{3,t}} + (1 - \delta_3) \right) \right], \\ L_{1,t}^{1+\eta_1} &= \varpi C_t^{-\gamma} (1 - \alpha_1) Y_{1,t}, \\ C_t + \sum_{k=1}^3 K_{k,t} &= \sum_{k=1}^3 Y_{k,t} + \sum_{k=1}^3 (1 - \delta_k) K_{k,t-1}, \\ 1 &= K_{2,t} = L_{2,t} = L_{3,t}. \end{aligned}$$

Finally, the above results also yield the steady state conditions:

$$\begin{aligned}1 &= \beta \left(\alpha_1 \frac{Y_1}{K_1} + (1 - \delta_1) \right), \\1 &= \beta \left(\alpha_3 \frac{Y_3}{K_3} + (1 - \delta_3) \right), \\L_1^{1+\eta_1} &= \varpi C^{-\gamma} (1 - \alpha_1) Y_1, \\C &= \sum_{k=1}^3 Y_k - \sum_{k=1}^3 \delta_k K_k.\end{aligned}$$

D Parameter values used

Table D.1: Parameter Values Used in Simulations: Deep Parameters.

Parameter	Description	Value
β	subjective discount factor/time preference rate	0.99
ϖ	(inverse) marginal disutility of labor	0.1
γ	coefficient of relative risk-aversion	2
$\boldsymbol{\eta}$	(inverse) Frisch elasticity of labor supply	[0.5; 0.5; 0.5]
$\boldsymbol{\alpha}$	output elasticity of capital/capital's share of output	[0.67; 1.00; 1.00]
δ	depreciation rate	[0.03; 0.00; 0.00]
$\check{\mathbf{Z}}$		[0; 1; 1]

Table D.2: Parameter Values Used in Simulations: Additional Parameters.

Parameter	Description	Value
$\boldsymbol{\rho}$	autoregressive coefficient for state $\boldsymbol{\mu}_t$	[0.9; -0.5; 0.0]
$\boldsymbol{\Sigma}_{\varepsilon\mu 0}$	upper var-cov submatrix of $\boldsymbol{\varepsilon}_{\mu,t}$ shock to $\boldsymbol{\mu}_t$	diag([0; 0; 0])
$\boldsymbol{\Sigma}_{\varepsilon\mu 1}$	lower var-cov submatrix of $\boldsymbol{\varepsilon}_{\mu,t}$ shock to $\boldsymbol{\mu}_t$	diag([0.02 ² ; 0.01 ² ; 0.00])
$\boldsymbol{\Sigma}_{\varepsilon z}$	var-cov matrix of $\boldsymbol{\varepsilon}_{z,t}$ shock to productivity \mathbf{z}_t	diag([0.02 ² ; 0.01 ² ; 0.00])
$\kappa_{\mathcal{T}}$	effective mathematical operations capacity bound	[19; 122]
$\kappa_{\mathcal{I}}$	effective stochastic evaluation capacity bound	-4.8

In the last Table, $\kappa_{\mathcal{T}}$ is calculated using formula (1) for $k = 1$: $(4+1+2+2+8)+(2)=19$; and for $k = 2$: $(32+2+16+4+64)+(4)=122$.

While $\kappa_{\mathcal{I}}$, which we have assumed to equal $\kappa_{\mathcal{I}}^{\#}$, is calculated using formula

$$\mathcal{I}(g(\mathbf{z}_{t+1}|\tilde{\boldsymbol{\mu}}_{t+1|t}); h(\hat{\mathbf{z}}_{t+1}|\hat{\boldsymbol{\mu}}_{t+1|t})) = \mathcal{E}(g(\mathbf{z}_{t+1}|\tilde{\boldsymbol{\mu}}_{t+1|t})) = \frac{1}{2} \ln \left| 2\pi e \left(\mathbf{A}_z \boldsymbol{\Upsilon}_{\mu,t+1|t} \mathbf{A}_z^{\top} + \boldsymbol{\Sigma}_{\varepsilon z} \right) \right| = -4.8 \text{ nats.}$$

E Baseline scenario simulations (for comparison)

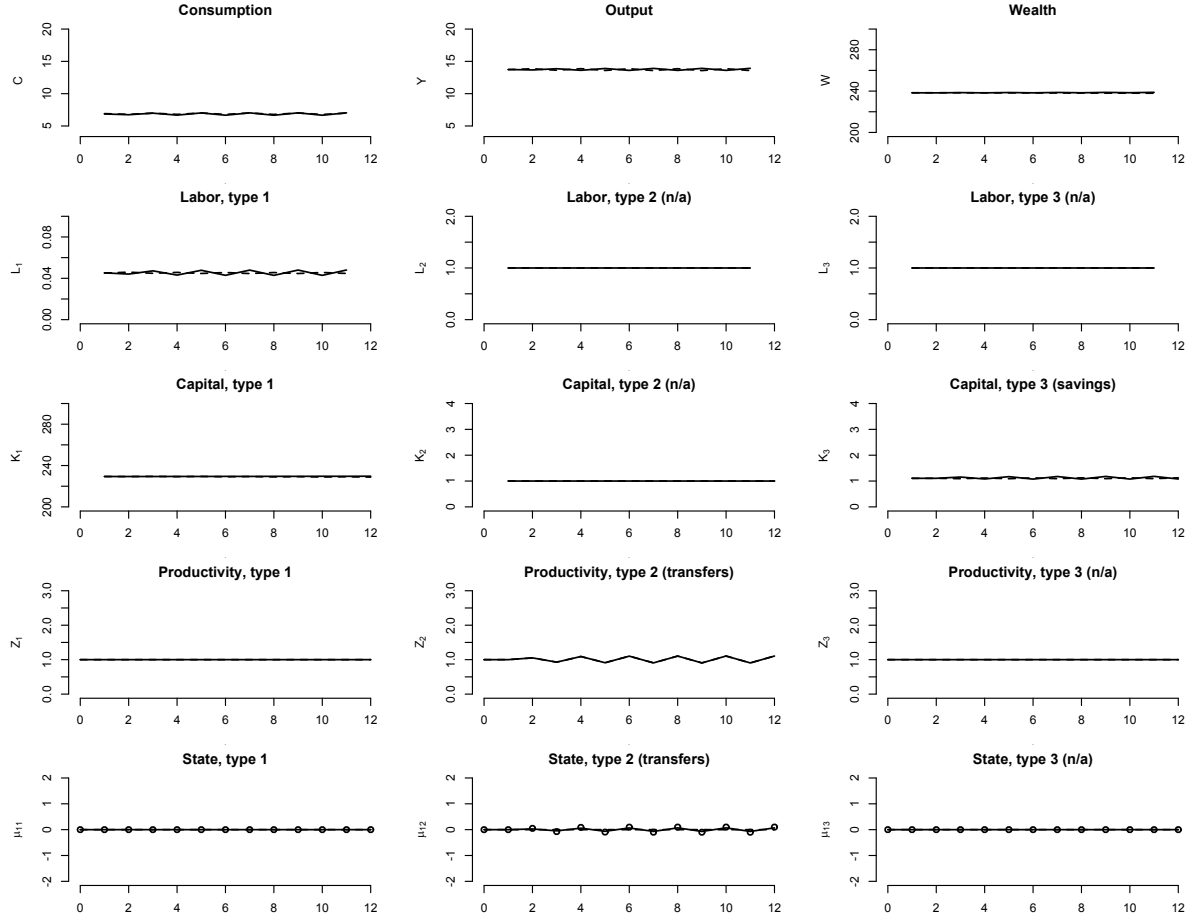


Figure E.1: Simulation of a baseline scenario (over a selected interval):

decisions and outcomes of unconstrained (solid) vs. constrained (dashed) agents (in the top three rows), agents' corresponding observed variables (fourth row), unobserved state variables (circles) with agents' corresponding inferences (fifth row) (parameters used given in Tables D.1 and D.2; remaining endogenous variables C_t , $L_{1,t}$, $K_{1,t}$ and $K_{3,t}$ initialized at steady-state values; exogenous random variables μ_t and z_t are 0).

F Proof of Proposition 1

Proof. The Proposition statement defined

$$\varsigma_{s:t-1} := \frac{1}{t-s} \sum_{l=s}^{t-1} \left| v(\mathbf{K}_l, \hat{\mathbf{z}}_{l+1}, \tilde{\boldsymbol{\mu}}_{l+1|l+1}) - \mathbb{E}_l^h \left[v(\mathbf{K}_l, \hat{\mathbf{z}}_{l+1}, \tilde{\boldsymbol{\mu}}_{l+1|l+1}) \right] \right|,$$

where both components of each summation term are fully observed. Then the probability of a large deviation can be bounded by

$$\Pr(\varsigma_{s:t-1} \geq \varsigma_a) \leq \frac{\mathbb{E}^g[\varsigma_{s:t-1}]}{\varsigma_a} \approx \frac{\varsigma_{0:s-1}}{\varsigma_a},$$

where we have utilized the Markov's inequality (also see Chernoff, 1952, bound), and approximated an unknown population moment with its sample counterpart that is based on observations from 0 to $s - 1$, i.e. those collected right after the model's last (re-) solution.

Hence, rejecting the hypothesis that the (null) model is still valid upon exceeding threshold ς_a will be an erroneous decision (Type-I error) with probability not higher than $\varsigma_{0:s-1}/\varsigma_a$. Increasing the threshold will reduce the chance of an erroneous decision. Specifically, a Type-I error rate of $a \times 100\%$ requires setting ς_a to $\varsigma_{0:s-1}/a$.

Finally, rejection of the hypothesis implies re-solution of problem $\mathcal{P}_{\mathcal{TL}}$. □

G Proof of Proposition 2

Proof. First, we are using the assumption on the validity of replacement of the value function with the wealth variable (as is the case when one is proportional to another). Wealth W_{t+1} here equals $(\mathbf{r}_{t+1} + (\mathbf{1} - \boldsymbol{\delta}))^\top \mathbf{K}_t + \mathbf{w}_{t+1}^\top \mathbf{L}_{t+1}$. Its stochastic component is $\mathbf{r}_{t+1}^\top \mathbf{K}_t + \mathbf{w}_{t+1}^\top \mathbf{L}_{t+1}$, i.e., aggregate output $\sum_k Y_{k,t+1}$.

Next, we express wealth in continuous time as follows (see Verstyuk, 2017; this approximation is related to Campbell and Viceira, 2002a). Let productivity evolve as

$$d\mathbf{Z}_t := d\exp(\mathbf{z}_t) := \text{diag}(\mathbf{Z}_t)(\mathbf{A}_z \boldsymbol{\mu}_t + \frac{1}{2} \text{diag}^{-1}(\boldsymbol{\sigma}_{\varepsilon z} \boldsymbol{\sigma}_{\varepsilon z}^\top)) dt + \text{diag}(\mathbf{Z}_t) \boldsymbol{\sigma}_{\varepsilon z} d\mathbf{B}_{\varepsilon z, t},$$

where $\boldsymbol{\sigma}_{\varepsilon z}$ is such that $\boldsymbol{\sigma}_{\varepsilon z} \boldsymbol{\sigma}_{\varepsilon z}^\top = \boldsymbol{\Sigma}_{\varepsilon z}$, and $\mathbf{B}_{\varepsilon z, t}$ is a standard K -dimensional Brownian motion.

Increasing time interval to $dt = 1$ and utilizing the fact that the mean of \mathbf{Z}_{t+1} is fixed at $t+1$ (it is a “predictable process”) produces the following continuous-time approximation to discrete-time stochastic dynamics:

$$\mathbf{Z}_{t+1} - \mathbf{Z}_t \approx \text{diag}(\mathbf{Z}_t)(\mathbf{A}_z \boldsymbol{\mu}_{t+1} + \frac{1}{2} \text{diag}^{-1}(\boldsymbol{\sigma}_{\varepsilon z} \boldsymbol{\sigma}_{\varepsilon z}^\top)) dt + \text{diag}(\mathbf{Z}_t) \boldsymbol{\sigma}_{\varepsilon z} (\mathbf{B}_{\varepsilon z, t+1} - \mathbf{B}_{\varepsilon z, t}).$$

Subtracting from the above expression its constrained counterpart

$$\hat{\mathbf{Z}}_{t+1} - \mathbf{Z}_t \approx \text{diag}(\mathbf{Z}_t)(\mathbf{A}_z \hat{\boldsymbol{\mu}}_{t+1} + \frac{1}{2} \text{diag}^{-1}(\hat{\boldsymbol{\sigma}}_{\varepsilon z} \hat{\boldsymbol{\sigma}}_{\varepsilon z}^\top)) dt + \text{diag}(\mathbf{Z}_t) \hat{\boldsymbol{\sigma}}_{\varepsilon z} (\mathbf{B}_{\varepsilon z, t+1} - \mathbf{B}_{\varepsilon z, t}).$$

gives

$$\begin{aligned} \mathbf{Z}_{t+1} - \hat{\mathbf{Z}}_{t+1} &\approx \text{diag}(\mathbf{Z}_t) \left(\mathbf{A}_z \boldsymbol{\mu}_{t+1} - \mathbf{A}_z \hat{\boldsymbol{\mu}}_{t+1} + \frac{1}{2} \text{diag}^{-1}(\boldsymbol{\sigma}_{\varepsilon z} \boldsymbol{\sigma}_{\varepsilon z}^\top) - \frac{1}{2} \text{diag}^{-1}(\hat{\boldsymbol{\sigma}}_{\varepsilon z} \hat{\boldsymbol{\sigma}}_{\varepsilon z}^\top) + \right. \\ &\quad \left. + \boldsymbol{\sigma}_{\varepsilon z} (\mathbf{B}_{\varepsilon z, t+1} - \mathbf{B}_{\varepsilon z, t}) - \hat{\boldsymbol{\sigma}}_{\varepsilon z} (\mathbf{B}_{\varepsilon z, t+1} - \mathbf{B}_{\varepsilon z, t}) \right) =: \\ &=: \text{diag}(\mathbf{Z}_t)(\mathbf{z}_{t+1} - \hat{\mathbf{z}}_{t+1} + \mathbf{A}_z \check{\boldsymbol{\mu}}_{t+1}). \end{aligned}$$

But of course, since $\boldsymbol{\mu}_t$ is unobserved, a more relevant expression for $d\mathbf{Z}_t$ is

$$d\mathbf{Z}_t := \text{diag}(\mathbf{Z}_t)(\mathbf{A}_z\boldsymbol{\mu}_{t|t} + \frac{1}{2}\text{diag}^{-1}(\boldsymbol{\sigma}_{\zeta\mu}\boldsymbol{\sigma}_{\zeta\mu}^\top) + \frac{1}{2}\text{diag}^{-1}(\boldsymbol{\sigma}_{\varepsilon z}\boldsymbol{\sigma}_{\varepsilon z}^\top))dt + \\ + \text{diag}(\mathbf{Z}_t)(\boldsymbol{\sigma}_{\zeta\mu}d\mathbf{B}_{\zeta\mu,t} + \boldsymbol{\sigma}_{\varepsilon z}d\mathbf{B}_{\varepsilon z,t}),$$

adding another layer of risks (and extending the notation in an obvious manner, except perhaps that $\boldsymbol{\sigma}_{\zeta\mu}\boldsymbol{\sigma}_{\zeta\mu}^\top = \boldsymbol{\Upsilon}_{\mu,t+1|t}$); and thus

$$\mathbf{Z}_{t+1} - \mathbf{Z}_t \approx \text{diag}(\mathbf{Z}_t)(\mathbf{A}_z\boldsymbol{\mu}_{t+1|t} + \frac{1}{2}\text{diag}^{-1}(\boldsymbol{\sigma}_{\zeta\mu}\boldsymbol{\sigma}_{\zeta\mu}^\top) + \frac{1}{2}\text{diag}^{-1}(\boldsymbol{\sigma}_{\varepsilon z}\boldsymbol{\sigma}_{\varepsilon z}^\top))dt + \\ + \text{diag}(\mathbf{Z}_t)(\boldsymbol{\sigma}_{\zeta\mu}(\mathbf{B}_{\zeta\mu,t+1} - \mathbf{B}_{\zeta\mu,t}) + \boldsymbol{\sigma}_{\varepsilon z}(\mathbf{B}_{\varepsilon z,t+1} - \mathbf{B}_{\varepsilon z,t})),$$

with the corresponding modifications to constrained counterparts, and so on.

Note that in the calculations that follow we use instead of $\boldsymbol{\mu}_{t+1|t}$ the object $\tilde{\boldsymbol{\mu}}_{t+1|t}$ (with the necessary corrections to volatilities). Technically, to allow the means to cancel out; economically, to avoid the complexity of computing both objects.

Then we have

$$W_{t+1} - \hat{W}_{t+1} = (\mathbf{Y}_{t+1} - \hat{\mathbf{Y}}_{t+1})^\top \mathbf{1} = \\ = \mathbf{Z}_{t+1}^\top (\mathbf{K}(\mathbf{K}_{t-1}, \mathbf{z}_t, \boldsymbol{\mu}_{t|t})^\alpha \odot \mathbf{L}(\mathbf{K}_t, \mathbf{z}_{t+1}, \boldsymbol{\mu}_{t+1|t+1})^{1-\alpha}) - \\ - \hat{\mathbf{Z}}_{t+1}^\top (\mathbf{K}(\mathbf{K}_{t-1}, \hat{\mathbf{z}}_t, \tilde{\boldsymbol{\mu}}_{t|t})^\alpha \odot \mathbf{L}(\mathbf{K}_t, \hat{\mathbf{z}}_{t+1}, \tilde{\boldsymbol{\mu}}_{t+1|t+1})^{1-\alpha}) \approx \\ \approx (\mathbf{Z}_{t+1} - \hat{\mathbf{Z}}_{t+1})^\top (\mathbf{K}(\mathbf{K}_{t-1}, \mathbf{z}_t, \boldsymbol{\mu}_{t|t})^\alpha \odot \mathbf{L}(\mathbf{K}_t, \mathbf{z}_{t+1}, \boldsymbol{\mu}_{t+1|t+1})^{1-\alpha}) \approx \\ \approx (\mathbf{z}_{t+1} - \hat{\mathbf{z}}_{t+1} + \mathbf{A}_z\check{\boldsymbol{\mu}}_{t+1})^\top (\mathbf{Z}_t \odot \mathbf{K}_t^\alpha \odot \mathbf{L}_{t+1}^{1-\alpha}).$$

In the above equation, the penultimate approximate equality used the second assumption that the differences in second-order effects of \mathbf{Z}_{t+1} and $\boldsymbol{\mu}_{t+1|t+1}$ as compared to $\hat{\mathbf{Z}}_{t+1}$ and $\tilde{\boldsymbol{\mu}}_{t+1|t+1}$ on wealth via labor (and similarly for capital) are “small”, and are dominated by the first-order effect ($\mathbf{Z}_{t+1} - \hat{\mathbf{Z}}_{t+1}$).

Using the third assumption about minimization of maximum loss for any possible combination of control and state variables entering the term $(\mathbf{Z}_t \odot \mathbf{K}_t^\alpha \odot \mathbf{L}_{t+1}^{1-\alpha})$ produces the claimed result (following an argument similar to Verstyuk, 2017):

$$(W_{t+1} - \hat{W}_{t+1})^2 \approx ((\mathbf{z}_{t+1} - \hat{\mathbf{z}}_{t+1} + \mathbf{A}_z\check{\boldsymbol{\mu}}_{t+1})^\top (\mathbf{Z}_t \odot \mathbf{K}_t^\alpha \odot \mathbf{L}_{t+1}^{1-\alpha}))^2 \propto \\ \propto (\mathbf{z}_{t+1} - \hat{\mathbf{z}}_{t+1} + \mathbf{A}_z\check{\boldsymbol{\mu}}_{t+1})^\top (\mathbf{z}_{t+1} - \hat{\mathbf{z}}_{t+1} + \mathbf{A}_z\check{\boldsymbol{\mu}}_{t+1})$$

Above, we have defined the conditional mean $\hat{\boldsymbol{\mu}}_{t+1|t}$ of a random variable $\hat{\mathbf{z}}_{t+1}$ as

$$\hat{\boldsymbol{\mu}}_{t+1|t} := \tilde{\boldsymbol{\mu}}_{t+1|t} + \check{\boldsymbol{\mu}}_{t+1},$$

where the bias term is

$$\check{\boldsymbol{\mu}}_{t+1} := \begin{bmatrix} \mathbf{0}_{K \times K} \\ \frac{1}{2}\text{diag}^{-1}(\boldsymbol{\Sigma}_{z,t+1|t} - \hat{\boldsymbol{\Sigma}}_{z,t+1|t}) \end{bmatrix}.$$

□

H Proof of Proposition 3

Proof. Complexity problem comprises

$$\min_{\tilde{\mathbf{B}}_t, f(\cdot, \cdot)} \mathbb{E}_t^f \left[\|\mathbf{z}_{t+1} - \hat{\mathbf{z}}_{t+1} + \mathbf{A}_z \check{\boldsymbol{\mu}}_{t+1}\|_2^2 \right] = \text{tr}(\boldsymbol{\Psi}_{z,t+1|t})$$

subject to constraints

$$\begin{aligned} \mathcal{T}(\tilde{\boldsymbol{\mu}}_{t+1|t}, \hat{\boldsymbol{\mu}}_{t+1|t} | \mathbf{z}_t, \tilde{\boldsymbol{\mu}}_{t|t-1}) &\leq \kappa_{\mathcal{T}}, \\ \mathcal{I}(g(\mathbf{z}_{t+1} | \tilde{\boldsymbol{\mu}}_{t+1|t}); h(\hat{\mathbf{z}}_{t+1} | \hat{\boldsymbol{\mu}}_{t+1|t})) &\leq \kappa_{\mathcal{I}}. \end{aligned}$$

Expanding the terms gives

$$\begin{aligned} \mathbf{z}_{t+1} - \hat{\mathbf{z}}_{t+1} + \mathbf{A}_z \check{\boldsymbol{\mu}}_{t+1} &= \mathbf{A}_z \boldsymbol{\mu}_{t+1} + \boldsymbol{\varepsilon}_{z,t+1} - (\mathbf{A}_z \hat{\boldsymbol{\mu}}_{t+1} + \hat{\boldsymbol{\varepsilon}}_{z,t+1}) + \mathbf{A}_z \check{\boldsymbol{\mu}}_{t+1} = \\ &= \mathbf{A}_z \boldsymbol{\mu}_{t+1|t} + \mathbf{A}_z \boldsymbol{\zeta}_{\mu,t+1|t} + \boldsymbol{\varepsilon}_{z,t+1} - (\mathbf{A}_z \tilde{\boldsymbol{\mu}}_{t+1|t} + \mathbf{A}_z \hat{\boldsymbol{\zeta}}_{\mu,t+1|t} + \hat{\boldsymbol{\varepsilon}}_{z,t+1}) = \\ &= \mathbf{A}_z \boldsymbol{\epsilon}_{\mu,t+1|t} + \boldsymbol{\epsilon}_{z,t+1}. \end{aligned}$$

Note that we replace $\boldsymbol{\mu}_{t+1|t}$ with $\tilde{\boldsymbol{\mu}}_{t+1|t}$ throughout (adding the necessary corrections to other related terms), which spares the agent from the complexity of computing both objects.

Then (conditionally on t) the variance-covariance matrix is

$$\begin{aligned} \mathbb{E}_t^f \left[(\mathbf{z}_{t+1} - \hat{\mathbf{z}}_{t+1} + \mathbf{A}_z \check{\boldsymbol{\mu}}_{t+1})(\mathbf{z}_{t+1} - \hat{\mathbf{z}}_{t+1} + \mathbf{A}_z \check{\boldsymbol{\mu}}_{t+1})^\top \right] &= \mathbf{A}_z \boldsymbol{\Upsilon}_{\mu,t+1|t} \mathbf{A}_z^\top + \boldsymbol{\Sigma}_{\varepsilon z} - \\ &\quad - (\mathbf{A}_z \hat{\boldsymbol{\Upsilon}}_{\mu,t+1|t} \mathbf{A}_z^\top + \hat{\boldsymbol{\Sigma}}_{\varepsilon z}) = \\ &= \mathbf{A}_z \boldsymbol{\Psi}_{\mu,t+1|t} \mathbf{A}_z^\top + \boldsymbol{\Psi}_{\varepsilon z} \end{aligned}$$

(this decomposition is verified by the subsequent solution to $\mathcal{P}_{\mathcal{I}}$), yielding

$$\min_{\tilde{\mathbf{B}}_t, f(\cdot, \cdot)} \mathbb{E}_t^f \left[\|\mathbf{z}_{t+1} - \hat{\mathbf{z}}_{t+1} + \mathbf{A}_z \check{\boldsymbol{\mu}}_{t+1}\|_2^2 \right] \iff \min_{\tilde{\mathbf{B}}_t, f(\cdot, \cdot)} \left\{ \text{tr}(\mathbf{A}_z \boldsymbol{\Psi}_{\mu,t+1|t} \mathbf{A}_z^\top + \boldsymbol{\Psi}_{\varepsilon z}) \right\}.$$

Now we get to the actual argument. On the one hand, taking $\tilde{\boldsymbol{\mu}}_{t+1|t}$ as given, minimizing the problem's criterion with respect to its second argument requires:

$$\min_{f(\cdot, \cdot)} \mathbb{E}_t^f \left[\|\mathbf{z}_{t+1} - \hat{\mathbf{z}}_{t+1} + \mathbf{A}_z \check{\boldsymbol{\mu}}_{t+1}\|_2^2 \right] = \text{tr}(\boldsymbol{\Psi}_{z,t+1|t}),$$

i.e. solving $\mathcal{P}_{\mathcal{I}}$.

On the other hand, taking the solution to $\mathcal{P}_{\mathcal{I}}$ (as well as the parameter values) as given, further reduction of $\mathbf{A}_z \boldsymbol{\Psi}_{\mu,t+1|t} \mathbf{A}_z^\top + \boldsymbol{\Psi}_{\varepsilon z}$ amounts to reducing $\mathbf{A}_z \boldsymbol{\Upsilon}_{\mu,t+1|t} \mathbf{A}_z^\top + \boldsymbol{\Sigma}_{\varepsilon z}$, that is the (conditional) variance of \mathbf{z}_{t+1} :

$$\min_{\tilde{\mathbf{B}}_t} \left\{ \text{tr}(\mathbf{A}_z \boldsymbol{\Upsilon}_{\mu,t+1|t} \mathbf{A}_z^\top + \boldsymbol{\Sigma}_{\varepsilon z}) \right\} \iff \min_{\tilde{\mathbf{B}}_t} \mathbb{E}_t^f \left[\|\mathbf{A}_z \boldsymbol{\zeta}_{\mu,t+1|t} + \boldsymbol{\varepsilon}_{z,t+1}\|_2^2 \right]$$

However, by orthogonality of $\boldsymbol{\zeta}_{\mu,t+1|t}$ and $\boldsymbol{\varepsilon}_{z,t+1}$ (as well as admitting any value of \mathbf{A}_z), this is equivalent to

$$\min_{\tilde{\mathbf{B}}_t} \mathbb{E}_t^f \left[\|\boldsymbol{\mu}_{t+1} - \tilde{\boldsymbol{\mu}}_{t+1|t}\|_2^2 \right] = \text{tr}(\boldsymbol{\Upsilon}_{\mu,t+1|t});$$

which in turn boils down to

$$\min_{\tilde{\mathbf{B}}_t} \mathbb{E}_t^f [\|\boldsymbol{\mu}_t - \tilde{\boldsymbol{\mu}}_{t|t}\|_2^2] = \text{tr}(\boldsymbol{\Upsilon}_{\mu,t|t}),$$

i.e. to solving $\mathcal{P}_{\mathcal{T}}$. □

I Proof of Proposition 4

Proof. Updates of the components of vector $\tilde{\boldsymbol{\mu}}_{t|t}$ (and eventually, $\tilde{\boldsymbol{\mu}}_{t+1|t}$) are produced by non-zero elements of $\tilde{\mathbf{B}}_t$. Exploiting the independence of shocks driving the stochastic process for productivity \mathbf{z}_t (as well as equal costs of a component's partial and full update), it is optimal to allocate the capacity $\kappa_{\mathcal{T}}$ starting from the largest diagonal element of matrix $\boldsymbol{\Upsilon}_{\mu,t|t}$. Thus in the optimum, the minimizer $\tilde{\mathbf{B}}_t$ is determined by the ordering of its elements' contributions to volatilities in matrix $\boldsymbol{\Upsilon}_{\mu,t|t}$.

Subsequent updates of the components of vector $\tilde{\boldsymbol{\mu}}_{t|t}$, which is the reason for updating $\tilde{\boldsymbol{\mu}}_{t|t}$ in the first place, similarly depends on the number of non-zero elements of $\tilde{\boldsymbol{\mu}}_t$ (eventually, $\hat{\boldsymbol{\mu}}_{t+1|t}$ and $\check{\boldsymbol{\mu}}_t$).

The number of operations required to update k elements of vector $\tilde{\boldsymbol{\mu}}_{t|t}$ (actually, $\tilde{\boldsymbol{\mu}}_{t+1|t}$) is $((2k)^2)k + k + (k^2)2k + 2k + (2k)^3$ (following the equations in $\mathcal{P}_{\rightsquigarrow}$; note that $\boldsymbol{\Upsilon}_{\mu,t|t-1}$ as well as $\boldsymbol{\Upsilon}_{\mu,t|t}$, and hence $\tilde{\mathbf{B}}_t$ are constants, hence do not require updating each period). Then, the number of operations required to update k elements of vector $\hat{\boldsymbol{\mu}}_{t|t}$ (actually, $\hat{\boldsymbol{\mu}}_{t+1|t}$) is $2k$ (using its definition, \mathcal{P}_{\vee} , from Proposition 2; note that $\boldsymbol{\Sigma}_{z,t+1|t}$, $\hat{\boldsymbol{\Sigma}}_{z,t+1|t}$ as well as $\boldsymbol{\Psi}_{z,t+1|t}$, and hence $\check{\boldsymbol{\mu}}_{t+1}$ are constants and do not require updating). The sum of these two quantities gives the number of operations required to update k largest elements of vectors $\tilde{\boldsymbol{\mu}}_{t|t}$ and $\hat{\boldsymbol{\mu}}_{t|t}$ (eventually, $\tilde{\boldsymbol{\mu}}_{t+1|t}$ and $\hat{\boldsymbol{\mu}}_{t+1|t}$). We want to make k as close to its upper bound K as resources $\kappa_{\mathcal{T}}$ allow, which gives the stated solution. □

References

- [1] Angeletos, George-Marios and Chen Lian. (2016) “Forward Guidance without Common Knowledge.” Manuscript.
- [2] Angeletos, George-Marios and Chen Lian. (2017) “Dampening General Equilibrium: From Micro to Macro.” Manuscript.
- [3] Arora, Sanjeev and Boaz Barak. (2009) *Computational Complexity: A Modern Approach*, New York, NY: Cambridge University Press.
- [4] Auerbach, Alan J. and Yuriy Gorodnichenko. (2012) “Measuring the Output Responses to Fiscal Policy”, *American Economic Journal: Economic Policy*, 4(2): 1–27.
- [5] Barro, Robert J. (1974) “Are Government Bonds Net Wealth?”, *Journal of Political Economy*, 82(6): 1095–1117.
- [6] Campbell, John Y. and John H. Cochrane. (1999) “By Force of Habit: A Consumption-Based Explanation of Aggregate Stock Market Behavior”, *Journal of Political Economy*, 107(2): 205–251.
- [7] Campbell, John and Angus Deaton. (1989) “Why is Consumption So Smooth?”, *Review of Economic Studies*, 56(3): 357–373.
- [8] Campbell, John Y. and N. Gregory Mankiw. (1989) “Consumption, Income, and Interest Rates: Reinterpreting the Time Series Evidence.” *NBER Macroeconomics Annual*, 4: 185–245.
- [9] Carroll, Christopher D., Jody Overland and David N. Weil. (2000) “Saving and Growth with Habit Formation”, *American Economic Review*, 90(3): 341–355.
- [10] Carroll, Christopher D. (2001) “A Theory of the Consumption Function, With and Without Liquidity Constraints (Expanded Version).” NBER Working Paper.
- [11] Congressional Budget Office. (2012) *Estimated Impact of the American Recovery and Reinvestment Act on Employment and Economic Output from October 2011 through December 2011*, Washington, DC: Congressional Budget Office.
- [12] Cochrane, John H. (2016) “Comments on ”A Behavioral New-Keynesian Model” by Xavier Gabaix.” Manuscript.
- [13] Cormen, Thomas H., Charles E. Leiserson, Ronald L. Rivest and Clifford Stein. (2009) *Introduction to Algorithms*, 3rd ed., Cambridge, MA: The MIT Press.
- [14] Cover, Thomas M. and Joy A. Thomas. (2006) *Elements of Information Theory*, 2nd ed., New York: Wiley-Interscience.
- [15] Del Negro, Marco, Marc Giannoni and Christina Patterson. (2015) “The Forward Guidance Puzzle”, *Federal Reserve Bank of New York Staff Report*, No. 574.
- [16] Farhi, Emmanuel and Iván Werning. (2017) “Monetary Policy, Bounded Rationality, and Incomplete Markets.” Manuscript.
- [17] Farhi, Emmanuel and Iván Werning. (2016) “Fiscal Multipliers: Liquidity Traps and Currency Unions.” In: John B. Taylor and Harald Uhlig (eds.), *Handbook of Macroeconomics*, edition 1, volume 2B, chapter 31, pages 2417–2492, Elsevier.
- [18] Faria-e-Castro, Miguel. (2017) “Fiscal Multipliers and Financial Crises.” Manuscript.
- [19] Fernáandez-Villaverde, J., J.F. Rubio-Ramírez and F. Schorfheide. (2016) “Solution and Estimation Methods for DSGE Models.” In: John B. Taylor and Harald Uhlig (eds.), *Handbook of Macroeconomics*, edition 1, volume 2A, chapter 9, pages 527–724, Elsevier.

- [20] Friedman, Milton. (1957) “The Permanent Income Hypothesis.” In: *A Theory of the Consumption Function*, Princeton University Press.
- [21] Fuhrer, Jeffrey C. (2000) “Habit Formation in Consumption and Its Implications for Monetary-Policy Models”, *American Economic Review*, 90(3): 367–390.
- [22] Gabaix, Xavier. (2016a) “A Behavioral New Keynesian Model.” Manuscript.
- [23] Gabaix, Xavier. (2016b) “Behavioral Macroeconomics Via Sparse Dynamic Programming.” Manuscript.
- [24] Galí, Jordi, J. David López-Salido and Javier Vallés. (2007) “Understanding the Effects of Government Spending on Consumption”, *Journal of the European Economic Association*, 5(1): 227–270.
- [25] George, Edward I. and Robert E. McCulloch. (1993) “Variable Selection Via Gibbs Sampling”, *Journal of the American Statistical Association*, 88(423): 881–889.
- [26] Leung-Yan-Cheong, Sik K. and Thomas M. Cover. (1978) “Some Equivalences Between Shannon Entropy and Kolmogorov Complexity”, *IEEE Transactions on Information Theory*, IT-24 (3): 331–338.
- [27] Lucas, Deborah J. (2016) “Credit Policy as Fiscal Policy”, *Brookings Papers on Economic Activity*: 1–41.
- [28] Madigan, David and Adrian E. Raftery. “Model Selection and Accounting for Model Uncertainty in Graphical Models Using Occam’s Window”, *Journal of the American Statistical Association*, 89(428): 1535–1546.
- [29] Mankiw, N. Gregory. (2000) “The Savers-Spenders Theory of Fiscal Policy”, *American Economic Review*, 90(2): 120–125.
- [30] Mitchell, T. J. and J. J. Beauchamp. (1988) “Bayesian Variable Selection in Linear Regression”, *Journal of the American Statistical Association*, 83(404): 1023–1032.
- [31] Parker, Jonathan A. (2017) “Why Don’t Households Smooth Consumption? Evidence from a \$25 million experiment”, *American Economic Journal: Macroeconomics*, forthcoming.
- [32] Pástor, Luboš and Pietro Veronesi. (2003) “Stock Valuation and Learning about Profitability”, *Journal of Finance*, 58(5): 1749–1789.
- [33] Rissanen, Jorma. (2007) *Information and Complexity in Statistical Modeling*, Springer.
- [34] Teixeira, Bruno Otávio Soares, Jaganath Chandrasekar, Harish J. Palanthandalam-Madapusi, Leonardo Antônio Borges Tôrres, Luis Antonio Aguirre and Dennis S. Bernstein. (2008) “Gain-constrained Kalman Filtering for Linear and Nonlinear Systems”, *IEEE Transactions on Signal Processing*, 56(9): 4113–4123.
- [35] Veronesi, Pietro. (1999) “Stock Market Overreaction to Bad News in Good Times: A Rational Expectations Equilibrium Model”, *Review of Financial Studies*, 12(5): 975–1007.
- [36] Veronesi, Pietro. (2000) “How Does Information Quality Affect Stock Returns?”, *Journal of Finance*, 55(2): 807–837.
- [37] Verstyuk, Sergiy. (2017) “Log, Stock and Two Simple Lotteries.” Manuscript.
- [38] Woodford, Michael. (2011) “Simple Analytics of the Government Expenditure Multiplier”, *American Economic Journal: Macroeconomics*, 3(1): 1–35.