

Macroeconomics with Learning and Misspecification: A General Theory and Applications*

Pooya Molavi[†]

MIT

Job Market Paper

January 27, 2019

The latest version can be found [here](#).

Abstract

This paper explores a form of bounded rationality where agents learn about the economy with possibly misspecified models. I consider a recursive general-equilibrium framework that nests a large class of macroeconomic models. Misspecification is represented as a constraint on the set of beliefs agents can entertain. I introduce the solution concept of constrained-rational-expectations equilibrium (CREE), in which each agent selects the belief from her constrained set that is closest to the endogenous distribution of observables in the Kullback–Leibler divergence. If the set of permissible beliefs contains the rational-expectations equilibria (REE), then the REE are CREE; otherwise, they are not. I show that a CREE exists, that it arises naturally as the limit of adaptive and Bayesian learning, and that it incorporates a version of the Lucas critique. I then apply CREE to a particular novel form of bounded rationality where beliefs are constrained to factor models with a small number of endogenously chosen factors. Misspecification leads to amplification or dampening of shocks and history dependence. The calibrated economy exhibits hump-shaped impulse responses and co-movements in consumption, output, hours, and investment that resemble business-cycle fluctuations.

*I am grateful to my advisors Iván Werning, Daron Acemoglu, Drew Fudenberg, and Marios Angeletos for their continuous guidance and support, to Arda Gitmez, Alireza Tahbaz-Salehi, and Olivier Wang for countless discussions, suggestions, and comments that helped me in every stage of this project, and to Abhijit Banerjee, Martin Beraja, Ricardo Caballero, Victor Chernozhukov, Arnaud Costinot, Michele Fornino, Joe Hazell, Ali Jadbabaie, Ali Kakhbod, Omer Karaduman, Amir Kermani, Danial Lashkari, Claire Lazar, Anna Mikusheva, Christina Patterson, Will Rafe, Karthik Sastry, Frank Schilbach, Alp Simsek, Phillip Strack, Andrea Vedolin, Daniel Waldinger, Alex Wolitzky, Muhamet Yildiz, and Nathan Zorzi for their helpful comments.

[†]Email: mpooya@mit.edu.

1 Introduction

The rational-expectations assumption underlies much of modern macroeconomic modeling. It makes expectations model-consistent and predicts how they vary with changes in policy. But the level of sophistication that it requires from agents may be unrealistically high. By putting a straitjacket on agents' expectations, the rational-expectations assumption also rules out many compelling biases that have been documented through lab experiments (e.g., by [Kahneman and Tversky, 1972, 1982a,b](#)) and survey evidence (e.g., by [Coibion and Gorodnichenko, 2012, 2015a,b](#)). While alternatives to rational expectations have been proposed (e.g., by [Marcet and Sargent, 1989a,b](#)), a unified and flexible framework has been lacking. Abandoning rational expectations leads a modeler into the “wilderness of bounded rationality,” as [Sargent \(1999\)](#) famously quipped. How might we improve the realism of our models while maintaining the endogeneity and discipline of rational expectations?

This paper proposes a flexible recursive general-equilibrium framework in which agents face cognitive and behavioral constraints but do the best they can within the confines of their constraints. It relaxes the rational-expectations assumption by breaking the link between agents' models of the economy and the economic models they inhabit. Each agent is endowed with a set of models for the economy, uses them to learn about her environment, and selects the model that best describes her observations. If the agents' models rule out the equilibrium distribution of observables, the models are misspecified. Model misspecification can be seen as a unified expression of bounded rationality: I show that many deviations from rational expectations proposed in the literature are instances of misspecification.

The paper makes three main contributions:

1. It introduces a general solution concept given arbitrary forms of misspecification and provides Bayesian and adaptive learning foundations for the solution concept.
2. It proposes a new deviation from rational expectations, where agents' models are constrained to low-dimensional hidden factor models.
3. It considers a medium-scale business-cycle model in which agents entertain one-dimensional factor models and studies the quantitative implications of bounded rationality.

In the first part of the paper, I introduce the solution concept of *constrained-rational-expectations equilibrium (CREE)* as the misspecified counterpart to the rational-expectations equilibrium (REE). In a CREE, (i) agents maximize their utilities given their expectations, (ii) the general-equilibrium consistency requirements are satisfied, and (iii) each agent's expectation is chosen endogenously to minimize the distance to the distribution of observables. I build on [Esponda and Pouzo \(2016a,b\)](#), extending their notion of Berk–Nash equilibrium to recursive

general-equilibrium macro settings and extending the scope of the theory to allow for non-convergence of beliefs. I show that a CREE exists under general conditions. I also prove that the economy converges to a CREE as agents learn about their environment—either adaptively or using Bayes’ rule. If the agents’ models are correctly specified, then the set of REE is a subset of the set of CREE; otherwise, it is not. In the second part of the paper, I use CREE to study the implications of a particular form of bounded rationality in a macroeconomic context. Agents’ models are given by hidden factor models with a small number of factors. I provide conditions under which the response of the CREE to shocks is amplified/dampened relative to the REE. I also show that the one-dimensional model with CREE exhibits outcomes that are qualitatively similar to noisy information models with Kalman filters (e.g., [Sims, 2003](#)). But in a CREE, agents conjure up the measurement noise to better fit their misspecified low-dimensional models to their observations.

In Section 2, I introduce the general framework used throughout the paper. I consider a recursive general-equilibrium economy with a representative agent. The representative agent solves a dynamic decision problem subject to a set of constraints on her choices. The agent’s preferences and her choice set are functions of *observables*, which may include variables such as prices, income, and taste shocks. The observables may be affected by the agent’s choices in equilibrium, but she takes them as given. The agent’s optimal decisions thus depend on her expectations about the future values of observables. The framework can nest a large class of macroeconomic models—including some with heterogeneous agents and incomplete information.

I begin Section 3 by adapting the temporary equilibrium of [Grandmont \(1977\)](#) and [Woodford \(2013\)](#) to the general framework of this paper. I define a *temporary equilibrium* as a mapping from the agent’s expectations and past values of equilibrium variables to the current distribution of equilibrium variables. In a temporary equilibrium, the agent’s decisions are optimal given her expectations, and the general-equilibrium consistency requirements such as market-clearing conditions are satisfied. A temporary equilibrium is a convenient device for summarizing the fundamentals of the economy, but it is silent on the determination of the agent’s expectations.

Instead of completing the description of equilibrium by imposing rational expectations, I assume that the agent is endowed with *a set of models* and forms her expectations by fitting her models to her observations. Each model defines a probability distribution over the path of observables. The set of models may or may not include the probability distribution that arises in equilibrium. If it does, then the agent’s collection of models is *correctly specified*; otherwise, it is *misspecified*. When the agent has a misspecified set of models, she selects a model from her misspecified set that minimizes a novel version of the Kullback–Leibler divergence from the equilibrium distribution of observables. I later provide two formal learning foundations for the Kullback–Leibler divergence as the appropriate notion of distance in the macro settings considered in this paper.

Section 3 goes on to formally define the solution concept of CREE. A CREE consists of a temporary equilibrium, a stationary probability distribution for equilibrium variables, and the set of models for the agent that best fit her observations. In a CREE, the evolution of equilibrium variables is determined by the temporary equilibrium; the agent updates her belief using Bayes' rule; the equilibrium variables are distributed according to the stationary distribution; and the agent puts positive probability only on the set of models that minimize the Kullback–Leibler divergence given the temporary equilibrium and the stationary distribution. I show that a CREE exists as long as some weak continuity and compactness conditions are satisfied.

In Section 4, I provide learning foundations for CREE. I first show that CREE arises as the limit of Bayesian learning. I consider an agent with a subjective prior who updates her belief over time using Bayes' rule and makes optimal choices given her belief—anticipating that she will update her belief in the future using Bayes' rule. This specification of beliefs leads to a recursive version of the internally-rational-expectations equilibrium of Adam and Marcet (2011). I show that—under a regularity condition called *asymptotic mean stationarity*—Bayesian learning leads the agent's belief to concentrate on the set of models that minimize the Kullback–Leibler divergence. It also leads the Bayesian equilibrium to converge to a CREE.

I then provide a second foundation for CREE by showing that adaptive learning also leads the agent's belief to concentrate on the set of models that minimize the Kullback–Leibler divergence. To study adaptive learning, I propose a generalization of the least-squares learning approach (e.g., Marcet and Sargent, 1989a,b, Evans and Honkapohja, 1995, 2012), in which the agent estimates her model using a quasi-maximum-likelihood estimator. The agent behaves optimally given her estimate and under the assumption that she will not update her belief. I show that under asymptotic mean stationarity of the adaptive equilibrium, the agent's belief concentrates on the set with the minimum Kullback–Leibler divergence from the equilibrium distribution of observables. I also provide sufficient conditions under which the equilibrium itself converges to a CREE as the result of adaptive learning.

The approach of CREE to the modeling of bounded rationality offers several conceptual and practical advantages over the existing alternatives. It unifies and connects various disparate approaches by casting them all as instances of misspecification. It facilitates the incorporation of new and rich behavioral biases in standard macro models. And it leads to tractable and portable models of expectation formation. The solution concept of CREE not only determines the agents' expectations in equilibrium but also predicts how those expectations vary with changes in the environment as agents learn. As such, it is particularly well-suited for counterfactual policy analysis: a policymaker who uses CREE to analyze the effects of alternative policies is not subject to the Lucas critique.

In Section 5, I use CREE to study how misspecified expectations change the response of an economy to shocks. I consider a simple linear economy with one choice variable, one state variable, and one shock. The agent's set of models is constrained to contain only hidden factor

models. I provide conditions under which the response of the economy to shocks in a CREE is amplified/dampened relative to the rational-expectations benchmark. I also show that while the REE does not exhibit history dependence, the CREE generically does so.

Section 6 describes the results of a quantitative exercise to study the role of bounded rationality in explaining business-cycle fluctuations. I consider a medium-scale economy with standard fundamentals. There are nominal price and wage rigidities, investment, neoclassical capital-adjustment costs, realistic monetary and fiscal policies, and a large number of nominal and real shocks. There is no habit formation in consumption, no wage or price indexation, no shock to the marginal product of investment, and no investment-adjustment cost.¹

I replace these frictions with a novel constraint on the agents' set of models: agents are constrained to entertain hidden factor models with d endogenously constructed factors. The agents' models are misspecified when d is small (relative to the size of the minimal state-space representation of the REE). I refer to the particular case of CREE where agents have factor models with d factors as CREE- d . A CREE- d has only one additional parameter compared to the REE: the number of factors d . The number of factors can be estimated jointly with the parameters of preferences, technology, and shocks using standard techniques. As a first pass, I set the number of factors equal to one and calibrate the other parameters to standard values from the DSGE literature.

The CREE-1 exhibits quantitatively different outcomes relative to the REE. The impulse responses of real variables to TFP and monetary-policy shocks exhibit hump shapes in contrast to the impulse responses in the REE. Demand shocks (i.e., government spending and discount-factor shocks) lead to co-movements in consumption, investment, hours, and output that match the salient features of business-cycle fluctuations. In the REE, in contrast, both government spending and discount-factor shocks lead to counterfactual co-movements. Bounded rationality in the form of CREE- d thus emerges as a parsimonious alternative to the battery of frictions needed in the DSGE literature to improve the empirical fit of standard models.

The CREE-1 model also gives rise to an *endogenous* measure of consumer confidence. Agents endogenously construct a hidden factor that they believe drives the movements in aggregate variables. I refer to the agents' estimate of the hidden factor as "consumer confidence." The consumer confidence increases the most with increases in investment, hours, the rental rate of capital, consumption, and income; it decreases the most with increases in taxes; and it is largely unaffected by changes in the values of shocks, inflation, the real wage, and the nominal interest rate. Positive supply and demand shocks both lead to increases in the consumer confidence and co-movements in aggregate variables. A change in consumer confidence thus works both as an

¹I assume a neoclassical *capital*-adjustment cost à la Hayashi (1982) and not an *investment*-adjustment cost like the one assumed in Christiano, Eichenbaum, and Evans (2005), Smets and Wouters (2007), and the DSGE literature that follows them.

amplification mechanism and as a propagation mechanism for the fundamental shocks that hit the economy.

Related Literature. This paper contributes to several lines of research in macroeconomics, game theory, behavioral economics, and econometrics. It builds on the literature on Bayesian inference under misspecification that goes back to the seminal work of Berk (1966).² In parallel to Berk's analysis of Bayesian inference, a literature going back to Huber (1967) has studied the asymptotic properties of the quasi-maximum-likelihood estimator (QMLE) under misspecification.³ My convergence results contribute to this literature by weakening the regularity assumption needed for convergence to that of asymptotic mean stationarity.

Esponda and Pouzo (2016a) adapt Berk's notion of misspecification to games. They introduce into an otherwise-standard static game the possibility that the players' models are misspecified and propose the solution concept of Berk–Nash equilibrium. In a Berk–Nash equilibrium, players play optimally given their beliefs, and beliefs minimize the weighted Kullback–Leibler divergence against the equilibrium distribution of play. Meanwhile, Esponda and Pouzo (2016b) adapt Berk–Nash equilibrium to Markov decision problems and show that Bayesian learning leads the agent's belief to concentrate on the set of models that minimize another version of the Kullback–Leibler divergence. In related work, Fudenberg, Romanyuk, and Strack (2017) consider active learning and information acquisition by an agent with a misspecified model, showing that whether the agent's belief converges may depend on her discount rate; Heidhues, Koszegi, and Strack (2018) study misspecified learning by an overconfident agent and show that overconfidence may systematically lead her away from the correct belief about the fundamental; and Gagnon-Bartsch, Rabin, and Schwartzstein (2018) propose channeled attention as a reason why an agent may not abandon her misspecified models in the face of mounting evidence that the models are misspecified.

This paper extends the analysis of Esponda and Pouzo to recursive dynamic general-equilibrium macro settings with a continuum of agents and continuous actions, observables, and state variables. What distinguishes general-equilibrium settings from games is the simultaneous determination of aggregate observables (such as prices) and the agents' choices (such as quantities consumed) through price-taking behavior and market clearing conditions.⁴ I adapt and generalize the notion of Berk–Nash equilibrium to macro settings, innovating along several dimensions. First, I exploit the recursive structure of the economy and the Markovian

²Berk showed that the posterior distribution of a Bayesian observer concentrates on the set of models that minimize the Kullback–Leibler divergence, where the Kullback–Leibler divergence given a candidate distribution is defined as the expected value (under the true distribution) of the log-likelihood ratio of the true distribution against the candidate distribution. Bunke and Milhaud (1998), Kleijn and Van Der Vaart (2006), and Shalizi (2009) extend the result of Berk by providing conditions for the weak convergence of posterior distributions and considering infinite-dimensional models and non-i.i.d. observations.

³Other important contributions to the analysis of QMLE include Pfanzagl (1969), White (1982), and Vuong (1989). A central finding of this literature is the convergence of QMLE to a model that minimizes the Kullback–Leibler divergence.

⁴It is important to note, however, that I do not assume that the economy is Walrasian. The flexible framework studied in the paper can embed frictions such as nominal rigidities and non-competitive behavior such as monopolistic competition.

structure of Bayes' rule to cast the dynamics of macro aggregates and beliefs, jointly, as a Markov chain over an infinite-dimensional space. Second, I generalize the definition of equilibrium to allow for ergodic distributions over aggregates and beliefs and weaken the notion of convergence to that of empirical distributions. Third, I use results from the theory of continuous Markov chains over general state spaces to establish the convergence of empirical distributions and concentration of beliefs under weak continuity and asymptotic mean stationarity conditions. Finally, I prove that learning in the macroeconomic settings considered in this paper leads the agents to minimize a new version of the weighted Kullback–Leibler divergence, where the weight is now given by the temporary equilibrium of the economy.

This paper also contributes to the literature that studies the implications of misspecification as a manifestation of bounded rationality in macroeconomics. This literature goes back to the work of [Bray \(1982\)](#) and [Bray and Savin \(1986\)](#) on restricted-perceptions equilibrium. In a restricted-perceptions equilibrium, misspecification takes the form of agents dropping some relevant variables from their regressions. [Hommes and Sorger \(1998\)](#) and [Branch and McGough \(2005\)](#) consider agents whose models are misspecified due to their use of linear equations to approximate non-linear equilibrium relationships. Agents in [Krusell and Smith \(1998\)](#) also have a misspecified model of the economy. They believe that current and future prices do not depend on anything but the first few moments of the wealth distribution. [Marcet and Sargent \(1989a,b\)](#), [Sargent \(1999\)](#), and [Marcet and Nicolini \(2003\)](#) consider agents who learn about their environment given a misspecified model according to which observations are generated from a stochastic process with a persistent hidden component.⁵ Besides unifying these models and nesting them as special cases, this paper endogenizes the parameters of the agents' learning rules: while the Kalman gain used by agents is a free parameter in the adaptive-learning literature, it is determined endogenously in a CREE and varies with changes in policy. The CREE thus incorporates a version of the Lucas critique.⁶ This paper also contributes to the adaptive-learning literature by proposing a generalization of least-squares learning in which agents update their beliefs using QMLE.

Agents with preferences that exhibit ambiguity aversion (as in [Gilboa and Schmeidler, 1989](#)) or robustness concerns (as in [Hansen and Sargent, 1995, 2001a,b](#)) also entertain a set of models, some of them misspecified.⁷ In the robustness approach of Hansen and Sargent, agents' set

⁵Other related papers in the learning literature include [Cho, Williams, and Sargent \(2002\)](#), [Bullard and Mitra \(2002\)](#), [Orphanides and Williams \(2005\)](#), [Preston \(2005\)](#), [Bullard, Evans, and Honkapohja \(2008\)](#), [Adam and Marcet \(2011\)](#), [Eusepi and Preston \(2011\)](#), [Malmendier and Nagel \(2016\)](#), [Adam, Marcet, and Beutel \(2017\)](#), and [Eusepi and Preston \(2018a\)](#). See [Eusepi and Preston \(2018b\)](#) for a recent survey of the literature.

⁶[Evans and Ramey \(2006\)](#) also provide conditions under which monetary policy is subject to the Lucas critique in a model with adaptive expectations.

⁷[Gilboa and Schmeidler \(1989\)](#) propose max-min expected utility as a formalization of ambiguity aversion and a possible resolution to the [Ellsberg \(1961\)](#) paradox. [Epstein and Schneider \(2003\)](#) extend the axiomatic approach of Gilboa and Schmeidler to dynamic settings. See [Maccheroni, Marinacci, and Rustichini \(2006\)](#) and [Strzalecki \(2011\)](#) for axiomatic foundations for the robust preferences of Hansen and Sargent and [Ilut and Schneider \(2014\)](#) for a business-cycle application with ambiguity-averse agents.

of models typically consists of a benchmark model and a surrounding set of models that are difficult to distinguish with finite data. Agents acknowledge the possibility of misspecification and are assumed to follow decision rules that perform well given any model in their set. As such, preference for robustness offers a substitute for high risk aversion in generating realistic risk premia.⁸ In contrast, agents in a CREE behave optimally assuming that the model that best describes their observations is the true model of the world. Ambiguity aversion may asymptotically vanish as agents learn—as shown by [Marinacci \(2002\)](#) and [Epstein and Schneider \(2007\)](#).⁹ But agents in a CREE do not settle on the correct model because of their misspecified prior, despite having access to an infinite sequence of observations. In short, while in the robustness literature the deviation from rational expectations captures a lack of data and a fear of misspecification, the deviation in a CREE is due to reality being more complex than what agents can perceive.

A close cousin of misspecification is lack of identification. A model is misspecified if it does not contain the true probability distribution; it is unidentified if it includes (wrong) probability distributions that are indistinguishable from the true probability distribution along the equilibrium path. [Fudenberg and Levine \(1993\)](#)'s self-confirming equilibrium is the leading way of thinking about the lack of identification in game theory and macroeconomics.¹⁰ While agents' models can be both misspecified and unidentified in this paper, my focus is model misspecification.

A number of behavioral biases such as sentiments ([Barberis, Shleifer, and Vishny, 1998](#)), analogy-based expectations ([Jehiel, 2005](#)), cursedness ([Eyster and Rabin, 2005](#)), and gambler's and hot-hand fallacies ([Rabin and Vayanos, 2010](#)) can also be viewed as particular cases of model misspecification. In Appendix C, I show how several examples of misspecification from the existing literature can be seen as special cases of the framework of this paper.

The idea behind CREE is related to the concept of rational inattention put forward by [Sims \(2003\)](#) and used to study business cycles by [Mankiw and Reis \(2002\)](#), [Maćkowiak and Wiederholt \(2009, 2015\)](#), [Alvarez, Lippi, and Paciello \(2015\)](#), and [Afrouzi \(2017\)](#), among others. Both CREE and rational inattention impose restrictions on the agents' expectations, and both notions propose selection criteria for choosing an expectation from the constrained set. The restriction in rational-inattention models is an information-flow constraint, whereas the selection criterion is often utility maximization or the minimization of discounted mean squared errors. This paper

⁸See [Hansen, Sargent, and Tallarini \(1999\)](#) for an application of robust preferences to the permanent income model and asset pricing. Other contributions to the robustness literature include [Anderson, Sargent, and Hansen \(2003\)](#), [Hansen and Sargent \(2005, 2007\)](#), [Hansen, Sargent, Turmuhambetova, and Williams \(2006\)](#), [Woodford \(2010\)](#), [Karantounias \(2013\)](#), [Bhandari, Borovicka, and Ho \(2017\)](#), [Hansen and Sargent \(2017\)](#), and [Hansen, Szoke, Han, and Sargent \(2018\)](#). See [Hansen and Sargent \(2008\)](#) for a thorough treatment.

⁹See also [Zimper and Ma \(2017\)](#) who show that ambiguity aversion may not vanish if the decision maker is "stubborn."

¹⁰See also [Fudenberg and Levine \(2009\)](#), who study the relationship between self-confirming beliefs and the Lucas critique, [Primiceri \(2006\)](#) for an explanation of the US inflation in the 1960s and 1970s using self-confirming equilibrium, and [Battigalli, Cerreia-Vioglio, Maccheroni, Marinacci, and Sargent \(2016\)](#), who propose a general framework for studying self-confirming policies in macroeconomics.

generalizes rational inattention by considering arbitrary constraints on the agents' set of models. It proposes the selection criterion of CREE as a micro-founded and tractable alternative to the selection criterion of rational inattention.¹¹

Finally, the applied contribution of the paper is closely connected to a literature that emphasizes how noisy information can be the source of inertia and hump-shaped dynamics in macroeconomic models (e.g., [Sims, 2003](#), [Woodford, 2003a](#), [Mankiw and Reis, 2002](#), [Nimark, 2008](#), [Angeletos and Huo, 2018](#)). Both in that literature and the medium-scale model studied in Section 6 of this paper, the agents solve a Kalman filtering problem with an unobservable state, and this is the source of persistence and hump-shaped dynamics. But while in that literature the measurement error is exogenously imposed in the form of noisy information, in my application it is endogenously chosen by the agents as a byproduct of their misspecification. Related is also a literature that uses incomplete information and/or ambiguity toward a theory of how confidence matters over the business cycle (e.g., [Angeletos, Collard, and Dellas, 2018](#), [Benhabib, Wang, and Wen, 2015](#), [Ilut and Saijo, 2016](#), [Bhandari, Borovicka, and Ho, 2017](#)).

Outline. The rest of the paper is organized as follows: Section 2 introduces the general framework. Section 3 introduces the notions of temporary equilibrium and CREE. Section 4 presents Bayesian and adaptive learning foundations for the solution concept of CREE. Section 5 studies a special case of the general framework in which agents entertain hidden factor models and provides conditions under which misspecification can lead to amplification or dampening of shocks and to history dependence. Section 6 proposes a business-cycle model with agents who have hidden factor models and studies the implications of misspecification for business-cycle fluctuations. Section 7 concludes. Appendix A presents examples of standard macro models and illustrates how they can be mapped to the general framework of Section 2. Appendix B provides additional mathematical details for the hidden factor model. Appendix C presents a list of models of bounded rationality from the literature and shows that they can be viewed as instances of misspecification. Appendix D presents additional details on the business-cycle model described in the body of the paper. Appendix E presents some mathematical definitions and the proofs of the main results. An Online Appendix contains the standard calculations for the business-cycle model as well as additional theoretical results.

2 General Framework

In this section, I present the general framework that is used in the remainder of the paper. The framework aims to strike a balance between generality and tractability. In Subsection 2.1, I

¹¹[Maćkowiak, Matejka, and Wiederholt \(2017\)](#) show that rational-inattention models can be represented as state-space models in which the variance of the “measurement error” is lower-bounded by a variable that depends on the agent’s amount of attention. In Appendix C, I use the result of [Maćkowiak et al. \(2017\)](#) to show that the rational-inattention restriction on expectations can be viewed as a special case of the general formulation of model misspecification in this paper.

describe the economic environment—consisting of preferences, general-equilibrium requirements, and the distribution of shocks. To simplify the notation and clarify the main contributions of the paper, I restrict my attention to representative-agent economies. In Appendix A, I give examples of workhorse macro models and show that they can be mapped to the framework of Subsection 2.1. I also illustrate how a heterogeneous-agent model can be cast as a special case of my baseline representative-agent framework by an appropriate redefinition of variables. Extensions to other heterogeneous-agent models are straightforward.

In Subsection 2.2, I formally define what I mean by a constrained set of models for the agent. The formulation is sufficiently general to allow for (almost) arbitrary constraints on the agent’s models. In Appendix C, I show how a wide range of deviations from rational expectations—including models of adaptive learning, sentiments, extrapolative expectations, gambler’s fallacy, and rational inattention—can be seen as particular examples of such constraints.

The following mathematical definitions are used throughout the paper: transition probability and invariant distribution. Given measurable spaces $(\Omega_1, \mathcal{F}_1)$ and $(\Omega_2, \mathcal{F}_2)$, function $K : \Omega_1 \times \mathcal{F}_2 \rightarrow [0, 1]$ is a *transition probability* from Ω_1 to Ω_2 if (i) the mapping $\omega_1 \mapsto K(B|\omega_1)$ is measurable for any set $B \in \mathcal{F}_2$ and (ii) $K(\cdot|\omega_1)$ is a probability distribution over $(\Omega_2, \mathcal{F}_2)$ for any $\omega_1 \in \Omega_1$. A probability distribution ρ over Ω is an *invariant distribution* for the transition probability K from Ω to itself if $\int K(B|\omega)\rho(d\omega) = \rho(B)$ for any measurable set $B \subseteq \Omega$.

2.1 Economic Environment

The representative agent chooses a sequence $\{x_t\}_{t=1}^\infty$ of *choice variables* to maximize the expected present-discounted value of her payoffs:

$$\max_{\{x_t\}_{t=1}^\infty} E \left[\sum_{t=1}^\infty \beta^{t-1} u(x_{t-1}, x_t, y_t) \right], \quad (1)$$

given x_0 and subject to constraints

$$x_t \in \Gamma(x_{t-1}, y_t) \quad t = 1, 2, \dots \quad (2)$$

where $\beta \in (0, 1)$ is the discount factor, u is the per-period utility function, and Γ is a correspondence that determines the agent’s choice set as a function of the choice x_{t-1} made by the agent in the previous period and an *observable* y_t . The observable y_t is a vector consisting of all the variables that are observed by the agent, can affect her optimal choices by shifting her payoffs or changing her choice set, and are outside of the agent’s control.¹²

Problem (1)–(2) would be identical to the canonical dynamic decision problem studied in Chapter 9 of Stokey and Lucas (1989) but for two important differences. First, the operator E

¹²The assumption that the agent’s choices do not affect the values of the observable variable implies that the agent has no incentive to engage in active learning or experimentation. This assumption distinguishes the paper from the literature on misspecified learning in games such as Fudenberg, Romanyuk, and Strack (2017).

represents the agent's *subjective* expectation about the path $\{y_t\}_{t=1}^{\infty}$ of the observables. It can in general be different than the model-implied (rational) expectation. The modeling and analysis of the evolution of the agent's subjective expectation is the focus of the theory part of this paper. The second difference concerns the determination of the observable y_t . I deviate from the textbook dynamic decision problem by allowing y_t to have endogenous elements (such as prices) that are determined in equilibrium and are required to be consistent with the agent's choices.

These consistency requirements are expressed parsimoniously by means of a function G of x and y that is required to be identically equal to zero in equilibrium. In any equilibrium and at all times, x_t and y_t are required to satisfy

$$G(x_t, y_t, z_t) \equiv 0, \quad (3)$$

where z_t is a *state variable* that is predetermined as of period t . The state variable in period t is distributed according to

$$z_t \sim \Pi(\cdot | y_{t-1}, z_{t-1}), \quad (4)$$

conditional on the values of the observable and the state variable in period $t - 1$. The state variable z_t may have exogenous and endogenous components that represent variables such as exogenous shocks or the capital stock of the economy. I assume that z_t is not observable to the agent. But note that this assumption is without loss of generality: any observable state variable can be included in the vector y_t .

The function G can be used to express two distinct types of consistency requirements. First, it can be used to express general equilibrium conditions such as market-clearing conditions. Second, it can be used to make the value of the observable y_t a function of the predetermined state variable. The latter consistency requirements can be used to model persistence in the value of (exogenous or endogenous) observable variables.

In Appendix A, I provide examples of textbook macroeconomic models and discuss how they can be mapped to the general framework of this section by appropriate choices of x , y , z , Π , Γ , u , β , and G .

2.2 The Agent's Models

I endow the representative agent with a set of models for the economy. A model consists of a probability distribution over the infinite sequence $\{y_t\}_{t=0}^{\infty}$ of observables. The agent's set of models may be *correctly specified*—containing the data-generating process that arises in equilibrium—or it may be *misspecified*.

I impose two weak assumptions on the set of models for the agent. First, I assume that the agent only entertains models according to which the sequence $\{y_t\}_{t=0}^{\infty}$ is a *Markov chain* (possibly over an infinite-dimensional space). Given the flexibility in choosing the domain of

y , this assumption is without any serious loss of generality.¹³ I also simplify the analysis by assuming that the agent's model is *parametric*, i.e., can be parametrized by a *finite* number of real-valued variables. This is in keeping with the overwhelming majority of the papers in economic theory and econometrics that study the implications of model misspecification.¹⁴ These two restrictions are sufficiently strong to enable me to prove general results while being sufficiently permissive to allow the framework to embed many interesting cases that arise in applications.

A set of parametric Markovian models for the agent is fully described by a pair (Θ, Q) , where Θ is a subset of a finite-dimensional Euclidean space and Q is a function that maps parameters to transition probabilities. Given parameter $\theta \in \Theta$, the agent believes that y_t is distributed according to

$$y_t \sim Q_\theta(\cdot | y_{t-1}), \quad (5)$$

conditional on y_{t-1} being the realization of the observable in period $t - 1$. Given any probability distribution μ_{y_0} for y_0 , the transition probability Q_θ induces a probability distribution over the space of sequences $\{y_t\}_{t=0}^\infty$. Whenever the mapping Q is clear from the context, I refer to θ as a *model* and to the set Θ as the agent's constrained *set of models* (or simply the agent's models).

In the next two sections, I remain agnostic about the nature of the constraints on the agent's set of models. Doing so allows me to establish general theoretical results that are valid given arbitrary specifications of the agent's models. Appendix C illustrates how a number of models of bounded rationality can be cast as particular specifications of the set Θ . In Sections 5 and 6, I specialize the agent's set of models to be induced by hidden factor models.

An economy is defined as an economic environment together with a set of parametric Markovian models for the agent.

Definition 1. An *economy* is a tuple $\mathfrak{E} = (X, Y, Z, \Pi, \Gamma, u, \beta, G, \Theta, Q)$, with X, Y , and Z denoting the domains of the choice variable x , the observable y , and the state variable z . An *initial condition* for the economy \mathfrak{E} consists of initial values (x_0, y_0, z_0) for the choice variable, the observable, and the state.

2.3 Technical Assumptions

I next state the weak technical assumptions that are maintained throughout the paper. More substantive assumptions are explicitly stated in the paper as numbered assumptions and are invoked in the results that use them. The sets X, Y , and Z are nonempty compact subsets of metric spaces with the corresponding Borel sigma-algebras denoted by \mathcal{X}, \mathcal{Y} , and \mathcal{Z} . The

¹³In particular, any probability distribution over $\Omega^\mathbb{N}$, where Ω denotes a Borel space, can be expressed as a Markov chain with state space $\Omega^\mathbb{N}$.

¹⁴Examples include White (1982), Vuong (1989), Esponda and Pouzo (2016a,b), and Fudenberg, Romanyuk, and Strack (2017). For a paper in statistics that deals with the complexities arising in Bayesian estimation of nonparametric misspecified models, see Kleijn and Van Der Vaart (2006).

function $u : X \times Y \times Y \rightarrow \mathbb{R}$ and the correspondence $\Gamma : X \times Y \rightrightarrows X$ are measurable. The function $\Pi : Y \times Z \times Z \rightarrow [0, 1]$ is a transition probability from $Y \times Z$ to Z . The function $G(x, y, z)$ takes values in \mathbb{R}^k , with k denoting the number of general-equilibrium consistency requirements. The set Θ is a nonempty compact subset of a finite-dimensional Euclidean space. The set $\Delta\Theta$ of probability distributions over Θ is endowed with the topology of weak convergence and the corresponding Borel sigma-algebra $\mathcal{B}(\Delta\Theta)$. For any $\theta \in \Theta$, the function $Q_\theta : Y \times \mathcal{Y} \rightarrow [0, 1]$ is a transition probability from Y to itself. For any $y_- \in Y$, there exists a sigma-finite measure ζ on Y such that, for all $\theta \in \Theta$, the probability distribution $Q_\theta(\cdot|y_-)$ is absolutely continuous with respect to ζ with the density denoted by $q_\theta(\cdot|y_-)$. The family of functions $\{\theta \mapsto q_\theta(y|y_-)\}_{y_-}$ is equicontinuous for any $y \in Y$, and the family of functions $\{q_\theta(\cdot|\cdot)\}_{\theta \in \Theta}$ is equicontinuous. The density $q_\theta(\cdot|\cdot)$ is strictly positive and bounded for all $\theta \in \Theta$.

3 Equilibrium

This section introduces the solution concept of constrained-rational-expectations equilibrium. As a stepping stone, I fix the agent's expectations and define a recursive temporary equilibrium.¹⁵

3.1 Recursive Temporary Equilibrium

It is more convenient to work with the recursive representation of problem (1)–(2). I assume that the principle of optimality holds so that the recursive and sequential problems are equivalent.¹⁶ Given model θ for the agent, her Bellman equation is given by

$$V(x_-, y, \theta) = \max_{x \in \Gamma(x_-, y)} \left[u(x_-, x, y) + \beta \int V(x, y', \theta) Q_\theta(dy'|y) \right]. \quad (6)$$

This is a standard Bellman equation—except for the fact that the expectation is taken with respect to the agent's subjective model for the transition probability of y . Let $\mathbf{x}(x_-, y, \theta)$ denote the policy correspondence

$$\mathbf{x}(x_-, y, \theta) \equiv \arg \max_{x \in \Gamma(x_-, y)} \left[u(x_-, x, y) + \beta \int V^*(x, y', \theta) Q_\theta(dy'|y) \right], \quad (7)$$

where $V^*(x_-, y, \theta)$ is the unique fixed point of the functional equation (6).

A recursive temporary equilibrium is defined as a transition probability that is consistent with the optimality condition (7), the general-equilibrium requirement (3), and the transition probability for the state variable given in (4).

¹⁵The idea of temporary equilibrium goes back to the writings of Hicks (1939) and Lindahl (1939). It has been extensively developed in the context of Arrow–Debreu economies by Grandmont (1977, 1982). See Woodford (2013) for a discussion of temporary equilibria in the context of modern monetary models and Farhi and Werning (2017) for an application in the context of a heterogeneous-agent new-Keynesian economy. The notion of recursive temporary equilibrium, defined in this paper, is tightly linked to the notion of the “Markov process of temporary equilibrium” introduced by Hildenbrand and Grandmont (1974). In particular, a recursive temporary equilibrium generates a random process for equilibrium variables that is a Markov process of temporary equilibrium.

¹⁶See Stokey and Lucas (1989, ch. 9) for conditions on the fundamentals that guarantee the principle of optimality to hold.

Definition 2. Given an economy \mathfrak{E} , a *recursive temporary equilibrium* is a transition probability T from $X \times Y \times Z \times \Theta$ to $X \times Y \times Z$ such that for any x_-, y_-, z_-, θ ,

(i) the agent's decisions are optimal:

$$x \in \mathbf{x}(x_-, y, \theta) \quad \text{for } T\text{-almost all } (x, y);$$

(ii) the general-equilibrium requirements are satisfied:

$$G(x, y, z) = 0 \quad \text{for } T\text{-almost all } (x, y, z);$$

(iii) and the distribution of the state variable is consistent with the given transition probability:

$$T(B|x_-, y_-, z_-, \theta) = \Pi(B|y_-, z_-) \quad \text{for any measurable set } B \subseteq Z.$$

It will prove useful to introduce a straightforward generalization of recursive temporary equilibrium that allows for the agent to assign positive probabilities to more than one model. When the agent's belief is non-degenerate, one has to take a stand on what the agent believes about the evolution of her own future expectations. I assume that the agent anticipates updating her belief using Bayes' rule. Due to the Markovian structure of Bayes' rule, the optimization problem of the agent can still be written as a Bellman equation. But now I also must keep track of the belief as a state variable (since it changes over time). Suppose the agent's belief in the current period is given by $\lambda \in \Delta\Theta$. The agent's optimization problem can be represented as the following Bellman equation:

$$V(x_-, y, \lambda) = \max_{x \in \Gamma(x_-, y)} \left[u(x_-, x, y) + \beta \int V(x, y', \lambda') Q_\theta(dy'|y) \lambda(d\theta) \right], \quad (8)$$

where $\lambda' = \phi(\lambda, y, y')$ is the Bayesian update of λ conditional on observing the transition from y to y' ; more formally, the mapping $\phi : (\lambda, y, y') \mapsto \lambda'$ is defined by letting, for any measurable set $B \subseteq \Delta\Theta$,

$$\lambda'(B) = \frac{\int_B q_\theta(y'|y) \lambda(d\theta)}{\int q_\theta(y'|y) \lambda(d\theta)}. \quad (9)$$

The continuity and full-support assumptions made on q guarantee that the Bayesian update $\phi(\lambda, y, y')$ is always well-defined and finite.¹⁷ I let $\mathbf{x}(x_-, y, \lambda)$ denote the policy correspondence given the Bellman equation (8). The following definition generalizes Definition 2 to accommodate non-degenerate beliefs and Bayesian updating:

Definition 2'. Given an economy \mathfrak{E} , a *Bayesian recursive temporary equilibrium* is a transition probability T from $X \times Y \times Z \times \Delta\Theta$ to itself such that for any x_-, y_-, z_-, λ_- ,

¹⁷See Lemma E.1 for a proof.

(i) the agent's decisions are optimal:

$$x \in \mathbf{x}(x_-, y, \lambda) \quad \text{for } T\text{-almost all } (x, y, \lambda);$$

(ii) the general-equilibrium requirements are satisfied:

$$G(x, y, z) = 0 \quad \text{for } T\text{-almost all } (x, y, z);$$

(iii) the distribution of the state variable is consistent with the given transition probability:

$$T(B|x_-, y_-, z_-, \lambda_-) = \Pi(B|y_-, z_-) \quad \text{for any measurable set } B \subseteq Z;$$

(iv) and the agent updates her belief using Bayes' rule:

$$\lambda = \phi(\lambda_-, y_-, y) \quad \text{for } T\text{-almost all } (y, \lambda).$$

For brevity, I refer to a Bayesian recursive temporary equilibrium simply as a *temporary equilibrium* whenever there is no risk of confusion.

A Bayesian recursive temporary equilibrium maps the value of equilibrium variables (x_-, y_-, z_-) and the agent's belief in the previous period λ_- to the distribution of equilibrium variables and the agent's belief in the current period. The distribution of equilibrium variables is consistent with the optimality and consistency requirements, and the agent's belief is updated by Bayes' rule.

Definition 2' reduces to Definition 2 when the prior λ_- is given by a degenerate distribution $\mathbf{1}_\theta$ for some model $\theta \in \Theta$. This is a simple consequence of the fact that the Bayesian update $\phi(\mathbf{1}_\theta, y, y')$ of a degenerate prior $\mathbf{1}_\theta$ is equal to the prior itself, regardless of the values of y and y' .

A temporary equilibrium is a mapping according to which the agent maximizes her utility given her expectations and the consistency requirements are satisfied. Existence of temporary equilibrium is, therefore, a necessary condition for the existence of equilibrium. I maintain the following assumption throughout the paper:

Assumption 0. There exists a temporary equilibrium T for the economy \mathfrak{E} .

I also assume that the temporary equilibrium is continuous to ensure that the economy is well-behaved. The following continuity assumption requires the mapping T to be continuous:

Assumption 1. The mapping $(\lambda_-, x_-, y_-, z_-) \mapsto T(\cdot|x_-, y_-, z_-, \lambda_-)$ is continuous, where the topology on $X \times Y \times Z \times \Delta\Theta$ is the product topology and the topology on $\Delta(X \times Y \times Z \times \Delta\Theta)$ is induced by the total variation norm.¹⁸

¹⁸More formally, $\Delta(X \times Y \times Z \times \Delta\Theta)$ is taken as a subset of $ca(X \times Y \times Z \times \Delta\Theta)$, the Banach space of countably additive signed measures over $X \times Y \times Z \times \Delta\Theta$ with the total variation norm. Assumption 1 is the counterpart of the continuity assumption made in Proposition 5.6 of [Futia \(1982\)](#).

The temporary equilibrium summarizes the fundamentals of the economy other than the agent's set of models.¹⁹ I take it as the starting point of my analysis in the rest of this section and in Section 4.

3.2 Constrained-Rational-Expectations Equilibrium

The description of equilibrium is complete once a temporary equilibrium is supplemented with a procedure for the determination of the agent's expectation. In a rational-expectations equilibrium, this is done by assuming that the agent's expectation coincides with the equilibrium distribution. I relax the rational-expectations assumption by requiring the agent's expectation to be the probability distribution among the ones allowed by her constrained set of models that comes closest to the distribution of observables. The notion of distance I use is a generalized version of the Kullback–Leibler divergence.²⁰ In Section 4, I provide two learning foundations for this choice.

Definition 3. Given a temporary equilibrium T and a probability distribution μ over $X \times Y \times Z \times \Delta\Theta$, the *Kullback–Leibler divergence* of model θ from T is defined as

$$H(Q_\theta, T, \mu) \equiv - \int \int \log(q_\theta(y|y_-)) T(dy|x_-, y_-, z_-, \lambda_-) \mu(dx_- \times dy_- \times dz_- \times d\lambda_-). \quad (10)$$

$H(Q_\theta, T, \mu)$ is a measure of the distance between the transition probability T and the transition probability Q_θ . The inner integral in (10) is a measure of the distance between distributions $Q_\theta(\cdot|y_-)$ and $T(\cdot|x_-, y_-, z_-, \lambda_-)$ over Y , taking as given the past value $(x_-, y_-, z_-, \lambda_-)$. The outer integral takes an average of the distance between $Q_\theta(\cdot|y_-)$ and $T(\cdot|x_-, y_-, z_-, \lambda_-)$ by integrating over $(x_-, y_-, z_-, \lambda_-)$ according to the probability measure μ . That $H(Q_\theta, T, \mu)$ is a measure of the distance between transition probabilities is best seen by considering the case where Θ contains a parameter θ^* for which $Q_{\theta^*}(\cdot|y_-) = T(\cdot|x_-, y_-, z_-, \lambda_-)$ almost everywhere and for μ -almost all $(x_-, y_-, z_-, \lambda_-)$. The parameter θ^* is then a minimizer of the Kullback–Leibler divergence. When Θ does not contain such a θ^* parameter, the Kullback–Leibler divergence is minimized by parameters θ for which Q_θ is the most similar to Q_{θ^*} . The following lemma establishes some basic properties of this measure of distance:

Lemma 1. $H(Q_\theta, T, \mu)$ is finite and well-defined. The mapping $\theta \mapsto H(Q_\theta, T, \mu)$ is continuous.

I can now define the solution concept of constrained-rational-expectations equilibrium.

¹⁹To be more precise, the temporary equilibrium *does* depend on the mapping Q but it does *not* depend on the constrained set of models Θ .

²⁰The Kullback–Leibler divergence is commonly used in statistics as a measure of the fit of a model. The Kullback–Leibler divergence of Q_θ from the true probability distribution P is often normalized by adding a constant. This is to ensure that the divergence is equal to zero when $Q_\theta = P$. I define the Kullback–Leibler divergence without the normalization. This ensures that the Kullback–Leibler divergence is well-defined and finite under weaker assumptions without affecting the set of models that minimize the Kullback–Leibler divergence. The non-normalized version of the Kullback–Leibler divergence is sometimes referred to as *cross entropy*. For more on the Kullback–Leibler divergence, cross entropy, and related topics, see Gray (2011).

Definition 4. Given an economy \mathfrak{E} and a temporary equilibrium T for \mathfrak{E} , a *constrained-rational-expectations equilibrium (CREE)* consists of a triple (T, μ^*, Θ^*) , where μ^* is a probability distribution over $X \times Y \times Z \times \Delta\Theta$ and $\Theta^* \subseteq \Theta$ is a closed set, such that

- (i) μ^* is an invariant distribution for T ;
- (ii) $\Theta^* = \arg \min_{\theta \in \Theta} H(Q_\theta, T, \mu^*)$;
- (iii) and μ^* is supported on $X \times Y \times Z \times \Delta\Theta^*$.

When Θ^* is a singleton, I refer to the equilibrium as a *pure CREE*.

The objects that appear in the definition of equilibrium have intuitive interpretations. The transition probability T describes the probability of transitioning from $(x_{t-1}, y_{t-1}, z_{t-1}, \lambda_{t-1})$ to $(x_t, y_t, z_t, \lambda_t)$ at any time t . The probability distribution μ^* describes the unconditional distribution of the tuple $(x_t, y_t, z_t, \lambda_t)$ of the choice, observable, and state variables and the agent's belief at any point in time. The set Θ^* is the set of models that the agent finds plausible.

The CREE describes a stationary environment in which the agent only considers a subset of models plausible and equilibrium variables follow a stationary process. The agent's belief is concentrated on the set of models that best describe her observations in the sense of minimizing the Kullback–Leibler divergence. And the distribution of equilibrium variables is consistent with the agent's belief, utility maximization by the agent, and the general equilibrium consistency requirements.

A CREE can be seen as the limit point in which the agent has learned all that can be learned from her observations, and the economy no longer experiences transitional dynamics originating from learning. In Section 4, I formalize this intuition by establishing that CREE arises in the limit as the agent learns about her environment.

To get more intuition for the solution concept of CREE, it is instructive to consider a world in which the agent's set of models is unconstrained; that is, any transition probability Q from Y to itself is a model considered possible by the agent. The minimizer of the Kullback–Leibler divergence is then the transition probability Q^* that agrees with the temporary equilibrium transition probability T almost everywhere and for μ^* -almost all $(x_-, y_-, z_-, \lambda_-)$. But this is exactly the requirement for a (recursive) rational-expectations equilibrium. CREE extends the definition of a rational-expectations equilibrium by allowing for the agent's model to rule out any such Q^* . The CREE selection then calls for the agent to choose a transition probability Q_θ that is closest to the equilibrium transition probability.

The solution concept of CREE is related to Berk–Nash equilibrium, proposed by [Esponda and Pouzo \(2016a\)](#) in the context of static games and extended by [Esponda and Pouzo \(2016b\)](#) to Markov decision problems. Both solution concepts require the agent's belief to be supported on the set of models that minimize a version of the Kullback–Leibler divergence from the equilibrium distribution of observables. But there are two major differences between these

solution concepts. First, while both CREE and Berk–Nash equilibrium define the Kullback–Leibler divergence as the weighted expected value of log-likelihoods, the weighting is different between the two solution concepts. It is by the endogenous distribution of actions in [Esponda and Pouzo \(2016a\)](#) and by an exogenous transition probability and an endogenous distribution in [Esponda and Pouzo \(2016b\)](#). In this paper by contrast, the log-likelihoods are weighted by the temporary equilibrium of the economy. These different choices reflect differences in the nature of the economic environment, but they all have foundations in Bayesian learning. Second, while each agent has a unique belief in a Berk–Nash equilibrium, the agent’s belief in a CREE is drawn from an ergodic distribution at any point in time. This relaxation enables me to establish the convergence of Bayesian equilibrium to a CREE under two weak regularity assumptions. Pure CREE comes closer to Berk–Nash equilibrium by restricting the agent’s belief to belong to a singleton.

The following theorem shows that Definition 4 is not vacuous by establishing that a CREE exists in any well-behaved economy.

Theorem 1. *If the economy \mathcal{E} has a temporary equilibrium T that satisfies Assumption 1, then a CREE exists.*

A CREE always exists if the economy has a continuous temporary equilibrium. The existence of a temporary equilibrium is clearly necessary for the existence of a CREE. Continuity of the temporary equilibrium, on the other hand, is a weak requirement that ensures that the mapping whose fixed point defines a CREE is well defined. A pure CREE exists if the agent’s model additionally satisfies a convexity requirement. The statement and proof of the result on the existence of a pure CREE has been relegated to the Online Appendix.

4 Learning Foundations of CREE

In this section, I provide two learning foundations for the solution concept of CREE. I first provide conditions under which Bayesian learning leads the economy to converge to a CREE. I then argue that adaptive learning also leads the agent’s belief to concentrate on the set of models that best describe her observations in the sense of minimizing the Kullback–Leibler divergence from the equilibrium distribution of observables. These results provide learning foundations for the solution concept of CREE and the choice of Kullback–Leibler divergence as the measure of the fit of a model. They also establish a formal connection between the Bayesian and adaptive approaches to learning. Readers interested in applications of CREE may directly skip to Section 5.

4.1 Bayesian Learning

A natural way of modeling learning is by assuming that the agent is Bayesian. A Bayesian agent has a subjective prior $\lambda_0 \in \Delta\Theta$ over her set of models and updates her belief over time using Bayes’

rule. Recall that a temporary equilibrium characterizes the evolution of equilibrium variables and the agent's belief. It readily leads to the following equilibrium definition

Definition 5. Given an economy \mathfrak{E} and a temporary equilibrium T for \mathfrak{E} , a *Bayesian equilibrium* with initial conditions (x_0, y_0, z_0) for the economy and prior λ_0 for the agent consists of a sequence of adapted random variables $\{x_t, y_t, z_t, \lambda_t\}_{t=0}^{\infty}$ with probability distribution \mathbb{P} such that

$$\mathbb{P} \left((x_t, y_t, z_t, \lambda_t) \in B \mid \{x_s, y_s, z_s, \lambda_s\}_{s=0}^{t-1} \right) = T(B \mid x_{t-1}, y_{t-1}, z_{t-1}, \lambda_{t-1})$$

for any measurable set B and all $t \geq 1$.

The definition of Bayesian equilibrium is quite natural. The agent behaves optimally, the consistency conditions are satisfied, the distribution of the state is consistent with the given transition probability, and the agent has a prior belief λ_0 over her set of models and updates her belief using Bayes' rule. The existence of a Bayesian equilibrium is a trivial corollary of the assumption that a temporary equilibrium exists. Moreover, since a Bayesian equilibrium is the Markov chain with transition probability T and initial state $(x_0, y_0, z_0, \lambda_0)$, it is unique almost surely.

Proposition 1. *Consider an economy \mathfrak{E} and a temporary equilibrium for \mathfrak{E} . For any initial condition (x_0, y_0, z_0) for the economy and any prior λ_0 for the agent, there exists a Bayesian equilibrium; the Bayesian equilibrium is unique up to sets of measure zero.*

Bayesian equilibrium is closely related to the internally-rational-expectations equilibrium (IREE) of [Adam and Marcet \(2011\)](#). Just as in [Adam and Marcet \(2011\)](#), I assume that agents have an internally-consistent belief systems about the path of y_t , choose their actions optimally given their beliefs, and update their beliefs using Bayes' rule. But I additionally assume that agents can only entertain Markovian and parametric models, and I focus on equilibria that have a recursive structure.²¹ The restriction to parametric and recursive models provides enough structure to prove general learning results for Bayesian equilibrium.

Despite being internally consistent and well-grounded in the Bayesian paradigm, the solution concept of Bayesian equilibrium has two practical drawbacks. First, the modeler has to keep track of the agent's belief as a state variable. This results in a large increase in the size of the state space. Second, the agent's prior over the set of models is a free parameter. This would be a minor nuisance if the equilibrium outcomes were not too sensitive to the choice of the prior. But, as [Adam and Marcet \(2011\)](#) argue, even small changes in the agents' priors can lead to large changes in the *short-run* economic outcomes.

The solution concept of CREE addresses these drawbacks while providing a good approximation to the long-run behavior of economy under Bayesian learning. As the agent updates

²¹Although IREE is defined without any reference to a parametric model, parametric models are often used in applications of the IREE. In particular, [Adam and Marcet \(2011\)](#) use a parametric model in their application of the IREE solution concept to asset pricing. Other applied papers that follow the IREE approach, e.g., [Adam, Kuang, and Marcet \(2012\)](#), [Adam, Marcet, and Nicolini \(2016\)](#), [Adam, Marcet, and Beutel \(2017\)](#), and [Gerko \(2017\)](#) also all use parametric models.

her belief about the model of the economy, the effect of the prior vanishes and the agent's belief concentrates on the set of models with minimal Kullback–Leibler divergence from the equilibrium distribution. When that set is a singleton (i.e., in a pure CREE), one can additionally drop the belief from the description of a CREE. The learning result presented next provides sufficient conditions for a Bayesian equilibrium to converge to a CREE.

The observation that enables a sharp learning result under relatively weak assumptions is that a Bayesian equilibrium is Markovian: a Bayesian equilibrium with initial condition (x_0, y_0, z_0) for the economy and prior λ_0 for the agent is the Markov chain over $X \times Y \times Z \times \Delta\Theta$ with initial value $(x_0, y_0, z_0, \lambda_0)$ and transition probability T . This is the consequence of the Markovian structure of the economy and the fact that Bayes' rule is Markovian.

I impose two weak regularity conditions on the transition probability T . The first condition is the continuity condition stated in Assumption 1. The second condition is a condition called asymptotic mean stationarity. Given a Bayesian equilibrium $\{x_t, y_t, z_t, \lambda_t\}_{t=0}^{\infty}$ with probability distribution \mathbb{P} , let μ_t denote the probability distribution over $X \times Y \times Z \times \Delta\Theta$ defined as

$$\mu_t(B) \equiv \mathbb{P}(\{(x_t, y_t, z_t, \lambda_t) \in B\}), \quad (11)$$

and define $\bar{\mu}_t$ as

$$\bar{\mu}_t(B) \equiv \frac{1}{t} \sum_{s=0}^{t-1} \mu_s(B), \quad (12)$$

where B denotes an arbitrary measurable subset of $X \times Y \times Z \times \Delta\Theta$. The probability distribution $\bar{\mu}_t$ is the time average of the distribution of equilibrium variables up to time t . Asymptotic mean stationarity requires this time average to stabilize.

Definition 6. A temporary equilibrium T for economy \mathfrak{E} is *asymptotically mean stationary* if for any initial condition (x_0, y_0, z_0) for the economy and any prior λ_0 for the agent, $\bar{\mu}_t$ converges weakly to some probability distribution $\bar{\mu}$ over $X \times Y \times Z \times \Delta\Theta$ and the mapping $(x_0, y_0, z_0, \lambda_0) \mapsto \bar{\mu}$ is continuous.

Asymptotic mean stationarity is a much weaker requirement than stationarity. The latter requires the probability distribution μ_t to be independent of time t . For generic initial conditions (x_0, y_0, z_0) for the economy and priors λ_0 for the agent, the Bayesian equilibrium will not be stationary. Asymptotic mean stationarity, on the other hand, only requires the time average of the sequence of probability distributions $\{\mu_t\}_t$ to eventually stabilize. It is consistent with a wide range of stochastic processes including those with transitional dynamics and deterministic and stochastic cycles—as long as the cycles themselves occur with some statistical regularity. It rules out cycles whose frequencies slowly decrease over time.

Whether asymptotic mean stationarity is satisfied only depends on the fundamentals of the economy as summarized by the temporary equilibrium T . In particular, it is equivalent to the requirement that the transition probability T satisfies a condition known as uniform mean

stability—which is automatically satisfied for transition probabilities on finite state spaces.²² The transition probability T does not need to satisfy any stronger conditions (such as aperiodicity, irreducibility, or ergodicity).

The main result of this subsection establishes that continuity and asymptotic mean stationarity are sufficient for a Bayesian equilibrium to converge to a CREE:

Theorem 2. *Consider an economy \mathfrak{E} and a temporary equilibrium T that satisfies Assumption 1 and is asymptotically mean stationary, let $\{x_t, y_t, z_t, \lambda_t\}_{t=0}^\infty$ be a Bayesian equilibrium with probability distribution \mathbb{P} , and let Θ_0 denote the support of the agent’s prior λ_0 . With \mathbb{P} -probability one, there exists a probability distribution μ^* over $X \times Y \times Z \times \Delta\Theta$ and a closed set $\Theta^* \subseteq \Theta_0$ such that*

(a) *for any sets $U_1 \subset K \subset U_2 \subseteq X \times Y \times Z \times \Delta\Theta$ such that U_1, U_2 are open and K is closed,*

$$\mu^*(U_1) \leq \liminf_{t \rightarrow \infty} \frac{1}{t} \sum_{s=0}^{t-1} \mathbb{1}\{(x_s, y_s, z_s, \lambda_s) \in K\} \leq \limsup_{t \rightarrow \infty} \frac{1}{t} \sum_{s=0}^{t-1} \mathbb{1}\{(x_s, y_s, z_s, \lambda_s) \in K\} \leq \mu^*(U_2);$$

(b) *for any open neighborhood U of Θ^* ,*

$$\lim_{t \rightarrow \infty} \lambda_t(U) = 1;$$

(c) *and the triple (T, μ^*, Θ^*) constitutes a CREE given the set of models Θ_0 for the agent.*

The theorem provides a learning foundation for the choice of Kullback–Leibler divergence as the notion of distance in the definition of a CREE. Part (a) shows that the empirical distribution of equilibrium variables eventually stabilizes with μ^* denoting the long-run empirical distribution. Parts (b) and (c) establish that the agent’s belief asymptotically concentrates on the set of models Θ^* that minimize the Kullback–Leibler divergence from the equilibrium distribution. Taken together, the three parts of the theorem establish that the long-run behavior of equilibrium variables in a Bayesian equilibrium are indistinguishable from the behavior observed in a CREE.

The theorem is related to the results of [Esponda and Pouzo \(2016a,b\)](#) on Bayesian learning foundations of Berk–Nash equilibrium. But while my result, Theorem 2 of [Esponda and Pouzo \(2016a\)](#), and Theorem 2 of [Esponda and Pouzo \(2016b\)](#) all provide learning foundations for the choice of Kullback–Leibler divergence as the notion of distance, the results are different in their assumptions. [Esponda and Pouzo \(2016a\)](#) establish that any stable (i.e., convergent) strategy profile is a Berk–Nash equilibrium, whereas [Esponda and Pouzo \(2016b\)](#) establish that *if* the strategy profile and the empirical distribution of the state-choice pair are both convergent, then the limit constitutes a Berk–Nash equilibrium. In contrast, I prove the convergence of the empirical distribution of equilibrium variables under the assumption of asymptotic

²²For a Markov chain over a finite-state space with transition matrix N and initial distribution π_0 , the sum $\frac{1}{t} \pi_0' \sum_{s=0}^{t-1} N^s$ converges to a limit, with the limit a continuous function of π_0 . But this conclusion does not necessarily hold for transition probabilities over infinite-dimensional spaces. Uniform mean stability is the assumption that it does. See [Jamison \(1965\)](#) or [Futia \(1982\)](#) for a discussion of uniform mean stability and related topics.

mean stationarity. Stability is an assumption on the convergence of random variables, while asymptotic mean stationarity is an assumption on the probability distributions from which those random variables are drawn.

Theorem 2 is also different from the results of [Esponda and Pouzo \(2016a,b\)](#) in its conclusion. It establishes the convergence of empirical distributions under asymptotic mean stationarity—a conclusion that has no counterpart in [Esponda and Pouzo \(2016a,b\)](#). The results are also different because of differences in the definition of CREE and Berk–Nash equilibrium. While CREE allows for the agent’s belief to fluctuate among the set of models that minimize the Kullback–Leibler divergence, the agent’s belief is constant in a Berk–Nash equilibrium. Consequently, the notion of convergence in Theorem 2 is weaker than that of Esponda and Pouzo: my result proves the convergence of the distribution of beliefs to an ergodic distribution. In the Online Appendix, I provide identifiability and convexity assumptions under which the agent’s belief converges to a degenerate distribution over some model θ^* and the Bayesian equilibrium converges to a pure CREE.

The proof relies on the observation that a Bayesian equilibrium can be represented as a Markov chain and a law-of-large-numbers (LLN) result for Markov chains due to [Breiman \(1960\)](#) and [Jamison \(1965\)](#). The result of [Jamison \(1965\)](#) requires the transition probability of a Markov chain to be continuous and uniformly mean stable. Continuity follows Assumption 1, while uniform mean stability is a consequence of asymptotic mean stationarity. The LLN of [Jamison \(1965\)](#) then implies that the time average of $f(x_t, y_t, z_t, \lambda_t)$ converges for any continuous function f . The first part of the theorem is proved by using Urysohn’s lemma to bound the indicator function by continuous functions and applying the LLN of Jamison. The second part follows a standard argument from Bayesian statistics. In particular, I extend [Bunke and Milhaud \(1998\)](#)’s proof of the concentration of measure on the set of minimal Kullback–Leibler divergence for i.i.d. observations to the case where observations follow a Markov chain. I use the LLN established previously to show that expected distance of θ at time t from the set of minimal Kullback–Leibler divergence converges to zero as t goes to infinity, where the expectation is taken with respect to the agent’s belief at t . This conclusion immediately implies the concentration-of-measure result in the second part of the theorem. The detailed proof is presented in Appendix E.

That the agent’s belief asymptotically concentrates on a set is not sufficient to guarantee the convergence of her belief. Rather, without additional assumptions, it is generally not possible to rule out the agent’s belief fluctuating among a set of models that are not identified given the agent’s observations.²³ In the Online Appendix, I show that under appropriate convexity and identifiability assumptions, the agent’s belief converges to a point mass and the Bayesian equilibrium converges to a pure CREE.

²³For an early example of the failure of convergence of the Bayesian posterior when the parameter is not identifiable, see [Diaconis and Freedman \(1986\)](#). For an example of the cycling of beliefs under misspecified learning in an economic context, see [Nyarko \(1991\)](#).

4.2 Adaptive Learning

An alternative to the Bayesian approach to learning is the adaptive-learning approach of [Marcet and Sargent \(1989a\)](#) and [Evans and Honkapohja \(2012\)](#). This approach assumes that agents estimate their models using a frequentist procedure (almost always ordinary least-squares estimation). In this subsection, I use the approach of adaptive learning in the context of the general framework of Section 2 and show that adaptive learning also leads the agent's belief to concentrate on the set of models with minimum Kullback–Leibler divergence from the equilibrium distribution.

In the adaptive-learning approach, one has to make several choices with no obvious answers.²⁴ First, when the agents' expectations about events in the far future matter for their decisions, one has to take a stand on what agents believe about their own future expectations. This issue is often dealt with using the anticipated-utility approach of [Kreps \(1998\)](#): agents behave as if they will never change their expectations—even though they do so in equilibrium.²⁵ I follow the anticipated-utility approach.

Second, one has to specify the procedure used by agents to map their observations to a point estimate for their models. The papers in the adaptive-learning literature often assume least-squares learning; that is, agents estimate the coefficients of their models using ordinary least squares (OLS) estimators. In the general framework studied in this paper, a model consists of a probability distribution over future values of observables—rather than being a set of coefficients in a linear regression. As a result, the agent cannot estimate her model using an OLS estimator. I instead make the natural assumption that the agent estimates her model using a quasi-maximum-likelihood estimator. When the agent's set of models only consists of linear models with normal error terms, the quasi-maximum-likelihood estimator coincides with the OLS estimator.

Definition 7. The function $\theta \mapsto L_t(\{y_s\}_{s=0}^t, \theta) \equiv \prod_{s=1}^t q_\theta(y_s | y_{s-1})$ is called the *quasi-likelihood function* given $\{y_s\}_{s=0}^\infty$. Any maximizer $\hat{\theta}_t \in \arg \max_{\theta \in \Theta} L_t(\{y_s\}_{s=0}^t, \theta)$ of the quasi-likelihood function is a *quasi-maximum-likelihood estimator (QMLE)* for θ given $\{y_s\}_{s=0}^t$.

With the definition of QMLE in hand, I can define an adaptive equilibrium:

Definition 8. Given an economy \mathfrak{E} and a temporary equilibrium T for \mathfrak{E} , an *adaptive equilibrium* with initial conditions (x_0, y_0, z_0) for the economy and initial estimate $\hat{\theta}_0$ for the agent's model consists of adapted random variables $\{x_t, y_t, z_t, \hat{\theta}_t\}_{t=0}^\infty$ with probability distribution \mathbb{P} such that

- (i) $\mathbb{P} \left((x_t, y_t, z_t) \in B \left\{ \{x_s, y_s, z_s, \hat{\theta}_s\}_{s=0}^{t-1} \right\} \right) = T \left(B \mid x_{t-1}, y_{t-1}, z_{t-1}, \mathbf{1}_{\hat{\theta}_{t-1}} \right)$ for any measurable set B and all $t \geq 1$;

²⁴For a discussion of the issues of arbitrariness arising in the adaptive-learning approach, see [Adam and Marcet \(2011\)](#).

²⁵The alternative to anticipated-utility learning is the so-called Euler-equation learning. See [Preston \(2005\)](#) for a discussion of the issues that arise in the adaptive-learning approach when the long-horizon expectations matter. For a discussion of anticipated utility and Euler-equation learning approaches, see the survey article of [Eusepi and Preston \(2011\)](#).

(ii) $\hat{\theta}_t$ is a measurable QMLE for θ given $\{y_s\}_{s=0}^t$ for all $t \geq 1$;

In an adaptive equilibrium, the agent acts at time t as if she is certain that model $\hat{\theta}_t$ is the right model. In other words, the agent's belief is always degenerate. As in a Bayesian equilibrium, the agent behaves optimally given her (now degenerate) belief, the consistency conditions are satisfied, and the distribution of the state is consistent with the given transition probability. The difference between an adaptive equilibrium and a Bayesian equilibrium comes from Condition (ii): in an adaptive equilibrium, the agent's belief at time t is given by the degenerate distribution over a QMLE $\hat{\theta}_t$ for θ given the history of her observations.

Note the cognitive dissonance exhibited by the agent in an adaptive equilibrium. Condition (i) requires the agent to act optimally under the assumption that she will update her belief using Bayes' rule, and a degenerate belief remains unchanged after the application of Bayes' rule. Yet Condition (ii) requires the agent to re-estimate her model at time $t + 1$ given the information that becomes newly available to her. This internally inconsistent behavior on the part of the agent—behaving as if she will not update her belief but updating her belief—is a quirk of the anticipated-utility approach to adaptive learning.

The following existence result is a corollary of the existence of a temporary equilibrium and standard results on the existence of a QMLE:²⁶

Proposition 2. *Consider an economy \mathfrak{E} and a temporary equilibrium for \mathfrak{E} . For any initial condition (x_0, y_0, z_0) for the economy and initial estimate $\hat{\theta}_0$ for the agent's model, an adaptive equilibrium exists.*

An appropriate notion of asymptotic mean stationarity is again sufficient for the concentration of beliefs on the set of models with minimal Kullback–Leibler divergence from the equilibrium distribution. But while Bayesian equilibria have a recursive structure, adaptive equilibria in general do not. This is due to the fact that a QMLE does not necessarily have a recursive representation.²⁷ Consequently, I can no longer appeal to the recursive structure of equilibrium to state the asymptotic-mean-stationarity assumption in terms of the temporary equilibrium of the economy. I instead have the following definition:

Definition 9. An adaptive equilibrium $\{x_t, y_t, z_t, \hat{\theta}_t\}_{t=0}^{\infty}$ with probability distribution \mathbb{P} is *asymptotically mean stationary* if the random process $\{x_t, y_t, z_t, \hat{\theta}_t\}_{t=0}^{\infty}$ is asymptotically mean stationary.²⁸

Asymptotic mean stationarity is the requirement that the long-run behavior of equilibrium variables exhibits some statistical regularity. It is much weaker than stationary and a necessary

²⁶The existence of a measurable QMLE follows from the compactness of Θ , the continuity of $q_\theta(y|y_-)$ in θ for all (y_-, y) , and the measurability of $q_\theta(y|y_-)$. See, for instance, White (1994, p. 16).

²⁷A special case of the QMLE that does have a recursive representation is the OLS estimator.

²⁸See Appendix E for a formal definition of asymptotic mean stationarity for general random processes.

condition for the law of large numbers to hold.²⁹ The following theorem establishes the convergence properties of adaptive equilibrium under asymptotic mean stationarity:

Theorem 3. *Consider an economy \mathcal{E} , a temporary equilibrium T for \mathcal{E} , and an asymptotically-mean-stationary adaptive equilibrium $\{x_t, y_t, z_t, \hat{\theta}_t\}_{t=0}^{\infty}$ with probability distribution \mathbb{P} . With \mathbb{P} -probability one, there exists a probability distribution μ^* over $X \times Y \times Z \times \Theta$ such that*

(a) *for any measurable set $B \subseteq X \times Y \times Z \times \Theta$,*

$$\lim_{t \rightarrow \infty} \frac{1}{t} \sum_{s=0}^{t-1} \mathbb{1}\{(x_s, y_s, z_s, \hat{\theta}_s) \in B\} = \mu^*(B) \quad (13)$$

(b) *the Kullback–Leibler divergence given the agent’s estimate satisfies*

$$\lim_{t \rightarrow \infty} H(Q_{\hat{\theta}_t}, T, \mu^*) = \min_{\theta \in \Theta} H(Q_{\theta}, T, \mu^*);$$

(c) *moreover, if $\min_{\theta \in \Theta} H(Q_{\theta}, T, \mu^*) = \{\theta^*\}$, then the triple $(T, \mu^*, \{\theta^*\})$ is a pure CREE given the set of models Θ for the agent.³⁰*

The theorem is the adaptive-learning counterpart to Theorem 2. The first part establishes the convergence of the empirical frequency of equilibrium variables. The second part proves that the agent’s point estimate converges to the set of models that minimize the Kullback–Leibler divergence from the long-run distribution of observables. But unlike in a Bayesian equilibrium, in an adaptive equilibrium, the long-run economic outcomes do not necessarily coincide with outcomes in a CREE. This is because in a CREE the agent updates her belief using Bayes’ rule, whereas in an adaptive equilibrium she uses QMLE to estimate her model. Still, the third part of the theorem shows that—when the model with minimal Kullback–Leibler divergence is unique—an adaptive equilibrium resembles a (pure) CREE in the long-run. In the Online Appendix, I provide convexity and identifiability assumptions that are sufficient to guarantee that the model with minimal Kullback–Leibler divergence is unique.

Theorems 2 and 3 suggest that the CREE is a natural extension of the rational-expectations equilibrium to scenarios where agents are constrained to choose a model θ belonging to some set Θ that may not contain the true equilibrium probability distribution. The theorems show that inference by such constrained agents leads them to choose models that minimize a novel version of the Kullback–Leibler divergence.

Theorems 2 and 3 also establish a formal link between Bayesian and adaptive approaches to learning. They show that, in the long run, both approaches select temporary equilibria that belong to the class of CREEs. Moreover, if the CREE is unique—as is the case in the application

²⁹See Corollary 8.2 of Gray (2009) for a formal result that establishes the necessity of asymptotic mean stationarity for the law of large numbers to hold.

³⁰Any probability distribution μ^* over $X \times Y \times Z \times \Theta$ is a probability distribution over $X \times Y \times Z \times \Delta\Theta$ that only assigns positive probability to degenerate beliefs over Θ .

I consider in Section 6—both approaches select the same temporary equilibrium. The long-run behavior of the economy is then independent of the details of the learning process; it only depends on the fundamentals of the economy and the models considered plausible by agents.

5 Misspecification in Hidden Factor Models

I next consider a special case of the general framework of Section 2 in which the agent entertains linear-Gaussian hidden factor models. I use this specification of the agent’s set of models to argue that bounded rationality in the form of misspecification can lead to amplification or dampening of shocks and to history dependence. In Section 6, I use a similar hidden factor model to quantitatively study the implications of bounded rationality for business-cycle fluctuations.

5.1 A Canonical Model

I consider a set of models for the agent according to which observables are linear functions of an unobservable state. The agent believes that there is an underlying state $\omega_t \in \mathbb{R}^d$ that follows the Markov process

$$\omega_t = A\omega_{t-1} + \epsilon_{\omega t}, \quad (14)$$

where $A \in \mathbb{R}^{d \times d}$ is a stable matrix and $\epsilon_{\omega t} \in \mathbb{R}^d$ is i.i.d. $\mathcal{N}(0, \Sigma_\omega)$.³¹ The vector of observables $o_t \in \mathbb{R}^n$ is a linear function of ω_t :

$$o_t = B'\omega_t + \epsilon_{ot}, \quad (15)$$

where $B \in \mathbb{R}^{d \times n}$, $\epsilon_{ot} \in \mathbb{R}^n$ is i.i.d. $\mathcal{N}(0, \Sigma_o)$, and B' denotes the transpose of B .³² Hidden factor models are also known as state-space models and hidden Markov models.³³ I use the terms hidden state and factor interchangeably when referring to ω_t . In Appendix C, I show how several behavioral models can be viewed as instances of hidden factor models in which matrices A , B , Σ_ω , and Σ_o have been constrained.

The agent’s set of models is parameterized by the tuple $\theta = (A, B, \Sigma_\omega, \Sigma_o)$. Matrices A and Σ_ω parametrize the agent’s view of the persistence and volatility of the factors and the correlations between different factors. Matrix B captures the agent’s view of the contribution of each factor to

³¹A square matrix is *stable* if its eigenvalues all lie inside the unit circle.

³²Since $\epsilon_{\omega t}$ and ϵ_{ot} are assumed to be distributed according to normal distributions, which do not have compact supports, the application considered in this section is not strictly speaking a special case of the framework of Section 2. But the proofs of the theoretical results can be extended to the non-compact case by assuming that the probability distributions that arise in equilibrium are *tight* and appealing to Prokhorov’s theorem. The assumptions maintained in this section (that the shocks follows stationary processes and that agents only entertain stationary models) guarantee that any probability measure that arises in equilibrium is tight.

³³Hidden factor models are extensively used in engineering and economics. See the handbook chapter of [Stock and Watson \(2016\)](#) for a thorough exposition on the theory and application of factor models in macroeconomics. They are also used in the Bayesian and adaptive learning literatures as the rationale for the use of constant-gain learning rules. A number of behavioral models such as [Barberis, Shleifer, and Vishny \(1998\)](#)’s model of sentiments, [Rabin and Vayanos \(2010\)](#)’s model of Gambler’s fallacy, and models of rational inattention (e.g., [Maćkowiak, Matejka, and Wiederholt, 2017](#)) can also be expressed as hidden factor models in which the agents’ models have been constrained by behavioral assumptions.

every observable. It can be seen as representing the “loading” of observables on different factors. The random variable ϵ_{ot} is often referred to as the “measurement error” in the literature. In the current framework, equations (14) and (15) are not structural equations—they only exist in the mind of the agent—so ϵ_{ot} is not a measurement error in the usual sense. But for the sake of consistency with existing works, I refer to ϵ_{ot} as the measurement error and to Σ_o as the variance-covariance matrix of the measurement error.

Hidden factor models can be seen as canonical representations of stationary stochastic processes. In particular, the Wold representation theorem implies that any stationary ARMA process with Gaussian innovations has a representation in the form of equations (14) and (15). But the representation is, in general, not unique. For instance, one can always scale both a factor and the loading of observables on the factor without changing the agent’s forecasts of observables.³⁴ I assume in the rest of this section that matrices $(A, B, \Sigma_\omega, \Sigma_o)$ are normalized in such a way that the representation in (14) and (15) is unique.³⁵

The agent’s models can be expressed in the recursive form of Section 2 by defining

$$y_t \equiv (o_t, o_{t-1}, \dots) \quad (16)$$

to be the history of realizations of o_t . The transition probability Q_θ is defined by a stationary Kalman filter. The agent believes that o_{t+1} is normally distributed with a variance-covariance matrix that is only a function of A, B, Σ_ω , and Σ_o and a conditional mean that is given by $E_t[o_{t+1}] = B'\hat{\omega}_t$. The vector $\hat{\omega}_t \equiv \hat{\omega}_{t+1|t} \in \mathbb{R}^d$ is the agent’s estimate of ω_t , defined recursively as

$$\hat{\omega}_t = (A - KB')\hat{\omega}_{t-1} + Ko_t, \quad (17)$$

where K is the matrix of “Kalman gains.” More generally, the agent’s conditional expectation of o_{t+s} at time t is given by

$$E_t[o_{t+s}] = B'A^{s-1}\hat{\omega}_t.$$

See Appendix B for details of how a hidden factor model can be mapped to the model of Section 2 and an explicit expression for the matrix of Kalman gains.

5.2 Amplification, Dampening, and History Dependence

I next study how replacing rational expectations with a hidden factor model for the agent changes the response of the economy to shocks. For the sake of concreteness, I consider a simple linear economy with one decision variable, one state variable, and one shock. The temporary

³⁴More generally, $\theta = (A, B, \Sigma_\omega, \Sigma_o)$ is observationally indistinguishable from $\bar{\theta} = (\bar{A}, \bar{B}, \bar{\Sigma}_\omega, \Sigma_o)$ whenever there exists an invertible matrix U such that $\bar{A} = UAU^{-1}$, $\bar{B} = (U^{-1})'B$, and $\bar{\Sigma}_\omega = U\Sigma_\omega U'$. See Gevers and Wertz (1984) for a discussion of identification of state-space models and a canonical parameterization.

³⁵One can, for instance, normalize A to be a diagonal stable matrix and normalize Σ_ω to be a correlation matrix.

equilibrium is described by the following equations:

$$x_t = Co_t + E_t \left[\sum_{s=1}^{\infty} \beta^s Co_{t+s} \right], \quad (18)$$

$$o_t = (x_t, z_t)', \quad (19)$$

$$z_t = \rho z_{t-1} + \varepsilon_t, \quad (20)$$

where x_t is the choice variable, z_t is the state variable, ε_t is an i.i.d. $\mathcal{N}(0, \sigma_\varepsilon^2)$ shock, $C = (c, 1)$ is a vector of constants that depends on the fundamentals of the economy with $c \in (0, 1)$, and $\rho \in [0, 1)$ is the persistence of the state variable. Equation (18) is the optimality condition for the agent, (19) is the general-equilibrium consistency requirement, and (20) is the equation describing the transition probability for the state variable.

The rational-expectations equilibrium of the economy is easy to characterize using standard techniques:

$$x_t^{\text{REE}} = \frac{1}{1 - c - \beta\rho} z_t. \quad (21)$$

The multiplier $1/(1 - c - \beta\rho)$ is the usual Keynesian multiplier, adjusted to account for the effect of the persistence of the state variable on the agent's expectations. Note that, in the rational-expectations equilibrium, x_t is independent of past values of the state variable conditional on its current value z_t .

The rational-expectations equilibrium can be represented as a hidden factor model of the form (14) and (15). Let $\omega_t^{\text{REE}} = z_t$, $A^{\text{REE}} = \rho$, $B^{\text{REE}} = (1/(1 - c - \beta\rho), 1)$, $\Sigma_\omega^{\text{REE}} = \sigma_\varepsilon^2$, and $\Sigma_o^{\text{REE}} = \mathbf{0}$, where $\mathbf{0}$ denotes the zero matrix. Any vector $K = (k_x, k_z)$ satisfying $k_x/(1 - c - \beta\rho) + k_z = \rho$ is a vector of (optimal) Kalman gains given $(A^{\text{REE}}, B^{\text{REE}}, \Sigma_\omega^{\text{REE}}, \Sigma_o^{\text{REE}})$. I set $K^{\text{REE}} = (\rho(1 - c - \beta\rho), 0)$. With these definitions, equations (14) and (15) are equivalent to equations (20) and (21), which describe the rational-expectations equilibrium. The following proposition is a trivial generalization of this observation:

Proposition 3. *Any stationary rational-expectations equilibrium of a linear economy with normally distributed shocks can be represented as a hidden factor model in the form of equations (14) and (15).*

This proposition establishes that the restriction to hidden factor models does *not* rule out rational-expectations equilibria by itself. Deviations from the REE benchmark only arise because of additional constraints on the agent's set of models. Those are direct restrictions on matrices A , B , Σ_ω , and Σ_o in learning, rational inattention, and behavioral models such as [Rabin and Vayanos \(2010\)](#). They arise from a constraint on the number of factors in the application that comes in the next section. In the rest of this subsection, I remain agnostic about the source of the deviation from the REE and study the response of the economy to shocks taking $\theta = (A, B, \Sigma_\omega, \Sigma_o)$ as given. The next subsection focuses on the endogenous determination of θ .

The response of the economy to shocks given the hidden factor model can be found by combining the agent's forecasts with the temporary equilibrium relationships. Substituting for $E_t[o_{t+s}]$ in (18) and using the definition o_t , I get

$$x_t^{\text{CREE}} = \frac{1}{1 - c - \beta CB'(I - \beta A)^{-1} k_x} \left[\beta CB'(I - \beta A)^{-1} (A - KB') \hat{\omega}_{t-1} + (1 + \beta CB'(I - \beta A)^{-1} k_z) z_t \right]. \quad (22)$$

Equation (22) is different from (21) in three ways. First, the part in the Keynesian multiplier that captures the general-equilibrium effects of changes in expectations is $\beta \rho$ in the REE and $\beta CB'(I - \beta A)^{-1} k_x$ in a CREE. In the latter, a unit increase in x_t leads the agent to increase her estimate of ω_t by k_x units. This leads to a $\beta CB'(I - \beta A)^{-1} k_x$ unit increase in x_t . Second, the coefficient capturing the direct effect of changes in z_t on the agent's choices is 1 in the REE and $1 + \beta CB'(I - \beta A)^{-1} k_z$ in a CREE. A unit change in the value of the state variable leads to a k_z -unit change in the agent's estimates of ω_t in a CREE. This in turn leads to an additional $\beta CB'(I - \beta A)^{-1} k_z$ unit change in the agent's action. Third, unlike equation (21), equation (22) generally has a backward-looking term: even conditional on the value of the state variable at time t , the agent's action in a CREE may depend on the past realizations of shocks as summarized by $\hat{\omega}_{t-1}$.

These deviations from the rational-expectations benchmark can lead to amplification or dampening of shocks. Whether shocks are amplified or dampened relative to the rational-expectations benchmark depends on the constraints on the agent's set of models and the model selected in the CREE.

Proposition 4. *Consider the linear economy described in equations (18)–(20) and a linear-Gaussian hidden factor model for agents as in (14) and (15). If*

$$CB'(I - \beta A)^{-1} \left(\frac{1}{1 - c - \beta \rho} k_x + k_z \right) \preceq \frac{\rho}{1 - c - \beta \rho} \quad (23)$$

in a CREE, then

$$\frac{\partial x_t^{\text{CREE}}}{\partial \varepsilon_t} \preceq \frac{\partial x_t^{\text{REE}}}{\partial \varepsilon_t}.$$

The right-hand side of equation (23) is a constant that only depends on the fundamentals of the economy. The left-hand side depends on the agent's selected model through its dependence on k_x , k_z , B , and A . It is increasing in all four. Larger values of k_x and k_z imply that the agent's estimates of the hidden states are more sensitive to changes in observables; a larger B implies that the agent's actions are more sensitive to changes in her estimates of the hidden states; and a larger A implies that the agent perceives the hidden states to be more persistent. All three contribute to larger responses of the economy to shocks.³⁶

The deviation from the rational-expectations benchmark can also lead to history dependence. Let $e_t \equiv x_t - \mathbb{E}[x_t | z_t]$ denote the part of the choice variable at time t that is not explained

³⁶Note however that K is not independent of matrices A and B . It is determined by Bayes' rule given parameters A , B , Σ_ω , and Σ_o of the agents' model. See Appendix B for details.

by the value of the state variable at time t , where \mathbb{E} denotes expectation with respect to the equilibrium probability distribution. One can think of e_t as the residual from a linear regression of x_t on z_t . I use e_t as the measure of history dependence. The following proposition characterizes the extent of history dependence as a function of the parameters of the agent's model that are selected in a CREE:

Proposition 5. *In the linear economy described in equations (18)–(20),*

$$e_t^{REE} \equiv 0.$$

If agents have a linear-Gaussian hidden factor model as in (14) and (15),

$$e_t^{CREE} = \frac{\beta CB'(I - \beta A)^{-1}}{1 - c - \beta CB'(I - \beta A)^{-1} k_x} \sum_{s=1}^{\infty} (A - KB')^s (k_x x_{t-s} + k_z z_{t-s}). \quad (24)$$

The first part of the proposition shows that the REE of the linear economy does not exhibit history dependence. The second part establishes that the economy exhibits history dependence for generic models for the agent—it is only when $A - KB' = 0$, as in the REE, that there is no history dependence. More formally, for almost any perturbation $(\tilde{A}, \tilde{B}, \tilde{\Sigma}_\omega, \tilde{\Sigma}_o)$ of the parameters of the agent's model away from their rational-expectations values, the temporary equilibrium in which the agent's expectations are induced by $(\tilde{A}, \tilde{B}, \tilde{\Sigma}_\omega, \tilde{\Sigma}_o)$ exhibits history dependence. The lack of history dependence in the REE is therefore not robust to the possibility that the agent's set of models is misspecified.

5.3 The CREE

A (pure) CREE is defined via the mapping from the agent's model to the distribution of the observable and the mapping from the distribution of the observable to the agent's model. I start by describing the first mapping. Suppose that the agent puts probability one on some model θ , and let $\zeta_t = (o'_t, \hat{w}'_t)'$ denote the vector consisting of the observable and the agent's estimate \hat{w}_t of the hidden state. Equations (17) and (22) can be combined to express ζ_t as a VAR(1) process:

$$\zeta_t = A_\zeta(\theta)\zeta_{t-1} + C_\zeta(\theta)\varepsilon_t. \quad (25)$$

Equation (25) and the assumption that the shock is normally distributed imply that ζ_t is normally distributed at all times. Therefore, the stationary distribution of ζ_t is fully described by the autocovariance matrices

$$\Xi_{\zeta,s}(\theta) = \mathbb{E}_\theta[\zeta_t \zeta'_{t-s}], \quad (26)$$

where \mathbb{E}_θ denotes the expectation with respect to the probability distribution over $\{\zeta_t\}_t$ induced by equation (25) and the distribution of the shock. Matrices $\{\Xi_{\zeta,s}(\theta)\}_s$ can be recursively computed in terms of matrices $A_\zeta(\theta)$ and $C_\zeta(\theta)$ and the variance σ_ε^2 of the shock. I let $\{\Xi_s(\theta)\}_s$,

where $\Xi_s(\theta) \equiv \mathbb{E}_\theta[o_t o'_{t-s}]$, denote the corresponding autocovariance matrices for the vector of observables.

The mapping from the distribution of observables to the agent's model is given by the minimization of the Kullback–Leibler divergence. In Appendix B, I show that the Kullback–Leibler divergence given model θ and autocovariance matrices $\{\Xi_s\}_s$ for observables is given by

$$H(\{\Xi_s\}_s, \theta) = -\frac{1}{2} \log \left(\det \left(\Omega^{-1} \right) \right) + \frac{n}{2} \log (2\pi) + \frac{1}{2} \text{tr} \left(\Omega^{-1} \Xi_0 \right) - \sum_{s=1}^{\infty} \text{tr} \left(\Omega^{-1} \Xi_s K' (A' - BK')^{s-1} B \right) \\ + \frac{1}{2} \sum_{s=1}^{\infty} \sum_{\tau=1}^{\infty} \text{tr} \left(\Omega^{-1} B' (A - KB')^{s-1} K \Xi_{\tau-s} K' (A' - BK')^{\tau-1} B \right), \quad (27)$$

where Ω is the variance-covariance matrix of the agent's perceived forecast errors. In a CREE, the agent's beliefs are concentrated on the set $\Theta(\{\Xi_s\}_s) \equiv \arg \min_{\theta \in \Theta} H(\{\Xi_s\}_s, \theta)$ of models that minimize the Kullback–Leibler divergence from the distribution of observables.

A pure CREE is defined by the fixed point of the mappings $\theta \mapsto \{\Xi_s(\theta)\}_s$ and $\{\Xi_s\}_s \mapsto \Theta(\{\Xi_s\}_s)$ defined in the previous two paragraphs. It consists of a model θ^* and autocovariance matrices $\{\Xi_s^*\}_s$ for observables such that $\theta^* \in \Theta(\{\Xi_s^*\}_s)$ and $\{\Xi_s^*\}_s = \{\Xi_s(\theta^*)\}_s$. The agent selects the model θ^* that minimizes the Kullback–Leibler divergence, and the distribution of observables is consistent with the agent's optimal actions given the fundamentals of the economy and her expectations. The CREE clearly depends on the agent's constrained set of models Θ . In Appendix C, I argue that models such as rational inattention, gambler's fallacy, and adaptive-learning models constrain the set Θ by imposing direct restrictions on matrices A , B , Σ_ω , and Σ_o .

The alternative I pursue in the remainder of this paper is a novel form of bounded rationality where the only constraint is on the number of factors entertained by agents. They are otherwise free to choose matrices A , B , Σ_ω , and Σ_o to minimize the Kullback–Leibler divergence from the distribution of their observations. This constraint captures an intuitive behavioral assumption: agents can only conceive low-dimensional models. It is inspired by research by [Stock and Watson \(2016\)](#) and [Angeletos, Collard, and Dellas \(2018\)](#), among others, who show that the business-cycle fluctuations in a large number of aggregate variables can be attributed to fluctuations in a small number of hidden factors.

I refer to the special case of CREE where the agents' models are given by d -factor hidden factor models as a CREE- d . The parameter d is a structural parameter that captures the complexity of the agents' models, with higher values of d corresponding to more complex models of the economy. I use CREE- d as the model of bounded rationality in the business-cycle application of the next section.

6 A Business-Cycle Model with CREE- d

In this section, I study a standard medium-scale new-Keynesian model in which the agents' models of the economy are constrained. I use this exercise to argue that bounded rationality

in the form of model misspecification improves the empirical fit of the model, that it is a micro-founded substitute for add-ons such as exogenous habit formation, price and wage indexation, and investment-adjustment costs, and that it leads to novel economic insights.³⁷

6.1 The Economic Environment

The economy is a new-Keynesian model with capital, neoclassical capital-adjustment costs, price and wage rigidities, and six shocks: a total factor productivity (TFP) shock, a discount-factor shock, a price-markup shock, a wage-markup shock, a monetary-policy shock, and a government-spending shock. The model can alternatively be seen as a DSGE model à la [Christiano, Eichenbaum, and Evans \(2005\)](#) and [Smets and Wouters \(2007\)](#) with the following bells and whistles dropped: (i) habit formation in consumption, (ii) wage and price indexation, (iii) the capital-utilization margin, and (iv) the shock to the marginal product of investment. I also replace investment-adjustment costs with a neoclassical capital adjustment cost. The specification of price and markup shocks is closest to [Justiniano, Primiceri, and Tambalotti \(2010\)](#), whose analysis I follow.

6.1.1 Final-good producers

The final good Y_t is produced by competitive firms by combining a continuum of intermediate goods, indexed by i , according to the CES production function

$$Y_t = \left[\int_0^1 Y_t(i)^{\frac{1}{1+\lambda_{pt}}} di \right]^{1+\lambda_{pt}}.$$

The elasticity λ_{pt} follows the following AR(1) process

$$\log(1 + \lambda_{pt}) = (1 - \rho_p) \log(1 + \lambda_p) + \rho_p \log(1 + \lambda_{p,t-1}) + \varepsilon_{pt},$$

where ε_{pt} is i.i.d. $\mathcal{N}(0, \sigma_p^2)$. Profit maximization and the zero-profit condition imply that the price of the final good is given by the price index

$$P_t = \left[\int_0^1 P_t(i)^{\frac{1}{\lambda_{pt}}} di \right]^{\lambda_{pt}},$$

where $P_t(i)$ denotes the price of the intermediate good i . The demand for good i is given by the isoelastic demand schedule

$$Y_t(i) = \left(\frac{P_t(i)}{P_t} \right)^{-\frac{1+\lambda_{pt}}{\lambda_{pt}}} Y_t. \quad (28)$$

³⁷A number of recent papers in macroeconomics have proposed deviations from the benchmark of full-information rational-expectations as ways of improving the predictions of the baseline new-Keynesian model. Examples include [Garcia-Schmidt and Woodford \(2015\)](#), [Gabaix \(2016\)](#), [Farhi and Werning \(2017\)](#), [Angeletos and Lian \(2018\)](#), [Woodford \(2018\)](#), and [Angeletos and Huo \(2018\)](#). In related work, [Acemoglu and Jensen \(2018\)](#) characterize the robust predictions of the neoclassical growth model in the presence of behavioral biases.

6.1.2 Intermediate-goods producers

A monopolist produces each intermediate good i according to the production function

$$Y_t(i) = \max \left\{ z_t K_t(i)^\alpha (\gamma^t L_t(i))^{1-\alpha} - \gamma^t F, 0 \right\},$$

where $K_t(i)$ and $L_t(i)$ denote the capital and labor input of the firm, F is a fixed cost of production, chosen so that profits are zero along the balanced growth path, γ is the exogenous rate of labor-augmenting technological progress, and z_t is the stationary part of TFP that follows the AR(1) process

$$\log z_t = \rho_z \log z_{t-1} + \varepsilon_{zt},$$

where ε_{zt} is i.i.d. $\mathcal{N}(0, \sigma_z^2)$.

Intermediate-good producers are subject to nominal frictions à la Calvo. Each period the price of a randomly-selected fraction ξ_p of intermediate goods is fixed. The remaining intermediate-goods producers choose their prices $P_t(i)$ optimally by maximizing the present-discounted value of future profits,

$$E_t \left[\sum_{s=0}^{\infty} \xi_p^s \frac{\beta^s \Lambda_{t+s}}{\Lambda_t} \left(P_t(i) Y_{t+s}(i) - W_{t+s} L_{t+s}(i) - r_{t+s} K_{t+s}(i) \right) \right],$$

subject to the demand curve (28), where Λ_t is the marginal utility of nominal income, W_t is the nominal wage, and r_t is the rental rate of capital. The operator E_t denote the expectation with respect to the intermediate-good producers' subjective expectation at time t about the path $\{\Lambda_{t+s}, W_{t+s}, r_{t+s}, P_{t+s}, Y_{t+s}, \lambda_{p,t+s}, z_{t+s}\}_{s \geq 1}$ of variables that the producers take as given.

6.1.3 Investment firms

The capital stock of the economy is owned by investment firms. The firms take the rental rate of capital and the price of the final good as given and maximize the present-discounted value of profits

$$E_t \left[\sum_{s=0}^{\infty} \beta^s \Lambda_{t+s} (r_t K_t - P_t I_t) \right],$$

subject to the physical capital accumulation equation

$$K_{t+1} = (1 - \delta)K_t + I_t - S \left(\frac{I_t}{K_t} \right) K_t. \quad (29)$$

I assume that the adjustment cost satisfies $S = S' = 0$ and $S'' > 0$ along the balanced growth path. Note that the adjustment cost is a neoclassical adjustment cost and not an investment-adjustment cost that is commonplace in the DSGE literature.³⁸ I assume that no spot market for capital exists.³⁹

³⁸The investment-adjustment-cost specification replaces $S(I_t/K_{t-1})K_{t-1}$ with $S(I_t/I_{t-1})I_{t-1}$. It leads to an Euler equation for investment that has a backward-looking term.

³⁹Since this is a representative-agent model, this assumption is immaterial in the rational-expectations equilibrium. But in a CREE, investment depends on expectations about the entire future path of returns to capital when there is no capital market, whereas it only depends on the rental rate of capital and its price in the next period with a spot market for capital.

6.1.4 Employment agencies

There is a continuum of households. Each household is a monopolistic supplier of a specialized type of labor, index by j . A competitive employment agency combines specialized labor into a homogeneous labor input using the CES function

$$L_t = \left[\int_0^1 L_t(j)^{\frac{1}{1+\lambda_{wt}}} dj \right]^{1+\lambda_{wt}},$$

where λ_{wt} is a wage-markup shock that follows the ARMA(1) process

$$\log(1 + \lambda_{wt}) = (1 - \rho_w) \log(1 + \lambda_w) + \rho_w \log(1 + \lambda_{w,t-1}) + \varepsilon_{wt}$$

with ε_{wt} i.i.d. $\mathcal{N}(0, \sigma_w^2)$.

Profit maximization by employment agencies and the zero-profit condition imply that the price of the homogeneous good is given by the wage index

$$W_t = \left[\int_0^1 W_t(j)^{\frac{1}{\lambda_{wt}}} dj \right]^{\lambda_{wt}},$$

and the demand for the labor of type j is given by the isoelastic curve

$$L_t(j) = \left(\frac{W_t(j)}{W_t} \right)^{-\frac{1+\lambda_{wt}}{\lambda_{wt}}} L_t. \quad (30)$$

6.1.5 Labor unions

Wages are set by a continuum of labor unions, also indexed by j , each representing a household. Each period a randomly-selected fraction ξ_w of unions cannot freely set the wage of the household they represent. The nominal wages of those households grow at the deterministic growth rate of TFP, γ .⁴⁰ The remaining fraction of labor unions set the optimal wage $W_t(j)$ by maximizing

$$E_t \left[\sum_{s=0}^{\infty} \xi_w^s \beta^s \left(-\psi_{t+s} \varphi \frac{L_{t+s}(j)^{1+\nu}}{1+\nu} + \Lambda_{t+s} W_t(j) L_{t+s}(j) \right) \right]$$

subject to (30), where Λ_t is the marginal utility of nominal income.

6.1.6 Households

I assume that there exists a competitive insurance agency that fully insures households against fluctuations in their labor income resulting from the inability of labor unions to reset the nominal wage. Therefore, the equilibrium labor income of each household is equal to $W_t L_t$, the average labor income in the economy.

⁴⁰Since there is technological progress, absent this assumption, there would be no balanced growth path without wage dispersion. Note that this is different than the assumption of wage indexation that is common in the DSGE literature: I do not assume the wage to be indexed to the current inflation rate.

Each household takes the labor income and the stream of profits from the ownership of firms as given and chooses consumption and saving in government bonds to maximize the utility function

$$E_t \left[\sum_{s=0}^{\infty} \beta^s \psi_{t+s} \left(\log(C_{t+s}) - \varphi \frac{L_{t+s}(j)^{1+\nu}}{1+\nu} \right) \right],$$

subject to a no-Ponzi condition and the nominal budget constraint

$$P_t C_t + T_t + B_t \leq R_{t-1} B_{t-1} + W_t L_t + \Pi_t, \quad (31)$$

where C_t is consumption, I_t is investment, T_t denotes lump-sum taxes, B_t is holding of one-period government bonds, R_t is the gross nominal interest rate, Π_t denotes profits from the ownership of firms, and ψ_t is a discount-factor shock that follows the AR(1) process

$$\log \psi_t = \rho_\psi \log \psi_{t-1} + \varepsilon_{\psi t},$$

with $\varepsilon_{\psi t}$ i.i.d. $\mathcal{N}(0, \sigma_\psi^2)$. The operator E_t denotes the expectation with respect to the household's subjective belief about the path $\{\psi_{t+s}, L_{t+s}, W_{t+s}, P_{t+s}, T_{t+s}, R_{t+s}, \Pi_{t+s}\}_{s \geq 1}$ of aggregate and idiosyncratic observables that enter her decision problem.

6.1.7 The government

The monetary policy sets the nominal interest rate following a Taylor rule

$$\frac{R_t}{R} = \left(\frac{R_{t-1}}{R} \right)^{\rho_R} \left(\frac{\pi_t}{\pi} \right)^{\phi_\pi (1-\rho_R)} \eta_{mt},$$

where R is the steady-state gross nominal interest rate.⁴¹ η_{mt} is a monetary policy shock that follows the AR(1) process

$$\log \eta_{mt} = \rho_m \log \eta_{m,t-1} + \varepsilon_{mt},$$

where ε_{mt} is i.i.d. $\mathcal{N}(0, \sigma_m^2)$.

Government spending G_t is exogenous, with $g_t \equiv G_t/\gamma^t$ given by the following AR(1) process:

$$\log g_t = (1 - \rho_g) \log g + \rho_g \log g_{t-1} + \varepsilon_{gt}$$

and ε_{gt} i.i.d. $\mathcal{N}(0, \sigma_g^2)$. Government finances spending by issuing short-term nominal bonds and levying lump-sum taxes on households. The nominal government budget constraint is given by

$$R_{t-1} B_{t-1} + P_t G_t - T_t = B_t, \quad (32)$$

where T_t denotes nominal taxes. Taxes follow a tax rule that ensures that the real value of public debt (i.e., $B_t/P_t \gamma^t$) remains constant over time.⁴²

⁴¹In the new-Keynesian literature, it is often assumed that the monetary authority responds both to changes in the inflation rate and to changes in the output gap. The right notion of an output gap in a CREE is not clear. The monetary authority may define the output gap relative to the flexible price allocation where agents re-estimate their model, the one in which the agents' models are unchanged, or the rational-expectations flexible price allocation. I bypass the question of how the output gap ought to be defined by assuming that the monetary authority only responds to changes in the inflation rate.

⁴²Ricardian equivalence does not necessarily hold in a CREE. The timing of taxes and the value of the outstanding public

6.1.8 Market clearing

The aggregate resource constraint is given by

$$C_t + I_t + G_t = Y_t.$$

6.2 Log-Linear Temporary Equilibrium

I start by characterizing the log-linearized temporary equilibrium of the economy under arbitrary specification of expectations. As the first step, I characterize the non-linear equilibrium conditions. I then characterize a balanced growth path (BGP) along which prices are constant and output, consumption, investment, government spending, capital stock, the real wage, and the public debt all grow at rate γ , the deterministic growth rate of labor productivity. Finally, I log-linearize the equilibrium conditions around the BGP. The calculations are tedious and are relegated to the Online Appendix.

Away from rational expectations, one agent's optimality conditions cannot be simplified using other equilibrium conditions that are not necessarily respected by the agent's expectation operator—conditions such as other agents' optimality conditions and resource constraints.⁴³ For instance, the optimality condition of firms resetting their prices cannot be combined with the evolution of the price index to obtain the standard recursive version of the Phillips curve. Likewise, the households' optimality conditions cannot be combined with the aggregate resource constraint to obtain the standard recursive consumption-Euler equation.

I instead combine the households' first-order optimality conditions with their budget constraints and the no-Ponzi condition to obtain the following version of the permanent-income hypothesis:

$$\begin{aligned} \hat{c}_t = & \beta \hat{\psi}_t + \frac{b}{c} \frac{1-\beta}{\beta} \left(\hat{R}_{t-1} + \hat{b}_{t-1} - \hat{\pi}_t \right) + (1-\beta) \left(\frac{x}{c} \hat{x}_t - \frac{\tau}{c} \hat{\tau}_t \right) - \beta \hat{R}_t \\ & + E_t \left[\sum_{s=1}^{\infty} \beta^s \left((1-\beta) \frac{x}{c} \hat{x}_{t+s} - (1-\beta) \frac{\tau}{c} \hat{\tau}_{t+s} - (1-\beta) \hat{\psi}_{t+s} - \beta \frac{x-\tau}{c} \hat{R}_{t+s} + \frac{x-\tau}{c} \hat{\pi}_{t+s} \right) \right]. \end{aligned} \quad (33)$$

where lowercase letters with hats denote log-deviations from the BGP and lowercase letters without hats denote the steady state values, π denotes the inflation rate, x denotes the households' total income from labor and the ownership of firms, and τ denotes their real tax burden.

Equation (33) has an intuitive interpretation. The first term is the direct effect of a discount-factor shock on consumption. The second term captures the wealth effect from changes in the public debt and the revaluation of government bonds. The third term is consumption out of

debt might therefore both affect the response of the economy to shocks. See also [Eusepi and Preston \(2018a\)](#), where the authors use an adaptive learning framework to study the effects of the level of public debt on the transmission of monetary policy.

⁴³[Preston \(2005\)](#) is the first to make this point in the context of adaptive-learning models. See, also, [Woodford \(2003b, p. 272\)](#) for a discussion.

changes in the current disposable income, with $1 - \beta$ the marginal propensity to consume (MPC). The fourth term is the intertemporal substitution effect from changes in the nominal interest rate. The term on the second line is consumption out of changes in the expected permanent income, with $1 - \beta$ the MPC.

Investment is given by

$$\hat{i}_t = \hat{k}_t + \frac{1}{\chi} (-\hat{\psi}_t + \hat{c}_t) + E_t \left[\sum_{s=1}^{\infty} \beta^s \left(\frac{1-\beta}{\chi\beta} \hat{\psi}_{t+s} - \frac{1-\beta}{\chi\beta} \hat{c}_{t+s} + \frac{1}{\chi} \left(\frac{1}{\beta} - \frac{1-\delta}{\gamma} \right) \hat{\rho}_{t+s} \right) \right], \quad (34)$$

where ρ denotes the rental rate of capital and χ is a constant that captures the convexity of the adjustment cost. Note that investment depends on the expected infinite presented-discounted value of the rental rate of capital, where the discount rate depends on consumption and the discount factor-shock. If there existed a market where firms could trade capital, investment would only depend on the investment firms' expectation of the price of capital in the next period. While the two expressions would coincide under rational expectations, under constrained-rational expectations, they only coincide in knife-edge cases.

Price and wage inflation are described by two Phillips curves that relate the current values of inflation rate and the wage index to the intermediate-goods' producers and the labor unions' expectations about the paths of observables. In particular, price inflation is described by

$$\begin{aligned} \hat{\pi}_t = & \kappa \left(\hat{\lambda}_{pt} + \alpha \hat{\rho}_t + (1 - \alpha) \hat{w}_t - \hat{z}_t \right) \\ & + E_t \left[\sum_{s=1}^{\infty} \xi_p^s \beta^s \left(\frac{1 - \xi_p}{\xi_p} \hat{\pi}_{t+s} + \kappa \left(\hat{\lambda}_{p,t+s} + \alpha \hat{\rho}_{t+s} + (1 - \alpha) \hat{w}_{t+s} - \hat{z}_{t+s} \right) \right) \right], \end{aligned} \quad (35)$$

where κ is a constant. The wage index is given by

$$\hat{w}_t = \xi_w (\hat{w}_{t-1} - \hat{\pi}_t) + \kappa_w \left(\hat{\lambda}_{wt} + v \hat{L}_t + \hat{c}_t \right) + \kappa_w (v_w - 1) \hat{w}_t \quad (36)$$

$$+ E_t \left[\sum_{s=1}^{\infty} \xi_w^s \beta^s \left((1 - \xi_w) \hat{\pi}_{t+s} + \kappa_w \left(\hat{\lambda}_{wt} + v \hat{L}_t + \hat{c}_t \right) + \kappa_w (v_w - 1) \hat{w}_{t+s} \right) \right], \quad (37)$$

where v_w and κ_w are constants. The remaining log-linearized equilibrium conditions are relatively standard. They are presented in Appendix D.1.

6.3 The Agents' Models

I assume that all agents have perfect foresight about the balanced growth path of the economy. So I only need to specify the expectations of intermediate-goods producers, investment firms, labor unions, and households about the paths $\{\hat{\pi}_s, \hat{\lambda}_{ps}, \hat{\rho}_s, \hat{w}_s, \hat{z}_s\}_s$, $\{\hat{\psi}_s, \hat{c}_s, \hat{\rho}_s\}_s$, $\{\hat{\pi}_s, \hat{\lambda}_{ws}, \hat{L}_s, \hat{c}_s, \hat{w}_s\}_s$, and $\{\hat{x}_s, \hat{\tau}_s, \hat{\psi}_s, \hat{R}_s, \hat{\pi}_s\}_s$, respectively, of log deviations from the BGP. For simplicity, I assume that all agents have the same model of the economy. Therefore, each model θ must determine the expectations about the paths $\{\hat{x}_s, \hat{\tau}_s, \hat{\psi}_s, \hat{R}_s, \hat{\pi}_s, \hat{\rho}_s, \hat{w}_s, \hat{z}_s, \hat{\lambda}_{ps}, \hat{\lambda}_{ws}, \hat{L}_s, \hat{c}_s\}_s$ of

observables that enter the problems of intermediate-goods producers, investment firms, labor unions, or households. I further enrich the set of agents' observations by assuming that they also observe aggregate investment and government spending. So agents are all assumed to observe and form expectations about the following vector of observables:

$$o_s \equiv \left(\hat{x}_s, \hat{\tau}_s, \hat{\psi}_s, \hat{R}_s, \hat{\pi}_s, \hat{\rho}_s, \hat{w}_s, \hat{z}_s, \hat{\lambda}_{ps}, \hat{\lambda}_{ws}, \hat{L}_s, \hat{c}_s, \hat{i}_s, \hat{g}_s \right)' \in \mathbb{R}^n. \quad (38)$$

Although the elements of o_s are interdependent in equilibrium, the agents may not be aware of equilibrium relationships, and so they may believe in combinations of variables that are inconsistent with equilibrium relationships. This is the manifestation of the possibility that the agents' expectations may not coincide with rational expectations.

I assume that the agents can only entertain d -factor linear-Gaussian hidden factor models as in Section 5. The agents believe that there are d underlying factors that determine the movements in the observable economic variables. The d -factor model fully describes the agents' constrained set of expectations. Agents can entertain any expectation induced by such a model, with different choices corresponding to different ways of constructing the factors. With $d = 2$, for instance, the agents can choose one factor to capture the movements in real variables and the other to capture movements in nominal variables. Or they can pick a factor to account for the long-run movements in the observables and one factor to account for the short-run movements. Or they can pick the factors by principal component analysis. The solution concept of CREE provides a systematic way of determining the equilibrium factors picked by the agents. In the rest of this section, I focus on the case where $d = 1$.

6.4 Calibration and the Solution Method

The model has two classes of parameters. The first class consists of structural parameters such as the discount factor β and the capital share in the Cobb–Douglas production function α . These parameters need to be estimated or calibrated by the economist analyzing the model. The second class consists of the parameters of the agents' model, namely, matrices A , B , Σ_ω , Σ_o . These are *not* free parameters that need to be estimated by the economist. Rather, they are determined endogenously in a CREE.

The parameters of the agents' model are not identified without further normalization. I ensure that the agents do not face any identification problem by normalizing A to be a diagonal matrix with diagonal elements between zero and one and normalizing Σ_ω to be a correlation matrix. These normalizations do not introduce any restrictions on the set Θ of models considered by the agents. I additionally assume that Σ_o is a multiple of the identity matrix I . That is, the agents believe the measurement errors to be i.i.d. across different observables. This assumption does restrict the set Θ . But doing so significantly reduces the number of parameters the agents need to estimate in the CREE- d , thus making the model more tractable. It also allows me to include variables in the vector of observables o_t that are co-linear in equilibrium. This

normalization leaves the agents with $d + dn + d(d - 1)/2 = 15$ parameters to estimate. It is important to note that, just as in REE, the agents estimate their model using an infinitely large sample generated by the model economy—and not the data used by the economist to calibrate the structural parameters.

The algorithm that finds a CREE- d consists of an inner loop and an outer loop. In the inner loop, the distribution of observables is fixed. The agents estimate the parameters of their model by minimizing the Kullback–Leibler divergence from the distribution of observables. The inner loop uses IPOPT, an interior-point optimization software, to minimize the Kullback–Leibler divergence. To speed up the computation, exact first and second derivatives are provided to the algorithm using an *automatic differentiation* software. The outer loop is a standard fixed-point iteration loop. It iterates on the parameters of the agents’ model and the distribution of observables until a fixed point corresponding to a CREE- d is reached.

The structural parameters of the model can in principle be estimated using standard techniques. Although CREE- d is more tractable than models of rational inattention, the computational cost of estimating a CREE- d is still hundreds of times higher than that of a rational-expectations equilibrium. So as a first pass, I calibrate the structural parameters of the model to standard values from the new-Keynesian and DSGE literatures. Doing so has the additional benefit of not giving CREE- d any advantage over the REE in matching the business-cycle moments.⁴⁴

The model is calibrated at a quarterly frequency. Following the textbook of Galí (2008), I set the discount factor to $\beta = 0.99$, the capital share in the production function to $\alpha = 1/3$, the inverse of the Frisch elasticity of labor supply to $\nu = 1$, the elasticity of substitution between intermediate goods to 6—this corresponds to a steady state price markup of $\lambda_p = 0.20$ —the elasticity of substitution between differentiated labor types to 4.5—corresponding to a steady state wage markup of about $\lambda_w = 0.29$ —the Calvo parameter for prices to $\xi_p = 2/3$, the Calvo parameter for wages to $\xi_w = 3/4$. The quarterly depreciation rate of capital is set to $\delta = 0.025$, the value commonly used in the literature. Two parameters of the model have no counterparts in the DSGE literature. The public debt-to-GDP ratio is irrelevant in rational-expectations models where Ricardian equivalence holds. I set it to $b/g = 0.50$. This is roughly equal to the corresponding value for the US economy in the period 1950–2007. The convexity of the capital-adjustment cost is set to $S'' = 20$, following the original estimate of Hayashi (1982). This value is on the higher end of the spectrum but within the range of values considered in the literature.⁴⁵ The remaining parameters are set to the state-of-the-art DSGE estimates obtained by Sala (2015) by performing Bayesian estimation in the frequency domain. Table 1 lists the values of all the structural parameters. Note that there are no parameters for the agents’ expectations.

⁴⁴Examples of estimated DSGE models with learning include Milani (2007) and Slobodyan and Wouters (2012).

⁴⁵See Cooper and Haltiwanger (2006) for a discussion of different estimates of the capital-adjustment cost.

parameter	description	value	source
<u>Preferences, technology, nominal rigidities</u>			
β	discount rate	0.99	Gali (2008)
ν	inverse Frisch elasticity of labor supply	1.0	Gali (2008)
α	capital share	1/3	Gali (2008)
ξ_p	price rigidity	2/3	Gali (2008)
ξ_w	wage rigidity	3/4	Gali (2008)
λ_p	steady state price markup	0.20	Gali (2008)
λ_w	steady state wage markup	0.29	Gali (2008)
γ	steady state growth rate	1.01	Sala (2015)
δ	depreciation rate	0.025	Sala (2015)
χ	convexity of capital adjustment cost	20	Hayashi (1982)
g/y	government spending as share of GDP	0.20	Sala (2015)
b/y	public debt to GDP ratio	0.50	FRED
<u>Monetary policy</u>			
ρ_R	persistence of monetary policy	0.80	Sala (2015)
ϕ_π	Taylor rule coefficient on inflation	2.05	Sala (2015)
<u>Persistence of shocks</u>			
ρ_z	TFP	0.21	Sala (2015)
ρ_ψ	discount-factor shock	0.67	Sala (2015)
ρ_p	price-markup shock	0.65	Sala (2015)
ρ_w	wage-markup shock	0.33	Sala (2015)
ρ_m	monetary-policy shock	0.30	Sala (2015)
ρ_g	government-spending shock	0.93	Sala (2015)
<u>Standard deviation of shocks</u>			
σ_z	TFP	0.90	Sala (2015)
σ_ψ	discount-factor shock	0.30	Sala (2015)
σ_p	price-markup shock	0.08	Sala (2015)
σ_w	wage-markup shock	0.13	Sala (2015)
σ_m	monetary-policy shock	0.23	Sala (2015)
σ_g	government-spending shock	0.40	Sala (2015)

Table 1. Calibrated parameters

6.5 Results

Figure 1 illustrates the *endogenous* vector of Kalman gains. A longer bar for an observable signifies that the agents' estimate of the hidden factor is more sensitive to changes in that observable. Note that only the *relative* magnitudes of the bars are meaningful—the hidden factor can be scaled without changing the agents' forecasts of observables. The agents' estimate of the hidden factor increases the most with changes in investment, hours, the rental rate of capital, consumption, and income; it decreases the most with increases in taxes; and it is largely unaffected by changes in the value of the shocks, the real wage, the inflation rate, and the nominal interest rate. The agents' estimate of the hidden factor is a summary of the values of real aggregate variables; it captures the agents' unidimensional view of the state of the economy. As such, I refer to the agents' estimate $\hat{\omega}_t$ as “consumer confidence.”⁴⁶ Note that no data on expectations have been used to construct this measure. Rather, it emerges as agents try to find the hidden factor that best describes their equilibrium observations.

Figures 2–7 plot the response of the economy to shocks in the REE and the CREE-1. The figures reveal several noteworthy facts about how the economy responds to shocks. First, unlike in the REE, in the CREE-1, shocks to TFP and the monetary policy lead to hump-shaped responses of real variables. The consumer confidence also exhibits a hump-shaped impulse response to TFP and monetary-policy shocks. Second, a positive shock to the discount factor leads to a simultaneous increase in income, output, consumption, hours, and investment in the CREE-1. In contrast, in the REE, investment falls. Third, an increase in government spending leads to a simultaneous increase in income, output, consumption, hours, and investment in the CREE-1. In the REE, in contrast, consumption falls following a positive government-spending shock. Therefore, neither of the demand shocks of the economy can generate business-cycle-like movements in the real variables in the REE, but they both do so in the CREE-1. Fourth, the response of the economy to a government-spending shock is more persistent in the CREE-1 than in the REE. Finally, the CREE-1 dampens the response of the economy to a wage-markup shock relative to the REE benchmark.

7 Conclusion

This paper proposes model misspecification as a unified expression of bounded rationality in macroeconomics. It argues that a number of models of bounded rationality previously studied in the literature can be represented as particular cases of model misspecification.

I propose the solution concept of CREE as the generalization of the rational-expectations equilibrium to the case where agents' models are constrained to a misspecified set. In a CREE, agents may not have the correct model of the economy, but they do the best they can given

⁴⁶Ilut and Schneider (2014) and Angeletos et al. (2018) also construct endogenous measures of consumer confidence.

their misspecified models and their observations. A CREE exists under weak conditions. It is reached in the limit as agents learn about their environment. It incorporates a version of the Lucas critique and is well-suited for counterfactual policy analysis.

I use a special case of CREE to study a business-cycle model in which agents can only entertain factor models with a small number of factors. The calibrated economy exhibits hump-shaped impulse responses to shocks and demand-driven co-movements in aggregate variables. The model also gives rise to an endogenous measure of consumer confidence with striking resemblance to measures of consumer confidence constructed using survey expectations.

Constrained rationality emerges from the quantitative exercise as a parsimonious and plausible alternative to the battery of frictions needed in the business-cycle literature to improve the empirical fit of standard models. The exercise illustrates the portability and tractability of the solution concept of CREE. It can be readily incorporated in existing macro models to enrich the dynamics of agents' expectations and improve the realism and empirical fit of our models.




Figure 1. The vector of Kalman gains

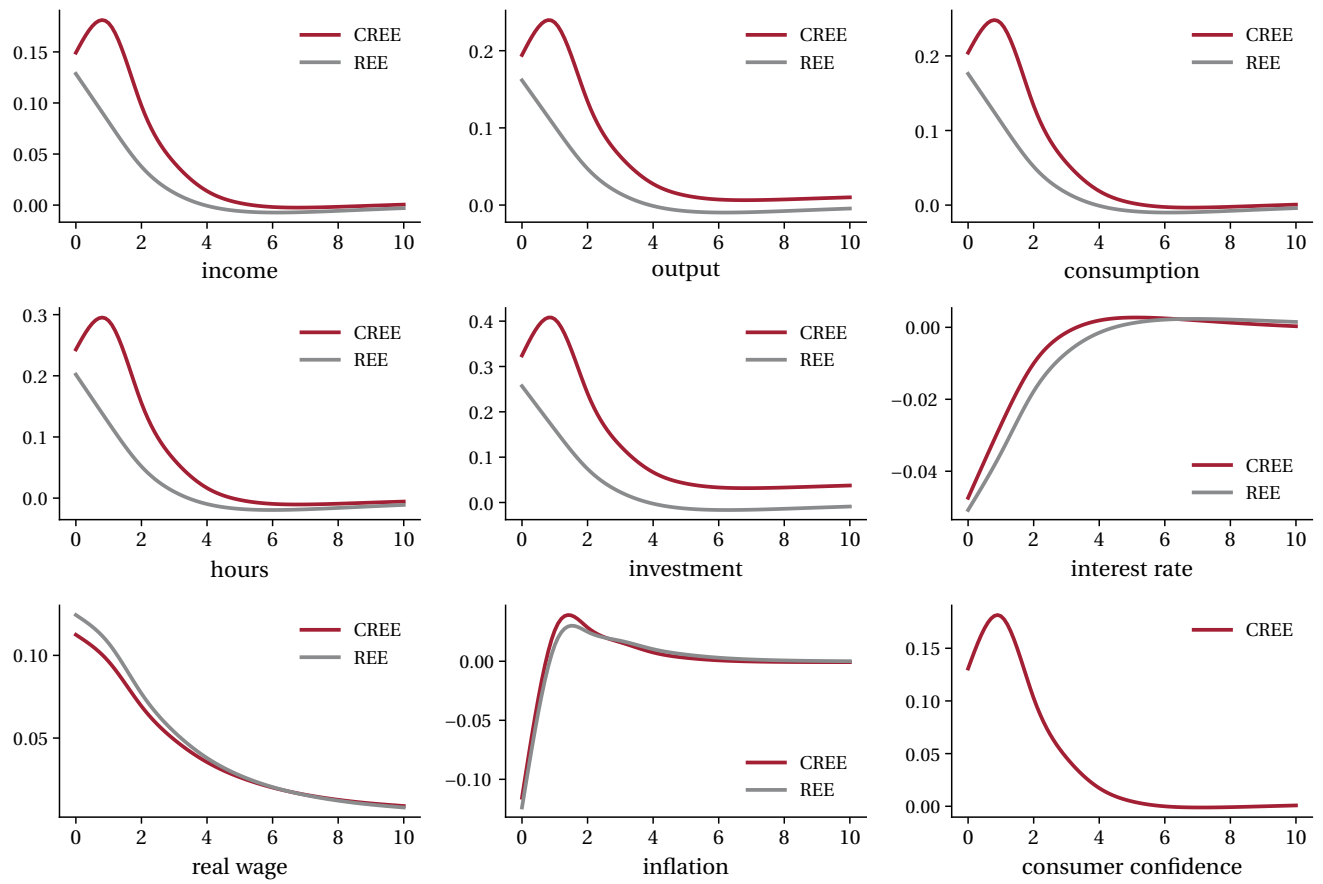


Figure 2. Impulse responses to the TFP shock

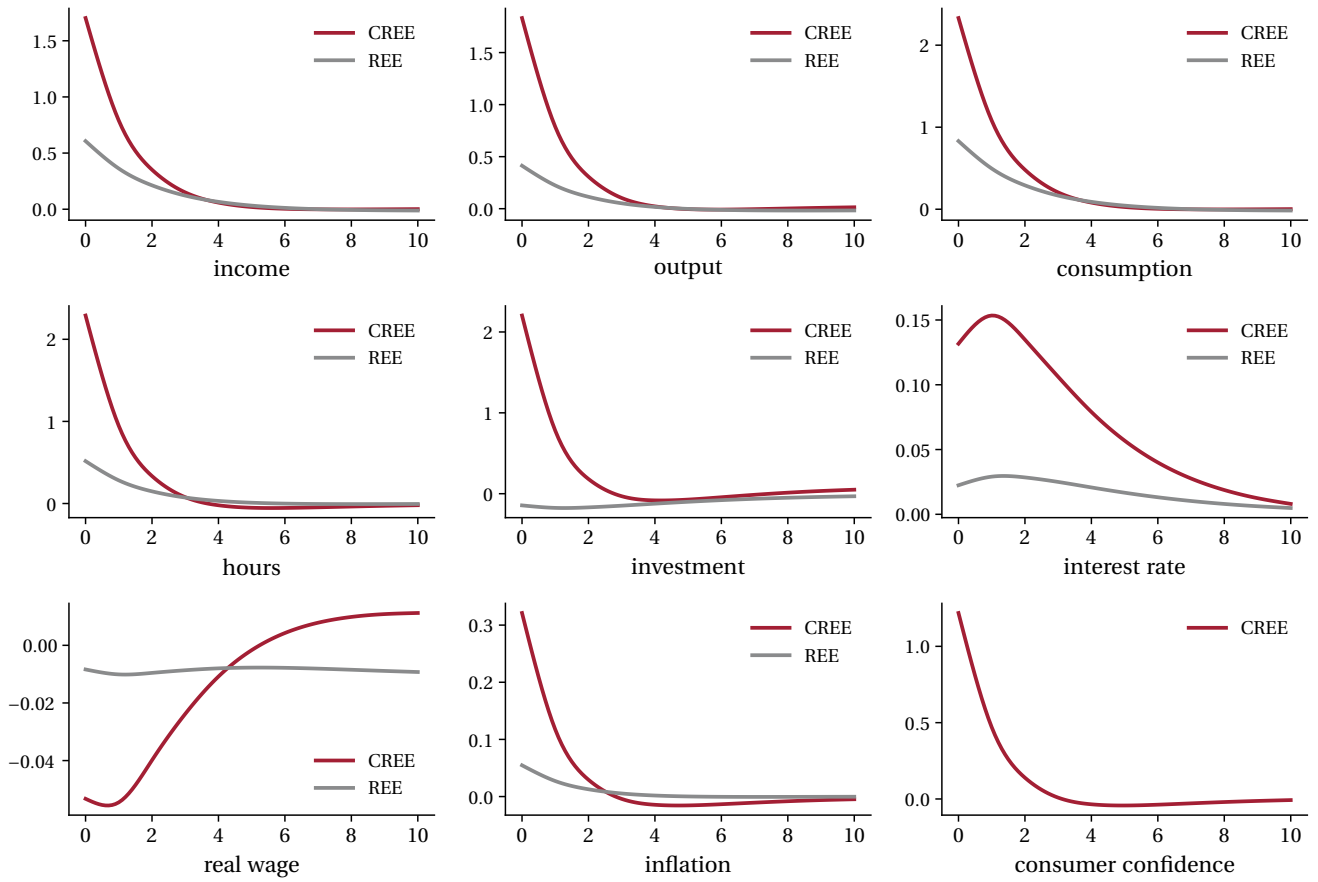


Figure 3. Impulse responses to the discount-factor shock

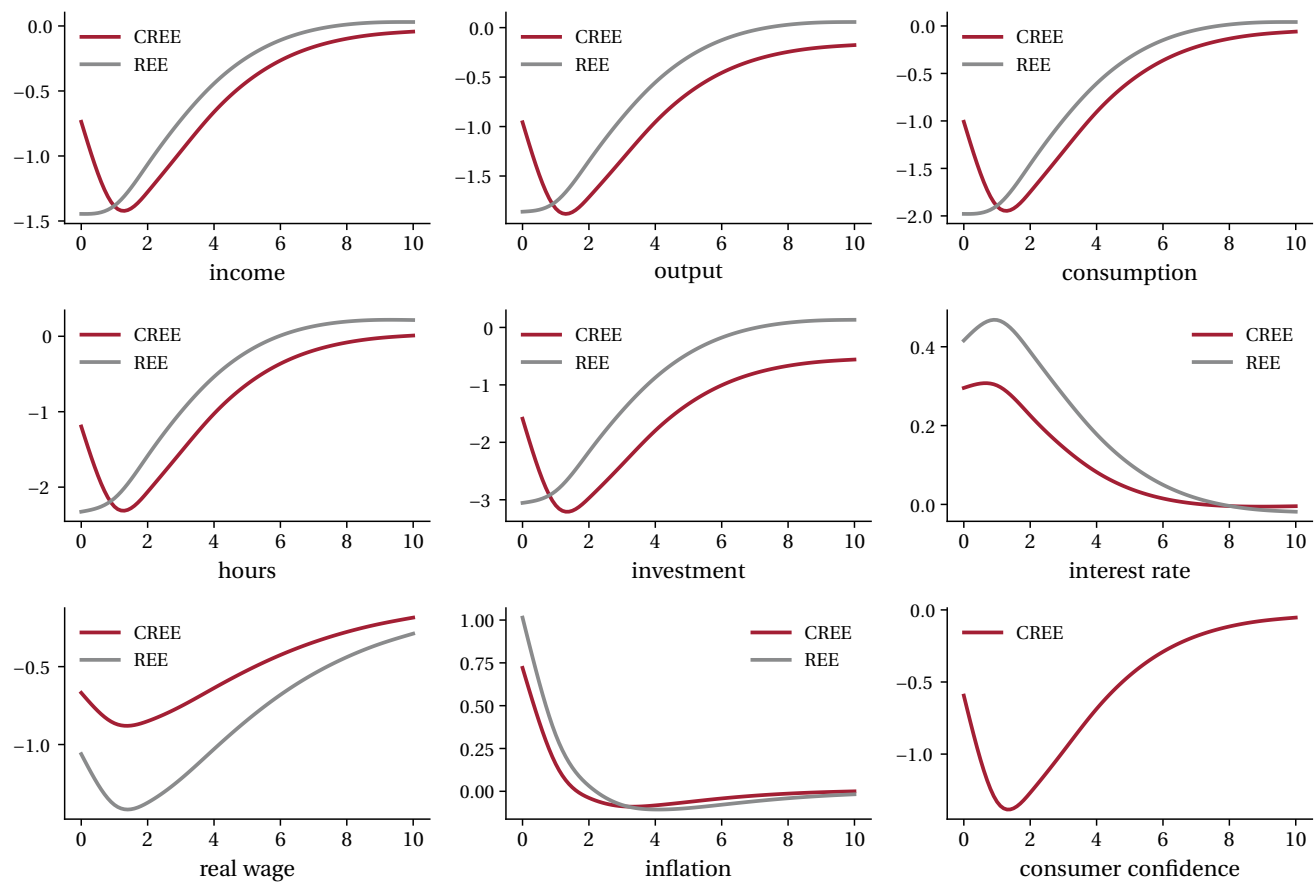


Figure 4. Impulse responses to the price-markup shock

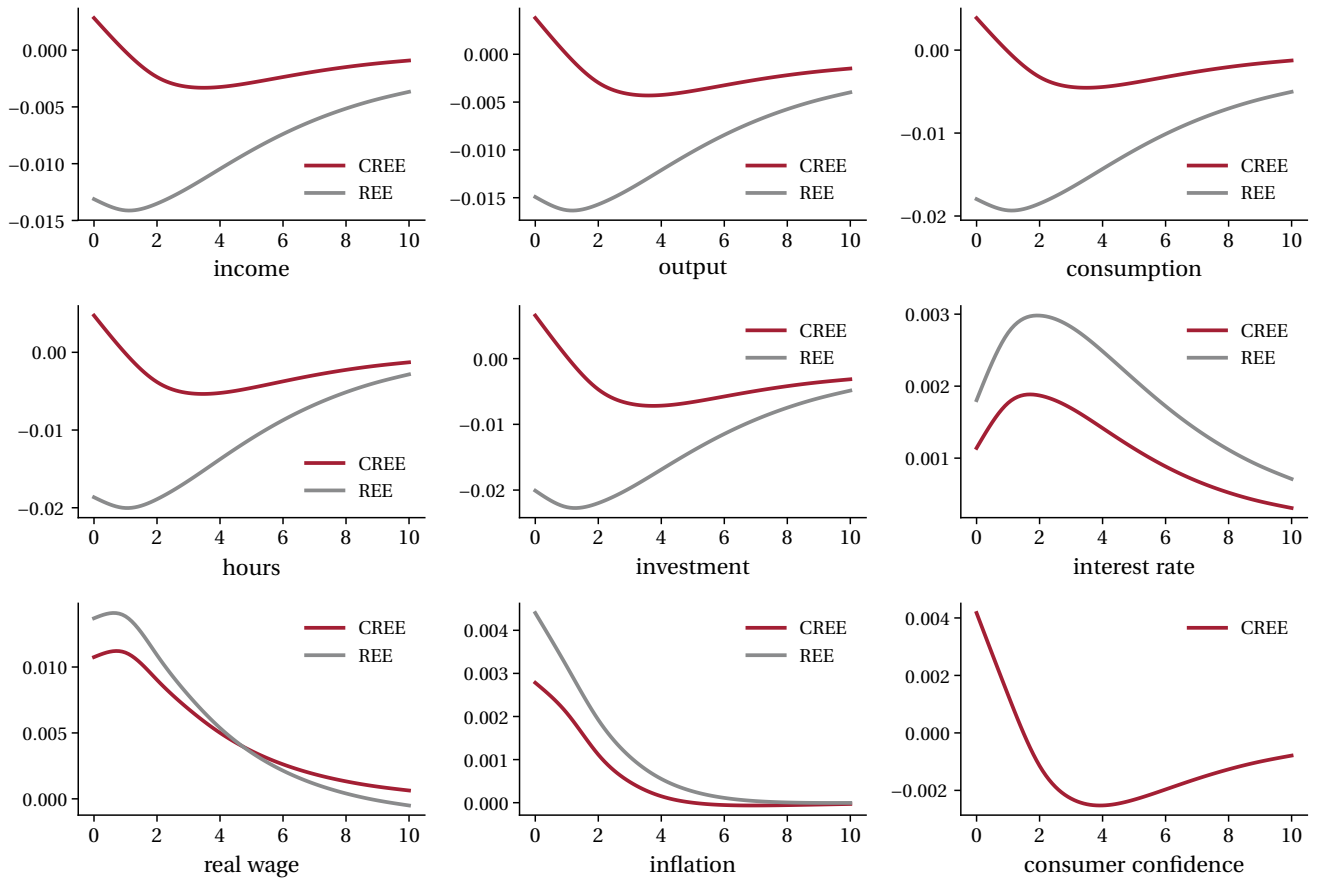


Figure 5. Impulse responses to the wage-markup shock




Figure 6. Impulse responses to the monetary-policy shock




Figure 7. Impulse responses to the government-spending shock

A Examples of the Economic Environment

In this appendix, I show how several benchmark macro models can be mapped to the abstract general-equilibrium economy introduced in Section 2.

A.1 The Lucas Asset-Pricing Model

Consider a representative-agent asset-pricing model with a single perishable consumption good in each period and an infinitely-lived Lucas tree. The representative agent has standard time-separable preferences with the flow utility in period t given by $u(c_t, \psi_t)$, where c_t denotes consumption and $\psi_t \in \mathbb{R}$ is a taste shifter. The tree is in fixed unit supply and yields a random dividend d_t in period t . The representative agent also receives an exogenously given stream of labor income, with the realized value of labor income in period t denoted by w_t . Let p_t denote the price of the tree in period t (in terms of the consumption good), and let s_t denote the agent's holding of the tree. I let $x_t = (c_t, s_t)'$ and $y_t = (\psi_t, d_t, w_t, p_t)'$. The constraints on the choice of the agent are represented by the correspondence Γ , defined as

$$\Gamma(x_{t-1}, y_t) = \{x_t : c_t + s_t p_t \leq s_{t-1}(p_t + d_t) + w_t\} \cap \{x_t : \underline{s} \leq s_t \leq \bar{s}\},$$

where \underline{s} and \bar{s} are lower and upper bounds on the agent's holding of the tree. The state variable z is an underlying shock that follows a Markov process with Markov kernel Π ; that is, z_t is distributed according to $\Pi(\cdot | z_{t-1})$ conditional on the history of shocks $\{z_s\}_{s=0}^{t-1}$. Note that since z is allowed to belong to a general Borel space Z , it can be used to represent any stationary ARMA process. The function $G : X \times Y \times Z \rightarrow \mathbb{R}^4$ is a vector-valued function that is given by:

$$\begin{aligned} G_1(x_t, y_t, z_t) &= \psi_t - \Psi(z_t), \\ G_2(x_t, y_t, z_t) &= d_t - D(z_t), \\ G_3(x_t, y_t, z_t) &= w_t - W(z_t), \\ G_4(x_t, y_t, z_t) &= s_t - 1, \end{aligned}$$

where Ψ , D , and W are functions that map the underlying shock to the values of the taste shifter, dividend, and labor income. The last element of G is used to represent the market clearing condition for the tree. The good market clears by Walras' law.

A.2 A Heterogeneous-Agent New-Keynesian Model

As a richer example, consider a heterogeneous-agent new-Keynesian model à la [Werning \(2015\)](#) in which the monetary authority controls the real interest rate and output is determined by aggregate demand.

There is a finite set of household types indexed by $i \in I$, with type i representing fraction α_i of the population. Types may differ in their flow utility functions, labor income, borrowing

constraints, and their models of the economy, but I assume different types to all have the same discount factor β . The flow utility function of households of type i is given by $u_i(c_{it}, \psi_{it})$, where c_{it} denotes consumption and $\psi_{it} \in \mathbb{R}$ is a taste shifter. Households of type i face the following budget constraints:

$$c_{it} + s_{it}p_t + \frac{b_{it}}{R_t} \leq s_{i,t-1}(p_t + d_t) + w_{it} + b_{i,t-1},$$

where s_{it} denotes the household's holding of a Lucas tree, p_t and d_t denote the price and dividend of the tree, respectively, b_{it} denotes the holding of a riskless one-period bond, R_t denotes the real interest rate paid on the bond between periods t and $t + 1$, and w_{it} denotes the labor income. I assume that households of a given type all have identical initial holdings of the tree and the riskless bond. Agents of type i are also subject to the following borrowing constraints:

$$b_{it} + s_{it}p_t \geq \underline{b}_{it},$$

where \underline{b}_{it} is a borrowing limit. I let C_t denote the aggregate output in the economy.

The information available to a household of type i is represented by some vector y_t^i that contains a subset of the elements of the vector

$$\left(\left(c_{jt}^i, s_{jt}^i, b_{jt}^i \right)_{j \neq i}, \left(\psi_{jt}^i, w_{jt}^i, \underline{b}_{jt}^i \right)_{j \in I}, d_t^i, p_t^i, R_t^i, C_t^i \right),$$

with the superscript i used to indicate that a variable is a ‘‘copy’’ that appears in the information set of households of type i . Copies of variables that directly enter the households' optimization problems (i.e., ψ_{it}^i , w_{it}^i , \underline{b}_{it}^i , d_t^i , p_t^i , and R_t^i) are always assumed to be included in y_t^i , and thus, to be observable to households of type i . The model considered by the households of type i is represented by a parameter set Θ_i and a mapping Q_i , where, for every $\theta_i \in \Theta_i$, $Q_{i\theta_i}$ is a Markov kernel over the space of all y^i with density $q_{i\theta_i}$ (with respect to the Lebesgue measure).

I can equivalently represent the decisions of households as being made by a representative agent who cares about all households but cannot transfer resources or information between households. Let $x_{it} = (c_{it}, s_{it}, b_{it})$, $x_t = (x_{it})_{i \in I}$, and $y_t = (y_t^i)_{i \in I}$. The representative agent's flow utility function is given by

$$u(x_t, y_t) = \sum_{i \in I} u_i(c_{it}, \psi_{it}^i).$$

Since the representative agent is unable to transfer resources from one household type to another, the weight assigned to different household types in the representative agent's utility function is irrelevant; I have simply chosen a utility function that weights all household types equally. The constraints faced by the representative agent are represented by the correspondence Γ defined as

$$\Gamma(x_{t-1}, y_t) = \{x_t : x_{it} \in \Gamma_i(x_{i,t-1}, y_t^i)\},$$

where

$$\Gamma_i(x_{i,t-1}, y_t^i) = \{x_{it} : c_{it} + s_{it}p_t^i + b_{it} \leq s_{i,t-1}(p_t^i + d_t^i) + w_{it}^i + R_{t-1}^i b_{i,t-1}\} \cap \{x_{it} : b_{it} + s_{it}p_t^i \geq \underline{b}_{it}^i\}.$$

The representative agent's model of the economy is given by the parameter space $\Theta = \prod_{i \in I} \Theta_i$ and the mapping $\theta \mapsto Q_\theta$, for $\theta = (\theta_i)_{i \in I}$, for which the density q_θ is given by

$$q_\theta(y_t | y_{t-1}) = \prod_{i \in I} q_{i\theta_i}(y_t^i | y_{t-1}^i).$$

The product structure entails that, from the point of view of the representative agent, realizations of y^i are statistically independent of those of y^j for all $j \neq i$ —although in equilibrium y_t^i and y_t^j include copies of the same set of variables, and so, are dependent. This is the manifestation of the assumption that households of different types do not share information with one another and thus face separate inference problems.

Finally, the state variable z is an underlying shock that belongs to a general Borel space Z and follows a Markov process with Markov kernel Π . The taste shifters, dividends, labor incomes, and borrowing constraints are all random variables that depend on the realized value of the shock. I assume that R_t , the real interest rate at time t , is also a random variable that only depends on z_t . In other words, the real interest rate is determined by the monetary authority as a function of monetary policy shocks that are included in z and whose values are determined outside of the model.

The consistency requirements are imposed by means of a function $G(x_t, y_t, z_t)$ that needs to be equal to zero for all t . I describe the function G simply by enumerating the expressions that are required to be equal to zero in any equilibrium. Consistency of different copies of choice and endogenous variables with one another and with the realized value of the shock requires

$$\begin{aligned} \psi_{jt}^i - \Psi_j(z_t) &= 0 & \forall i, j \in I, \\ w_{jt}^i - W_j(C_t, z_t) &= 0 & \forall i, j \in I, \\ \underline{b}_{jt}^i - B_j(C_t, z_t) &= 0 & \forall i, j \in I, \\ d_t^i - D(C_t) &= 0 & \forall i \in I, \\ R_t^i - R(z_t) &= 0 & \forall i \in I, \\ c_{jt}^i - c_{jt} &= 0 & \forall i, j \in I, \\ s_{jt}^i - s_{jt} &= 0 & \forall i, j \in I, \\ b_{jt}^i - b_{jt} &= 0 & \forall i, j \in I, \end{aligned}$$

where Ψ_j , W_j , B_j , D , and R are functions. By definition, aggregate income, C_t , must satisfy

$$\sum_{i \in I} \alpha_i W_i(C_t, z_t) + D(C_t) - C_t = 0.$$

And market clearing for the consumption good, the Lucas tree, and bonds require that

$$\sum_{i \in I} \alpha_i c_{it} - C_t = 0,$$

$$\sum_{i \in I} \alpha_i s_{it} - 1 = 0,$$

$$\sum_{i \in I} \alpha_i b_{it} = 0.$$

This completes the mapping of the economy to the general framework of Section 2.

B Hidden Factor Models and the Kalman Filter

The expressions presented in this section are instances of standard results on Kalman filtering that can be found, among other places, in [Anderson and Moore \(2005\)](#) and [Hamilton \(1994, ch. 13\)](#). In the first subsection, I characterize the Kullback–Leibler divergence in the general case. In the second subsection, I consider a special case where agents consider one-dimensional factor models with i.i.d. measurement errors and simplify the Kullback–Leibler divergence.

B.1 The General Case

Agents are assumed to have a linear-Gaussian hidden-factor model of the economy that is described by the following equations:

$$\omega_t = A\omega_{t-1} + \epsilon_{\omega t},$$

$$o_t = B'\omega_t + \epsilon_{o t},$$

where $\omega_t, \epsilon_{\omega t} \in \mathbb{R}^d$, $o_t, \epsilon_{o t} \in \mathbb{R}^n$, $A \in \mathbb{R}^{d \times d}$, $B \in \mathbb{R}^{d \times n}$, $\epsilon_{\omega t}$ is i.i.d. $\mathcal{N}(0, \Sigma_\omega)$, and $\epsilon_{o t}$ is i.i.d. $\mathcal{N}(0, \Sigma_o)$. All of the eigenvalues of matrix A are assumed to be inside the unit circle.

Conditional on $\{o_s\}_{s=-\infty}^t$, agents believe that the hidden state variable ω_{t+1} is normally distributed with mean $\hat{\omega}_t \equiv \hat{\omega}_{t+1|t}$ and variance-covariance matrix $\hat{\Sigma}_\omega$, where $\hat{\Sigma}_\omega$ is the unique positive semidefinite symmetric matrix that satisfies the following algebraic Riccati equation

$$\hat{\Sigma}_\omega = A \left(\hat{\Sigma}_\omega - \hat{\Sigma}_\omega B (B' \hat{\Sigma}_\omega B + \Sigma_o)^{-1} B' \hat{\Sigma}_\omega \right) A' + \Sigma_\omega, \quad (\text{B.1})$$

$\hat{\omega}_t$ is defined recursively as

$$\hat{\omega}_t = (A - KB')\hat{\omega}_{t-1} + K o_t, \quad (\text{B.2})$$

$K \in \mathbb{R}^{d \times n}$ is the Kalman gain defined as

$$K \equiv A \hat{\Sigma}_\omega B (B' \hat{\Sigma}_\omega B + \Sigma_o)^{-1}, \quad (\text{B.3})$$

and A' denotes the transpose of matrix A .⁴⁷

Since the eigenvalues of A are strictly smaller than 1 in magnitude by assumption, equation (B.2) can be solved backward to get $\hat{\omega}_t = \sum_{s=0}^{\infty} (A - KB')^s K o_{t-s}$.⁴⁸ Thus, conditional on $\{o_s\}_{s=-\infty}^t$,

⁴⁷See, for instance, equations (1.2), (4.2), and (4.4) in Chapter 4 of [Anderson and Moore \(2005\)](#).

⁴⁸See, for instance, [Anderson and Moore \(2005, p. 77\)](#).

agents believe that the observable o_{t+1} is normally distributed with mean $\hat{o}_{t+1|t} = B'\hat{\omega}_t$ and variance-covariance matrix

$$\Omega \equiv B'\hat{\Sigma}_\omega B + \Sigma_o. \quad (\text{B.4})$$

More generally, conditional on $y_t = (o_t, o_{t-1}, \dots)$, agents believe that the mean of o_{t+s} is given by

$$\hat{o}_{t+s|t} = B'A^{s-1}\hat{\omega}_t = B'A^{s-1} \sum_{\tau=0}^{\infty} (A - KB')^\tau K o_{t-\tau} \quad (\text{B.5})$$

for any $s \geq 1$.

The agents' model can be cast in the canonical form of Section 2 by defining $y_t \equiv (o_t, o_{t-1}, \dots)$. The vector y_t consists of the entire history of realizations of the vector of observables o_s for $s \leq t$. According to the agents' model, the sequence $\{y_t\}$ is a (first-order) time-homogeneous Markov chain. In particular, conditional on $y_t = (o_t, o_{t-1}, \dots)$, agents believe that $y_{t+1} = (o_{t+1}, y_t)$, where o_{t+1} is normally distributed with mean $\hat{o}_{t+1|t} = \sum_{s=0}^{\infty} B'(A - KB')^s K o_{t-s}$ and variance-covariance matrix Ω .

The agents' model can be described by a parameter set Θ and a mapping Q . Let $\tilde{\Theta}$ denote the set of all tuples $\theta = (A, B, \Sigma_\omega, \Sigma_o)$ such that A has all of its eigenvalues in the unit circle and Σ_ω and Σ_o are variance-covariance matrices. The parameter set Θ is a (possibly strict) subset of $\tilde{\Theta}$. Any parameter $\theta \in \Theta$ defines a Markov kernel Q_θ over $Y \equiv \mathbb{R}^N$ as follows. Given any $y_t = (o_t, o_{t-1}, \dots)$, let ν_{y_t} denote the measure over Y , defined by setting, for all measurable sets $E_1 \subseteq \mathbb{R}^n$ and $E_2 \subseteq Y$,

$$\nu_{y_t}(E_1 \times E_2) = \tilde{\nu}(E_1) \mathbb{1}\{y_t \in E_2\},$$

where $\tilde{\nu}$ denotes the n -dimensional Lebesgue measure. $Q_\theta(\cdot|y_t)$ is absolutely continuous with respect to the sigma-finite measure ν_{y_t} with the log-density given by

$$\begin{aligned} \log(q_\theta(y_{t+1}|y_t)) &= -\frac{1}{2} \log(\det(\Omega)) - \frac{n}{2} \log(2\pi) \\ &\quad - \frac{1}{2} \left(o_{t+1} - \sum_{s=0}^{\infty} B'(A - KB')^s K o_{t-s} \right)' \Omega^{-1} \left(o_{t+1} - \sum_{s=0}^{\infty} B'(A - KB')^s K o_{t-s} \right), \end{aligned}$$

where K and Ω are defined in equations (B.3) and (B.4), respectively, and $\det(\Omega)$ denotes the determinant of matrix Ω .

The Kullback–Leibler divergence of model $\theta = (A, B, \Sigma_\omega, \Sigma_o)$ from a distribution \mathbb{P} for the path $\{o_t\}_{t=-\infty}^{\infty}$ of observables is equal to

$$\begin{aligned} H(\mathbb{P}, \theta) &= -\frac{1}{2} \log(\det(\Omega^{-1})) + \frac{n}{2} \log(2\pi) \\ &\quad + \frac{1}{2} \text{tr} \left(\Omega^{-1} \mathbb{E} \left[\left(o_{t+1} - \sum_{s=0}^{\infty} B'(A - KB')^s K o_{t-s} \right) \left(o_{t+1} - \sum_{s=0}^{\infty} B'(A - KB')^s K o_{t-s} \right)' \right] \right), \quad (\text{B.6}) \end{aligned}$$

where \mathbb{E} denotes the expectation with respect to \mathbb{P} and tr denotes the trace of a matrix (i.e., the sum of its diagonal elements). Expanding the last term in equation (B.6), I get

$$\begin{aligned} H(\mathbb{P}, \theta) = & -\frac{1}{2} \log \left(\det \left(\Omega^{-1} \right) \right) + \frac{n}{2} \log(2\pi) + \frac{1}{2} \text{tr} \left(\Omega^{-1} \Xi_0 \right) \\ & - \sum_{s=1}^{\infty} \text{tr} \left(\Omega^{-1} \Xi_s K' (A' - BK')^{s-1} B \right) \\ & + \frac{1}{2} \sum_{s=1}^{\infty} \sum_{\tau=1}^{\infty} \text{tr} \left(\Omega^{-1} B' (A - KB')^{s-1} K \Xi_{\tau-s} K' (A' - BK')^{\tau-1} B \right), \end{aligned} \quad (\text{B.7})$$

where I am using the linearity of the trace operator and

$$\Xi_s \equiv \mathbb{E} \left[o_t o_{t-s}' \right]$$

is the autocovariance matrix of vector o_t at lag $s \in \mathbb{Z}$.

The calculation in the paragraph establishes that the Kullback–Leibler divergence given the distribution \mathbb{P} for the sequence $\{o_t\}_{t=-\infty}^{\infty}$ of observables only depends on the autocovariance matrices of o_t at different lags. With some abuse of notation, I let $H(\{\Xi_s\}_{s \in \mathbb{Z}}, \theta)$ denote the Kullback–Leibler divergence given the sequence of autocovariance matrices $\{\Xi_s\}_{j \in \mathbb{Z}}$.

B.2 The One-Factor Model

I next consider a special case of the hidden-factor model of the previous subsection. Agents believe that the hidden factor is one dimensional and the measurement errors are independent and identically distributed across observables. The matrix A is then equal to some scalar $a \in [-1, 1]$ and the matrix Σ_o is equal to $\sigma_o^2 I$ for some $\sigma_o > 0$, where I denotes the identity matrix. Recall by equation (B.7) that the Kullback–Leibler divergence is given by

$$\begin{aligned} H(\mathbb{P}, \theta) = & -\frac{1}{2} \log \det \left(\Omega^{-1} \right) + \frac{n}{2} \log(2\pi) + \frac{1}{2} \text{tr} \left(\Omega^{-1} \Xi_0 \right) \\ & - \sum_{s=1}^{\infty} \text{tr} \left(\Omega^{-1} \Xi_s K' (A' - BK')^{s-1} B \right) \\ & + \frac{1}{2} \sum_{s=1}^{\infty} \sum_{\tau=1}^{\infty} \text{tr} \left(\Omega^{-1} B' (A - KB')^{s-1} K \Xi_{\tau-s} K' (A' - BK')^{\tau-1} B \right). \end{aligned}$$

When the hidden factor is one dimensional, the above expression can be simplified further. Define

$$\begin{aligned} M \equiv & D' \Xi_{\zeta_0} D - 2D' (I - (A - KB') A_{\zeta})^{-1} A_{\zeta} \Xi_{\zeta_0} D K' B + \left(1 - (A - KB')^2 \right)^{-1} B' K D' \Xi_{\zeta_0} D K' B \\ & + 2(A - KB') \left(1 - (A - KB')^2 \right)^{-1} B' K D' (I - (A - KB') A_{\zeta})^{-1} A_{\zeta} \Xi_{\zeta_0} D K' B. \end{aligned}$$

Then,

$$H(\mathbb{P}, \theta) = \frac{1}{2} \log \det(\Omega) + \frac{n}{2} \log(2\pi) + \frac{1}{2} \text{tr} \left(\Omega^{-1} M \right).$$

Without loss of generality I can normalize the hidden factor to have $\hat{\Sigma}_\omega = 1$. The matrix Ω then simplifies to

$$\Omega = B'B + \sigma_o^2 I.$$

The matrix-determinant lemma implies that

$$\log \det(\Omega) = \log \left(1 + B\Sigma_o^{-1}B' \right) + 2n \log(\sigma_o),$$

where recall that n denotes the number of observables. Define

$$v \equiv \frac{1}{\sigma_o} B'.$$

Then,

$$\frac{1}{2} \log \det(\Omega) = \frac{1}{2} \log(1 + v'v) + n \log(\sigma_o).$$

The Sherman–Morrison formula implies that

$$\Omega^{-1} = \Sigma_o^{-1} - \frac{\Sigma_o^{-1}B'B\Sigma_o^{-1}}{1 + B\Sigma_o^{-1}B'} = \frac{1}{\sigma_o^2} \left(I - \frac{vv'}{1 + v'v} \right).$$

Therefore, the Kullback–Leibler divergence can be written as

$$H(\mathbb{P}, \theta) = \frac{1}{2} \log(1 + v'v) + n \log(\sigma_o) + \frac{n}{2} \log(2\pi) + \frac{1}{2} \frac{1}{\sigma_o^2} \text{tr} \left(\left(I - \frac{vv'}{1 + v'v} \right) M \right).$$

The vector of Kalman gains is given by

$$K = AB\Omega^{-1} = \frac{av'}{\sigma_o} \left(I - \frac{vv'}{1 + v'v} \right) = \frac{av'}{\sigma_o(1 + v'v)}.$$

Therefore,

$$B'K = K'B = \frac{a}{1 + v'v} vv',$$

and

$$A - KB' = A' - BK' = \frac{a}{1 + v'v}.$$

Minimizing the Kullback–Leibler divergence with respect to a , B , and σ_o is equivalent to minimizing it with respect to a , v , and σ_o . The matrix M only depends on v . Therefore, the first-order condition with respect to σ_o is given by

$$n\sigma_o^{-1} - \sigma_o^{-3} \text{tr} \left(\left(I - \frac{vv'}{1 + v'v} \right) M \right) = 0.$$

It is easy to verify that the second-order condition is also satisfied. Solving the above equation for σ_o , I get

$$(\sigma_o^*)^2 = \frac{1}{n} \text{tr} \left(\left(I - \frac{vv'}{1 + v'v} \right) M \right).$$

Substituting in the Kullback–Leibler divergence, I get

$$H(\mathbb{P}, \theta) = \frac{1}{2} \log(1 + v'v) + \frac{n}{2} \log \operatorname{tr} \left(\left(I - \frac{vv'}{1 + v'v} \right) M \right) + \text{constant}.$$

Using the expression for M , the above expression can be simplified to

$$\begin{aligned} H(\mathbb{P}, \theta) &= \frac{1}{2} \log(1 + v'v) \\ &+ \frac{n}{2} \log \left(\operatorname{tr} (D' \Xi_{\zeta_0} D) - \frac{1 + v'v - a^2}{(1 + v'v)^2 - a^2} v' D' \left[I + 2 \left(I - \frac{a}{1 + v'v} A_{\zeta} \right)^{-1} \frac{a}{1 + v'v} A_{\zeta} \right] \Xi_{\zeta_0} D v \right) \\ &+ \text{constant}. \end{aligned}$$

Define

$$\begin{aligned} \gamma &= \frac{1}{1 + v'v}, \\ u &= \frac{v}{\sqrt{v'v}}. \end{aligned}$$

Then,

$$\begin{aligned} H(\mathbb{P}, \theta) &= -\frac{1}{2} \log(\gamma) + \frac{n}{2} \log \left(\operatorname{tr} (D' \Xi_{\zeta_0} D) - \frac{(\gamma^{-1} - a^2)(\gamma^{-1} - 1)}{\gamma^{-2} - a^2} u' D' \left[I + 2 (I - \gamma a A_{\zeta})^{-1} \gamma a A_{\zeta} \right] \Xi_{\zeta_0} D u \right) \\ &+ \text{constant}. \end{aligned}$$

Minimizing the Kullback–Leibler divergence with respect to a and v is equivalent to minimizing the following expression with respect to γ , a , and u subject to the constraints that $a, \gamma \in [0, 1]$ and $u'u = 1$:

$$-\log(\gamma) + n \log \left(\operatorname{tr} (D' \Xi_{\zeta_0} D) - \frac{(1 - a^2 \gamma)(1 - \gamma)}{1 - a^2 \gamma^2} u' D' \left[I + 2 (I - \gamma a A_{\zeta})^{-1} \gamma a A_{\zeta} \right] \Xi_{\zeta_0} D u \right).$$

This is the expression that is used in the quantitative exercise.

C Special Cases from the Literature on Bounded Rationality

In this appendix, I show how several boundedly-rational models of expectation formation can be viewed as imposing particular constraints on the set of models entertained by the agent. For the sake of exposition, in this appendix, I assume that the observable y_t follows an exogenous Markov process.⁴⁹

C.1 Covariance-Stationary VARs

Consider an agent who believes that the observable variable follows a covariance-stationary VAR(p) process with normal innovations. Let $\tilde{y}_t \in \mathbb{R}^l$ denote the finite-dimensional vector of

⁴⁹This can be done formally by letting $\Pi(\cdot|y_-, z_-) = \Pi(\cdot|z_-)$ for all y_-, z_- and $G(x, y, z) = y - z$; that is, the state variable z follows a Markov process and the observable is equal to the state variable at all times.

observables at time t . The agent's set of models is parametrized by the VAR coefficients $\theta_1, \theta_2, \dots, \theta_p \in \mathbb{R}^{l \times l}$ and the variance-covariance matrix of innovations, which I denote by $\theta_0 \in \mathbb{R}^{l \times l}$. Given parameter $\theta = (\theta_0, \theta_1, \dots, \theta_p)$, the probability that \tilde{y}_t belongs to a set $\tilde{B} \subseteq \mathbb{R}^l$ is given by

$$\tilde{Q}_\theta(\tilde{B}|\tilde{y}_{t-1}, \dots, \tilde{y}_{t-p}) = \Phi_{\theta_0}(\{\varepsilon : \theta_1 \tilde{y}_{t-1} + \dots + \theta_p \tilde{y}_{t-p} + \varepsilon \in \tilde{B}\}),$$

where ε denotes the innovation and Φ_{θ_0} denotes the l -dimensional normal probability distribution with mean zero and variance-covariance matrix θ_0 . I assume that the agent only considers parameters θ for which θ_0 belongs to some compact subset Θ_0 of the set of symmetric and positive-definite $l \times l$ matrices. This assumption guarantees that any θ_0 considered by the agent is a proper variance-covariance matrix.

The set of models considered by the agent can be expressed in the Markovian form of Section 2. Let $y_t = (\tilde{y}_t, \tilde{y}_{t-1}, \dots, \tilde{y}_{t-p+1}) \in Y \equiv \mathbb{R}^{l \times p}$, and let $B_r \subseteq \mathbb{R}^l$ be an arbitrary set for $r = 1, \dots, p$. Under model θ for the agent, the probability that y_t belongs to the set $B = B_1 \times \dots \times B_p$ conditional on $y_{t-1} = (\tilde{y}_{t-1}, \dots, \tilde{y}_{t-p})$ is given by

$$Q_\theta(B|y_{t-1}) = \tilde{Q}_\theta(B_1|\tilde{y}_{t-1}, \dots, \tilde{y}_{t-p}) \mathbb{1}\{(\tilde{y}_{t-1}, \dots, \tilde{y}_{t-p+1}) \in B_2 \times \dots \times B_p\},$$

where $\mathbb{1}$ denotes the indicator function. The expression above defines a Markov kernel over Y for any θ . The set of parameters is given by $\Theta = \Theta_0 \times \tilde{\Theta}$, where Θ_0 is the set of variance-covariance matrices specified above and $\tilde{\Theta}$ is the closure of the set of all tuples $(\theta_1, \dots, \theta_p)$ for which the VAR(p) process with autoregressive coefficients $\theta_1, \dots, \theta_p$ is covariance stationary.⁵⁰ The set Θ and the mapping $\theta \mapsto Q_\theta$ defined above constitute a well-defined set of models for the agent.

The agent's set of models may or may not contain the true data-generating process. When \tilde{y}_t can be represented as a VAR(p) process with normally distributed innovations, the agent's set of models is correctly specified; if \tilde{y}_t does not have a VAR(p) representation or if the innovations are not normally distributed, then the agent's set of models is misspecified.⁵¹

More generally, the agent's set of models is said to be misspecified whenever Θ does not contain a parameter θ for which the probability distribution induced by Q_θ coincides with the true data-generating process. In the rest of this appendix, I show that a number of deviations from the rational-expectations benchmark that have been previously proposed in the literature can be viewed as examples of misspecification by imposing appropriate constraints on the agent's set of models.

C.2 Restricted Perceptions Equilibrium

Suppose that the agent in the previous subsection only considers a strict subset Θ_{RP} of Θ to be plausible. This is the starting point of the literature on *restricted-perceptions equilibrium (RPE)*

⁵⁰The parameters on the boundary of $\tilde{\Theta}$ do not correspond to covariance stationary VAR processes. If the agent believes that the process for \tilde{y}_t is strictly covariance stationary, then one has to replace $\tilde{\Theta}$ with some compact subset contained in its interior.

⁵¹Note, however, that if the innovations are not normally distributed but the agent's optimal choices only depend on her beliefs about the first moment of observables, then misspecification does not affect the agent's optimal choices.

that goes back to [Bray \(1982\)](#).⁵² The set Θ_{RP} is used to express the analyst's a priori knowledge of the type of expectations that may be reasonably entertained by agents.

A particular case that is the focus of much of the RPE literature is one in which the agent does not make use of some elements of vectors $\tilde{y}_{t-1}, \dots, \tilde{y}_{t-p} \in \mathbb{R}^l$ when forecasting \tilde{y}_t . For the sake of argument, let me assume that $\tilde{y}_t = (\tilde{y}_{1t}, \dots, \tilde{y}_{lt})$ and that the agent believes the values of $\{\tilde{y}_{js}\}_{s < t}$ for $j \in J$ to have no predictive power for $(\tilde{y}_{kt})_{k \notin J}$. The observables $(\tilde{y}_{js})_{j \in J}$, may, for example, represent shocks (such as TFP) or state variables (such as the capital stock) that do not directly appear in the agent's optimization problem—even though they may affect other variables that do. The constraint on the agent's model can be expressed by setting

$$\Theta_{\text{RP}} = \left\{ \theta = (\theta_0, \dots, \theta_p) \in \Theta : [\theta_r]_{kj} = 0 \quad \forall k \notin J, \forall r = 1, \dots, p \right\},$$

where $[\theta_r]_{kj}$ denotes the element of matrix θ_r that is in row k and column j . The mapping $\theta \mapsto Q_\theta$ is simply the restriction of the mapping from the previous example to the set Θ_{RP} .

C.3 Extrapolative Expectations

Another special case of the restricted-perceptions equilibrium is the model of extrapolative expectations à la [Fuster, Laibson, and Mendel \(2010\)](#) and [Fuster, Hebert, and Laibson \(2012\)](#).⁵³ The assumption in those papers is that agents form their forecasts using fewer lags of the vector of observables than what is needed in a rational-expectations equilibrium. In particular, it is often assumed that agents use a single lag to form their expectations. Extrapolative expectations can be cast as a special case of the framework of Section 2 by constraining the set of parameters to belong to

$$\Theta_{\text{EE}} = \left\{ \theta = (\theta_0, \dots, \theta_p) \in \Theta : \theta_r = 0 \quad \forall r > 1 \right\},$$

and defining Q_θ to be the restriction of the mapping $\theta \mapsto Q_\theta$, defined in the context of the unrestricted VAR(p) model in subsection C.1, to the constrained set Θ_{EE} .

C.4 Consistent Expectations

In the consistent-expectations equilibrium of [Hommes and Sorger \(1998\)](#) and [Branch and McGough \(2005\)](#), agents' perceptions of economic variables are based on linear models, even though the true model of the economy may be nonlinear. Consistent perceptions are easily embedded in the general framework of Section 2. Suppose that the observable variable y_t belongs to some compact Banach space Y . Consistent perceptions are modeled by assuming that, according to the agent's model, the probability that the observable y_t belongs to set B is

⁵²Other notable contribution to the literature on RPE include [Bray and Savin \(1986\)](#) and [Branch \(2004\)](#). For more on the RPE solution concept, see Chapter 13 of [Evans and Honkapohja \(2012\)](#).

⁵³For related models and applications of extrapolative expectations, see [Hirshleifer \(2001\)](#), [Hirshleifer, Li, and Yu \(2015\)](#), and [Barberis, Greenwood, Jin, and Shleifer \(2015\)](#).

given by

$$Q_\theta(y_{t-1}, B) = P_{\theta_0}(\{\varepsilon : A_{\theta_1}y_{t-1} + \varepsilon \in B\}),$$

where P_{θ_0} is a parametric probability distribution for the shock ε , A_{θ_1} is a linear operator, which is parameterized by some parameter θ_1 , and $\theta = (\theta_0, \theta_1)$ is the parameter of the agent's model.⁵⁴

C.5 Analogy-Based Expectations, Cursedness, and Correlation Misperception

Another class of deviations from rational expectations are those capturing the misperception of correlations among observables. In the analogy-based-expectations equilibrium (ABEE) of [Jehiel \(2005\)](#) and cursed equilibrium of [Eyster and Rabin \(2005\)](#), players in a game are modeled as neglecting the dependence of their opponents' actions on their information. [Ellis and Piccione \(2017\)](#) generalize this idea by allowing for arbitrary patterns of correlation misperception, while [Spiegler \(2016\)](#) generalizes it by using Bayesian networks to model more general errors of statistical inference.⁵⁵

Correlation misperception on the part of agents can also be viewed as particular constraints on the agents' set of models.⁵⁶ For the sake of exposition, I simplify other aspects of the model by considering an agent who believes the vector of observables $y_t \in \mathbb{R}^l$ to be i.i.d. over time. By [Sklar \(1959\)](#)'s theorem, any distribution for y_t can be expressed in terms of marginal distributions of y_{tj} for $j = 1, \dots, l$ and a *copula* $C : [0, 1]^l \rightarrow [0, 1]$ that captures the dependence relationships among components of y_t .⁵⁷ Suppose that the agent believes that the marginal distribution of y_{tj} is given by Q_{θ_j} and the copula describing the dependence relationships is given by C_{θ_0} , where $\theta_0, \theta_1, \dots, \theta_l$ are parameters that belong to subsets $\Theta_0, \Theta_1, \dots, \Theta_l$ of Euclidean spaces.

This formulation allows me to restrict the dependence relationships among components of y_t without restricting the marginal distributions. In particular, it may be the case that, for all j , the set Θ_j contains a parameter θ_j for which Q_{θ_j} coincides with the true distribution of y_{tj} , while there exists no $\theta_0 \in \Theta_0$ for which the copula C_{θ_0} describes the true dependence relationship among elements of y_t . Fully-cursed expectations are modeled by assuming that Θ_0 is a singleton $\{\theta_0^*\}$ such that $C_{\theta_0^*}$ is the independence copula, i.e., $C_{\theta_0^*}(y_1, \dots, y_l) = \prod_{j=1}^l y_j$ for all $y = (y_1, \dots, y_l) \in Y$.

Constraints on the agent's view of the dependence relationships among variables can also be expressed in a number of other ways. If the agent believes y_t to be normally distributed, for example, then the agent's model can be parameterized using parameters $\theta_l = (\mu_l, \sigma_l^2)$, which capture the mean and variance of the marginals, and parameter θ_0 , which is the $l \times l$ correlation

⁵⁴Since y_t is assumed to belong to a compact Banach space, the Markovian structure is without loss of generality as y_t can be defined as $y_t = (\tilde{y}_t, \tilde{y}_{t-1}, \dots)$, the infinite history of realizations of \tilde{y}_t .

⁵⁵See also [Eyster and Piccione \(2013\)](#) for an asset-pricing application in a model that is formally close to the ABEE and [Eyster, Madarasz, and Michaillat \(2017\)](#) for an application of the cursed equilibrium in a new-Keynesian model.

⁵⁶The observation that correlation misperception is a form of model misspecification is not new. [Esponda and Pouzo \(2016a\)](#) formulate analogy-based expectations as a special type of misspecification in a game-theoretical context and formally establish that the ABEE is a special case of their more general Berk–Nash Equilibrium concept.

⁵⁷For more on copula theory, see [Nelsen \(2006\)](#).

matrix. Fully cursed expectations are then captured by assuming that $\Theta_0 = \{\theta_0^*\}$, where θ_0^* is the identity matrix.

One can also use this parametrization to study more subtle forms of correlation misperception. Suppose, for instance, that the agent believes y_t to be a normally-distributed vector and the correlation between any two components of y_t to be either 0, +1, or -1.⁵⁸ This belief can be represented by assuming that the set Θ_0 from the previous paragraph is constrained to be the set of $l \times l$ matrices with diagonal elements equal to +1 and off-diagonal elements belonging to the set $\{0, +1, -1\}$. In a CREE, the correlations perceived by the agent are determined as endogenous functions of the economic environment and the rules followed by policymakers.

C.6 Sentiments, Law of Small Numbers, and Gambler’s Fallacy

Another class of behavioral biases considers agents who believe that an underlying state exhibits more mean reversion than it actually does. Barberis, Shleifer, and Vishny (1998)’s theory of sentiments, Rabin (2002)’s theory of the “law of small numbers,” and Rabin and Vayanos (2010)’s theory of gambler’s fallacy are related works that all consider variations on this general theme. For the sake of concreteness, I focus on Rabin and Vayanos (2010), which can be most easily mapped to the linear-Gaussian hidden factor model of Appendix B.

Rabin and Vayanos (2010) consider an agent who observes an observable o_t whose distribution depends on an underlying state z_t . The observable is given by $o_t = z_t + \varepsilon_t$, where z_t follows an AR(1) process and ε_t is a normally distributed i.i.d. measurement error. While ε_t does not exhibit any mean reversion given the true data-generating process, the agent only considers a set of misspecified models according to which ε_t is mean reverting. This misspecification leads the agent to underreact to short streaks in the observable, overreact to longer ones, and underreact to very long ones.

Rabin and Vayanos show that the agent’s models can be written as linear-Gaussian hidden factor models. In particular, the agent’s set of models can be written as a special case of the set of models in Appendix B, in which the hidden state ω_t is a two-dimensional vector, the matrix A is given by

$$A = \begin{pmatrix} \rho & 0 \\ 0 & \delta\rho - \alpha\rho \end{pmatrix},$$

the vector C is given by $C = (\rho, -\alpha\rho)'$, and innovations $\epsilon_{\omega t}$ and ϵ_{ot} are jointly normal with $\mathbb{E}[\epsilon_{\omega t}\epsilon'_{\omega t}] = \Sigma_{\omega}$, $\mathbb{E}[\epsilon_{ot}\epsilon'_{ot}] = \Sigma_o$, and $\mathbb{E}[\epsilon_{\omega t}\epsilon'_{ot}] = \Sigma_{\omega o}$.⁵⁹ Parameters α , δ , ρ , Σ_{ω} , Σ_o , and $\Sigma_{\omega o}$ parameterize the agent’s model.

Gambler’s fallacy is modeled as a constraint on the parameter α . In the true data-generating

⁵⁸This form of bounded rationality is in the spirit of intuitive heuristics put forward by Kahneman and Tversky as ways people seem to deal with complex probabilistic situations. Kahneman and Tversky (1982a) argue that “people rely on heuristic principles which reduce the complex tasks of assessing probabilities and predicting values to simpler judgmental operations. In general, these heuristics are quite useful, but sometimes they lead to severe and systematic errors.” See, also, Kahneman and Tversky (1972) for the results of lab experiments that document some of these heuristics.

⁵⁹Note that in contrast to the models of Appendix B, $\epsilon_{\omega t}$ and ϵ_{ot} may be correlated in Rabin and Vayanos (2010).

process, α is equal to zero. But an agent who is subject to gambler’s fallacy has a doctrinaire belief that $\alpha > 0$. [Rabin and Vayanos \(2010\)](#) assume that the remaining parameters of the agent’s models are chosen endogenously using a (quasi)-maximum-likelihood estimator given an infinitely long sequence of observations. Their selection rule thus coincides with the CREE selection criterion.

C.7 Constant-Gain Learning

Constant-gain learning rules, such as [Marcet and Sargent \(1989a\)](#), can also be seen as instances of model misspecification.⁶⁰ An agent who performs constant-gain learning cares about some unknown state ω_t and updates her estimate $\hat{\omega}_t$ of the state using the following update rule:

$$\hat{\omega}_t = \hat{\omega}_{t-1} + k(o_t - \hat{o}_t), \tag{C.1}$$

where k is the gain parameter, o_t is an observable, and \hat{o}_t is the agent’s estimate of the observable based on her information at time $t - 1$. In adaptive-learning models, the gain k is a free parameter that is calibrated by the economist. So the updating rule (C.1) can be viewed as an instance of the “update step” of a Kalman filter, described in equation (B.2), in which parameters A , B , and K are given particular values.

A constant-gain learning rule with a calibrated gain does two things from the perspective of this paper. First, it constrains the agents’ set of models to the set of hidden factor models. Second, it selects a model from the constrained set that corresponds to the calibrated value of the gain. The solution concept of CREE- d introduced in Section 5 and used in the application of Section 6 also does the former. But rather than calibrating the gain parameter, it allows the agent to choose it endogenously to best fit the distribution of observables.

While calibrating the gain parameter could improve the fit of a model, it may lead to policy prescriptions that are subject to the Lucas critique: changes in policy may lead the agents to change the gains they use. The solution concept of CREE- d does not face this difficulty. As the parameters of policy vary, agents endogenously change the Kalman gains they use to update their expectations as they learn about the new policy regime.

C.8 Rational Inattention

[Sims \(2003\)](#)’s rational inattention can also be seen as an instance of model misspecification. Rationally inattentive agents can condition their actions only on noisy signals of the aggregate observables—with an information-flow constraint limiting the informativeness of the agents’ signals. More formally, rational inattention requires the average mutual information between an agent’s stream of signals and the stream of observables to be upper-bounded by a parameter

⁶⁰Other notable examples of models where agents use constant-gain learning rules include [Sargent \(1993\)](#), [Sargent \(1999\)](#), [Cho, Williams, and Sargent \(2002\)](#), [Marcet and Nicolini \(2003\)](#), [Adam, Kuang, and Marcet \(2012\)](#), [Adam, Marcet, and Nicolini \(2016\)](#), and [Adam, Marcet, and Beutel \(2017\)](#).

κ , which measures the agent’s limited amount of attention.⁶¹ The agent selects her signals to optimize an objective function. The objective function is either the agent’s expected payoff (as in [Maćkowiak and Wiederholt, 2015](#)) or the present-discounted value of the mean squared errors (as in [Maćkowiak, Matejka, and Wiederholt, 2017](#)).

I use the results of [Maćkowiak et al. \(2017\)](#) to cast rational inattention as an instance of misspecification. They show that, when the observable variable follows a stationary ARMA process with normally distributed innovations, one can restrict attention to signals that have a state space representation of the form (14) and (15). They also show that the information-follow constraint can be expressed as follows:

$$\frac{1}{2} \log_2 \det(\Omega) - \frac{1}{2} \log_2 \det(\hat{\Sigma}_\omega) \leq \kappa,$$

where \log_2 denotes base 2 logarithm and Ω and $\hat{\Sigma}_\omega$ are matrices defined in equations (B.4) and (B.1), respectively. By restricting matrices Ω and $\hat{\Sigma}_\omega$, the information-flow constraint restricts the parameters of the agents’ models.

In addition to constraining the agents’ set of models, rational attention also offers a selection criterion given the agents’ constrained sets of models. The selection criterion of rational inattention is generically different than that of CREE. The latter is computationally more tractable than the former—modern DSGE models can be easily modified to incorporate CREE. It also has Bayesian and adaptive learning foundations. The rational-inattention constraint on the set of models can be combined with the CREE selection criterion to develop a model that captures the intuitive idea of inattention while maintaining the analytical tractability and learning foundations of CREE.

C.9 Non-Bayesian Models of Updating

Another literature in behavioral economics, going back to the representativeness heuristic of [Kahneman and Tversky \(1972\)](#), considers deviations from Bayesian updating. Notable examples include base-rate neglect ([Kahneman and Tversky, 1982a](#)), confirmation bias ([Rabin and Schrag, 1999](#)), local thinking ([Gennaioli and Shleifer, 2010](#)), and diagnostic expectations ([Bordalo, Gennaioli, and Shleifer, 2018](#)). Models of non-Bayesian updating are strictly speaking not special cases of the framework considered in this paper. This is due to the fact that agents have internally consistent (i.e., Bayesian) belief systems in a CREE. However, [Molavi \(2018\)](#) establishes that (almost) any seemingly non-Bayesian sequence of beliefs is observationally indistinguishable from a sequence of beliefs that is generated through the application of Bayes’ rule over a larger state space.⁶²

⁶¹Rational inattention has several closely related formulations. Here, I follow [Maćkowiak et al. \(2017\)](#).

⁶²This observational indistinguishability is *not* a finite-sample phenomenon. In particular, it holds even if the external observer directly observes the underlying *population* distribution of beliefs. The observer can only rule out Bayesian rationality if there exist a pair (μ_{t-1}, μ_t) , where μ_t denotes the agent’s time- t belief about the payoff-relevant variable, such

D Additional Details for the Business-Cycle Model of Section 6

Standard calculations have been relegated to the Online Appendix.

D.1 Temporary-Equilibrium Relationships

The log-linear permanent income equation is given by

$$\begin{aligned} \hat{c}_t = & \beta \hat{\psi}_t + \frac{b}{c} \frac{1-\beta}{\beta} (\hat{R}_{t-1} - \hat{\pi}_t) + (1-\beta) \left(\frac{x}{c} \hat{x}_t - \frac{\tau}{c} \hat{\tau}_t \right) - \beta \hat{R}_t \\ & + E_t \left[\sum_{s=1}^{\infty} \beta^s \left((1-\beta) \frac{x}{c} \hat{x}_{t+s} - (1-\beta) \frac{\tau}{c} \hat{\tau}_{t+s} - (1-\beta) \hat{\psi}_{t+s} - \beta \frac{x-\tau}{c} \hat{R}_{t+s} + \frac{x-\tau}{c} \hat{\pi}_{t+s} \right) \right]. \end{aligned} \quad (\text{D.1})$$

Households' income is given by

$$\hat{x}_t = \frac{y}{x} \hat{y}_t - \frac{i}{x} \hat{i}_t. \quad (\text{D.2})$$

The discount-factor shock follows the exogenous process

$$\hat{\psi}_t = \rho_\psi \hat{\psi}_{t-1} + \varepsilon_{\psi t}. \quad (\text{D.3})$$

Investment is given by

$$\hat{i}_t = \hat{k}_t + \frac{1}{\chi} (-\hat{\psi}_t + \hat{c}_t) + E_t \left[\sum_{s=1}^{\infty} \beta^s \left(\frac{1-\beta}{\chi \beta} \hat{\psi}_{t+s} - \frac{1-\beta}{\chi \beta} \hat{c}_{t+s} + \frac{1}{\chi} \left(\frac{1}{\beta} - \frac{1-\delta}{\gamma} \right) \hat{\rho}_{t+s} \right) \right]. \quad (\text{D.4})$$

Capital evolves according to

$$\hat{k}_t = \frac{1-\delta}{\gamma} \hat{k}_{t-1} + \left(1 - \frac{1-\delta}{\gamma} \right) \hat{i}_{t-1}. \quad (\text{D.5})$$

Government spending follows the exogenous process

$$\hat{g}_t = \rho_g \hat{g}_{t-1} + \varepsilon_{g t}, \quad (\text{D.6})$$

and output is given by

$$\hat{y}_t = \frac{c}{y} \hat{c}_t + \frac{i}{y} \hat{i}_t + \frac{g}{y} \hat{g}_t. \quad (\text{D.7})$$

Inflation is given by

$$\begin{aligned} \hat{\pi}_t = & \kappa \left(\hat{\lambda}_{p t} + \alpha \hat{\rho}_t + (1-\alpha) \hat{w}_t - \hat{z}_t \right) \\ & + E_t \left[\sum_{s=1}^{\infty} \xi_p^s \beta^s \left(\frac{1-\xi_p}{\xi_p} \hat{\pi}_{t+s} + \kappa \left(\hat{\lambda}_{p, t+s} + \alpha \hat{\rho}_{t+s} + (1-\alpha) \hat{w}_{t+s} - \hat{z}_{t+s} \right) \right) \right], \end{aligned} \quad (\text{D.8})$$

where

$$\kappa \equiv \frac{(1-\xi_p)(1-\xi_p \beta)}{\xi_p}$$

that (i) the pair (μ_{t-1}, μ_t) is realized with positive probability according to the population distribution of beliefs and (ii) μ_t is not absolutely continuous with respect to μ_{t-1} . See [Shmaya and Yariv \(2016\)](#) for a related “anything goes” result.

is a constant, TFP follows the exogenous process

$$\hat{z}_t = \rho_z \hat{z}_{t-1} + \varepsilon_{zt}, \quad (\text{D.9})$$

and the markup shock follows the exogenous process

$$\hat{\lambda}_{pt} = \rho_p \hat{\lambda}_{p,t-1} + \varepsilon_{pt}. \quad (\text{D.10})$$

The real wage is given by

$$\hat{w}_t = \xi_w (\hat{w}_{t-1} - \hat{\pi}_t) + \kappa_w (\hat{\lambda}_{wt} + \nu \hat{L}_t + \hat{c}_t) + \kappa_w (\nu_w - 1) \hat{w}_t \quad (\text{D.11})$$

$$+ E_t \left[\sum_{s=1}^{\infty} \xi_w^s \beta^s \left((1 - \xi_w) \hat{\pi}_{t+s} + \kappa_w (\hat{\lambda}_{wt} + \nu \hat{L}_t + \hat{c}_t) + \kappa_w (\nu_w - 1) \hat{w}_{t+s} \right) \right]. \quad (\text{D.12})$$

where

$$\kappa_w \equiv \frac{(1 - \xi_w)(1 - \xi_w \beta)}{\nu_w}$$

is a constant, and the wage markup shock $\hat{\lambda}_{wt}$ follows the exogenous process

$$\hat{\lambda}_{wt} = \rho_w \hat{\lambda}_{w,t-1} + \varepsilon_{wt}. \quad (\text{D.13})$$

Hours are given by

$$\hat{L}_t = \frac{1}{1 - \alpha} \left[\frac{y}{y + F} \hat{y}_t - \alpha \hat{k}_t - \hat{z}_t \right], \quad (\text{D.14})$$

and the rental rate of capital by

$$\hat{\rho}_t = \hat{w}_t + \hat{L}_t - \hat{k}_t. \quad (\text{D.15})$$

The nominal interest rate follows the interest-rate rule

$$\hat{R}_t = \rho_R \hat{R}_{t-1} + (1 - \rho_R) \phi_\pi \hat{\pi}_t + \hat{\eta}_{mt}, \quad (\text{D.16})$$

where monetary-policy shock follows the exogenous process

$$\hat{\eta}_{mt} = \rho_m \hat{\eta}_{m,t-1} + \varepsilon_{mt}. \quad (\text{D.17})$$

Finally, taxes follow the tax rule

$$\hat{\tau}_t = \frac{g}{\tau} \hat{g}_t + \frac{b}{\beta \tau} (\hat{R}_{t-1} - \hat{\pi}_t). \quad (\text{D.18})$$

D.2 Constrained-Rational-Expectations Equilibrium

I next characterize the CREE when the agents' models are given by factor models with d factors. Equation (D.1) can be written in vector form as

$$\hat{c}_t = \beta \hat{\psi}_t + \frac{b}{c} \frac{1 - \beta}{\beta} (\hat{R}_{t-1} + \hat{b}_{t-1} - \hat{\pi}_t) + (1 - \beta) \left(\frac{x}{c} \hat{x}_t - \frac{\tau}{c} \hat{\tau}_t \right) - \beta \hat{R}_t + \sum_{s=1}^{\infty} \beta^s \nu'_c E_t [o_{t+s}],$$

where $v_c \in \mathbb{R}^n$ is a vector of constants. Using equation (B.5) to substitute for $E_t [o_{t+s}]$, I get

$$\hat{c}_t = \beta \hat{\psi}_t + \frac{b}{c} \frac{1-\beta}{\beta} \left(\hat{R}_{t-1} + \hat{b}_{t-1} - \hat{\pi}_t \right) + (1-\beta) \left(\frac{x}{c} \hat{x}_t - \frac{\tau}{c} \hat{\tau}_t \right) - \beta \hat{R}_t + \beta v_c' B' (I - \beta A)^{-1} \hat{\omega}_t, \quad (\text{D.19})$$

where I denotes the identity matrix. Likewise, equations (D.4), (D.8), and (D.12) can be written as

$$\hat{i}_t = \hat{k}_t + \frac{1}{\chi} (-\hat{\psi}_t + \hat{c}_t) + \beta v_i' B' (I - \beta A)^{-1} \hat{\omega}_t, \quad (\text{D.20})$$

$$\hat{\pi}_t = \kappa \left(\hat{\lambda}_{pt} + \alpha \hat{\rho}_t + (1-\alpha) \hat{w}_t - \hat{z}_t \right) + \xi_p \beta v_\pi' B' (I - \xi_p \beta A)^{-1} \hat{\omega}_t, \quad (\text{D.21})$$

$$\hat{w}_t = \xi_w (\hat{w}_{t-1} - \hat{\pi}_t) + \kappa_w \left(\hat{\lambda}_{wt} + \nu \hat{L}_t + \hat{c}_t \right) + \kappa_w (\nu_w - 1) \hat{w}_t + \xi_w \beta v_w' B' (I - \xi_w \beta A)^{-1} \hat{\omega}_t, \quad (\text{D.22})$$

where $v_i, v_\pi, v_w \in \mathbb{R}^n$ are constants.

$$\zeta_t \equiv \left(\hat{x}_t, \hat{y}_t, \hat{c}_t, \hat{i}_t, \hat{L}_t, \hat{k}_t, \hat{R}_t, \hat{\eta}_{mt}, \hat{g}_t, \hat{\tau}_t, \hat{\pi}_t, \hat{w}_t, \hat{\rho}_t, \hat{z}_t, \hat{\lambda}_{pt}, \hat{\lambda}_{wt}, \hat{\psi}_t, \hat{\omega}_t \right)'$$

Equations (17), (D.2), (D.3), (D.4)–(D.9), (D.10), (D.5)–(D.18), and (D.19)–(D.22) can be written in vector form as

$$\zeta_t = \Phi_\zeta \zeta_t + \tilde{A}_\zeta \zeta_{t-1} + \tilde{C}_\zeta \varepsilon_t, \quad (\text{D.23})$$

where

$$\varepsilon_t \equiv (\varepsilon_{zt}, \varepsilon_{\psi t}, \varepsilon_{pt}, \varepsilon_{wt}, \varepsilon_{mt}, \varepsilon_{gt})' \quad (\text{D.24})$$

is the vector of the 6 shocks that hit the economy. Assuming that $I - \Phi_\zeta$ is invertible, I can solve (D.23) for ζ_t to get

$$\zeta_t = A_\zeta \zeta_{t-1} + C_\zeta \varepsilon_t, \quad (\text{D.25})$$

where

$$A_\zeta \equiv (I - \Phi_\zeta)^{-1} \tilde{A}_\zeta,$$

$$C_\zeta \equiv (I - \Phi_\zeta)^{-1} \tilde{C}_\zeta.$$

Note that matrices A_ζ and C_ζ are functions of the fundamentals of the economy and the parameters of the agents' model as summarized by (A, B, K) .

Equation (D.23) can be used to compute the autocovariance matrices of ζ_t . Let

$$\Xi_{\zeta s} \equiv \mathbb{E} [\zeta_t \zeta_{t-s}']$$

denote the autocovariance matrix of ζ_t at lag $s \in \mathbb{Z}$, where \mathbb{E} is the probability distribution induced on paths of ζ_t by equation (D.23) and the stationary distribution of $\{\varepsilon_t\}_t$. Assuming that ζ_t is a stationary process (and so the eigenvalues of A_ζ are inside the unit circle), $\Xi_{\zeta 0}$ is the unique solution to the following equation:

$$\Xi_{\zeta 0} = A_\zeta \Xi_{\zeta 0} A_\zeta' + C_\zeta \Sigma C_\zeta',$$

where

$$\Sigma \equiv \text{diag} \left(\sigma_z^2, \sigma_\psi^2, \sigma_p^2, \sigma_w^2, \sigma_m^2, \sigma_g^2, \sigma_\tau^2 \right)$$

is the covariance matrix of shocks; Ξ_{ζ_s} is given, for $s > 0$, as

$$\Xi_{\zeta_s} = A_\zeta^s \Xi_{\zeta_0};$$

and Ξ_{ζ_s} is given, for $s < 0$, as

$$\Xi_{\zeta_s} = \Xi'_{\zeta, -s}.$$

I can use the above expressions to compute the autocovariance matrices of o_t . The vector o_t can be written as

$$o_t = D' \zeta_t,$$

where $D \in \mathbb{R}^{(21+d) \times n}$ is a vector of zeros and ones, with $D_{ij} = 1$ whenever the j th element of o_t is equal to the i th element of ζ_t . It is easy to see that

$$\Xi_s \equiv \mathbb{E} [o_t o'_{t-s}] = D' \Xi_{\zeta_s} D.$$

It is now easy to characterize a CREE. The equilibrium consists of the tuple $(A^*, B^*, \Sigma_\omega^*, \Sigma_o^*)$, which parameterize the agents' model, and the stationary distribution of endogenous random variables as summarized by the sequence $\{\Xi_s^*\}_{s \in \mathbb{Z}}$ of autocovariance matrices. Agents take $\{\Xi_s^*\}_{s \in \mathbb{Z}}$ as given and chooses $(A^*, B^*, \Sigma_\omega^*, \Sigma_o^*)$ to solve

$$(A^*, B^*, \Sigma_\omega^*, \Sigma_o^*) \in \arg \min_{(A, B, \Sigma_\omega, \Sigma_o)} H(\{\Xi_s^*\}_{s \in \mathbb{Z}}, \theta),$$

where H is the Kullback–Leibler divergence, defined in (B.7), and $\{\Xi_s^*\}_{s \in \mathbb{Z}}$ is the sequence of autocovariance matrices given $(A^*, B^*, \Sigma_\omega^*, \Sigma_o^*)$.

E Proofs

I first present some mathematical definitions that are used in the proofs of the results.

Dynamical Systems and Ergodic Theory

The definitions presented in this subsection can be found in [Gray \(2009\)](#). A discrete-time abstract *dynamical system* is a tuple $(\Omega, \mathcal{B}, \rho, \varphi)$ where (Ω, \mathcal{B}) is measurable space, ρ is a probability measure on (Ω, \mathcal{B}) , and $\varphi : \Omega \rightarrow \Omega$ is a measurable mapping. Given a dynamical system $(\Omega, \mathcal{B}, \rho, \varphi)$, a measurable function $f : \Omega \rightarrow \mathbb{R}$ is *invariant* if $f(\varphi(\omega)) = f(\omega)$ for all $\omega \in \Omega$; a measurable set $B \in \mathcal{B}$ is *invariant* if its indicator function $\mathbb{1}_B$ is invariant.

Given a dynamical system $(\Omega, \mathcal{B}, \rho, \varphi)$ and a measurable set $B \in \mathcal{B}$, the set $\varphi^{-1}(B)$ is the measurable set defined as $\varphi^{-1}(B) \equiv \{\omega \in \Omega : \varphi(\omega) \in B\}$. Likewise, the measurable sets $\varphi^{-t}(B)$ are defined recursively as $\varphi^{-t}(B) \equiv \varphi^{-1}(\varphi^{-t+1}(B))$ for all t . The dynamical system is *asymptotically*

mean stationary if, for every measurable set $B \in \mathcal{B}$, the limit $\bar{\rho}(B) = \lim_{t \rightarrow \infty} \frac{1}{t} \sum_{s=0}^{t-1} \rho(\varphi^{-s}(B))$ exists. The function $\bar{\rho} : \mathcal{B} \rightarrow [0, 1]$, when well-defined, is a probability distribution over (Ω, \mathcal{B}) , which is called the *stationary distribution* of the dynamical system.

A discrete-time *random process* is a sequence of random variables ξ_0, ξ_1, \dots all defined on some common probability space (O, \mathcal{O}, P) and taking values in some measurable space (Ξ, \mathcal{F}) . Any random process defines a dynamical system as follows. Let $(\Omega, \mathcal{B}) = (\Xi^{\mathbb{N}}, \mathcal{F}^{\mathbb{N}})$. Any element of Ω is an infinitely-long sequence $\omega = (\omega_0, \omega_1, \dots)$, where $\omega_t \in \Xi$ for all t . Let ρ denote the probability distribution of the random process ξ_0, ξ_1, \dots ; that is, for any measurable set $B \in \mathcal{B}$, define $\rho(B)$ as $\rho(B) \equiv \rho(\{o \in O : (\xi_0(o), \xi_1(o), \dots) \in B\})$. And let $\varphi : \Omega \rightarrow \Omega$ be the measurable mapping defined as $\varphi : (\omega_0, \omega_1, \dots) \mapsto (\omega_1, \omega_2, \dots)$. In other words, the mapping φ shifts the sequence “to the left” in time. The tuple $(\Omega, \mathcal{F}, \rho, \varphi)$ constructed above is the dynamical system corresponding to the random process ξ_0, ξ_1, \dots . A random process is *asymptotically mean stationary* if the corresponding dynamical system is asymptotically mean stationary. A probability distribution $\bar{\rho}$ over $(\Omega, \mathcal{B}) = (\Xi^{\mathbb{N}}, \mathcal{F}^{\mathbb{N}})$ is a *stationary distribution* of the random process if it is the stationary distribution of the corresponding dynamical system defined above.

Proof of Lemma 1

That the Kullback–Leibler divergence is well-defined and finite is due to the assumption that $q_\theta(\cdot|\cdot)$ is positive and bounded. The continuity of the mapping $\theta \mapsto H(Q_\theta, T, \mu)$ is a consequence of the assumption that the mapping $\theta \mapsto \log(q_\theta(y|y_-))$ is continuous for all y_-, y and the dominated convergence theorem. \square

Proof of Theorem 1

I prove the theorem using the Fan–Glicksberg fixed-point theorem. Define $\Psi \equiv \Delta(X \times Y \times Z \times \Delta\Theta)$. The space Ψ is given the topology induced by the total variation norm. Since X, Y, Z , and Θ are nonempty compact subsets of metric spaces, Ψ is a nonempty compact convex subset of a locally convex Hausdorff space. In what follows I define a correspondence $F : \Psi \rightarrow \Psi$ such that any fixed point μ^* of F defines a CREE and show that F has closed graph and nonempty convex values. The theorem then immediately follows the Fan–Glicksberg fixed-point theorem. Throughout the proof, I fix a temporary equilibrium T satisfying Assumption 1.

I start by defining an auxiliary correspondence. Given any closed set $\tilde{\Theta} \subseteq \Theta$, let $ba(\tilde{\Theta}) \equiv ba(X \times Y \times Z \times \Delta\tilde{\Theta})$ denote the Banach space of bounded and countably additive measures on $X \times Y \times Z \times \Delta\tilde{\Theta}$. The transition probability T defines a linear operator $\mathbf{T} : ba(\tilde{\Theta}) \rightarrow ba(\tilde{\Theta})$ as follows:

$$(\mathbf{T}\mu)(B) = \int_{X \times Y \times Z \times \Delta\tilde{\Theta}} T(B|x_-, y_-, z_-, \lambda_-) \mu(dx_- \times dy_- \times dz_- \times d\lambda_-),$$

where B denotes an arbitrary measurable subset of $X \times Y \times Z \times \Delta\tilde{\Theta}$. That \mathbf{T} maps any $\mu \in ba(\tilde{\Theta})$ to an element of $ba(\tilde{\Theta})$ is by the definition of T and the fact that Bayes’ rule always maps a belief

supported on $\tilde{\Theta}$ to a belief supported on $\tilde{\Theta}$. Since T is continuous by Assumption 1 and $X \times Y \times Z \times \Delta\tilde{\Theta}$ is a compact metric space, for any $\tilde{\Theta}$, there exists a nonempty convex compact set of invariant probabilities $\Xi(\tilde{\Theta}) \subseteq \Delta(X \times Y \times Z \times \Delta\tilde{\Theta})$ such that $\mathbf{T}\mu = \mu$ for all $\mu \in \Xi(\tilde{\Theta})$.⁶³

Define the correspondence $F : \Psi \rightarrow \Psi$ as

$$F(\mu) = \Xi \left(\arg \min_{\theta \in \Theta} H(Q_\theta, T, \mu) \right).$$

Suppose μ^* is a fixed point of the correspondence F . It is easy to see that μ^* , T , and $\Theta^* \equiv \arg \min_{\theta \in \Theta} H(Q_\theta, T, \mu^*)$ constitute a CREE. Moreover, by the argument above, F has nonempty convex values. So to prove the theorem, I only need to show that F has closed graph.

I do so by first proving that $H(Q_\theta, T, \mu)$ is continuous in μ and θ . Since (μ, θ) belongs to a metric space, to establish the continuity of $H(Q_\theta, T, \mu)$, it is sufficient to show that whenever (μ_k, θ_k) converges to (μ, θ) for a sequence $\{(\mu_k, \theta_k)\}_k$, the sequence $H(Q_{\theta_k}, T, \mu_k)$ converges to $H(Q_\theta, T, \mu)$. By definition,

$$H(Q_\theta, T, \mu) = - \int \int \log(q_\theta(y|y_-)) T(dy|x_-, y_-, z_-, \lambda_-) \mu(dx_- \times dy_- \times dz_- \times d\lambda_-).$$

Since μ_k converges to μ in total variation distance and $q_\theta(y|y_-)$ is bounded, $H(Q_\theta, T, \mu_k) \rightarrow H(Q_\theta, T, \mu)$ uniformly in θ . On the other hand, since the mapping $\theta \mapsto q_\theta(y_-|y)$ is continuous for all y_- , $y \in Y$, the function sequence $q_{\theta_k}(\cdot|\cdot)$ converges to $q_\theta(\cdot|\cdot)$ pointwise. Therefore, by the dominated convergence theorem, $H(Q_{\theta_k}, T, \mu) \rightarrow H(Q_\theta, T, \mu)$ for any μ . The uniform convergence of $H(Q_\theta, T, \mu_k)$ to $H(Q_\theta, T, \mu)$ and the convergence of $H(Q_{\theta_k}, T, \mu)$ to $H(Q_\theta, T, \mu)$ together imply that $H(Q_{\theta_k}, T, \mu_k)$ converges to $H(Q_\theta, T, \mu)$.

I can now show that F has closed graph. Let $\{\mu_k\}_k$ and $\{\mu'_k\}_k$ be two convergent sequences in Ψ such that $\mu'_k \in F(\mu_k)$ for all k , and let $\mu = \lim_{k \rightarrow \infty} \mu_k$ and $\mu' = \lim_{k \rightarrow \infty} \mu'_k$. To prove that F has closed graph, I need to show that $\mu' \in F(\mu)$. That $\mu' = \mathbf{T}\mu'$ is a trivial consequence of the assumption that $\mu'_k = \mathbf{T}\mu'_k$ for all k . So I only need to show that μ' is supported on $X \times Y \times Z \times \Delta \left(\arg \min_{\theta \in \Theta} H(Q_\theta, T, \mu) \right)$. Continuity of H established above and the Berge maximum theorem imply that the mapping $\mu \mapsto \arg \min_{\theta \in \Theta} H(Q_\theta, T, \mu)$ is upper hemicontinuous. Fix an arbitrary open neighborhood U of $\arg \min_{\theta \in \Theta} H(Q_\theta, T, \mu)$. By upper hemicontinuity of $\arg \min_{\theta \in \Theta} H(Q_\theta, T, \mu)$, there exists some K such that, if $k \geq K$, then $\arg \min_{\theta \in \Theta} H(Q_\theta, T, \mu_k) \subseteq U$. Since $\mu'_k \in F(\mu_k)$ for all k ,

$$\mu'_k \in \Delta \left(X \times Y \times Z \times \Delta \left(\arg \min_{\theta \in \Theta} H(Q_\theta, T, \mu_k) \right) \right).$$

Therefore, $\mu'_k \in \Delta(X \times Y \times Z \times \Delta U)$, and so $\mu'_k(X \times Y \times Z \times \Delta U) = 1$, for all $k \geq K$. Since $\mu'_k \rightarrow \mu'$ in total variation norm, $\mu'(X \times Y \times Z \times \Delta U) = 1$. The assumption that U is an arbitrary neighborhood of $\arg \min_{\theta \in \Theta} H(Q_\theta, T, \mu)$ then implies that μ' is supported on $X \times Y \times Z \times$

⁶³See, for instance, Theorem 19.18 of Aliprantis and Border (2006).

$\Delta \left(\arg \min_{\theta \in \Theta} H(Q_\theta, T, \mu) \right)$. This establishes that $\mu' \in F(\mu)$, thus proving that F has closed graph. Appealing to Fan–Glicksberg fixed-point theorem completes the proof of the theorem. \square

Statement and Proof of Lemma E.1

Lemma E.1. $\phi(\lambda_-, y_-, y)$ is well-defined and finite and the mapping $(\lambda_-, y_-) \mapsto \phi(\lambda_-, y_-, y)$ is continuous for all y .

Proof. Since $q_\theta(y|y_-) > 0$ for all (θ, y_-, y) , the Bayesian update $\phi(\lambda_-, y_-, y)$ is well-defined and finite. To show the continuity of ϕ , it is sufficient to show that $\int_B q_\theta(y|y_-)\lambda_-(d\theta)$ and $\int q_\theta(y|y_-)\lambda_-(d\theta)$ are continuous functions of (λ_-, y_-) . I only show the continuity of the latter as the proof of the continuity of the former is identical. Since $Y \times \Delta\Theta$ is a metric space, to show that the mapping $(\lambda_-, y_-) \mapsto \int q_\theta(y|y_-)\lambda_-(d\theta)$ is continuous for all y , it is sufficient to show that whenever $(\lambda_{k,-}, y_{k,-}) \rightarrow (\lambda_-, y_-)$ as $k \rightarrow \infty$, it is the case that $\int q_\theta(y|y_{k,-})\lambda_{k,-}(d\theta) \rightarrow \int q_\theta(y|y_-)\lambda_-(d\theta)$. Since the mapping $y_- \mapsto q_\theta(y|y_-)$ is continuous for all θ and y , the sequence $q_\theta(y|y_{k,-})$ converges to $q_\theta(y|y_-)$ pointwise. Since the family $\{\theta \mapsto q_\theta(y|y_-)\}_{y_- \in Y}$ is equicontinuous, Arzelà–Ascoli theorem implies that $q_\theta(y|y_{k,-})$ converges to $q_\theta(y|y_-)$ uniformly in θ . On the other hand, by the definition of weak convergence, for any $\tilde{y}_-, \tilde{y} \in Y$, the sequence $\int q_\theta(\tilde{y}|\tilde{y}_-)\lambda_{k,-}(d\theta)$ converges to $\int q_\theta(\tilde{y}|\tilde{y}_-)\lambda_-(d\theta)$ as k goes to infinity. Therefore, $\int q_\theta(y|y_{k,-})\lambda_{k,-}(d\theta) \rightarrow \int q_\theta(y|y_-)\lambda_-(d\theta)$. This proves that $\phi(\lambda_-, y_-, y)$ is a continuous function of (λ_-, y_-) . \square

Statement and Proof of Lemma E.2

Lemma E.2. Consider an economy \mathfrak{E} and a temporary equilibrium T that satisfies Assumption 1 and is asymptotically mean stationary. In any Bayesian equilibrium, with \mathbb{P} -probability one, there exists an invariant probability distribution μ^* for T such that, for any continuous function $f : (X \times Y \times Z \times \Delta\Theta)^2 \rightarrow \mathbb{R}$,

$$\begin{aligned} & \lim_{t \rightarrow \infty} \frac{1}{t} \sum_{s=1}^t f(x_{s-1}, y_{s-1}, z_{s-1}, \lambda_{s-1}, x_s, y_s, z_s, \lambda_s) \\ &= \int f(x_-, y_-, z_-, \lambda_-, x, y, z, \lambda) T(dx \times dy \times dz \times d\lambda | x_-, y_-, z_-, \lambda_-) \mu^*(dx_- \times dy_- \times dz_- \times d\lambda_-). \end{aligned} \tag{E.1}$$

Proof. Given any function $f : (X \times Y \times Z \times \Delta\Theta)^2 \rightarrow \mathbb{R}$, define the function $\bar{f} : X \times Y \times Z \times \Delta\Theta \rightarrow \mathbb{R}$ as

$$\bar{f}(x, y, z, \lambda) \equiv \int f(x, y, z, \lambda, x', y', z', \lambda') T(dx' \times dy' \times dz' \times d\lambda' | x, y, z, \lambda).$$

I first show that \bar{f} is a continuous function for any continuous f . Since $(X \times Y \times Z \times \Delta\Theta)^2$ is a metric space, to show that \bar{f} is a continuous function of (x, y, z, λ) , it is sufficient to show that whenever

$(x_k, y_k, z_k, \lambda_k) \rightarrow (x, y, z, \lambda)$ as $k \rightarrow \infty$,

$$\begin{aligned} & \int f(x_k, y_k, z_k, \lambda_k, x', y', z', \lambda') T(dx' \times dy' \times dz' \times d\lambda' | x_k, y_k, z_k, \lambda_k) \\ & \rightarrow \int f(x, y, z, \lambda, x', y', z', \lambda') T(dx' \times dy' \times dz' \times d\lambda' | x, y, z, \lambda). \end{aligned}$$

Since f is a continuous function over a compact space, the family of functions $\{f(x_k, y_k, z_k, \lambda_k, x', y', z', \lambda')\}_k$ is uniformly bounded. On the other hand, by Assumption 1, the mapping $(x, y, z, \lambda) \rightarrow T(\cdot | x, y, z, \lambda)$ is continuous with respect to the total variation norm on the space of countably additive measures over $X \times Y \times Z \times \Delta\Theta$. Therefore,

$$\begin{aligned} & \int f(x_k, y_k, z_k, \lambda_k, x', y', z', \lambda') T(dx' \times dy' \times dz' \times d\lambda' | x_k, y_k, z_k, \lambda_k) \\ & - \int f(x_k, y_k, z_k, \lambda_k, x', y', z', \lambda') T(dx' \times dy' \times dz' \times d\lambda' | x, y, z, \lambda) \rightarrow 0. \end{aligned}$$

Consequently,

$$\begin{aligned} & \lim_{k \rightarrow \infty} \int f(x_k, y_k, z_k, \lambda_k, x', y', z', \lambda') T(dx' \times dy' \times dz' \times d\lambda' | x_k, y_k, z_k, \lambda_k) \\ & = \lim_{k \rightarrow \infty} \int f(x_k, y_k, z_k, \lambda_k, x', y', z', \lambda') T(dx' \times dy' \times dz' \times d\lambda' | x, y, z, \lambda). \end{aligned}$$

But since f is continuous by assumption,

$$f(x_k, y_k, z_k, \lambda_k, x', y', z', \lambda') \rightarrow f(x, y, z, \lambda, x', y', z', \lambda')$$

for all x', y', z' , and λ' . The fact that f is bounded and the dominated convergence theorem then imply that

$$\begin{aligned} & \lim_{k \rightarrow \infty} \int f(x_k, y_k, z_k, \lambda_k, x', y', z', \lambda') T(dx' \times dy' \times dz' \times d\lambda' | x, y, z, \lambda) \\ & = \int f(x, y, z, \lambda, x', y', z', \lambda') T(dx' \times dy' \times dz' \times d\lambda' | x, y, z, \lambda). \end{aligned}$$

This proves that \bar{f} is continuous whenever f is continuous.

Next note that, since $\{x_s, y_s, z_s, \lambda_s\}_{s=0}^{\infty}$ is a Bayesian equilibrium with probability distribution \mathbb{P} ,

$$\bar{f}(x_{t-1}, y_{t-1}, z_{t-1}, \lambda_{t-1}) = \mathbb{E}[f(x_{t-1}, y_{t-1}, z_{t-1}, \lambda_{t-1}, x_t, y_t, z_t, \lambda_t) | \{x_s, y_s, z_s, \lambda_s\}_{s=0}^{t-1}]$$

for all t . Furthermore, since $X \times Y \times Z \times \Delta\Theta$ is compact, f and \bar{f} are bounded whenever f is continuous. Therefore, by Loève (1960, p. 387), with \mathbb{P} -probability one,⁶⁴

$$\lim_{t \rightarrow \infty} \frac{1}{t} \sum_{s=1}^t [f(x_{s-1}, y_{s-1}, z_{s-1}, \lambda_{s-1}, x_s, y_s, z_s, \lambda_s) - \bar{f}(x_{s-1}, y_{s-1}, z_{s-1}, \lambda_{s-1})] \rightarrow 0.$$

⁶⁴The result in Loève (1960) establishes that $\frac{1}{t} \sum_{s=0}^{t-1} (X_t - \mathbb{E}[X_t | X_{t-1}, \dots, X_0]) = 0$ with \mathbb{P} -probability one for any sequence of bounded random variables $\{X_t\}$.

So to prove the lemma, it is sufficient to show that, with \mathbb{P} -probability one, there exists an invariant probability distribution μ^* for T such that, for any continuous function $\bar{f} : X \times Y \times Z \times \Delta\Theta \rightarrow \mathbb{R}$,

$$\lim_{t \rightarrow \infty} \frac{1}{t} \sum_{s=0}^{t-1} \bar{f}(x_s, y_s, z_s, \lambda_s) = \int \bar{f}(x, y, z, \lambda) \mu^*(dx \times dy \times dz \times d\lambda).$$

I do so by showing that assumptions of Theorem 3.2 of [Jamison \(1965\)](#) are satisfied and applying the theorem.⁶⁵ By an argument identical to the one presented in the first paragraph of the proof, the transition probability T maps any continuous function to a continuous function, that is, T is Feller. The transition probability T is also uniformly stable in mean. To see this, note that the assumption that the temporary recursive equilibrium T is asymptotically mean stationary immediately implies that Condition (d) in Theorem 1.1 of [Jamison \(1965\)](#) is satisfied. So Theorem 1.1 of Jamison, which is a result establishing the equivalence of asymptotic mean stationarity and uniform mean stability, guarantees that T is uniformly stable in mean. Lemma [E.2](#) is then an immediate consequence of part (a) of Lemma 1.3 and Theorem 3.2 in Jamison. \square

Proof of Part (a) of Theorem 2

Since $X \times Y \times Z \times \Delta\Theta$ is a metric space, by Urysohn's Lemma, there exists a continuous function $f : X \times Y \times Z \times \Delta\Theta \rightarrow \mathbb{R}$ such that

$$\mathbb{1}\{(x, y, z, \lambda) \in K\} \leq f(x, y, z, \lambda) \leq \mathbb{1}\{(x, y, z, \lambda) \in U_2\}$$

for all $(x, y, z, \lambda) \in X \times Y \times Z \times \Delta\Theta$. Therefore,

$$\frac{1}{t} \sum_{s=0}^{t-1} \mathbb{1}\{(x_s, y_s, z_s, \lambda_s) \in K\} \leq \frac{1}{t} \sum_{s=0}^{t-1} f(x_s, y_s, z_s, \lambda_s).$$

So by Lemma [E.2](#),

$$\begin{aligned} \limsup_{t \rightarrow \infty} \frac{1}{t} \sum_{s=0}^{t-1} \mathbb{1}\{(x_s, y_s, z_s, \lambda_s) \in K\} &\leq \lim_{t \rightarrow \infty} \frac{1}{t} \sum_{s=0}^{t-1} f(x_s, y_s, z_s, \lambda_s) \\ &= \int f(x, y, z, \lambda) \mu^*(dx \times dy \times dz \times d\lambda). \end{aligned}$$

On the other hand,

$$\int f(x, y, z, \lambda) \mu^*(dx \times dy \times dz \times d\lambda) \leq \int \mathbb{1}\{(x, y, z, \lambda) \in U_2\} \mu^*(dx \times dy \times dz \times d\lambda) = \mu^*(U_2).$$

This establishes the rightmost inequality. By an identical argument,

$$\limsup_{t \rightarrow \infty} \frac{1}{t} \sum_{s=0}^{t-1} \mathbb{1}\{(x_s, y_s, z_s, \lambda_s) \in \bar{K}\} \leq \mu^*(\bar{U}_1),$$

⁶⁵Theorem 3.2 of [Jamison \(1965\)](#) is a law-of-large-numbers result. Together with Jamison's Lemma 1.3, it establishes that for any transition probability M that is Feller and uniformly stable in mean and for any initial distribution μ_0 and almost all realizations of the Markov chain with initial distribution μ_0 and transition probability M , there exists an invariant distribution μ for M such that the time-average of any continuous function of the chain converges to its expectation under μ .

where \bar{K} and \bar{U}_1 denote the complements of K and U_1 , respectively, in $X \times Y \times Z \times \Delta\Theta$. This establishes the leftmost inequality. \square

Proof of Part (b) of Theorem 2

To simplify the expressions, I introduce some notation. Fix an arbitrary $(x_0, y_0, z_0, \lambda_0)$, a Bayesian equilibrium $\{x_t, y_t, z_t, \lambda_t\}_{t=0}^\infty$ with probability distribution \mathbb{P} and initial condition $(x_0, y_0, z_0, \lambda_0)$, and a realization of $\{x_t, y_t, z_t, \lambda_t\}_{t=1}^\infty$ for which (E.1) holds. By Lemma E.2, the set of such $\{x_t, y_t, z_t, \lambda_t\}_{t=1}^\infty$ has \mathbb{P} -probability one. Let μ^* be the corresponding invariant probability distribution for T , and define

$$\begin{aligned}\Theta^* &\equiv \arg \min_{\theta \in \Theta} H^*(\theta), \\ d^*(\theta) &\equiv \min_{\vartheta \in \Theta^*} \|\theta - \vartheta\|,\end{aligned}$$

where I use $H^*(\theta)$ as a shorthand for $H(Q_\theta, T, \mu^*)$ and $\|\cdot\|$ denotes a norm on the Euclidean space. For any $\epsilon > 0$, define

$$\Theta_\epsilon^* \equiv \{\theta \in \Theta : d^*(\theta) \geq \epsilon\},$$

and let

$$\begin{aligned}H_\epsilon^* &\equiv \inf_{\theta \in \Theta_\epsilon^*} H^*(\theta), \\ H_0^* &\equiv \min_{\theta \in \Theta} H^*(\theta).\end{aligned}$$

It is easy to see that Θ^* is nonempty and compact, $d^*(\theta)$ is well-defined, the mapping $\theta \mapsto d^*(\theta)$ is continuous, the set Θ_ϵ^* is compact, and $H_\epsilon^* > H_0^*$.

I start by establishing that, for any $\epsilon > 0$,

$$\liminf_{t \rightarrow \infty} \frac{1}{t} \inf_{\theta \in \Theta_\epsilon^*} \left\{ \sum_{s=1}^t -\log(q_\theta(y_s|y_{s-1})) \right\} \geq H_\epsilon^*. \quad (\text{E.2})$$

Since the mapping $\theta \mapsto \log(q_\theta(y|y_-))$ is continuous for all y_- , y and Y is compact, by the dominated convergence theorem, for any $\theta \in \Theta_\epsilon^*$, there exists an open ball $B_\eta(\theta)$ of radius $\eta = \eta(\theta, \epsilon)$ centered at θ such that

$$\begin{aligned}&\int \inf_{\vartheta \in B_\eta(\theta)} \left\{ -\log(q_\vartheta(y|y_-)) \right\} T(dy|x_-, y_-, z_-, \lambda_-) \mu^*(dx_- \times dy_- \times dz_- \times d\lambda_-) \\ &\geq - \int \log(q_\theta(y|y_-)) T(dy|x_-, y_-, z_-, \lambda_-) \mu^*(dx_- \times dy_- \times dz_- \times d\lambda_-) - \epsilon \\ &= H^*(\theta) - \epsilon,\end{aligned}$$

where the equality is by the definition of H^* . Since Θ_ϵ^* is compact, it can be covered by a finite number J of such balls: $B_{\eta_j}(\theta_j)$, where $j = 1, \dots, J$. Define

$$\gamma_{sj} \equiv \inf_{\theta \in \Theta_\epsilon^* \cap B_{\eta_j}(\theta_j)} \left\{ -\log(q_\theta(y_s|y_{s-1})) \right\}.$$

Given that γ_{sj} is the infimum of a family of equicontinuous functions, it is continuous. An application of Lemma E.2 then establishes that, for all $j = 1, \dots, J$,

$$\lim_{t \rightarrow \infty} \frac{1}{t} \sum_{s=1}^t \gamma_{sj} \geq H^*(\theta_j) - \varepsilon.$$

On the other hand,

$$H^*(\theta_j) \geq \min_j H^*(\theta_j) \geq H_\varepsilon^*.$$

Finally, note that

$$\frac{1}{t} \inf_{\theta \in \Theta_\varepsilon^*} \left\{ \sum_{s=1}^t -\log(q_\theta(y_s | y_{s-1})) \right\} \geq \frac{1}{t} \min_j \sum_{s=1}^t \gamma_{sj} = \min_j \frac{1}{t} \sum_{s=1}^t \gamma_{sj}.$$

Thus,

$$\liminf_{t \rightarrow \infty} \frac{1}{t} \inf_{\theta \in \Theta_\varepsilon^*} \left\{ \sum_{s=1}^t -\log(q_\theta(y_s | y_{s-1})) \right\} \geq H_\varepsilon^* - \varepsilon.$$

Noting that $\varepsilon > 0$ is arbitrary establishes (E.2).

I now proceed to prove part (b) of the theorem. Without loss of generality, assume that $\text{supp } \lambda_0 = \Theta$. To prove the theorem, it is sufficient to show that

$$\lim_{t \rightarrow \infty} \int_{\Theta} d^*(\theta) \lambda_t(d\theta) = 0. \tag{E.3}$$

By Bayes' rule,

$$\int_{\Theta} d^*(\theta) \lambda_t(d\theta) = \frac{\int_{\Theta} d^*(\theta) \prod_{s=1}^t q_\theta(y_s | y_{s-1}) \lambda_0(d\theta)}{\int_{\Theta} \prod_{s=1}^t q_\theta(y_s | y_{s-1}) \lambda_0(d\theta)}.$$

I can write the above equation as

$$\int_{\Theta} d^*(\theta) \lambda_t(d\theta) = \frac{\int_{\Theta} d^*(\theta) L_t(\theta) \lambda_0(d\theta)}{\int_{\Theta} L_t(\theta) \lambda_0(d\theta)},$$

where $L_t(\theta)$ is defined as

$$L_t(\theta) = \prod_{s=1}^t q_\theta(y_s | y_{s-1}) = \exp \left(\sum_{s=1}^t \log(q_\theta(y_s | y_{s-1})) \right).$$

Recall that

$$\Theta_\varepsilon^* = \{\theta \in \Theta : d^*(\theta) \geq \varepsilon\},$$

and define

$$\overline{\Theta}_\delta^* \equiv \{\theta \in \Theta : d^*(\theta) < \delta\}.$$

For any $\varepsilon > 0$ and $\delta > 0$,

$$\int_{\Theta} d^*(\theta) \lambda_t(d\theta) \leq \varepsilon + C \frac{\int_{\Theta_\varepsilon^*} L_t(\theta) \lambda_0(d\theta)}{\int_{\overline{\Theta}_\delta^*} L_t(\theta) \lambda_0(d\theta)},$$

where $C \equiv \max_{\theta \in \Theta} d^*(\theta)$ is finite due to the compactness of Θ and the continuity of $d^*(\theta)$ in θ . Let $\alpha \equiv (H_\epsilon^* - H_0^*) / 2 > 0$ and multiply both the numerator and the denominator of the fraction on the right-hand side of the above equation by $\exp(t(H_0^* + \alpha))$ to get

$$\int_{\Theta} d^*(\theta) \lambda_t(d\theta) \leq \epsilon + C \frac{e^{t(H_0^* + \alpha)} \int_{\Theta_\epsilon^*} L_t(\theta) \lambda_0(d\theta)}{e^{t(H_0^* + \alpha)} \int_{\overline{\Theta}_\delta^*} L_t(\theta) \lambda_0(d\theta)}. \quad (\text{E.4})$$

Since $\epsilon > 0$ is arbitrary, it is sufficient to show that the numerator of the fraction on the right-hand side goes to zero for any $\epsilon > 0$ and its denominator goes to infinity for some $\delta > 0$.

I first consider the numerator. Define

$$\ell_t(\theta) \equiv -\log(L_t(\theta)) = \sum_{s=1}^t -\log(q_\theta(y_s | y_{s-1})),$$

and note that

$$\begin{aligned} e^{t(H_0^* + \alpha)} \int_{\Theta_\epsilon^*} L_t(\theta) \lambda_0(d\theta) &= \int_{\Theta_\epsilon^*} \exp\left(t \left[H_0^* + \alpha - \frac{1}{t} \ell_t(\theta) \right]\right) \lambda_0(d\theta) \\ &\leq \lambda_0(\Theta_\epsilon^*) \sup_{\theta \in \Theta_\epsilon^*} \exp\left(t \left[H_0^* + \alpha - \frac{1}{t} \ell_t(\theta) \right]\right) \\ &\leq \lambda_0(\Theta_\epsilon^*) \exp\left(t \left[H_0^* + \alpha - \frac{1}{t} \inf_{\theta \in \Theta_\epsilon^*} \ell_t(\theta) \right]\right). \end{aligned}$$

Equation (E.2) then implies that

$$\begin{aligned} \limsup_{t \rightarrow \infty} e^{t(H_0^* + \alpha)} \int_{\Theta_\epsilon^*} L_t(\theta) \lambda_0(d\theta) \\ \leq \lambda_0(\Theta_\epsilon^*) \exp\left(t \left[H_0^* + \alpha - \liminf_{t \rightarrow \infty} \frac{1}{t} \inf_{\theta \in \Theta_\epsilon^*} \ell_t(\theta) \right]\right) = 0. \end{aligned}$$

This proves that the numerator goes to zero.

Next consider the denominator. By Fatou's lemma,

$$\begin{aligned} \liminf_{t \rightarrow \infty} e^{t(H_0^* + \alpha)} \int_{\overline{\Theta}_\delta^*} L_t(\theta) \lambda_0(d\theta) \\ \geq \int_{\overline{\Theta}_\delta^*} \liminf_{t \rightarrow \infty} \exp\left(t \left[H_0^* + \alpha - \frac{1}{t} \ell_t(\theta) \right]\right) \lambda_0(d\theta) \\ \geq \lambda_0(\overline{\Theta}_\delta^*) \inf_{\theta \in \overline{\Theta}_\delta^*} \left\{ \liminf_{t \rightarrow \infty} \exp\left(t \left[H_0^* + \alpha - \frac{1}{t} \ell_t(\theta) \right]\right) \right\}. \end{aligned}$$

For any $\delta > 0$, the set $\overline{\Theta}_\delta^*$ is open. Furthermore, by assumption, λ_0 has full support over Θ , so $\lambda_0(\overline{\Theta}_\delta^*) > 0$ for any $\delta > 0$. Therefore, to prove that the above expression goes to infinity, it is sufficient to show that $\delta > 0$ can be chosen such that

$$\inf_{\theta \in \overline{\Theta}_\delta^*} \left\{ \liminf_{t \rightarrow \infty} \exp\left(t \left[H_0^* + \alpha - \frac{1}{t} \ell_t(\theta) \right]\right) \right\} = +\infty. \quad (\text{E.5})$$

Since $q_\theta(y|y_-)$ is continuous in (y_-, y) , by Lemma E.2, for any $\theta \in \overline{\Theta}_\delta^*$,

$$\liminf_{t \rightarrow \infty} \exp \left(t \left[H_0^* + \alpha - \frac{1}{t} \ell_t(\theta) \right] \right) = \liminf_{t \rightarrow \infty} \exp \left(t \left[H_0^* + \alpha - H^*(\theta) \right] \right).$$

Therefore,

$$\inf_{\theta \in \overline{\Theta}_\delta^*} \left\{ \liminf_{t \rightarrow \infty} \exp \left(t \left[H_0^* + \alpha - \frac{1}{t} \ell_t(\theta) \right] \right) \right\} = \liminf_{t \rightarrow \infty} \exp \left(t \inf_{\theta \in \overline{\Theta}_\delta^*} \left[H_0^* + \alpha - H^*(\theta) \right] \right).$$

Since $H^*(\theta)$ is continuous in θ and $\alpha > 0$ by Lemma 1, I can choose δ to be sufficiently small that

$$\inf_{\theta \in \overline{\Theta}_\delta^*} \left[H_0^* + \alpha - H^*(\theta) \right] > 0.$$

Picking such a $\delta > 0$ establishes (E.5) and completes the proof of part (b). \square

Proof of Part (c) of Theorem 2

Let μ^* and Θ^* be as in the proof of the first two parts. The triple (T, μ^*, Θ^*) trivially satisfies Conditions (i) and (ii) for a CREE. I next show that it also satisfies Condition (iii). Let U_1 be an arbitrary open neighborhood of Θ^* , and U_1 be an open set and K be a closed set such that $\Theta^* \subset U_2 \subset K \subset U_1$. By part (a) of the theorem,

$$\begin{aligned} \mu^*(X \times Y \times Z \times \Delta U_1) &\geq \limsup_{t \rightarrow \infty} \frac{1}{t} \sum_{s=0}^{t-1} \mathbb{1}\{(x_s, y_s, z_s, \lambda_s) \in X \times Y \times Z \times \Delta K\} \\ &\geq \limsup_{t \rightarrow \infty} \frac{1}{t} \sum_{s=0}^{t-1} \mathbb{1}\{\lambda_s \in \Delta U_2\}. \end{aligned}$$

Since U_2 is an open neighborhood of Θ^* , part (b) of the theorem implies that $\lambda_t(U_2) \rightarrow 0$ as $t \rightarrow \infty$. Therefore, $\limsup_{t \rightarrow \infty} \frac{1}{t} \sum_{s=0}^{t-1} \mathbb{1}\{\lambda_s \in \Delta U_2\} = \lim_{t \rightarrow \infty} \frac{1}{t} \sum_{s=0}^{t-1} \mathbb{1}\{\lambda_s \in \Delta U_2\} = 1$. Since U_1 is an arbitrary open neighborhood of Θ^* , the probability distribution μ^* is supported on $X \times Y \times Z \times \Delta \Theta^*$. \square

Statement and Proof of Lemma E.3

Lemma E.3. *Consider an economy \mathfrak{E} and an asymptotically mean stationary adaptive equilibrium $\{x_t, y_t, z_t, \hat{\theta}_t\}_{t=0}^\infty$ for \mathfrak{E} with probability distribution \mathbb{P} . With \mathbb{P} -probability one, there exists a probability distribution μ^* over $X \times Y \times Z \times \Theta$ such that, for any bounded measurable function $f : (X \times Y \times Z)^2 \times \Theta \rightarrow \mathbb{R}$,*

$$\begin{aligned} &\lim_{t \rightarrow \infty} \frac{1}{t} \sum_{s=1}^t f(x_{s-1}, y_{s-1}, z_{s-1}, \theta_{s-1}, x_s, y_s, z_s) \\ &= \int f(x_-, y_-, z_-, \theta_-, x, y, z) T(dx \times dy \times dz | x_-, y_-, z_-, \theta_-) \mu^*(dx_- \times dy_- \times dz_- \times d\theta_-). \end{aligned} \quad (\text{E.6})$$

Proof. The proof relies on the ergodic decomposition theorem for asymptotically mean stationary dynamical systems as presented in part (g) of Theorem 8.3 of [Gray \(2009\)](#). Let $\Omega = (X \times Y \times Z \times \Theta)^\mathbb{N}$ and let $\mathcal{F} = (\mathcal{X} \times \mathcal{Y} \times \mathcal{Z} \times \mathcal{B}(\Theta))^\mathbb{N}$, where $\mathcal{B}(\Theta)$ denotes the Borel sigma-algebra on Θ . Since $X \times Y \times Z \times \Theta$ is a compact metric space and \mathcal{X} , \mathcal{Y} , \mathcal{Z} , and $\mathcal{B}(\Theta)$ are Borel sigma-algebras, \mathcal{F} is a countably-generated sigma-algebra. Furthermore, the dynamical system defined on (Ω, \mathcal{F}) by the adaptive equilibrium is asymptotically mean stationary by assumption. Therefore, by Theorem 8.3 of [Gray \(2009\)](#), for \mathbb{P} -almost all $\{x_t, y_t, z_t, \hat{\theta}_t\}_{t=0}^\infty$, there exists a probability distribution ρ over $(X \times Y \times Z)^2 \times \Theta$ such that

$$\begin{aligned} & \lim_{t \rightarrow \infty} \frac{1}{t} \sum_{s=1}^t f(x_{s-1}, y_{s-1}, z_{s-1}, \hat{\theta}_{s-1}, x_s, y_s, z_s) \\ &= \int f(x_-, y_-, z_-, \theta_-, x, y, z) \rho(dx_- \times dy_- \times dz_- \times d\theta_- \times dx \times dy \times dz) \end{aligned} \quad (\text{E.7})$$

for any bounded measurable function f .

Let μ^* denote the marginal of ρ over $(x_-, y_-, z_-, \theta_-)$. Since $(X \times Y \times Z)^2 \times \Theta$ is a compact metric space, there exists a μ^* -almost everywhere unique transition probability M from $X \times Y \times Z \times \Theta$ to $X \times Y \times Z$ such that $\rho(B_- \times B) = \int_{B_-} M(B|x_-, y_-, z_-, \theta_-) \mu(dx_- \times dy_- \times dz_- \times d\theta_-)$ for all measurable sets $B_- \subseteq X \times Y \times Z \times \Theta$ and $B \subseteq X \times Y \times Z$.⁶⁶ So to prove the lemma, I only need to show that $M(B|x_-, y_-, z_-, \theta_-) = T(B|x_-, y_-, z_-, \theta_-)$ for any measurable set $B \subseteq X \times Y \times Z$ and μ^* -almost all $(x_-, y_-, z_-, \theta_-)$. I do so in the rest of the proof.

Fix some arbitrary $B \subseteq X \times Y \times Z$ and $B_- \subseteq X \times Y \times Z \times \Theta$. By equation (E.7), with \mathbb{P} -probability one,

$$\frac{1}{t} \sum_{s=1}^t \mathbb{1} \{ (x_{s-1}, y_{s-1}, z_{s-1}, \hat{\theta}_{s-1}, x_s, y_s, z_s) \in B_- \times B \} - \rho(B_- \times B) \rightarrow 0. \quad (\text{E.8})$$

Since $\{x_s, y_s, z_s, \hat{\theta}_s\}_{s=0}^\infty$ is an adaptive equilibrium, by definition,

$$\begin{aligned} \bar{\mathbb{1}} \left(x_{t-1}, y_{t-1}, z_{t-1}, \hat{\theta}_{t-1} \right) &\equiv \mathbb{E} \left[\mathbb{1} \{ (x_{t-1}, y_{t-1}, z_{t-1}, x_t, y_t, z_t) \in B_- \times B \} \mid \{x_s, y_s, z_s, \hat{\theta}_s\}_{s=0}^{t-1} \right] \\ &= T \left(B \mid x_{t-1}, y_{t-1}, z_{t-1}, \hat{\theta}_{t-1} \right) \mathbb{1} \{ (x_{t-1}, y_{t-1}, z_{t-1}, \hat{\theta}_{t-1}) \in B_- \}, \end{aligned} \quad (\text{E.9})$$

where the first equality is a definition. By [Loève \(1960, p. 387\)](#), with \mathbb{P} -probability one,

$$\frac{1}{t} \sum_{s=1}^t \left[\mathbb{1} \{ (x_{s-1}, y_{s-1}, z_{s-1}, \hat{\theta}_{s-1}, x_s, y_s, z_s) \in B_- \times B \} - \bar{\mathbb{1}} \left(x_{s-1}, y_{s-1}, z_{s-1}, \hat{\theta}_{s-1} \right) \right] \rightarrow 0. \quad (\text{E.10})$$

⁶⁶This is a consequence of the disintegration theorem. See, [Faden \(1985\)](#) for a statement and proof.

On the other hand, equations (E.7) and (E.9) imply that, with \mathbb{P}_- -probability one,

$$\begin{aligned}
& \frac{1}{t} \sum_{s=1}^t \mathbb{1} \left(x_{s-1}, y_{s-1}, z_{s-1}, \hat{\theta}_{s-1} \right) \\
&= \int T(B|x_-, y_-, z_-, \theta_-) \mathbb{1} \{(x_-, y_-, z_-, \theta_-) \in B_-\} \rho(dx_- \times dy_- \times dz_- \times d\theta_-) \\
&= \int_{B_-} T(B|x_-, y_-, z_-, \theta_-) \rho(dx_- \times dy_- \times dz_- \times d\theta_-). \tag{E.11}
\end{aligned}$$

Equations (E.8), (E.10), and (E.11) imply that

$$\begin{aligned}
\rho(B_- \times B) &= \int_{B_-} T(B|x_-, y_-, z_-, \theta_-) \rho(dx_- \times dy_- \times dz_- \times d\theta_-) \\
&= \int_{B_-} T(B|x_-, y_-, z_-, \theta_-) \mu^*(dx_- \times dy_- \times dz_- \times d\theta_-),
\end{aligned}$$

where the second equality is by the definition of μ^* . But by construction,

$$\rho(B_- \times B) = \int_{B_-} M(B|x_-, y_-, z_-, \theta_-) \mu^*(dx_- \times dy_- \times dz_- \times d\theta_-).$$

Since B_- is arbitrary, the last two equalities guarantee that $M(B|x_-, y_-, z_-, \theta_-) = T(B|x_-, y_-, z_-, \theta_-)$ for μ^* -almost all $(x_-, y_-, z_-, \theta_-)$. Noting that $B \subseteq X \times Y \times Z$ was an arbitrary measurable set completes the proof of the lemma. \square

Proof of Part (a) of Theorem 3

Fix an arbitrary $(x_0, y_0, z_0, \hat{\theta}_0)$, an adaptive equilibrium $\{x_t, y_t, z_t, \hat{\theta}_t\}_{t=0}^\infty$ with probability distribution \mathbb{P} and initial condition $(x_0, y_0, z_0, \hat{\theta}_0)$, and a realization of $\{x_t, y_t, z_t, \hat{\theta}_t\}_{t=1}^\infty$ for which (E.6) holds. By Lemma E.3, the set of such $\{x_t, y_t, z_t, \hat{\theta}_t\}_{t=1}^\infty$ has \mathbb{P} -probability one. An application of Lemma E.3 with $f(x_-, y_-, z_-, \theta_-, x, y, z) = \mathbb{1}\{(x_-, y_-, z_-, \theta_-) \in B\}$ leads to

$$\lim_{t \rightarrow \infty} \frac{1}{t} \sum_{s=0}^{t-1} \mathbb{1}\{(x_s, y_s, z_s, \hat{\theta}_s) \in B\} = \mu^*(B).$$

\square

Proof of Part (b) of Theorem 3

Throughout the proof, fix an arbitrary $(x_0, y_0, z_0, \hat{\theta}_0)$, an adaptive equilibrium $\{x_t, y_t, z_t, \hat{\theta}_t\}_{t=0}^\infty$ with probability distribution \mathbb{P} and initial condition $(x_0, y_0, z_0, \hat{\theta}_0)$, and a realization of $\{x_t, y_t, z_t, \hat{\theta}_t\}_{t=1}^\infty$ for which (E.6) holds. By Lemma E.3, the set of such $\{x_t, y_t, z_t, \hat{\theta}_t\}_{t=1}^\infty$ has \mathbb{P} -probability one.

I start by proving that a uniform law-of-large-numbers result holds. By assumption, the mapping $\theta \mapsto \log(q_\theta(y|y_-))$ is continuous for all y_-, y and $\log(q_\theta(y|y_-))$ is bounded. Therefore, by the dominated convergence theorem, for all $\theta \in \Theta$ and $\varepsilon > 0$, there exists an open ball $B_\eta(\theta)$ of radius $\eta = \eta(\theta, \varepsilon)$ centered at θ such that

$$\int \left(\sup_{\vartheta \in B_\eta(\theta)} \left\{ -\log(q_\vartheta(y|y_-)) \right\} - \inf_{\vartheta \in B_\eta(\theta)} \left\{ -\log(q_\vartheta(y|y_-)) \right\} \right) \rho(dy_- \times dy) \leq \varepsilon, \tag{E.12}$$

where ρ is the probability distribution over $(X \times Y \times Z)^2 \times \Theta$ defined as

$$\rho(B_- \times B) \equiv \int_{B_-} T(B|x_-, y_-, z_-, \theta_-) \mu^*(dx_- \times dy_- \times dz_- \times d\theta_-).$$

Since Θ is compact, it can be covered by a finite number J of such balls: $B_{\eta_j}(\theta_j)$, where $j = 1, \dots, J$. Define

$$\gamma_{sj} \equiv \sup_{\vartheta \in B_{\eta_j}(\theta_j)} \left\{ -\log(q_{\vartheta}(y_s|y_{s-1})) \right\}.$$

Since γ_{sj} is bounded, by Lemma E.3, for all $j = 1, \dots, J$,

$$\lim_{t \rightarrow \infty} \frac{1}{t} \sum_{s=1}^t \gamma_{sj} = \int \sup_{\vartheta \in B_{\eta_j}(\theta_j)} \left\{ -\log(q_{\vartheta}(y|y_-)) \right\} \rho(dy_- \times dy). \quad (\text{E.13})$$

For any $t \geq 1$, define

$$H_t(\theta) \equiv \frac{1}{t} \sum_{s=1}^t -\log(q_{\theta}(y_s|y_{s-1}))$$

to be the quasi-log-likelihood function. Note that

$$\begin{aligned} & \sup_{\theta \in \Theta} \left\{ H_t(\theta) - H(Q_{\theta}, T, \mu^*) \right\} \\ &= \max_j \sup_{\vartheta \in B_{\eta_j}(\theta_j)} \left\{ H_t(\vartheta) - H(Q_{\vartheta}, T, \mu^*) \right\} \\ &\leq \max_j \left\{ \frac{1}{t} \sum_{s=1}^t \gamma_{sj} - \int \inf_{\vartheta \in B_{\eta_j}(\theta_j)} \left\{ -\log(q_{\vartheta}(y|y_-)) \right\} \rho(dy_- \times dy) \right\}, \end{aligned} \quad (\text{E.14})$$

where I am using the fact that $\rho(dy_- \times dy) = T(dy|y_-) \mu^*(dy_-)$ by definition. By equation (E.13), for any $\epsilon > 0$ and $j = 1, \dots, J$, there exist some time $\tau_j = \tau_j(\epsilon)$ such that for all $t \geq \tau_j$,

$$\left| \frac{1}{t} \sum_{s=1}^t \gamma_{sj} - \int \sup_{\vartheta \in B_{\eta_j}(\theta_j)} \left\{ -\log(q_{\vartheta}(y|y_-)) \right\} \rho(dy_- \times dy) \right| \leq \epsilon.$$

Therefore, for all $t \geq \tau \equiv \max_j \tau_j$,

$$\begin{aligned} & \max_j \left\{ \frac{1}{t} \sum_{s=1}^t \gamma_{sj} - \int \inf_{\vartheta \in B_{\eta_j}(\theta_j)} \left\{ -\log(q_{\vartheta}(y|y_-)) \right\} \rho(dy_- \times dy) \right\} \\ &\leq \int \left(\sup_{\vartheta \in B_{\eta_j}(\theta_j)} \left\{ -\log(q_{\vartheta}(y|y_-)) \right\} - \inf_{\vartheta \in B_{\eta_j}(\theta_j)} \left\{ -\log(q_{\vartheta}(y|y_-)) \right\} \right) \rho(dy_- \times dy) + \epsilon \\ &\leq \epsilon + \epsilon, \end{aligned}$$

where the second inequality is by (E.12). Thus, by (E.14), for all $t \geq \tau$,

$$\sup_{\theta \in \Theta} \left\{ H_t(\theta) - H(Q_{\theta}, T, \mu^*) \right\} \leq \epsilon + \epsilon.$$

By an identical argument, there exists a time τ' such that for all $t \geq \tau'$,

$$\sup_{\theta \in \Theta} \left\{ H(Q_\theta, T, \mu^*) - H_t(\theta) \right\} \leq \varepsilon + \epsilon.$$

Therefore, since $\varepsilon, \epsilon > 0$ were arbitrary,

$$\limsup_{t \rightarrow \infty} \sup_{\theta \in \Theta} |H_t(\theta) - H(Q_\theta, T, \mu^*)| = 0.$$

This establishes the uniform law-of-large-numbers result.

Since $H_t(\theta) - H(Q_\theta, T, \mu^*) \rightarrow 0$ uniformly in θ , for any $\epsilon > 0$, there exists some τ such that for all $t \geq \tau$,

$$H_t(\theta) - H(Q_\theta, T, \mu^*) + \frac{\epsilon}{2} \quad \forall \theta \in \Theta.$$

In particular, for any $\theta^* \in \arg \min H(Q_\theta, T, \mu^*)$,

$$H_t(\theta^*) < H(Q_{\theta^*}, T, \mu^*) + \frac{\epsilon}{2}.$$

Since $\hat{\theta}_t$ minimizes $H_t(\theta)$,

$$H_t(\hat{\theta}_t) \leq H_t(\theta^*).$$

Thus, for all $t \geq \tau$,

$$H_t(\hat{\theta}_t) < H(Q_{\theta^*}, T, \mu^*) + \frac{\epsilon}{2}.$$

On the other hand, since $H_t(\theta) - H(Q_\theta, T, \mu^*) \rightarrow 0$ uniformly in θ , for all $t \geq \tau$,

$$H(\hat{\theta}_t, T, \mu^*) < H_t(\hat{\theta}_t) + \frac{\epsilon}{2}.$$

The last two displays imply that

$$H(\hat{\theta}_t, T, \mu^*) - H(Q_{\theta^*}, T, \mu^*) < \epsilon$$

Noting that ϵ is arbitrary completes the proof of this part of the theorem. \square

Proof of Part (c) of Theorem 3

I need to show that the triple $(T, \mu^*, \{\theta^*\})$ satisfies Conditions (i)–(iii) for being a CREE. Condition (ii) is trivially satisfied. Part (b) of the theorem—in conjunction with the fact that $H(Q_\theta, T, \mu^*)$ is continuous in θ , established in Lemma 1, and the assumptions that $\arg \min_{\theta \in \Theta} H(Q_\theta, T, \mu^*) = \{\theta^*\}$ and Θ is compact—implies that $\hat{\theta}_t \rightarrow \theta^*$. Part (a) of the theorem then implies that μ^* is supported on $X \times Y \times Z \times \{\theta^*\}$. This conclusion establishes Condition (iii).

So I only need to show that μ^* is an invariant distribution for T . Let $B_1 \subseteq X \times Y \times Z$ and $B_2 \subseteq \Delta\Theta$ be arbitrary measurable sets. Since μ^* is supported on $X \times Y \times Z \times \{\theta^*\}$ and the Bayesian update of the degenerate belief on θ^* is equal to the degenerate belief on θ^* ,

$$\begin{aligned} & \int T(B_1 \times B_2 | x_-, y_-, z_-, \theta_-) \mu^*(dx_- \times dy_- \times dz_- \times d\theta_-) \\ &= \{\mathbf{1}_{\theta^*} \in B_2\} \int T(B_1 | x_-, y_-, z_-, \theta^*) \mu^*(dx_- \times dy_- \times dz_-). \end{aligned}$$

On the other hand, by part (a) of the theorem,

$$\begin{aligned} \int T(B_1|x_-, y_-, z_-, \theta^*)\mu^*(dx_- \times dy_- \times dz_-) &= \lim_{t \rightarrow \infty} \frac{1}{t} \sum_{s=1}^t \mathbb{1}\{(x_s, y_s, z_s) \in B_1\} \\ &= \lim_{t \rightarrow \infty} \frac{1}{t} \sum_{s=1}^t \mathbb{1}\{(x_{s_1}, y_{s-1}, z_{s-1}) \in B_1\} = \mu^*(B_1), \end{aligned}$$

where the last equality is again by part (a). But since μ^* is supported on $X \times Y \times Z \times \{\theta^*\}$,

$$\mu^*(B_1 \times B_2) = \mu^*(B_1)\{\mathbf{1}_{\theta^*} \in B_2\}.$$

Therefore,

$$\int T(B_1 \times B_2|x_-, y_-, z_-, \theta_-)\mu^*(dx_- \times dy_- \times dz_- \times d\theta_-) = \mu^*(B_1 \times B_2).$$

Since B_1 and B_2 are arbitrary, the last display establishes that μ^* is an invariant distribution for T . □

References

- Acemoglu, Daron and Martin Kaae Jensen (2018), "Equilibrium Analysis in the Behavioral Neoclassical Growth Model." Working paper.
- Adam, Klaus, Pei Kuang, and Albert Marcet (2012), "House Price Booms and the Current Account." In *NBER Macroeconomics Annual*, volume 26, 77–122.
- Adam, Klaus and Albert Marcet (2011), "Internal Rationality, Imperfect Market Knowledge and Asset Prices." *Journal of Economic Theory*, 146, 1224–1252.
- Adam, Klaus, Albert Marcet, and Johannes Beutel (2017), "Stock Price Booms and Expected Capital Gains." *American Economic Review*, 107, 2352–2408.
- Adam, Klaus, Albert Marcet, and Juan Pablo Nicolini (2016), "Stock Market Volatility and Learning." *Journal of Finance*, 71, 33–82.
- Afrouzi, Hassan (2017), "Strategic Inattention, Inflation Dynamics and the Non-Neutrality of Money." Working paper.
- Aliprantis, Charalambos D. and Kim C. Border (2006), *Infinite Dimensional Analysis: A Hitchhiker's Guide*. Springer Berlin Heidelberg.
- Alvarez, Fernando E., Francesco Lippi, and Luigi Paciello (2015), "Monetary Shocks in Models with Inattentive Producers." *The Review of economic studies*, 83, 421–459.
- Anderson, Brian D.O. and John B. Moore (2005), *Optimal Filtering*. Dover Publications, Mineola, N.Y.
- Anderson, Evan W., Thomas J. Sargent, and Lars Peter Hansen (2003), "A Quartet of Semigroups for Model Specification, Robustness, Prices of Risk, and Model Detection." *Journal of the European Economic Association*, 1, 68–123.
- Angeletos, George-Marios, Fabrice Collard, and Harris Dellas (2018), "Quantifying Confidence." *Econometrica*, 86, 1689–1726.
- Angeletos, George-Marios and Zhen Huo (2018), "Myopia and Anchoring." Working paper.
- Angeletos, George-Marios and Chen Lian (2018), "Forward Guidance Without Common Knowledge." *American Economic Review*, 108, 2477–2512.
- Barberis, Nicholas, Robin Greenwood, Lawrence Jin, and Andrei Shleifer (2015), "X-CAPM: An Extrapolative Capital Asset Pricing Model." *Journal of Financial Economics*, 115, 1–24.
- Barberis, Nicholas, Andrei Shleifer, and Robert W. Vishny (1998), "A Model of Investor Sentiment." *Journal of Financial Economics*, 49, 307–343.
- Battigalli, Pierpaolo, Simone Cerreia-Vioglio, Fabio Maccheroni, Massimo Marinacci, and Thomas J. Sargent (2016), "A Framework for the Analysis of Self-Confirming Policies." Working paper.
- Benhabib, Jess, Pengfei Wang, and Yi Wen (2015), "Sentiments and Aggregate Demand Fluctuations." *Econometrica*, 83, 549–585.
- Berk, Robert H. (1966), "Limiting Behavior of Posterior Distributions When the Model is Incorrect." *Annals of Mathematical Statistics*, 37, 51–58.
- Bhandari, Anmol, Jaroslav Borovicka, and Paul Ho (2017), "Survey Data and Subjective Beliefs in Business Cycle Models." Working Paper.
- Bordalo, Pedro, Nicola Gennaioli, and Andrei Shleifer (2018), "Diagnostic Expectations and Credit Cycles." *Journal of Finance*, 73, 199–227.
- Branch, William A. (2004), "Restricted Perceptions Equilibria and Learning in Macroeconomics." Working paper.
- Branch, William A. and Bruce McGough (2005), "Consistent Expectations and Misspecification in Stochastic Non-Linear Economies." *Journal of Economic Dynamics and Control*, 29, 659–676.
- Bray, M. M. and N. E. Savin (1986), "Rational Expectations Equilibria, Learning, and Model Specification." *Econometrica*, 54, 1129–1160.

- Bray, Margaret (1982), "Learning, Estimation, and the Stability of Rational Expectations." *Journal of Economic Theory*, 26, 318–339.
- Breiman, Leo (1960), "The Strong Law of Large Numbers for a Class of Markov Chains." *Annals of Mathematical Statistics*, 31, 801–803.
- Bullard, James, George W. Evans, and Seppo Honkapohja (2008), "Monetary Policy, Judgment, and Near-Rational Exuberance." *American Economic Review*, 98, 1163–1177.
- Bullard, James and Kaushik Mitra (2002), "Learning About Monetary Policy Rules." *Journal of Monetary Economics*, 49, 1105–1129.
- Bunke, Olaf and Xavier Milhaud (1998), "Asymptotic Behavior of Bayes Estimates Under Possibly Incorrect Models." *Annals of Statistics*, 26, 617–644.
- Cho, In Koo, Noah Williams, and Thomas J. Sargent (2002), "Escaping Nash Inflation." *Review of Economic Studies*, 69, 1–40.
- Christiano, Lawrence J., Martin Eichenbaum, and Charles L. Evans (2005), "Nominal Rigidities and the Dynamic Effects of a Shock to Monetary Policy." *Journal of Political Economy*, 113, 1–45.
- Coibion, Olivier and Yuriy Gorodnichenko (2012), "What Can Survey Forecasts Tell us About Information Rigidities?" *Journal of Political Economy*, 120, 116–159.
- Coibion, Olivier and Yuriy Gorodnichenko (2015a), "Information Rigidity and the Expectations Formation Process: A Simple Framework and new Facts." *American Economic Review*, 105, 2644–78.
- Coibion, Olivier and Yuriy Gorodnichenko (2015b), "Is the Phillips Curve Alive and Well After All? Inflation Expectations and the Missing Disinflation." *American Economic Journal: Macroeconomics*, 7, 197–232.
- Cooper, Russell W. and John C. Haltiwanger (2006), "On the Nature of Capital Adjustment Costs." *The Review of Economic Studies*, 73, 611–633.
- Diaconis, Persi and David Freedman (1986), "On Inconsistent Bayes Estimates of Location." *Annals of Statistics*, 68–87.
- Ellis, Andrew and Michele Piccione (2017), "Correlation Misperception in Choice." *American Economic Review*, 107, 1264–1292.
- Ellsberg, Daniel (1961), "Risk, Ambiguity, and the Savage Axioms." *The Quarterly Journal of Economics*, 75, 643.
- Epstein, Larry and Martin Schneider (2007), "Learning Under Ambiguity." *Review of Economic Studies*, 74, 1275–1303.
- Epstein, Larry G. and Martin Schneider (2003), "Recursive Multiple-Priors." *Journal of Economic Theory*, 113, 1–31.
- Esponda, Ignacio and Demian Pouzo (2016a), "Berk-Nash Equilibrium: A Framework for Modeling Agents with Misspecified Models." *Econometrica*, 84, 1093–1130.
- Esponda, Ignacio and Demian Pouzo (2016b), "Equilibrium in Misspecified Markov Decision Processes." Working paper.
- Eusepi, Stefano and Bruce Preston (2011), "Expectations, Learning, and Business Cycle Fluctuations." *American Economic Review*, 101, 2844–2872.
- Eusepi, Stefano and Bruce Preston (2018a), "Fiscal Foundations of Inflation: Imperfect Knowledge." *American Economic Review*, 108, 2551–89.
- Eusepi, Stefano and Bruce Preston (2018b), "The Science of Monetary Policy: An Imperfect Knowledge Perspective." *Journal of Economic Literature*, 56, 3–59.
- Evans, George W. and Seppo Honkapohja (1995), "Local Convergence of Recursive Learning to Steady States and Cycles in Stochastic Nonlinear Models." *Econometrica*, 63, 195–206.
- Evans, George W. and Seppo Honkapohja (2012), *Learning and Expectations in Macroeconomics*. Princeton University Press.
- Evans, George W. and Garey Ramey (2006), "Adaptive Expectations, Underparameterization and the Lucas Critique." *Journal of Monetary Economics*, 53, 249–264.

- Eyster, Erik, Kristof Madarasz, and Pascal Michailat (2017), "Pricing When Customers Care about Fairness but Misinfer Markups." Working paper.
- Eyster, Erik and Michele Piccione (2013), "An Approach to Asset Pricing Under Incomplete and Diverse Perceptions." *Econometrica*, 81, 1483–1506.
- Eyster, Erik and Matthew Rabin (2005), "Cursed Equilibrium." *Econometrica*, 73, 1623–1672.
- Faden, Arnold M. (1985), "The Existence of Regular Conditional Probabilities: Necessary and Sufficient Conditions." *Annals of Probability*, 13, 288–298.
- Farhi, Emmanuel and Iván Werning (2017), "Monetary Policy, Bounded Rationality and Incomplete Markets."
- Fudenberg, Drew and David K. Levine (1993), "Self-Confirming Equilibrium." *Econometrica*, 61, 523–545.
- Fudenberg, Drew and David K. Levine (2009), "Self-Confirming Equilibrium and the Lucas Critique." *Journal of Economic Theory*, 144, 2354–2371.
- Fudenberg, Drew, Gleb Romanyuk, and Philipp Strack (2017), "Active Learning with a Misspecified Prior." *Theoretical Economics*, 12, 1155–1189.
- Fuster, Andreas, Benjamin Hebert, and David Laibson (2012), "Investment Dynamics with Natural Expectations." *International Journal of Central Banking*, 8, 243–265.
- Fuster, Andreas, David Laibson, and Brock Mendel (2010), "Natural Expectations and Macroeconomic Fluctuations." *Journal of Economic Perspectives*, 24, 67–84.
- Futia, Carl A. (1982), "Invariant Distributions and the Limiting Behavior of Markovian Economic Models." *Econometrica*, 50, 377–408.
- Gabaix, Xavier (2016), "A Behavioral New Keynesian Model." Working paper.
- Gagnon-Bartsch, Tristan, Matthew Rabin, and Joshua Schwartzstein (2018), "Channeled Attention and Stable Errors." Working paper.
- Galí, Jordi (2008), *Monetary Policy, Inflation, and the Business Cycle*. Princeton University Press.
- Garcia-Schmidt, Mariana and Michael Woodford (2015), "Are Low Interest Rates Deflationary? A Paradox of Perfect-Foresight Analysis." Working paper.
- Gennaioli, Nicola and Andrei Shleifer (2010), "What Comes to Mind." *Quarterly Journal of Economics*, 125, 1399–1433.
- Gerko, Elena (2017), "Expectations and Monetary Policy." Working paper.
- Gevers, Michel and Vincent Wertz (1984), "Uniquely Identifiable State-Space and ARMA Parametrizations for Multivariable Linear Systems." *Automatica*, 20, 333–347.
- Gilboa, Itzhak and David Schmeidler (1989), "Maxmin Expected Utility with Non-Unique Prior." *Journal of Mathematical Economics*, 18, 141–153.
- Grandmont, Jean-Michel (1977), "Temporary General Equilibrium Theory." *Econometrica*, 45, 535–572.
- Grandmont, Jean-Michel (1982), "Temporary General Equilibrium Theory." *Handbook of Mathematical Economics*, 2, 879–922.
- Gray, Robert M. (2009), *Probability, Random Processes, and Ergodic Properties*, second edition. Springer Publishing Company, Incorporated.
- Gray, Robert M. (2011), *Entropy and Information Theory*. Springer Science & Business Media.
- Hamilton, James D. (1994), *Time Series Analysis*. Princeton University Press.
- Hansen, Lars Peter and Thomas J. Sargent (1995), "Discounted Linear Exponential Quadratic Gaussian Control." *IEEE Transactions on Automatic control*, 40, 968–971.

- Hansen, Lars Peter and Thomas J. Sargent (2001a), "Acknowledging Misspecification in Macroeconomic Theory." *Review of Economic Dynamics*, 4, 519–535.
- Hansen, Lars Peter and Thomas J. Sargent (2001b), "Robust Control and Model Uncertainty." *American Economic Review*, 91, 60–66.
- Hansen, Lars Peter and Thomas J. Sargent (2005), "Robust Estimation and Control Under Commitment." *Journal of Economic Theory*, 124, 258–301.
- Hansen, Lars Peter and Thomas J. Sargent (2007), "Recursive Robust Estimation and Control Without Commitment." *Journal of Economic Theory*, 136, 1–27.
- Hansen, Lars Peter and Thomas J. Sargent (2008), *Robustness*. Princeton University Press.
- Hansen, Lars Peter and Thomas J. Sargent (2017), "Prices of Macroeconomic Uncertainties with Tenuous Beliefs."
- Hansen, Lars Peter, Thomas J. Sargent, and Thomas D. Jr. Tallarini (1999), "Robust Permanent Income and Pricing." *Review of Economic Studies*, 66, 873–907.
- Hansen, Lars Peter, Thomas J. Sargent, Gauhar Turmuhambetova, and Noah Williams (2006), "Robust Control and Model Misspecification." *Journal of Economic Theory*, 128, 45–90.
- Hansen, Lars Peter, Balint Szoke, Lloyd S. Han, and Thomas J. Sargent (2018), "Twisted Probabilities, Uncertainty, and Prices."
- Hayashi, Fumio (1982), "Tobin's Marginal Q and Average Q: A Neoclassical Interpretation." *Econometrica*, 213–224.
- Heidhues, Paul, Botond Koszegi, and Philipp Strack (2018), "Unrealistic Expectations and Misguided Learning." *Econometrica*, 86, 1159–1214.
- Hicks, John R. (1939), *Value and Capital*. Oxford: Clarendon.
- Hildenbrand, Werner and Jean-Michel Grandmont (1974), "Stochastic Processes of Temporary Equilibria." *Journal of Mathematical Economics*, 1, 247–277.
- Hirshleifer, David (2001), "Investor Psychology and Asset Pricing." *The Journal of Finance*, 56, 1533–1597.
- Hirshleifer, David, Jun Li, and Jianfeng Yu (2015), "Asset Pricing in Production Economies With Extrapolative Expectations." *Journal of Monetary Economics*, 76, 87–106.
- Hommes, Cars and Gerhard Sorger (1998), "Consistent Expectations Equilibria." *Macroeconomic Dynamics*, 2, 287–321.
- Huber, Peter J. (1967), "The Behavior of Maximum Likelihood Estimates Under Nonstandard Conditions." In *Proceedings of the Fifth Berkeley Symposium on Mathematical Statistics and Probability*, volume 1, 221–233, University of California Press.
- Ilut, Cosmin L and Hikaru Saijo (2016), "Learning, Confidence, and Business Cycles." Working paper.
- Ilut, Cosmin L. and Martin Schneider (2014), "Ambiguous Business Cycles." *American Economic Review*, 104, 2368–99.
- Jamison, Benton (1965), "Ergodic Decompositions Induced by Certain Markov Operators." *Transactions of the American Mathematical Society*, 117, 451–468.
- Jehiel, Philippe (2005), "Analogy-Based Expectation Equilibrium." *Journal of Economic Theory*, 123, 81–104.
- Justiniano, Alejandro, Giorgio E. Primiceri, and Andrea Tambalotti (2010), "Investment Shocks and Business Cycles." *Journal of Monetary Economics*, 57, 132–145.
- Kahneman, Daniel and Amos Tversky (1972), "Subjective Probability: A Judgment of Representativeness." *Cognitive Psychology*, 3, 430–454.
- Kahneman, Daniel and Amos Tversky (1982a), "Judgment under Uncertainty: Heuristics and Biases." *Science*, 185, 1124–1131.
- Kahneman, Paul Slovic, Daniel and Amos Tversky (1982b), *Judgment under Uncertainty: Heuristics and Biases*. Cambridge University Press.

- Karantounias, Anastasios G. (2013), "Managing Pessimistic Expectations and Fiscal Policy." *Theoretical Economics*, 8, 193–231.
- Kleijn, B. J. K. and A. W. Van Der Vaart (2006), "Misspecification in Infinite-Dimensional Bayesian Statistics." *Annals of Statistics*, 34, 837–877.
- Kreps, David M. (1998), "Anticipated Utility and Dynamic Choice." *Econometric Society Monographs*, 29, 242–274.
- Krusell, Per and Anthony A. Smith, Jr. (1998), "Income and Wealth Heterogeneity in the Macroeconomy." *Journal of Political Economy*, 106, 867–896.
- Lindahl, Erik (1939), *Studies in the Theory of Money and Capital*. London: Allen & Unwin.
- Loève, Michel (1960), *Probability Theory*, second edition. D. Van Nostrand, Princeton, NJ.
- Maccheroni, Fabio, Massimo Marinacci, and Aldo Rustichini (2006), "Ambiguity Aversion, Robustness, and the Variational Representation of Preferences." *Econometrica*, 74, 1447–1498.
- Maćkowiak, Bartosz, Filip Matejka, and Mirko Wiederholt (2017), "The Rational Inattention Filter." Working paper.
- Maćkowiak, Bartosz and Mirko Wiederholt (2009), "Optimal Sticky Prices under Rational Inattention." *American Economic Review*, 99, 769–803.
- Maćkowiak, Bartosz and Mirko Wiederholt (2015), "Business Cycle Dynamics under Rational Inattention." *Review of Economic Studies*, 82, 1502–1532.
- Malmendier, Ulrike and Stefan Nagel (2016), "Learning From Inflation Experiences." *Quarterly Journal of Economics*, 131, 53–87.
- Mankiw, N. Gregory and Ricardo Reis (2002), "Sticky Information versus Sticky Prices: A Proposal to Replace the New Keynesian Phillips Curve." *Quarterly Journal of Economics*, 117, 1295–1328.
- Marcet, Albert and Juan P. Nicolini (2003), "Recurrent Hyperinflations and Learning." *American Economic Review*, 93, 1476–1498.
- Marcet, Albert and Thomas J. Sargent (1989a), "Convergence of Least-Squares Learning in Environments with Hidden State Variables and Private Information." *Journal of Political Economy*, 97, 1306–1322.
- Marcet, Albert and Thomas J. Sargent (1989b), "Convergence of Least Squares Learning Mechanisms in Self-Referential Linear Stochastic Models." *Journal of Economic Theory*, 48, 337–368.
- Marinacci, Massimo (2002), "Learning from Ambiguous Urns." *Statistical Papers*, 43, 143–151.
- Milani, Fabio (2007), "Expectations, Learning and Macroeconomic Persistence." *Journal of Monetary Economics*, 54, 2065–2082.
- Molavi, Pooya (2018), "Tests of Subjective Rationality." Working paper.
- Nelsen, Roger B. (2006), *An Introduction to Copulas (Springer Series in Statistics)*, second edition. Springer-Verlag, Berlin, Heidelberg.
- Nimark, Kristoffer (2008), "Dynamic Pricing and Imperfect Common Knowledge." *Journal of Monetary Economics*, 55, 365–382.
- Nyarko, Yaw (1991), "Learning in Mis-specified Models and the Possibility of Cycles." *Journal of Economic Theory*, 55, 416–427.
- Orphanides, Athanasios and John C. Williams (2005), "Inflation Scares and Forecast-Based Monetary Policy." *Review of Economic Dynamics*, 8, 498–527.
- Pfanzagl, J. (1969), "On the Measurability and Consistency of Minimum Contrast Estimates." *Metrika*, 14, 249–272.
- Preston, Bruce (2005), "Learning about Monetary Policy Rules when Long-Horizon Expectations Matter." *International Journal of Central Banking*.

- Primiceri, Giorgio E. (2006), "Why Inflation Rose and Fell: Policy-Makers' Beliefs and U.S. Postwar Stabilization Policy." *Quarterly Journal of Economics*, 121, 867–901.
- Rabin, Matthew (2002), "Inference By Believers in the Law of Small Numbers." *Quarterly Journal of Economics*, 775–816.
- Rabin, Matthew and Joel L. Schrag (1999), "First Impressions Matter: A Model of Confirmatory Bias." *Quarterly Journal of Economics*, 114, 37–82.
- Rabin, Matthew and Dimitri Vayanos (2010), "The Gambler's and Hot-Hand Fallacies: Theory and Applications." *Review of Economic Studies*, 77, 730–778.
- Sala, Luca (2015), "DSGE Models in the Frequency Domains." *Journal of Applied Econometrics*, 30, 219–240.
- Sargent, Thomas J. (1993), *Bounded Rationality in Macroeconomics*. Oxford University Press.
- Sargent, Thomas J. (1999), *The Conquest of American Inflation*. Princeton University Press.
- Shalizi, Cosma Rohilla (2009), "Dynamics of Bayesian Updating with Dependent Data and Misspecified Models." *Electronic Journal of Statistics*, 3, 1039–1074.
- Shmaya, Eran and Leeat Yariv (2016), "Experiments on Decisions Under Uncertainty: A Theoretical Framework." *American Economic Review*, 106, 1775–1801.
- Sims, Christopher A. (2003), "Implications of Rational Inattention." *Journal of Monetary Economics*, 50, 665–690.
- Sklar, A. (1959), "Fonctions de Répartition À N Dimensions Et Leurs Marges." *Publ. Inst. Statist. Univ. Paris*, 8, 229–231.
- Slobodyan, Sergey and Raf Wouters (2012), "Learning in an Estimated Medium-Scale DSGE Model." *Journal of Economic Dynamics and Control*, 36, 26–46.
- Smets, Frank and Rafael Wouters (2007), "Shocks and Frictions in US Business Cycles: A Bayesian DSGE Approach." *American Economic Review*, 97, 586–606.
- Spiegler, Ran (2016), "Bayesian Networks and Boundedly Rational Expectations." *Quarterly Journal of Economics*, 131, 1243–1290.
- Stock, James H. and Mark W. Watson (2016), *Dynamic Factor Models, Factor-Augmented Vector Autoregressions, and Structural Vector Autoregressions in Macroeconomics*, first edition, volume 2. Elsevier B.V.
- Stokey, Nancy and Robert E. Jr. Lucas (1989), *Recursive Methods in Economic Dynamics*. Harvard University Press.
- Strzalecki, Tomasz (2011), "Axiomatic Foundations of Multiplier Preferences." *Econometrica*, 79, 47–73.
- Vuong, Quang H. (1989), "Likelihood Ratio Tests for Model Selection and Non-Nested Hypotheses." *Econometrica*, 57, 307–333.
- Werning, Iván (2015), "Incomplete Markets and Aggregate Demand." Working paper.
- White, Halbert (1982), "Maximum Likelihood Estimation of Misspecified Models." *Econometrica*, 50, 1–25.
- White, Halbert (1994), *Estimation, Inference and Specification Analysis*. Econometric Society Monographs, Cambridge University Press.
- Woodford, Michael (2003a), "Imperfect Common Knowledge and the Effects of Monetary Policy." *Knowledge, Information, and Expectations in Modern Macroeconomics: In Honor of Edmund S. Phelps*, 25.
- Woodford, Michael (2003b), *Interest and Prices: Foundations of a Theory of Monetary Policy*. Princeton University Press.
- Woodford, Michael (2010), "Robustly Optimal Monetary Policy with Near-Rational Expectations." *American Economic Review*, 100, 274–303.
- Woodford, Michael (2013), "Macroeconomic Analysis Without the Rational Expectations Hypothesis." *Annual Review of Economics*, 5, 303–346.
- Woodford, Michael (2018), "Monetary Policy Analysis When Planning Horizons Are Finite." Working paper.
- Zimper, Alexander and Wei Ma (2017), "Bayesian Learning with Multiple Priors and Nonvanishing Ambiguity." *Economic Theory*, 64, 409–447.