

Local rationality*

David Evans
University of Oregon

Jungang Li
University of Oregon

Bruce McGough
University of Oregon

November 1, 2022

Abstract

A new behavioral concept, local rationality, is developed within the context of a simple heterogeneous-agent model with incomplete markets. To make savings decisions, agents forecast the shadow price of asset holdings. Absent aggregate uncertainty, *locally rational agents* forecast shadow prices rationally, and thereby make optimal state-contingent decisions. They use adaptive learning to extend their forecasts to accommodate aggregate uncertainty. Over time the state evolves to an ergodic distribution near the economy's restricted perceptions equilibrium. In a partial equilibrium environment we develop intuition for locally rational decision making, documenting an important hysteresis effect. General equilibrium dynamics are examined via a calibration exercise. Calibrated representative-agent RBC models induce low consumption volatility relative to the data. Extending the model by either incorporating adaptive learning or heterogeneous agents fails to alter this conclusion. Via the hysteresis effect, local rationality, which interacts heterogeneity and adaptive learning, significantly improves the model's fit along this dimension.

JEL Classifications: E31; E32; E52; D84; D83

Key Words: Bounded rationality; real business cycles; heterogenous agent models; adaptive learning

1 Introduction

Aiyagari (1994) introduced uninsurable idiosyncratic risk into a real economy with capital in his work on precautionary savings motives; in doing so, he illustrated the potential of Bewley models

*We thank the Editor and two referees, as well as George Evans and participants at the University of Connecticut Macroeconomics Seminar, the University of California - Riverside Macroeconomics seminar, and the University of Oregon Macro Group for comments. We thank Giorgi Nikolaishvili for research assistance.

to serve as laboratories for the study of incomplete markets in general equilibrium environments. By developing the needed technical machinery to incorporate aggregate risk into Aiyagari's model, Krusell and Smith (1998) (KS) realized this illustrated potential, and the era of *heterogeneous agent* (HA) models had arrived. HA models now figure prominently in the standard canon of first-year courses in macroeconomics and in the standard toolkit of working macroeconomists.

HA models allow economists to relax the rigid representative agent (RA) assumption and thereby consider the dynamics of wealth distributions, and the models have had some success explaining the distributional dynamics observed in the data. However, HA models generally fare no better than RA models when confronting stylized business cycle facts. In fact, in many cases, an HA model calibrated to the same long-run moments as an RA model will feature nearly identically business cycle moments. For example, in a standard RBC environment both RA and HA models are qualitatively successful in their prediction of consumption smoothing, but they are quantitatively off: fully rational agents, regardless of the economic environment, smooth consumption more than is evidenced in the data.¹

Theoretical concerns also challenge the heterogeneous-agent modeling paradigm. HA models are almost ubiquitously anchored to the rational expectations (RE) hypothesis, and the myriad criticisms leveled at the assumption of rational expectations as a behavioral primitive apply with magnified vigor when models include heterogeneity of the type under consideration here. A singularly damning criticism involves optimal decision making in the HA environment. To solve their dynamic program, and thereby make fully rational choices, an agent must understand the transition dynamics of their state, which in an HA model includes the wealth distribution. This distribution is an infinite dimensional object with transition dynamics whose very existence remains only speculative.² How are we to take seriously a model that presupposes full knowledge of a complex object that the modelers can't even prove exists?

Bounded rationality, broadly interpreted to include boundedly optimal decision making, provides a natural behavioral alternative to rational expectations. Economic agents are modeled as small players in a big world, adhering to a collection of behavioral primitives governing decision making that are designed to be more realistic than their rational counterparts. Agents' attendant actions are coordinated each period via market clearing; these temporary equilibrium outcomes then determine the dynamics of economic aggregates; and the short and long run patterns exhibited by these aggregate dynamics comprise the implications of the model.

Our goal, in this paper, is to develop a behavioral approach that addresses the theoretical and empirical challenges to HA models populated by rational agents. We study bounded rationality in an HA environment using a standard heterogeneous-agent model along the lines of Krusell and Smith (1998) by introducing a new behavioral concept – *local rationality*. Informally, a locally rational agent knows (i.e. has already learned) how to respond optimally to idiosyncratic (local)

¹Stadler (1994) provides a nice survey of various representative-agent RBC implementations and their empirical successes and failures. Krusell and Smith (1998) established the inability of RBC-type HA models to overcome the corresponding inability of RA models to match certain aggregate moments.

²We know of no global general existence results for HA models; however, a recent and exciting contribution of Cao (2020) demonstrates the existence of REE in the Krusell-Smith model.

shocks, but must learn how to account for aggregate (global) shocks. The rationale for this is that it's reasonable to assume agents know how to forecast and respond well to their individual states, but they are less certain about how the aggregate state vector evolves and how its evolution should inform their decisions. We use the *shadow price approach*, developed in Evans and McGough (2020b), to model how agents respond to variation in aggregates.

The advantage of the shadow price approach is the simplicity it affords agents: they make decisions based on perceived trade offs which they measure using shadow prices, and they form expectations of future shadow prices using simple linear forecast rules, or *perceived laws of motion* (PLMs), which they update over time using learning algorithms. Agents are not required to understand the evolution of the economy's states nor are they required to solve complex, nonlinear dynamic programs. We follow the literature on adaptive learning by assuming our agents use *recursive least squares* (RLS), which adjusts coefficient estimates in accordance with the most recent forecast error. The weight placed on the forecast error is controlled by the algorithm's *gain*, which may be taken as constant over time, or decreasing at a suitable rate.

The use of a recursive least squares motivates a natural long-run solution concept for this environment: a *restricted perceptions equilibrium* (RPE). An RPE is a natural generalization of a *rational expectations equilibrium* (REE) to environments in which agents use misspecified linear forecast models when making decisions. Under *restricted perceptions*, agents are assumed to restrict consideration to forecast models having the same functional form as their perceived laws of motion. In an RPE, agents are *not* updating their forecast models; instead, each agent is using the optimal forecast model among those under consideration, where optimality here is measured in terms of long-run expected squared forecast error. We note that, in a linearized, representative-agent modeling environment, if agents are using a PLM that has the same functional form as the model's rational expectations equilibrium, then the RPE is the REE. In this sense, a rational expectations equilibrium is a special case of a restricted perceptions equilibrium.

A restricted perceptions equilibrium is a Nash concept: the optimality of a given agent's forecast model depends on the forecast models being used by other agents. The RPE thus serves as a natural anchor for long run equilibrium behavior in models under adaptive learning. If the model's RPE is *stable* then decreasing gain algorithms result in convergence of agents' beliefs to the RPE. Constant-gain algorithms, which remain alert to changing economic conditions, are natural alternatives for agents concerned with forecast-model misspecification. Learning algorithms of this type cannot, even asymptotically, lead to constant beliefs; however, even when constant-gain algorithms are used, an RPE a natural object to study as the long ergodic distribution of beliefs is generally near the RPE. While an RPE does not a priori impose that all agents use the same forecast model, we find that, at least for small shocks, homogeneity is a feature of any RPE of our model.

To flesh out the behavioral implications of local rationality, we consider a *locally rational* (LR) agent acting in a partial equilibrium environment. To engender tractability we adopt CARA preferences in consumption. The agent owns a stochastic income flow and makes consumption/savings decisions in face of a (stationary) stochastic return. In this environment, the manifestation of local rationality is that the agent knows how to make optimal decisions if the return is held permanently

at its mean, but must learn how to respond to variations in the return.

Because of our simplifying assumptions, we are able to obtain a closed form solution to the agent's decisions when returns are constant; and using this form, together with first-order analysis of the agent's behavior when returns are stochastic, we are able to carefully compare the behavior of our LR agent to that of a fully rational agent. A particularly important feature of LR behavior is a hysteresis effect. The response of fully rational agents to changes in return are sensitive to their contemporaneous asset holdings. On the other hand, the response of an LR agent to a change in returns is determined by beliefs, which are relatively insensitive to contemporaneous asset holdings, instead reflecting the agent's accumulated experience. Even when LR agents are wealthy, they remember being poor, and their behavior reflects it.

To assess the general equilibrium implications of locally rationality, we study a calibrated heterogeneous agent economy. To isolate the interaction of bounded rationality and heterogeneous agents, we compare outcomes to those of a representative agent model populated with either rational or learning agents, and a heterogeneous agent model populated with fully rational agents. All models are calibrated to match the same long run moments.³ To simplify our exposition we focus on a single business cycle moment: the relative volatility of consumption and output. A well known issue with RA real business cycle models is that they fail to match this moment.

Our results confirm this finding in the literature: in the data $\text{std}(C)/\text{std}(Y) = 0.50$ whereas in the RA calibration with a rational agent that ratio is 0.32. Neither bounded rationality nor heterogeneous agents can alter this ratio much on their own: in the RA model with shadow price learning⁴ the ratio varied from 0.32 to 0.34 while in the HA model under RE it was 0.36. However, when bounded rationality and heterogeneity are combined in the locally rational model it is possible to exactly match this ratio with a gain consistent with the literature.

To understand the interaction of bounded rationality and heterogeneity, recall that a key feature of the RPE is that all agents have the same beliefs when forecasting their future shadow price of savings. In a representative agent model, this feature is innocuous as all agents are ex-post identical. The same is not true in a heterogeneous agent model: the effect of an aggregate shock differs across agents since agents differ both in their exposure to, and in their ability to smooth, aggregate shocks. Poorer agents are more exposed to recessions than their richer counterparts, and as a consequence, their consumption is more responsive to aggregate shocks. Under local rationality, this leads to the hysteresis effect mentioned above. Rich agents remember what it was like to be poor during a recession, and they adjust their forecasts and behaviors accordingly: the consumption of

³An advantage of our approach is that absent aggregate shocks the behavior of our agents are identical to those of the rational expectations model which results in the same calibrated parameters for preferences and technology.

⁴We focus on shadow price learning in part to maintain comparability with our locally rational model. In fact, in the limit as the size of the idiosyncratic shocks approaches zero the locally rational model reduces to the RA model with shadow price learning. Other learning assumptions produce similar results: Williams (2003) introduced CGL into a real business cycle model and concluded that it was largely ineffective as a resolution to moment discrepancies. In part because of his findings, most of the subsequent work on matching moments using models with learning agents was conducting either in new Keynesian environments or in asset pricing models. Eusepi and Preston (2011) is an important exception, though they too were unable to match consumption volatility relative to output volatility as obtained from the data.

wealthy locally rational agents is more volatile than their wealthy rational counterparts. Since aggregate consumption is driven by the consumption patterns of wealthy agents, this higher volatility is also realized on aggregate.

Under constant gain learning, agents adjust their beliefs over time. In the extreme limit when the gain is very small, and the locally rational economy is approximately at the RPE, agents place near equal weight on all of their past experience, which results in excess consumption volatility relative to the data: $\text{std}(C)/\text{std}(Y) = 0.70$. As the gain parameter is increased, agents forget their past experience at a faster rate and better adapt their behavior to their current context, which brings the behavior of the locally rational model more in line with the rational expectations calibration. As a result, the gain parameter has some surprising effects on business cycle moments. In a representative agent learning model, increasing the gain parameter has the predictable effect of increasing the volatility of endogenous variables through the additional volatility of beliefs. As we have noted this effect is often small relative to the baseline volatility of the RE model. In our locally rational model, through the mechanism described above, increasing the gain parameters significantly *reduces* the volatility of some endogenous variables (consumption) and raises the volatility of others (investment). While we focused on a simple RBC style model, we expect this mechanism to apply more generally.

1.1 Related Literature

To our knowledge, implementation of bounded rationality in HA models is limited to Giusto (2014). To explain his findings and how they bear on our efforts, it helps to first review the KS computational technique, which, as it happens, uses an RPE to approximate the model's REE. The state vector of an HA model includes the economy's wealth distribution, a high-dimensional object that is not feasibly tracked. KS postulate that it is necessary only to track a finite number of moments of this distribution, thus simplifying the analysis. In fact, they argue that the first moment is enough, i.e. it is sufficient for agents in the economy to forecast the aggregate capital stock, and to do this, agents use a simple, linear forecast model. For a fixed parameterization of this forecast model, agents' behavior, and thus the implied dynamics of aggregate capital, can be computed; and, via linear projection, the corresponding *optimal linear forecast model* can be determined. KS declare victory when agents' forecast model aligns with the optimal one, i.e. when the economy is in an RPE.

Giusto (2014) adds adaptive learning to the methodology of KS. In Giusto's world, agents update their linear forecast model as new data become available. Other than that, agents act exactly as they do in KS: each period agents fully solve their dynamic programming (DP) problem taking their beliefs as given. This behavioral premise is known as the *anticipated utility approach*, originally due to Kreps (1998), and is similar in spirit to the long horizon approach of Preston (2005). Giusto shows that the KS RPE is stable under adaptive learning, and that learning implemented using a decreasing gain algorithm allows the model to better match the dynamics of wealth distributions.

Our approach departs from Giusto in two important ways. First, our agents are not assumed

to be able to solve DP problems with many aggregate states. Instead, they follow the cognitively less demanding primitives laid out above to make decisions. Of particular note is that our agents, in effect, repeatedly solve simple two-period problems based one-step-ahead forecasts; Guisto's agents solve DP problems based on forecasts at arbitrarily long horizons, and they must resolve these programs every period. Second, we are not analyzing the stability of the REE, or, more accurately, the KS RPE. The RPE we analyze is distinct from the KS RPE, and, as discussed in Section 6, this distinction has important implications for moment matching.

As mentioned above, our implementation of boundedly optimal decision making is based on the shadow-price approach, which is one of several mechanisms in the literature that link boundedly rational forecasting and boundedly optimal decision-making. Others include Euler-equation learning found in Evans and Honkapohja (2006), the long horizon approach emphasized in Preston (2005), and the sparse programming approach found in Gabaix (2017) and Gabaix (2020).⁵

Restricted perceptions equilibria have a venerated history in macroeconomics, particularly in relation to boundedly rational behavior. The nomenclature was introduced in Evans and Honkapohja (2001) but the concept is older: see Branch (2006) for a survey. Early work on the topic involved linear environments in which the forecast model misspecification involved omitted variables: Marcet and Sargent (1989), Sargent (1991), Evans, Honkapohja, and Sargent (1993) and Bullard, Evans, and Honkapohja (2008) study forecast rules that omit informative lags; and in Branch and Evans (2006a), Branch and Evans (2007) and Adam (2007) the forecast models omit relevant explanatory variables. Some recent work has linear forecast models in non-linear environments, which is more closely related to the concepts pursued in this paper. In a non-linear real business cycle environment, Evans, Evans, and McGough (2021a) demonstrate the existence of RPE associated with linear forecast models.⁶ Hommes and Zhu (2014) introduce the closely related concept of *behavioral learning equilibria*, which casts agents as using simple AR(1) forecast models in complex economic environments. RPE have also been central in a number of empirical DSGE models, e.g. Slobodyan and Wouters (2012).

There has been considerable research done on heterogeneous expectations in macroeconomic models. Early work includes Honkapohja and Mitra (2006), who consider the impact of expectations heterogeneity on equilibrium stability in a complete markets model, and apply their results to policy considerations in a new Keynesian model. Branch and McGough (2009) develop a tractable new Keynesian model with heterogeneous expectations; and Gasteiger (2018) and Anufriev, Assenza, Hommes, and Massaro (2013) explore the policy ramifications of heterogeneous expectations in neo-classical economies. See Branch and McGough (2018) for a survey.

The paper is organized as follows. Section 2 develops with care the modeling environment under rationality. Section 3 modifies the modeling environment to allow for local rationality, and

⁵See Branch, Evans, and McGough (2013) and Woodford (2018) for approaches that involve finite planning horizons. See Hommes (2013) for a broad exposition on behavioral models of the macroeconomy.

⁶See also Evans and McGough (2020c) and Evans and McGough (2020a). Evans and McGough (2018) consider the case in which exogenous variables are unobserved and use autoregressions or VARs as forecast models. Branch, McGough, and Zhu (2021) combine non-observability of exogenous shocks with the presence of observable sunspot processes to demonstrate the existence of stable RPE even in models that are determinate under RE.

includes a detailed discussion of restricted perceptions equilibrium. Section 4 conducts an in-depth examination of locally rational behavior in a partial equilibrium environment. Section 5 provides the calibration details and the methods used for our numerical work. Section 6 presents our computational evidence for existence and stability of the model's RPE, and discusses the results of our calibration exercise. Section 7 concludes.

2 The rational model

To construct our concept of local rationality we use a standard heterogeneous agent environment in the style of Aiyagari (1994), which we augment to include endogenous labor choice, as well as aggregate shocks in the spirit of Krusell and Smith (1998). In this section we adopt the usual behavioral assumption that agents are fully rational. In Section 3 we use this development as a platform to introduce and motivate local rationality as an alternative behavioral assumption.

Heterogeneous agent models with rational agents are, by now, so commonplace in the literature that their presentation is often high-level and brief, with emphasis placed only on the novelty under examination. The reader typically is assumed sufficiently familiar with the many technical details that they can either proceed with confidence of the model's internal consistency or they can work through the analysis themselves. As our work here re-imagines the agents' behavioral primitives, it is ground-level and necessarily detailed. To motivate our modified primitives and to facilitate comparison to the benchmark case, we develop the well-known rational model in a manner that is better aligned with the boundedly rational assumptions.⁷

2.1 The household problem

The household's decision problem is recursive, and under rationality it can be naturally framed using a time-invariant Bellman system; however, to motivate the behavioral primitives adopted in the boundedly rational case, it is more natural to characterize agent behavior sequentially via their first-order conditions.

Time is discrete. There is a unit mass of agents who are identical up to idiosyncratic wage shocks. Each agent is endowed with one unit of labor/leisure per period and measures their flow utility as a function of their current consumption c and leisure l with utility function $u(c, l)$. Different agents have different efficiency units of labor per hour worked. In return for supplying labor, each agent receives a wage that can be separated into two parts: an aggregate component w that is the same across all agents; and an idiosyncratic efficiency component ε that is independent and identically distributed across all agents. We assume that $\{\varepsilon\}$ is a Markov process with time-invariant transition function Π . An agent cannot fully insure against this idiosyncratic risk, but

⁷See Krusell and Smith (1998) for an early, detailed development, and Krueger, Mitman, and Perri (2016) for more details.

in each period they can trade one-period claims to capital up to an exogenously given borrowing constraint \underline{a} , for net return r . Goods and factor markets are assumed competitive.

Given a stochastic process for factor prices, $\{r_t, w_t\}$, the decision problem for an agent can be summarized as follows. In period t , a given agent finds themselves holding claims a , experiencing idiosyncratic efficiency ε , and facing current prices r_t and w_t . Additionally, the agent has at their disposal a host of additional data and information useful for forming forecasts and making decisions.⁸ We will use the subscript t to denote dependence on this time t information set. The agent proceeds to make period t decisions by choosing values $c_t(a, \varepsilon)$, $l_t(a, \varepsilon)$ and $a_t(a, \varepsilon)$ to satisfy

$$u_c(c_t(a, \varepsilon), l_t(a, \varepsilon)) \geq \beta E_t \left[\int \lambda_{t+1}(a_t(a, \varepsilon), \varepsilon') \Pi(d\varepsilon' | \varepsilon) \right] \quad (1)$$

$$\text{and } a_t(a, \varepsilon) \geq \underline{a}, \text{ with c.s.}$$

$$u_l(c_t(a, \varepsilon), l_t(a, \varepsilon)) = u_c(c_t(a, \varepsilon), l_t(a, \varepsilon)) w_t \quad (2)$$

$$a_t(a, \varepsilon) = (1 + r_t) a + w_t \cdot \varepsilon \cdot (1 - l_t(a, \varepsilon)) - c_t(a, \varepsilon) \quad (3)$$

$$\lambda_t(a, \varepsilon) = (1 + r_t) u_c(c_t(a, \varepsilon), l_t(a, \varepsilon)). \quad (4)$$

Here λ_t is the period t shadow price of an additional unit of claims held from period $t - 1$ to period t . The inequality pair (1) is the standard Euler condition and balances the agent's inter-temporal trade-off between consumption and savings. Equation (2) balances their intra-temporal trade-off between labor and leisure.

The right-hand side of equation (1) is the period t forecast of the period $t + 1$ shadow price of savings. This forecast is taken over both idiosyncratic risk faced by the agent, which is captured through the integral over next period's productivity ε' , as well as aggregate risk, which is summarized by the dependence of λ_{t+1} on the period $t + 1$ information set. The expectations operator, E_t , represents expectations over the aggregate shocks conditional on the time t information set. To emphasize per-period decision making, which will be useful when connecting the rational case to our locally rational implementation below, let $\lambda_t^e(a', \varepsilon)$ represent the agent's period t forecast of their period $t + 1$ shadow price given their savings for next period and current productivity. In the rational case under examination here

$$\lambda_t^e(a', \varepsilon) = E_t \left[\int \lambda_{t+1}(a', \varepsilon') \Pi(d\varepsilon' | \varepsilon) \right],$$

and the period t Euler equation can be written more succinctly as

$$u_c(c_t(a, \varepsilon), l_t(a, \varepsilon)) \geq \beta \lambda_t^e(a_t(a, \varepsilon), \varepsilon) \text{ and } a_t(a, \varepsilon) \geq \underline{a}, \text{ with c.s.} \quad (5)$$

In the rational model, one of the elements of the time t information set is the current joint distribution, μ_t , over agent states (a, ε) . When making forecasts, the agents must both know this

⁸In the rational model, for example, the agent must know the distribution of shocks and claims across agents and understand its evolution over time.

distribution as well as its transition dynamics. In equilibrium, these transition dynamics must be consistent with the decision rules of agents, $a_t(a, \varepsilon)$, as well as the transition density Π of the idiosyncratic shocks. By construction, the dynamics of μ_t and the optimal decisions rules of the agents, $a_t(a, \varepsilon)$, must be jointly determined in equilibrium.

2.2 The firm problem

The representative firm rents capital k_t at real rental rate $r_t + \delta$, hires effective labor n_t at real wage w_t , and produces output under perfect competition using CRTS technology $\theta f(k, n)$, where δ is the capital depreciation rate. We take $\{\theta_t\}$ to be a stationary process that affects total factor productivity, with dynamics given by $\theta_{t+1} = v_t \theta_t^\rho$, $|\rho| < 1$, and $\{v_t\}$ is iid having log-normal distribution. There are no capital installation costs. Profit maximizing behavior by the firm implies factors earn their marginal products:

$$w_t = \theta_t f_n(k_t, n_t) \quad \text{and} \quad r_t + \delta = \theta_t f_k(k_t, n_t). \quad (6)$$

2.3 Dynamic recursive equilibrium

We define a *dynamic recursive equilibrium* as a collection of stochastic processes consisting of agent decision rules $\{c_t, l_t, a_t\}$, agent forecasts $\{\lambda_t^e\}$, factor prices $\{r_t, w_t\}$, and the joint distribution of individual states $\{\mu_t\}$, satisfying

- *Agent optimality*: For all t and (a, ε) , the choices $c_t(a, \varepsilon)$, $l_t(a, \varepsilon)$, and $a_t(a, \varepsilon)$ satisfy (2), (3), and (5) given forecasts λ_t^e and current prices r_t, w_t .
- *Agent rationality*: For all t and (a', ε)

$$\lambda_t^e(a', \varepsilon) = E_t \left[\int \lambda_{t+1}(a', \varepsilon') \Pi(d\varepsilon' | \varepsilon) \right] \quad (7)$$

where $\lambda_t(a, \varepsilon)$ is the period t shadow price given by (4).

- *Market clearing*: $k_t = \int a \cdot \mu_t(da, d\varepsilon)$ and $n_t = \int (1 - l_t(a, \varepsilon)) \cdot \mu_t(da, d\varepsilon)$.
- *Firm optimality*: Prices r_t and w_t satisfy (6).
- *State dynamics*: μ_{t+1} evolves consistent with a_t and Π , and $\theta_{t+1} = v_{t+1} \theta_t^\rho$.

Observe that, given any initial aggregate state (μ_0, θ_0) , a dynamic recursive equilibrium, together with a sequence of innovation draws $\{v_t\}$, uniquely determines a time path of agent-state distributions $\{\mu_t\}$ and prices $\{r_t, w_t\}$. Most of the components in the definition of the dynamic recursive equilibrium are standard in the literature. The one exception to this is that we have explicitly decoupled agent optimality and rationality. When we extend our analysis to boundedly rational agents we will only need to change the forecasting rules used by agents.

2.4 The representative agent model: dynamic equilibrium

We will want to compare the dynamics of our model to those obtained under the representative agent (RA) analog, and to facilitate this comparison we highlight the natural sense in which the RA model is a special case of the HA model under examination.

Consider the model developed above, but with the cross-sectional variation in productivity shut down: $\varepsilon_t = 1$ for all agents. Assuming also that agents are initially endowed with the same wealth holdings, per period consumption/savings and labor/leisure decisions will be the same across agents, thus eliminating the need to track agent-state distributions and the dependence of policies on individual states. Equations (1) – (4) and (6) still hold, and by identifying agent-specific variables with corresponding aggregates, the model's dynamics are quite simple to characterize:

- *Agent optimality:* Aggregate wealth a_t , consumption c_t , and leisure l_t satisfy

$$u_c(c_t, l_t) \geq \beta \lambda_t^e, \quad a_t \geq \underline{a}, \quad \text{with c.s.},$$

$$u_l(c_t, l_t) = u_c(c_t, l_t) w_t, \quad \text{and} \quad a_t = (1 + r_t)a + w_t \cdot (1 - l_t) - c_t,$$

given forecasts λ_t^e .

- *Agent rationality:* $\lambda_t^e = E_t((1 + r_{t+1})u_c(c_{t+1}, l_{t+1}))$.
- *Market clearing:* $k_t = a_{t-1}$ and $n_t = 1 - l_t$.
- *Firm optimality:* Prices r_t and w_t satisfy (6).
- *State dynamics:* Capital evolves as $k_{t+1} = \theta_t f(k_t, n_t) + (1 - \delta)k_t - c_t$.

2.5 Stationary recursive equilibrium

The need to track the dynamics of the infinite dimensional aggregate state is a serious impediment, both to the modeler and to the model's agents. The suppression of aggregate risk, together with a focus on a stationary equilibrium, i.e. a steady-state distribution of agent-specific states, greatly simplifies matters. Because this simplification will feature prominently in our implementation of local rationality, we discuss it in detail here.

Setting $\theta = \nu = 1$ and assuming the distribution of agent-states is constant, the time subscript may be dropped: no information other than the agent-state is needed to make decisions. Using over-bars to distinguish this special case, we define a *stationary recursive equilibrium* (SRE) as a tuple $(\bar{c}, \bar{l}, \bar{a}, \bar{\lambda}^e, \bar{r}, \bar{w}, \bar{\mu})$ satisfying

- *Agent optimality:* For all (a, ε) , the choices $\bar{c}(a, \varepsilon)$, $\bar{l}(a, \varepsilon)$, and $\bar{a}(a, \varepsilon)$ satisfy (2), (3) and (5) given $\bar{\lambda}^e$.

- *Agent rationality*: For all (a', ε) ,

$$\bar{\lambda}^e(a', \varepsilon) = \int \bar{\lambda}(a', \varepsilon') \Pi(\varepsilon, d\varepsilon'), \quad (8)$$

where $\bar{\lambda}(a, \varepsilon) = (1 + \bar{r})u_c(\bar{c}(a, \varepsilon), \bar{l}(a, \varepsilon))$.

- *Market clearing*: $\bar{k} = \int \bar{a}(a, \varepsilon) \cdot \bar{\mu}(da, d\varepsilon)$ and $\bar{n} = \int (1 - \bar{l}(a, \varepsilon)) \cdot \bar{\mu}(da, d\varepsilon)$.
- *Firm optimality*: Prices \bar{r} and \bar{w} satisfy (6) given \bar{k} and \bar{n} .
- *State dynamics*: $\bar{\mu}$ is stationary under \bar{a} and Π .

3 Local rationality

The difficulty faced both by the modeler and by the model's agents, when attempting to determine, or even approximate, fully rational decision making, lies in the fact that policy rules and, hence the law of motion, depend on the distribution μ , which is a high dimensional object. Multiple approaches have been used in the literature to approximate the REE of these models. Broadly speaking they can be categorized into two types of approaches. The first type uses projection methods along the lines of Krusell and Smith (1998) to summarize the distribution with a finite set of moments. The exact method can vary, but generally faces the problem that each additional moment adds an additional dimension to the state space. Thus, the curse of dimensionality is quickly faced. The second approach, first introduced by Reiter (2009), instead linearizes policy rules around the REE.

Both of the approaches are appropriately viewed as addressing the *modeler's problem*, the assumption being that the model's agents are fully rational, whereas the modeler must rely on numerical methods to approximate their behavior. The supposition of fully rational agents is a common and natural benchmark; however, it strains the model's realism to imbue its agents with such sophistication. Said differently, the assumption of agent rationality in this model conflicts with the *cognitive consistency principle*, which has been emphasized by Evans and Honkapohja, and asserts that a model's agents should not be much more sophisticated, nor much less sophisticated, than the agents' modeler. In this section we develop a bounded rationality approach that navigates this cognitive conflict while also mitigating technical challenges faced by the modeler. Our implementation of bounded rationality, which borrows from both the RE literature mentioned above, and from the representative agent learning literature, is termed *local rationality*.

3.1 Locally rational agents

We begin with a description of the behavior of individual agents and then discuss equilibrium dynamics in Section 3.2. In an REE, the model's agents know not only the current distribution

of agent-states but also its law of motion and the associated effect on prices; further, they know how to use this knowledge to fully solve their decision problem. The RE model is silent on how agents came to acquire this knowledge and these skills. In contrast, we adopt the *agent-level learning* view, advanced by Evans and McGough (2020b), that agents may not have access to the full aggregate state, that they forecast aggregates using linear models which are updated over time as new data become available, and that they make decisions based on perceived tradeoffs that are informed by these forecasts.

In period t , a given agent is identified by their state (a, ε) , and by their beliefs, ψ , which comprise the coefficients of the forecast model used to form expectations of next period's shadow price. Together with all other agents, they are assumed to observe some common vector of aggregates $X_t \in \mathbb{R}^n$, and they condition their forecasts, $\hat{\lambda}_t^e$, on these aggregates. Given current prices r_t and w_t , they then use this forecast rule to determine their period t decisions $\hat{c}_t(a, \varepsilon, \psi)$, $\hat{l}_t(a, \varepsilon, \psi)$ and $\hat{a}_t(a, \varepsilon, \psi)$, which satisfy the following system of equations:

$$u_c(\hat{c}_t(a, \varepsilon, \psi), \hat{l}_t(a, \varepsilon, \psi)) \geq \beta \hat{\lambda}_t^e(\hat{a}_t(a, \varepsilon, \psi), \varepsilon, \psi) \text{ and } \hat{a}_t(a, \varepsilon, \psi) \geq \underline{a}, \text{ with c.s.} \quad (9)$$

$$u_l(\hat{c}_t(a, \varepsilon, \psi), \hat{l}_t(a, \varepsilon, \psi)) = u_c(\hat{c}_t(a, \varepsilon, \psi), \hat{l}_t(a, \varepsilon, \psi)) w_t \quad (10)$$

$$\hat{a}_t(a, \varepsilon, \psi) = (1 + r_t)a + w_t \cdot \varepsilon \cdot (1 - \hat{l}_t(a, \varepsilon, \psi)) - \hat{c}_t(a, \varepsilon, \psi). \quad (11)$$

Importantly, equations (9) – (11) are taken as *behavioral primitives*: they are imposed assumptions on the behavior households. Equation (9) balances the agent's inter-temporal consumption/savings trade off, and equation (10) balances their intra-temporal labor/leisure trade off. Equation (11) is the agent's budget constraint.

It remains to specify how the expectation $\hat{\lambda}_t^e$ is formed. In a heterogeneous agent economy, agents must learn how to forecast optimally in response to both idiosyncratic *and* aggregate shocks. In this paper we focus on learning how to forecast in the presence of aggregate shocks: our *local rationality* assumption is that, absent aggregate risk, an agent knows how to form forecasts optimally; in the presence of aggregate risk, the agent forms expectations *relative* to the rational forecasts they would have made in a stationary environment. Our reasons for assuming this are two fold. First, idiosyncratic shocks are larger and, thus, agents would learn how to optimally forecast in response to idiosyncratic shocks faster. Second, this implementation provides a clean comparison to the rational model as most solution techniques approximate decisions rules around a stationary recursive equilibrium. By having the benchmark model be that same stationary recursive equilibrium, we can ensure the differences in behavior under local rationality are driven by aggregate shocks.

Operationally, we assume that, to form forecasts, agents use the following forecast model, or *perceived law of motion* (PLM):

$$\log \hat{\lambda}_t = \log \bar{\lambda}_t + \langle \psi, X_{t-1} \rangle, \quad (12)$$

where $\psi \in \mathbb{R}^n$ is a vector of *beliefs* and $X_t \in \mathbb{R}^n$ is a vector of observable aggregates, and $\langle \cdot, \cdot \rangle$ is the standard inner product on \mathbb{R}^n . Using this PLM, the agent forms expectations as

$$\hat{\lambda}_t^e(a', \varepsilon, \psi) = \bar{\lambda}^e(a', \varepsilon) \cdot \exp(\langle \psi, X_t \rangle), \quad (13)$$

where $\bar{\lambda}^e$ is as defined in (8).⁹ In this way, $\langle \psi, X_t \rangle$ represents the agent's forecasted log deviation of their shadow-price from their stationary forecast, $\bar{\lambda}^e$.

Equation (13) is the key behavioral primitive of our model, and reflects the cognitive consistency principle mentioned above. Our agents are assumed able to behave optimally in absence of aggregate risk: in effect, we are assuming that agents have *already learned* how to behave optimally in the absence of aggregate uncertainty.¹⁰ The introduction of aggregate uncertainty greatly increases the complexity of the agent's problem, and here we incorporate bounded rationality: agents are not assumed to know how to forecast the evolution of aggregates optimally, nor how to solve their dynamic decision problem in the face of aggregate risk. Instead they use linear models to form forecasts and they make decisions based on the trade-offs these forecasts impart.

An agent's beliefs evolve as new data are observed, and here we follow the adaptive learning literature's emphasis on recursive least squares algorithms. These algorithms take new estimates (in our case, beliefs) to be a combination of prior estimates and the forecast error adjusted to account for the relative magnitudes and variations of the regressors. The weight placed on the adjusted forecast error is called *the gain* – denoted by g_t – and may be taken as decreasing or constant over time.

To update their beliefs, an agent with state (a, ε) and beliefs ψ regresses log deviations of the realized shadow price

$$\hat{\lambda}_t(a, \varepsilon, \psi) = (1 + r_t)u_c(\hat{c}_t(a, \varepsilon, \psi), \hat{l}_t(a, \varepsilon, \psi)) \quad (14)$$

from its stationary counterpart $\bar{\lambda}(a, \varepsilon)$ on to the previous period's observables X_{t-1} . Letting R_t measure the estimate of the second-moments of X , the recursive formulation of the updating rule for beliefs is given by

$$\hat{\psi}_t(a, \varepsilon, \psi) = \psi + g_t \cdot R_{t+1}^{-1} X_{t-1} (\log(\hat{\lambda}_t(a, \varepsilon, \psi)/\bar{\lambda}(a, \varepsilon)) - \langle \psi, X_{t-1} \rangle), \quad (15)$$

where $R_{t+1} = R_t + g_t \cdot (X_{t-1} \otimes X_{t-1} - R_t)$. Note that the term R_{t+1}^{-1} depends only on aggregates and so may be taken as common across agents.

Some of our results below feature algorithms with decreasing gains and so we retain the more general notation; however, local rationality is most naturally modeled by assuming agents use constant-gain learning (CGL), and this warrants further comment. The decreasing gain $g_t = 1/t$ results in ordinary least squares: see Ch. 2 of Evans and Honkapohja (2001); further, almost sure convergence to the RPE or REE in general requires decreasing gains in which $g_t \rightarrow 0$ at a suitable rate like t^{-1} . In applied work it is common to assumed the gain is a (small) constant: $g_t = g \in (0, 1)$.

⁹An alternative forecasting rule decomposes ψ into two components, one used to forecast the future aggregate state and the other used to specify the relationship between the aggregate state and the shadow price relative to the stationary case. Forecasting the shadow price then requires computing the product of these components. We opt for the simpler method of estimating this product directly.

¹⁰It can be shown, in a stationary environment with only idiosyncratic risk, that if agents are provided forecasting models for $\bar{\lambda}$ that depend on higher-order terms the economy's asymptotic behavior will approximate the REE arbitrarily well.

CGL algorithms discount older data at a geometric rate $1 - g$, and, in stable systems, result in weak convergence to a distribution near the RPE or REE when the gain is small enough.

Several reasons for using constant gain algorithms have been advanced in the literature: see Evans, Evans, and McGough (2021b) for a discussion. Our preference for CGL reflects the concern agents might have about model misspecification. Our agents use simple linear forecast models, but also recognize that these models may not capture the full complexity of the decision-making environments. To account for this, the agents reason that more recent data might be more informative about the current forecasting problem, and thus they discount past data. See Williams (2019) for more on the use of CGL as a robust procedure in the face of model misspecification.

Before turning to equilibrium dynamics it is worth reflecting on the simple nature of our agent's behavior. They enter the period with individual states (a, ε) and individual beliefs ψ . They observe the aggregate X_t and prices (r_t, w_t) , use their beliefs ψ to make forecasts which results in choices \hat{c}_t, \hat{l}_t , and \hat{a}_t for each agent. They go to work, get their wage, go to their broker to trade claims, and stop by the store on the way home to collect their consumables. Finally, they measure their realized shadow price, $\hat{\lambda}_t$, and update their beliefs. Now they're ready to relax, no more decisions or actions being needed until tomorrow.

3.2 Locally rational dynamics

Given SRE behavior $\bar{\lambda}$, agent-specific states and beliefs (a, ε, ψ) , and observable aggregates X_t , the conditions (9) - (11) determine agents' decision schedules in terms of prices (r_t, w_t) . The realized values of prices and other endogenous aggregates are determined by market clearing, i.e. *temporary equilibrium*. Mechanically, this determination requires tracking the evolving distribution of agent-specific states *and* agent-specific beliefs.

Let μ_t be the contemporaneous distribution of agent-states and beliefs. Then temporary equilibrium imposes that $r_t = \theta_t f_k(k_t, n_t) - \delta$ and $w_t = \theta_t f_n(k_t, n_t)$, where k_t and n_t are determined by the market clearing conditions

$$k_t = \int a \cdot \mu_t(da, d\varepsilon, d\psi) \quad \text{and} \quad n_t = \int (1 - \hat{l}_t(a, \varepsilon, \psi)) \mu_t(da, d\varepsilon, d\psi), \quad (16)$$

and θ_t is the realized TFP shock. The n_t in (16) depends on the policy rules $\hat{l}_t(a, \varepsilon, \psi)$, which, in turn, depend implicitly on current factor prices (r_t, w_t) . All must be jointly determined in the temporary equilibrium as solutions to a system of non-linear equations. Note that, just as in the RE model, prices are determined by the distribution μ_t and the productivity shock. The difference is that here the distribution μ_t is over states (a, ε) *and* beliefs ψ , which evolves consistent with \hat{a}_t and $\hat{\psi}_t$.

Denote by Γ the map that takes the full aggregate state $\xi_t = (\mu_t, \theta_t, R_t, X_{t-1})$ to the aggregate observables X_t , i.e. $X_t = \Gamma(\xi_t)$. The dynamics of the economy, which we refer to as *locally rational dynamics* (LR-dynamics), are given in recursive causal ordering as follows:

1. $X_t = \Gamma(\xi_t)$
2. Find (r_t, w_t, k_t, n_t) that solve $r_t = \theta_t f_k(k_t, n_t) - \delta$ and $w_t = \theta_t f_n(k_t, n_t)$ and (16).
3. $R_{t+1} = R_t + g_t \cdot (X_{t-1} \otimes X_{t-1} - R_t)$
4. $\theta_{t+1} = v_{t+1} \theta_t^\rho$
5. μ_{t+1} evolves consistent with \hat{a}_t and $\hat{\psi}_t$

The recursive causality of this dynamic system simplifies the computational burden faced by the modeler: it is no longer necessary to search for (an approximation of) a distributional transition dynamic that is consistent with rational expectations on the part of agents. This simplification comes at a cost: resolution of the temporary equilibrium (item 2) and approximation of the distributional dynamics (item 5) require analysis of an agent-specific state-space that has been expanded to include beliefs. Under rationality, beliefs are homogeneous among agents and consistent with the equilibrium dynamics: lovely in terms of parsimony but difficult to compute. Under local rationality, beliefs vary across agents and are updated recursively: less parsimonious but more computationally tractable.

A final observation: as in the rational case, given initial aggregate state $\xi_0 = (\mu_0, \theta_0, R_0, X_{-1})$, the LR-dynamics, together with a sequence of innovation draws $\{v_t\}$, uniquely determines a time path of aggregate states $\{\xi_t\}$, and thus of agent-state distributions $\{\mu_t\}$, as well as a time path of prices $\{r_t, w_t\}$.

3.3 Restricted perceptions equilibria

In a rational expectations equilibrium, agents make forecasts optimally conditional on the economy's data generating process (DGP). Importantly, the DGP is an endogenous object: it is determined by the actions, and hence the forecasts, of agents. A *restricted perceptions equilibrium* (RPE) is the analogous solution concept under the additional restriction that agents are constrained in their choice of forecast models. Each agent is assumed to choose a forecast model from a pre-determined class. In an RPE, agents make forecasts optimally conditional on the economy's DGP *and* conditional on their restricted class of forecast models. As with an REE, the data generating process arising from a RPE is an endogenous object: the optimality of a given agent's forecast model is conditional on the forecast models used by other agents.

A restricted perceptions equilibrium identifies a natural long-run solution concept in models with learning agents. In this case, the class of forecast models under consideration is typically taken as same for all agents; and the forecast models themselves are assumed linear in parameters, arising from the learning agents' PLMs. If the economy is stable under adaptive learning then its long run behavior, with a decreasing gain learning algorithm, is well-captured by the associated RPE.

Recall that, as a learning agent, his PLM is given by (12), i.e. $\log \hat{\lambda}_t = \log \bar{\lambda}_t + \langle \psi, X_{t-1} \rangle$, where the vector ψ is naturally interpreted as the agent's beliefs. To define a restricted perceptions

equilibrium under local rationality, we require that agents restrict attention to forecast models of the form (12). Conceptually, in an RPE, each agent's beliefs are fixed at a forecast model that is optimal in the sense that, conditional on the behaviors of all other agents (and thus on the economy's aggregate dynamics), it minimizes long-run mean-square forecast errors.

Formally (and computationally) the definition of an RPE in our environment begins with the behavior of a given agent in isolation. Thus imagine the agent holding beliefs ψ fixed for all time, and facing an economic environment summarized by the exogenous (to the agent) processes $\Xi_t = (X_t, r_t, w_t, \varepsilon_t)$. Equations (9)-(11), together with

$$\hat{\lambda}_t / \bar{\lambda}_t = (1 + r_t) u_c(\hat{c}_t, \hat{l}_t) / \bar{\lambda}(\hat{a}_{t-1}, \varepsilon_t),$$

induce a stochastic dynamic system in $\eta_t(\psi) = \left(\log \left(\hat{\lambda}_t / \bar{\lambda}_t \right), X_{t-1} \right)$. Under appropriate conditions this system has a stable ergodic distribution $v_{\Xi}(\psi)$, thus $\eta_t(\psi) \xrightarrow{\mathcal{D}} v_{\Xi}(\psi)$. We define $T_{\Xi}(\psi)$ to be the projection of $\log \left(\hat{\lambda}_t / \bar{\lambda}_t \right)$ onto the span of X_{t-1} using the distribution $v_{\Xi}(\psi)$. Importantly, the distribution $v_{\Xi}(\psi)$, and hence the map T_{Ξ} , depend on the exogenous processes, which is the reason for the subscript Ξ .

The map T_{Ξ} , which takes beliefs to projected values, is defined for a given agent, taking Ξ as exogenous. To extend this map to the general equilibrium environment, define a *beliefs profile* Ψ as an assignment of beliefs to each agent conditional on their initial state. Now note that a beliefs profile, together with the exogenous (to the economy) aggregate and idiosyncratic productivity processes, determines the time-path of the economy, and thus the stochastic process $\Xi_t(\Psi) = (X_t, r_t, w_t, \varepsilon_t)$. Furthermore, because a given agent is small, their beliefs do not impact $\Xi_t(\Psi)$, and so we may define a map $T_{\Xi(\Psi)}$, just as above. This construction provides the following definition:

Definition 1. *A beliefs profile Ψ is a restricted perceptions equilibrium if each agent's beliefs vector ψ , as determined by the profile Ψ , is a fixed point of $T_{\Xi(\Psi)}$.*

This definition of an RPE does not a-priori impose that all agents hold the same beliefs; however, homogeneity is implied in our model, at least for small shocks:

Proposition 1. *If the supports of the aggregate shocks are sufficiently small then agents hold the same beliefs in any restricted perceptions equilibrium that is local to the stationary recursive equilibrium.*

All proofs are in the Appendix. This proposition simplifies considerably our computational work. By appealing to homogeneity of beliefs, when computing an RPE we may replace the beliefs profile Ψ with a beliefs vector ψ , thus greatly reducing the dimension of the problem: see Section 6.1 for implementation details.

It is important to emphasize that the deck here is stacked in favor of homogeneity: after all, we assumed agents use forecast models from the same restricted class. This assumption is not necessary: indeed much of the literature on heterogeneous expectations involves agents-types being

distinguished by the forecast models they use, and the equilibrium concept adopted in these studies is often some version of an RPE.¹¹

In Section 4, which considers the behavior of a locally rational agent in partial equilibrium, we demonstrate existence and uniqueness of an RPE. And while, for our general equilibrium model, no analytic results are reachable with current technology, it is expected (based on a wealth of findings in the learning literature) that for appropriate aggregate observables there is a unique RPE; that for appropriately decreasing gains it is (locally) stable under adaptive learning; and that for small constant gains, agents' beliefs converge weakly to an ergodic distribution very near it. Each of these properties is demonstrated numerically for our model in Section 6.

Finally, it warrants returning to our view that benchmark behavior should be taken as CGL. The notion of optimality used to identify an RPE is predicated on the idea that agents will hold their beliefs constant over time; however, due to model misspecification, it may be advantageous for a given agent to allow their beliefs to vary over time. This line of reasoning is consistent with the views espoused by Williams (2018), and suggests that, in the environment under consideration, constant gain learning may be superior to learning algorithms that induce almost sure convergence to the RPE. In fact, we will see that a constant gain learning algorithm will allow our model to better fit moments of the data.

3.4 Special cases

The behavior of locally rational agents has two interesting limits. The first natural limit is when the size of the aggregate shocks approaches zero. It's clear from the definition that in the absence of aggregate shocks the model's SRE is an RPE with $\psi^* = 0$. With small aggregate shocks, the behavior of a locally rational equilibrium, therefore, inherits properties from the stationary recursive equilibrium such as the wealth distribution and level of precautionary savings. This allows us to isolate how agents learn in the presence of aggregate shocks.

In the other direction, we can take the limit as the size of idiosyncratic shocks ε approaches zero, with the initial distribution μ being a point mass on homogeneous initial conditions for wealth and beliefs. In this limit, the distribution of agents will remain a point mass throughout time, and we recover RA behavior similar to shadow price learning of Evans and McGough (2020b).

Because we will analyze local rationality in the RA model for comparison with the HA case, we elaborate here on some details. In a representative agent environment, locally rational agents, in effect, know the non-stochastic steady-state value of $\log \bar{\lambda}$, and scale it in response to aggregate conditions, just as in the HA case: $\log \lambda_t^e = \langle \psi_t, X_t \rangle \cdot \log \bar{\lambda}$, where ψ_t capture common beliefs. Analogs to (9) – (11) are used to form decision schedules, and competitive factor prices and market clearing result in realized values for the economy's aggregates. Finally, agent's beliefs are then updated using (15), just as in the HA case.

¹¹See Branch and McGough (2018) for more on expectations heterogeneity. In Appendix D, we demonstrate how to incorporate heterogeneous forecasting rules by allowing for types of agents who differ in the aggregates X_t that they condition on.

4 Local rationality in partial equilibrium

To explore the mechanisms underlying local rationality, we consider the behavior of an LR agent in isolation, i.e. absent general equilibrium effects. Additionally, we abstract from labor/leisure considerations and assume the agent supplies a unit of labor inelastically, with the aggregate component of wage set to one and idiosyncratic productivity ε_t taken to be iid and normally distributed about one.

We assume that the agent has *constant absolute risk aversion* (CARA) preferences over per-period consumption: $u(c) = -\gamma^{-1} \exp(-\gamma c)$, and receives net return on their savings given by r_t , which is taken as a stationary AR(1) process in logs:

$$r_{t+1} = (1 + \bar{r})^{1-\rho} (1 + r_t)^\rho \exp(\sigma_\eta \eta_{t+1}) - 1, \quad (17)$$

where η_t is an iid random variable with compact support, mean zero, and standard deviation one. Finally, we assume that agent's face a natural borrowing limit in the form of a transversality condition on wealth. The CARA utility form, together with the natural borrowing limit and the assumption that labor income is normally distributed and iid, allows for a closed form solution to the agent's stationary problem, i.e. when returns are held constant at $r_t = \bar{r}$, and thus lends transparency to locally rational behavior. On the other hand, the agent's savings, in this case, is a unit root process; to induce stationarity, we assume the agent has a probability ϕ of dying at the end of every period, leading to an effective discount rate of $\beta\phi$. If the agent dies we assume they are replaced with a new agent with zero assets, which will ensure a stationary distribution of assets.

As in Section 2.1, rational decisions can be written recursively as $a_t(a, \varepsilon)$, $c_t(a, \varepsilon)$ and $\lambda_t(a, \varepsilon)$. In the stationary case, the agent's decisions rules are independent of t , and we will denote them as $\bar{a}(a, \varepsilon)$, $\bar{c}(a, \varepsilon)$, and $\bar{\lambda}(a, \varepsilon)$, just as in Section 2.5. The results for the stationary case are well-known, see for example Acharya and Dogra (2020), and summarized in the following Proposition 2.

Proposition 2. *The optimal consumption and savings behavior of the agent absent aggregate risk is linear in cash in hand:*

$$\begin{aligned} \bar{c}(a, \varepsilon) &= \bar{\mathcal{C}} + \bar{\mu} ((1 + \bar{r})a + \varepsilon) \\ \bar{a}(a, \varepsilon) &= (1 - \bar{\mu}) ((1 + \bar{r})a + \varepsilon) - \bar{\mathcal{C}} \end{aligned}$$

with $\bar{\mu} = \frac{\bar{r}}{1 + \bar{r}}$ and $\bar{\mathcal{C}} = \frac{-\frac{1}{\gamma} \log(\beta\phi(1 + \bar{r})) - \bar{\mu} - \frac{1}{2} \gamma \bar{\mu}^2 \sigma_\varepsilon^2}{\bar{r}}$.

For completeness we provide a proof of this proposition in the Appendix. Proposition 2 tells us that rational behavior absent aggregate risk follows a simple linear structure. In fact, as $(1 - \bar{\mu})(1 + \bar{r}) = 1$ we can conclude that savings follow a random walk with drift. Moreover, as the shadow value of savings is log linear in consumption,

$$\log \bar{\lambda}(a, \varepsilon) = \log(1 + \bar{r}) - \gamma \bar{c}(a, \varepsilon),$$

we can also conclude that the log shadow price is linear in cash in hand.

For the remainder of this section we will assume that the model is calibrated such that there is no drift: $(1 + \bar{r})\bar{c} = 1$. When displaying plots we will assume the following quarterly calibration. We set the discount factor $\beta = 0.99$, the risk aversion parameter $\gamma = 2$, and the standard deviation of the income shock to be $\sigma_\varepsilon = 0.25$. We assume a 0.5% probability of dying every period so $\phi = 0.995$. Finally, we assume that a one standard deviation shock to returns is 25 basis points with an auto-correlation of $\rho = 0.85$. None of our results are sensitive to these parameters values.

4.1 Local rationality

The behavior of the rational agent is contingent on knowing the stochastic process for r_t and knowing how to respond optimally to its realizations. Under local rationality, the agent is assumed to know how to behave optimally absent aggregate risk, but they must learn how to respond to variation in returns. Operationally this means the agent knows the functions \bar{c} , \bar{a} , and $\bar{\lambda}$, and forecasts the future shadow value of savings using the PLM

$$\log \hat{\lambda}_t = \log \bar{\lambda}_t + \psi \log \left((1 + r_{t-1}) / (1 + \bar{r}) \right), \quad (18)$$

which is the analog to (12), and thus forms expectations via

$$\hat{\lambda}_t^e(a', \varepsilon, \psi) = \bar{\lambda}^e(a', \varepsilon) \cdot \exp(\psi \log \left((1 + r_t) / (1 + \bar{r}) \right)),$$

where $\bar{\lambda}^e(a', \varepsilon) = \int \bar{\lambda}(a', \varepsilon') d\text{Pr}(\varepsilon')$, and $d\text{Pr}$ is the normal density.¹²

Conditional on this forecast, the equations characterizing the behavior of the LR agent are analogous to equations (9), (11), and (14), and given by

$$\begin{aligned} \exp(-\gamma \hat{c}_t(a, \varepsilon, \psi)) &= \beta \phi \hat{\lambda}_t^e(\hat{a}_t(a, \varepsilon, \psi), \varepsilon, \psi) \\ \hat{a}_t(a, \varepsilon, \psi) &= (1 + r_t)a + y - \hat{c}_t(a, \varepsilon, \psi) \\ \hat{\lambda}_t(a, \varepsilon, \psi) &= (1 + r_t) \exp(-\gamma \hat{c}_t(a, \varepsilon, \psi)). \end{aligned}$$

Finally, as before, we assume that agents beliefs are updated using the following learning algorithm:

$$\begin{aligned} R_{t+1} &= R_t + g_t \left(\log \left((1 + r_{t-1}) / (1 + \bar{r}) \right)^2 - R_t \right) \\ \hat{\psi}_t(a, \varepsilon, \psi) &= \psi + g_t R_{t+1}^{-1} \log \left(\frac{1 + r_{t-1}}{1 + \bar{r}} \right) \left(\log \left(\frac{\hat{\lambda}_t(a, \varepsilon, \psi)}{\bar{\lambda}(a, \varepsilon)} \right) - \psi \log \left(\frac{1 + r_{t-1}}{1 + \bar{r}} \right) \right), \end{aligned}$$

where R_{t+1} is the agent's estimate for the covariance of the $\log \left((1 + r_{t-1}) / (1 + \bar{r}) \right)$.

¹²The dependence of $\bar{\lambda}^e$ on ε is not needed here since the process ε_t is iid; however, we retain it for notational symmetry.

4.2 First-Order Analysis

A standard tool for understanding the behavior of heterogenous agent models in the presence of aggregate risk is to linearize around the no aggregate risk policy rules, and this same technique provides useful insights when applied to the dynamics characterizing our locally rational agent in this section. We will proceed in this manner by taking a first order expansion of the policy rules with respect to the level of interest rate risk σ_η . As notation, we'll use $\mathcal{O}(\sigma_\eta^2)$ to denote errors that arise from a first order approximation and therefore scale with σ_η^2 . We will use $\mathcal{O}(\sigma_\eta)$ to denote errors that arise from the zeroth order approximation and therefore scale with σ_η . Our first result concerns the behavior of the rational agent.

Proposition 3 (Rational behavior). *There exists constants $\psi_0 > 0, \psi_a < 0$, and ψ_ε such that*

$$c_t(a, \varepsilon) = \hat{c}_t(a, \varepsilon, \psi^{RE}(a, \varepsilon)) + \mathcal{O}(\sigma_\eta^2)$$

where

$$\psi^{RE}(a, \varepsilon) = \psi_0 + \psi_a a + \psi_\varepsilon (\varepsilon - 1) \quad (19)$$

Proposition 3 tells us that, for any level of the idiosyncratic state, there exists beliefs such that the locally rational agent would behave in the same manner as the rational agent. To understand why beliefs take the form in (19) consider the response of a rational agent to an increase in the returns to wealth. The consumption response of the rational agent depends on the present value budget constraint where the expected present value of consumption must equal the human wealth of the agent: their financial wealth plus the expected present value of labor income. When the agent has no financial wealth an increase in returns decreases both the price of future consumption and the present value of wage income leading to a decrease in consumption. As returns are persistent, higher returns in the current period predict both higher future returns and lower future consumption. This, in turn implies a higher expected marginal value of wealth, λ_t^e , and thus is rationalized by $\psi_0 > 0$.

This logic is altered when the agent has some financial wealth in addition to their wage income. An increase in returns to wealth increases the value of that wealth.¹³ Agents with higher financial wealth have a correspondingly smaller decrease in consumption when returns increase. This is captured by $\psi_a < 0$, which implies that agents with higher financial wealth will have a smaller decrease in their expected future shadow price of wealth than agents with less financial wealth.

Next we turn to the behavior of the locally rational agent. Our first result concerns the existence of an RPE as defined in Section 3.3.

Theorem 1. *For σ_η small enough there exists a unique RPE associated with the PLM in (18) given by*

$$\psi^{RPE} = \frac{\rho}{1 - \frac{\rho}{1+\bar{r}}} + \mathcal{O}(\sigma_\eta).$$

¹³This occurs directly via an increase in the present returns to wealth, and indirectly through a decrease in the price of consumption




Figure 1: Best linear predictor of beliefs conditional on financial wealth

This theorem guarantees the existence and uniqueness of this RPE for small enough levels of risk, and provides an approximation for that RPE with an error that scales linearly with the size of the shock.¹⁴ Assuming the RPE is locally stable, we would expect the long run beliefs of agents with a decreasing gain learning algorithm to converge this RPE. We verify this fact numerically in the Appendix B.4. The nature of the PLM in (18) implies that unlike the rational beliefs in Proposition 3, the RPE beliefs are independent of the agent’s idiosyncratic states. This is due to agents drawing upon their entire experience when forming expectations. This includes periods when they are financially well off as well as periods when they are in dire financial straights.

To illustrate this last point, in Figure 1 we plot both the RPE beliefs (blue) and rational beliefs (black) as a function of financial wealth. In addition to those lines, we also plot the best linear predictor of beliefs ψ as a function of wealth using the ergodic joint distribution of beliefs and wealth constructed by simulating the behavior of the locally rational agent for various values of the gain parameter. We note that the rational beliefs cross the RPE beliefs almost exactly at the long run expected assets of the agent. This reflects that the agent learns on average to behave optimally over time. However, the RPE beliefs are far away from the rational beliefs when the agent either has very positive or very negative levels of financial wealth. Recall that agents with

¹⁴Solving for the RPE requires regressing the individual’s shadow price of savings on the interest rate. We use a first order approximation to approximate the ergodic variance of both with an error which scales with $\mathcal{O}(\sigma_\eta^3)$. Computing the regression coefficient requires dividing by the variance of r_t which leaves an error which scales as $\mathcal{O}(\sigma_\eta)$.




Figure 2: Impulse response of consumption to a one standard deviation increase in returns conditional on low wealth (left) and high wealth (right).

different levels of financial wealth experience the return shock differently. The RPE, being the long run limit of the decreasing gain learning algorithm, averages over those experiences so when agents are forecasting their future shadow price of savings they are remembering what it is like to be poor even when they are currently rich.

A similar effect is true when the agent employs a constant gain learning algorithm. Studying Figure 1 we see that for low values of the gain parameters the expectation of beliefs conditional on financial wealth aligns almost exactly with the RPE value. This is because when the gain is small enough, the agent is essentially averaging over their entire past experience. However, as the gain parameter increases the conditional expectation of beliefs tilts towards the rational value. With higher values of the gain parameter, the agent places more weight on recent experience when learning to how to respond optimally to variation in returns. This results in the agent heavily discounting past experiences when they had very different financial positions. Essentially, they learn how to behave optimally in response to their current circumstances.

These differences in beliefs can result in different responses to changes in returns. In Figure 2 we plot impulse responses to a one-time one standard deviation increase in returns. The right-hand side of the figure plots the percentage deviation of the consumption of high financial wealth individuals as a result of this shock while the left-hand side plots the response of low financial wealth individuals. In both cases consumption falls on impact with the shock, but it falls much further for low wealth individuals since they are far more reliant their labor income and must pay additional interest on their debt.¹⁵ Starting with the high wealth agents, we see that the fall in consumption is less for rational agents than for their boundedly rational counterparts. Those with the RPE beliefs have the greatest fall in consumption as they remember and put equal weight on their experiences when they were poor. As we increase the gain parameter the impulse responses approach rational behavior as the agent places more weight on recent experiences. The reverse

¹⁵Individual wealth in this model follows a near random walk. As a result, the time paths of the consumption IRFs do not return to zero.

is true for the low wealth individuals. In that case, the initial fall in consumption is larger for the rational agent since those with RPE beliefs place equal weight on their experiences when they were wealthy. This hysteresis effect will figure prominently in the behavior of the fully calibrated GE model below.

5 Calibration and numerical methods

In this section we outline the baseline calibration as well as the numerical methods used to simulate both the locally rational dynamics and the rational expectations equilibrium.

5.1 Functional forms and calibrations

We use the following standard calibration for the heterogeneous agent economy which follows closely the calibration in Boppart, Krusell, and Mitman (2018). Agents are assumed to have a utility function over consumption and leisure given by

$$u(c, l) = \frac{1}{1 - \sigma} (c^{1 - \sigma} - 1) - \eta \frac{(1 - l)^{1 + \varphi}}{1 + \varphi}.$$

The production function is assumed to be Cobb-Douglas: $f(k, n) = k^\alpha n^{1 - \alpha}$.

We begin by specifying the parameters common to both the rational expectations and boundedly rational model. We assume that the length of a period is one quarter and, therefore, assume a long run capital to output ratio of 10.26 (see Den Haan, Judd, and Juillard (2010)). The parameter α is chosen to be 0.36 to match the capital share of income. The depreciation rate, δ , is set to match an annualized steady state real interest rate of 4% per year. Given the long run capital to output ratio and production function this implies a value of $\delta = 0.025$. We assume logarithmic utility from consumption ($\sigma = 1$) as a benchmark value in the literature. We choose $\varphi = 1$ to target a Frisch elasticity of 1. For the TFP process we use a standard parameterization, setting the serial correlation coefficient to 0.95 and letting the standard deviation of the innovation be 0.007. To capture idiosyncratic efficiency, we follow Krueger, Mitman, and Perri (2016) who estimate a process for log earnings after taxes and transfers using the PSID. They estimated a quarterly persistence for innovations, ρ , to be 0.9923 with a standard deviation, σ_ε , of 0.0983. We use a finite state approximation to this AR(1) process using Rouwenhorst's method (see Kopecky and Suen (2010)) with 11 grid points. We assume that households cannot borrow, $\underline{a} = 0$.

The final two parameters β and η are internally calibrated and chosen to match moments for the stationary distribution. We set $\beta = 0.985$ to ensure that the steady state capital to output ratio matches the aforementioned target of 10.26. The parameter η is set to 7.8 to target an average supply of hours by households to 1/3.

For learning models, it remains to specify the aggregate observables and to calibrate the gain. Concerning the former, we follow the inspiration of Krusell and Smith (1998) and take

$X = (1, \log(k/\bar{k}), \log(\theta))$. Assuming agents observe \bar{k} is innocuous: after all, they are regressing on a constant. The assumption that agents observe aggregate capital and aggregate productivity, while less natural, is also harmless – realized prices r and w contain the same information – and the computational simplicity the assumption affords makes it standard in the literature: see, for example, Krusell and Smith (1998), Eusepi and Preston (2011) and Branch and McGough (2011).

We set our benchmark gain at $g = 0.035$ to match our preferred moment, the ratio of consumption to output volatility. Noting that the gain discounts past data at rate $1 - g$, this value implies a half-life of approximately 5 years, based on quarterly measures, i.e. $0.965^{20} \approx 0.5$. Our value of g is consistent with those used in the literature for calibration exercises and applied analysis.¹⁶

5.2 Wealth Inequality in the Model Economy

As argued by Krueger, Mitman, and Perri (2016) it is crucial, when using an HA model to study aggregate fluctuations, to have a model-implied cross-sectional wealth distribution that is consistent with the empirically observed concentrations, and it is especially relevant to match the the share of the bottom 40% being close to zero. In table 1 we document how well the stationary distribution of wealth in our model matches moments of the wealth distribution observed in the data. We focus on the wealth held by the quintiles of the wealth distribution as well as those at the very top.

% Share Held By	Data		Model
	PSID, 06	SCF, 07	
Q1	-0.90	-0.20	0.00
Q2	0.80	1.20	0.01
Q3	4.40	4.60	2.66
Q4	13.00	11.90	15.50
Q5	82.70	82.50	81.83
90-95	13.70	11.10	19.40
95-99	22.80	25.30	27.23
T1%	30.90	33.50	13.86

Table 1: Wealth distribution in the data and the model. The data columns are from table 6 of Krueger, Mitman, and Perri (2016).

From the table we observe that the benchmark model fits the empirical wealth distribution well and, specifically, is able to match the fact that the bottom 40% have essentially no wealth. This is despite these being un-targeted moments. This is unsurprising as our model uses the same idiosyncratic productivity process as Krueger, Mitman, and Perri (2016), and there they document how that highly persistent process is key for generating enough dispersion in wealth. The only

¹⁶For example, using quarterly data on US aggregates, Milani (2007) estimates a gain of 0.018; in their influential paper on monetary policy, Orphanides and Williams (2003) set $g = 0.05$; Branch and Evans (2006b) find that for quarterly GDP and inflation data, a range of 0.02 – 0.05 works well for both forecasting and for matching the Survey of Professional Forecasters; and Eusepi and Preston (2011) use an optimizing procedure to select a gain of 0.0029.

problematic moment for the model is the very top of the wealth distribution. In the data the top 1% of the wealth holders account for over 30% of overall net worth whereas the corresponding moment in the model is only 14%.

5.3 Numerical methods

For both the locally rational model as well as the rational expectations model, the first step is to approximate the stationary equilibrium. We carry this out by first solving the consumer’s problem, given fixed prices, using the endogenous grid method of Carroll (2006). The decision rules for each productivity level are approximated using cubic interpolation with 150 non-linearly spaced grid points. With the household decisions in hand, the stationary distribution of assets and productivities are approximated using a histogram over income and assets defined on a finer grid with 5000 points per productivity level. From the household policy rules, we construct a transition matrix between individual states and compute the associated invariant distribution.

To approximate the rational expectations equilibrium, we compute an impulse response to a one-time unexpected shock to productivity assuming perfect foresight. Boppart, Krusell, and Mitman (2018) demonstrated that, for a small enough shock, dividing the impulse response by the size of the shock constructs a numeric derivative which is isomorphic to linearizing the model’s dynamics with respect the productivity shock. We compute this impulse response by assuming that the economy is initially at the long run steady state. We then assume that log TFP receives a one time increase in productivity that mean reverts back to steady state level at rate ρ . By assuming that after $T = 350$ periods the economy has returned to the steady state, we can solve for the path of the capital to labor ratio¹⁷ that represents the perfect foresight equilibrium.

Once the impulse response has been recovered it is possible to simulate the time series of aggregates as follows. For a given aggregate variable, z , let $\{z_{\theta,t}\}$ be the impulse response of that variable to a one time, unanticipated productivity shock normalized such that $z_{\theta,t}\sigma_V$ is the response to a one standard deviation shock. The time series of z_t generated by a sequence of shocks v_t is then constructed by aggregating the effects of all past shocks:

$$z_t = \sum_{k=0}^T z_{\theta,k} v_{t-k}.$$

To simulate an economy with locally rational agents we need, at any given period, the joint distribution of assets, productivities, and beliefs. We approximate this distribution, μ_t , each period using 100,000 agents. Every period, given the current productivity level, θ_t , and distribution of agent characteristics, μ_t , we solve for the temporary equilibrium and update the aggregate state based on the decision rules \hat{a}_t and $\hat{\psi}_t$. See Appendix C for full details.

¹⁷For a given path of TFP, the capital/labor ratio pins down the path of prices that are inputs for the agents problem.

6 Results

In this section, we study the behavior of the calibrated economy populated by agents who are locally rational. We use numerical methods to show the existence of an RPE and demonstrate that the RPE is stable under learning with decreasing gain. Turning to simulations, we study the business cycle properties of the locally rational model and contrast them with the properties of both the rational expectations equilibrium and its representative agent counterpart.

6.1 Existence and Stability of Restricted Perceptions Equilibrium

To verify the existence of a restricted perceptions equilibrium, we appeal to Proposition 1 and assume agents hold common beliefs; we proceed to find the fixed point of the finite sample analogue of the T -map, denoted \hat{T} . To compute this map, we begin with a distribution of N agents drawn from the distribution of assets and productivities present in the stationary recursive equilibrium. We endow all agents with the same initial beliefs ψ and simulate the resulting locally rational dynamics for $S + 1$ periods assuming $g = 0$, which implies that beliefs are fixed at these initial beliefs. Let $\hat{\lambda}_{i,t}(\psi)$ and $\bar{\lambda}_{i,t}(\psi) = (1 + \bar{r})\bar{\lambda}(a_{i,t-1}(\psi), \varepsilon_{i,t})$ be the resulting path of the shadow price of wealth for agent i as well as the no aggregate risk counterpart. Similarly, let $X_t(\psi)$ be the path of observables. The map $\psi \rightarrow \hat{T}(\psi)$ is defined implicitly via

$$\frac{1}{SN} \sum_{t=2}^{S+1} \sum_{i=1}^N \left(\log \left(\frac{\hat{\lambda}_{i,t}(\psi)}{\bar{\lambda}_{i,t}(\psi)} \right) - \langle \hat{T}(\psi), X_{t-1}(\psi) \rangle \right) \cdot X_{t-1}(\psi) = 0.$$

This map converges point-wise to T as N and S approach infinity.¹⁸ We numerically verify the existence of an RPE by finding the fixed point of $\hat{T}(\psi)$ when $N = 100,000$ and $S = 1,000$.¹⁹

To verify the stability of the RPE under learning we simulate the dynamics of the locally rational economy from two different initial conditions with a decreasing gain learning algorithm.²⁰ In both experiments, we initialize agents from the distribution of wealth and productivities in the stationary recursive equilibrium and endow all agents with homogeneous beliefs. In the first experiment, all agents begin with the RPE beliefs, while in the second all agents start at beliefs consistent with the SRE, i.e. $\psi = 0$. We plot the path of the average beliefs, across the distribution of agents, over the simulation in Figure 3. The black line represents the path of average beliefs initialized at the RPE and the blue line represents the path initialized at $\psi = 0$. Theory predicts that as long as the RPE is locally stable there will be a basin of attraction around the RPE such that the beliefs of agents will converge to RPE beliefs in probability when using a decreasing gain learning algorithm. This is born out in our numerical simulations numerical simulations as both lines converge to the RPE levels over the course of 80,000 periods. The blue line illustrates that the basis of attraction of the RPE is large, with the average beliefs converging to the RPE values by

¹⁸More precisely, for each ψ , $\hat{T}(\psi)$ converges almost surely to $T(\psi)$ as N and S approach infinity.

¹⁹Increasing both N and S does not appreciably change the value of the fixed point.

²⁰We used $g_t = 0.5t^{-0.8}$ as the specification for the gain function




Figure 3: Time path of average beliefs when agents are initialized with the RPE beliefs (black) and $\psi = \mathbf{0}$ (blue). The RPE beliefs are represented with a solid red line.

the end of the simulation even if beliefs are initialized far away from the RPE. We provide further evidence for this in Appendix D, as well as stability analysis when agents have heterogeneous forecasting rules. We should emphasize that the slow rate convergence of beliefs to the RPE is indicative of the decreasing gain learning algorithm. Constant gain learning algorithms impart faster convergence but, as we will emphasize in the next section, result in different long run values.

6.2 Statistical Properties

Next, we evaluate the business cycle properties of the locally rational model and compare with the corresponding behavior of the heterogeneous agent rational expectations equilibrium as well as the representative agent economy under both rational expectations and shadow price learning. In all cases, we simulate the economy for 50,000 periods to construct an ergodic distribution of the relevant state variables. Drawing from the ergodic distribution, each model is simulated for 240 periods and moments are constructed after HP-filtering the log of all relevant variables.²¹ The same procedure is applied to the U.S. data which runs 240 quarters from 1948Q1 to 2007Q4.

	Representative Agent					Heterogeneous Agent			
	Data	RE	$g = 0.001$	$g = 0.01$	$g = 0.035$	RE	$g = 0.001$	$g = 0.01$	$g = 0.035$
$\frac{\text{std}(C)}{\text{std}(Y)}$	0.50	0.32	0.32	0.33	0.34	0.36	0.70	0.63	0.50
$\frac{\text{std}(I)}{\text{std}(Y)}$	2.73	3.10	3.09	3.08	3.07	2.91	1.88	2.10	2.50

Table 2: Business Cycle Statistics

Table 2 reports standard deviations for consumption and investment relative to the standard deviation of output for all models and the data. As has been well documented in the literature (see Romer (2012)), the benchmark real business cycle model both overstates the variation of investment and, correspondingly, understates the variation of consumption relative to the data. Neither the introduction of bounded rationality through shadow price learning nor the introduction of heterogeneous agents with rational expectations is able to significantly change any of these moments. However, the interaction of bounded rationality and agent heterogeneity leads to substantially different second moments, bringing them closer to the data by increasing the standard deviation of consumption while decreasing the standard deviation of investment.

Focusing on the last 4 columns of table 2, we observe that increasing the gain appears to bring heterogeneous agent model closer in line with the rational expectations equilibrium. This observation is born out when inspecting the impulse responses to one standard deviation productivity shock plotted in Figure 4. The black line in figure 4 plots the impulse response of rational expectations equilibrium constructed from a one-time unanticipated increase in productivity under perfect foresight. The colored lines are the responses of the locally rational economy. We construct these impulse responses by repeatedly drawing an initial distribution of assets, productivities and beliefs

²¹We construct 5000 simulations for each model and average over all simulations.

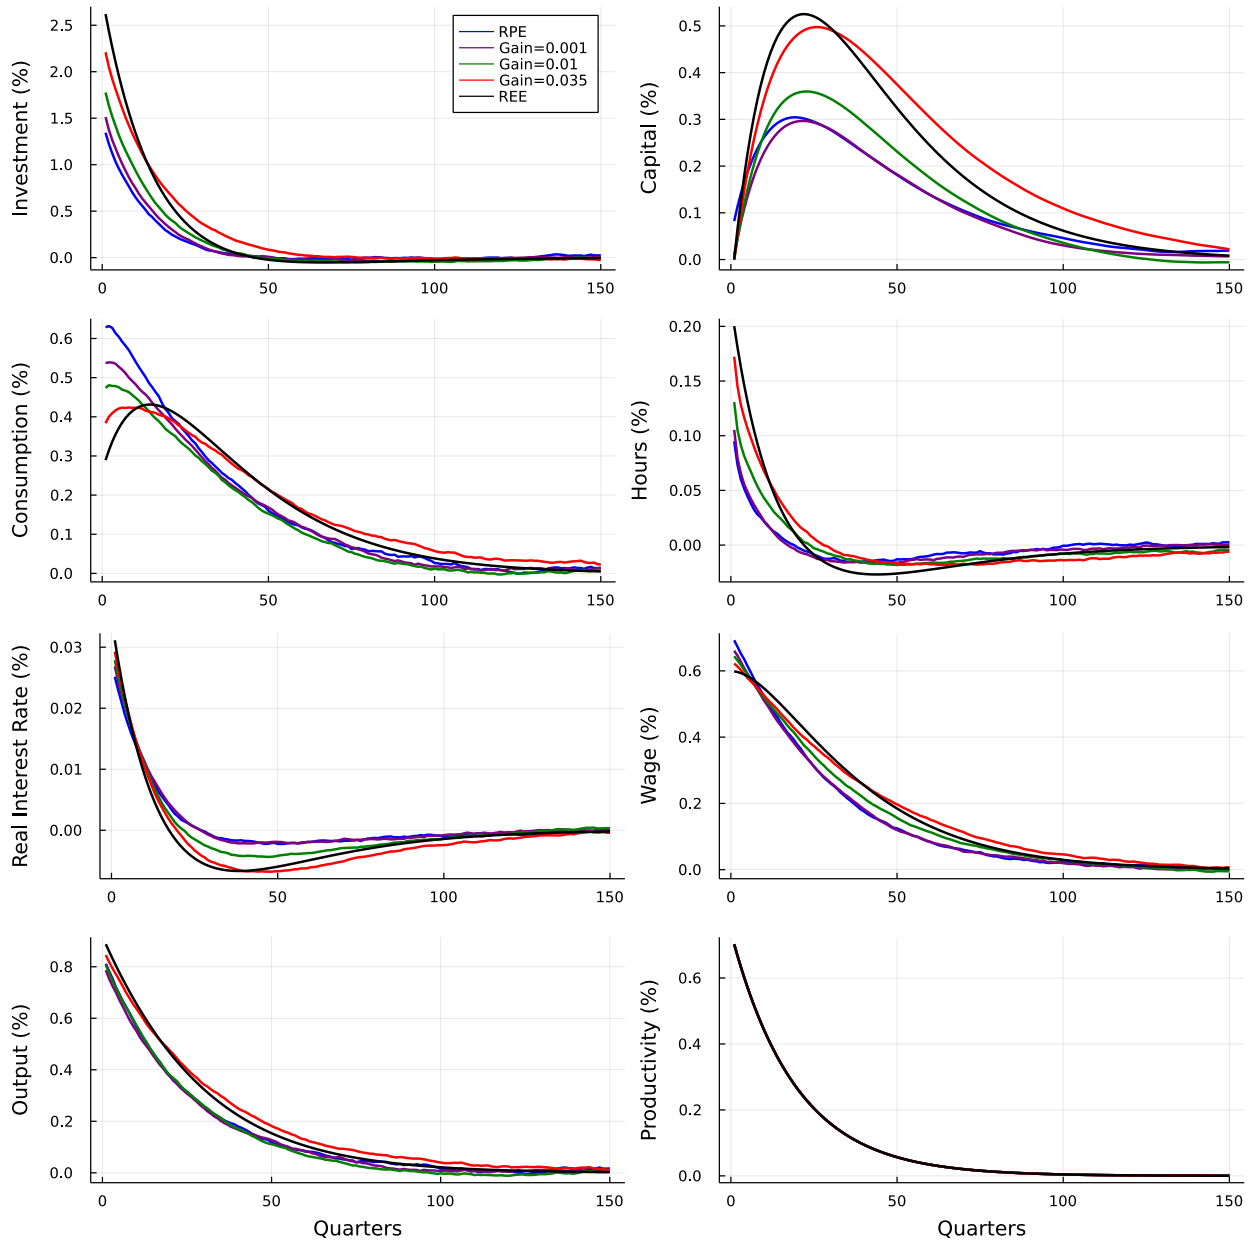


Figure 4: Impulse responses to a one-standard deviation increase in productivity. The black line refers to the linearized rational expectations equilibrium. The blue line refers to the RPE. The purple, green and red lines refer to the mean response of the locally rational dynamics with gains equal to 0.001, 0.01, and 0.035 respectively.

from the ergodic distribution generated by long simulation. We then record the impulse responses to a one-standard-deviation productivity shock from those initial starting points and plot the mean response of all variables as a percentage deviation from the path which would prevail in absence of the shock.

In all cases we observe the familiar humped shape responses of capital and consumption but, under local rationality, the response of capital is muted while the response of consumption is amplified. The locally rational agents appear to be smoothing consumption less than their rational counterparts. This is especially apparent in the response of the RPE, as seen in the blue line, which has very muted responses of investment and capital accumulation but amplified responses of consumption. Unsurprisingly, the smallest gain economy ($g = 0.001$), represented by the purple line, tracks the response of the RPE very closely. As the gain increases, the response of the locally rational economy converges towards rational expectations with the closest being the red line ($g = 0.035$). To gain a better understanding of this behavior it is necessary to explore the endogenous distribution of beliefs that arises in these economies.

We begin by constructing beliefs that rationalize the rational expectations equilibrium. Following the procedure of Appendix D.3, we construct beliefs $\psi^{RE}(a, \varepsilon)$ such that if agents use these beliefs to forecast their future shadow price of savings then, to first order, the economy behaves identically to the RE economy. Note that the rational expectations beliefs vary based on individual states, which codifies that different agents have different experiences in recessions (or booms) depending on their current situation. These different experiences are what give rise to the endogenous distribution of beliefs present in the locally rational model.

To study this distribution of beliefs, we follow the analysis of Section 4 and construct the best linear predictor of beliefs conditional on wealth, $E^*[\psi|a]$.²² As our interest is the response to a TFP shock we focus on $E^*[\psi_\theta|a]$, the best linear predictor of the TFP coefficient conditional on individual wealth, which we plot in Figure 5. The lines in the figure represent the long run ergodic distributions associated with the REE (black), three different constant gain learning algorithms (purple, green and red for gains 0.001, 0.01 and 0.035 respectively), and the RPE associated with the decreasing gain learning algorithm (blue). Finally, as a point of reference, we plot the average beliefs of the REE with the orange line. In all cases, observe that the coefficient on TFP is, on average, negative. This reflects that an increase in TFP raises average consumption thus reducing the marginal value of savings.²³ Starting with the black line representing beliefs in the REE, observe that the belief coefficient on TFP is less negative for agents with more wealth. This positive relationship captures the ability of wealthier agents to use their wealth to buffer themselves against business cycle fluctuations. Poorer agents are more exposed to aggregate shocks: their consumption falls, proportionally, more when TFP is low and increases more when TFP is high. We contrast that with the blue line representing the beliefs in the RPE, which differs from the REE line both in slope and level. To understand the differences in slope, recall that the RPE is the long run limit of the decreasing gain learning algorithm, where agents are using the entirety of their lifetime experience to construct forecasts. Wealthy agents remember (and place near equal weight)

²²See Appendix D.4 for details

²³Higher TFP also raises interest rates but in this calibration the other effect dominates.

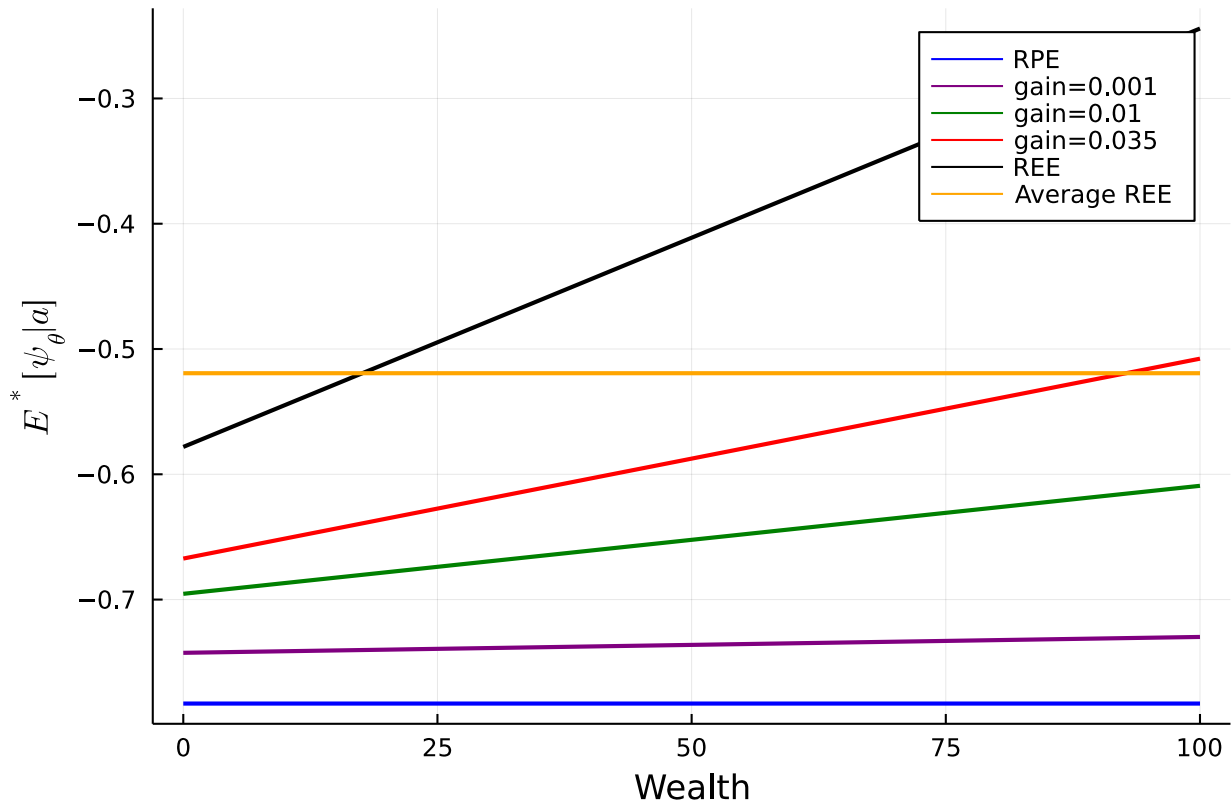


Figure 5: Best linear predictor of belief coefficient on TFP conditional on individual wealth for the ergodic joint distribution of wealth, productivities and beliefs.

on their experiences when they were poor. Similarly, poorer agents remember their experiences when they were rich. A result of this is that, on average, all agents have the same beliefs, which is represented by a horizontal line, similar to what we observed in Section 4.

Unlike the partial equilibrium case studied in Section 4, beliefs also differ in their level: the uniform beliefs associated with the RPE do not correspond the average beliefs of the REE. This difference in average levels is a result of general equilibrium effects and helps partially explain the differences in dynamics between the RPE and REE observed in Figure 4. To isolate the relative contribution of the slope and level effects, consider a *hypothetical economy* populated by agents with homogeneous beliefs at the average beliefs of the REE.²⁴ This hypothetical economy removes the level effect and focuses entirely on the dynamics driven homogeneous beliefs. The orange line in Figure 6 plots the impulse responses of consumption, capital, hours and the real interest rate in this hypothetical economy compared to those of REE (black) and the RPE (blue). The hypothetical economy lies roughly halfway between those two lines.




Figure 6: Impulse responses to a one-standard deviation increase in productivity. Black line refers to the linearized rational expectations equilibrium. The blue line refers to the RPE. The orange line refers to IRF of the hypothetical economy where all agents have homogeneous beliefs at the average value of the REE beliefs.

To understand this behavior, consider the beliefs of a wealthy agent in the hypothetical economy represented by the orange line in Figure 5. Wealthy agents in the hypothetical economy have a more negative coefficient on TFP than their counterparts in the REE, which implies that they believe themselves to be relatively more exposed to business cycles. As a result, wealthier agents

²⁴This is also the long run limit of beliefs in partial equilibrium: of a measure zero locally rational agent using a decreasing gain learning algorithm embedded in the REE model.

over-consume in booms and under-consume in recessions, which explains the amplified response of consumption observed in Figure 6.²⁵ Along with their higher consumption, the wealthier, more productive agents also supply less labor in booms relative to rational expectations, which causes a smaller increase in hours and interest rates relative to REE.

If the agents of this hypothetical economy could update their beliefs to improve their forecasts, they would want to have a more negative coefficient of TFP since the smaller increase in interest rates, in response to a positive TFP shock, would result even lower shadow prices of savings. This effect is amplified in equilibrium, as a more negative coefficient on TFP results in even lower forecasts of the shadow price of savings when TFP is high, which, in turn, results in even higher consumption and even lower labor and interest rates. In equilibrium, this process lowers the coefficient on TFP to the RPE value observed in Figure 5. The interaction of the level and slope effects is needed to fully explain the increased response of consumption observed in the RPE. Examining Figure 6, we conclude that roughly half the difference between the REE and RPE is explained by homogeneous beliefs and the remaining differences are due to general equilibrium.

Similar to what we observed in partial equilibrium, when agents update their beliefs using a constant gain learning algorithm they start discounting their past experiences and, as a result, their beliefs depend more on their recent experiences. Increasing the constant gain places more weight on the recent experiences, and hence their current state, which brings the resulting distribution of beliefs more in line with the rational expectations beliefs both in slope and level. We observe this clearly in Figure 5 as the slope and level of the purple, green, and red lines are increasing in their respective gains. This shift in beliefs is reflected in the Figure 4 impulse responses with higher gains being closer to the rational expectations paths. While the corresponding moments in Table 2 are also closer to rational expectations, they are not identical.²⁶ The model with a gain of $g = 0.035$ has the best fit as wealthier agents better respond to their current circumstances, but also remember what it was like to be poor. This hysteresis effect is not merely a theoretical construct, it parallels many results documented in the empirical literature. For example, see the seminal paper by Malmendier and Nagel (2011).

7 Conclusion

By providing a modeling environment that engenders tractable distributional dynamics, the heterogeneous agent literature has greatly expanded the reach of DSGE models; however, to an extent even greater than their RA counterparts, these modeling environments place unrealistically extreme demands on the cognitive capacity of agents. Local rationality provides a behavioral paradigm that

²⁵Note that poorer agents will have the reverse effect: under-consuming in booms and over-consuming in recessions. Aggregate consumption dynamics are determined by the behavior of the richer agents.

²⁶A curious reader might be interested in the dynamics if the forecast rule were expanded to include interaction with idiosyncratic states. We explore this in Section D.2.2 of the appendix. The extended learning rule brings the learning dynamics more in line with rational expectations but features a far more complicated forecasting model and instability not present in the baseline model.

mitigates this criticism: locally rational agents are very good at understanding themselves and their behaviors, but are less certain about how their behaviors interact with the behaviors of others and the attendant aggregate consequences; thus, instead of taking the RE view that agents understand the endogenously determined evolution of the economy's wealth distribution, locally rational agents simply estimate the evolution of certain aggregates over time, as well as the relationship between these aggregates and their own behavior.

Local rationality adheres to the cognitive consistency principle, which improves a model's realism. Interestingly, in the heterogeneous-agent environment, this improved realism benefits the modeler: because this principle puts the modeler and agents on equal footing, it is not necessary to solve for a time-invariant transition dynamic over an infinite dimensional state space – the modeler can work recursively exactly as the agents do.

An economy populated with locally rational agents has associated with it a restricted perceptions equilibrium that is homogeneous in beliefs, and that serves as a disciplined benchmark; however, it is natural to assume locally rational agents use constant gain learning algorithms when updating their beliefs, as this allows them to adjust their responses to aggregate conditions as local conditions vary. Under this assumption, agents' beliefs converge over time to an ergodic distribution that is centered near, but not directly on the economy's RPE.

Under low gain the distribution of beliefs is tightly centered near the RPE: agents respond only slowly to, e.g., changes in their wealth; for larger gains the distribution is more widely spread as agents' response times quicken and their attendant behaviors more closely approximate those of rational agents. This feature provides a nice avenue through which the gain can be used as a tuning device to match models to data. Using the Krusell-Smith environment, and in contrast to RE, we found that for reasonable gain levels the model under local rationality could reproduce the volatility of consumption relative to output found in US data. This is explained by the slow adjustment of agents' beliefs: under local rationality the beliefs of newly rich agents are clouded by the recent experiences with poverty which, in effect, amplifies their optimism and thus raises their consumption response to positive TFP shocks relative to their rational counterparts.

References

- ACHARYA, S., AND K. DOGRA (2020): “Understanding HANK: Insights from a PRANK,” *Econometrica*, 88(3), 1113–1158.
- ADAM, K. (2007): “Experimental Evidence on the Persistence of Output and Inflation,” *Economic Journal*, 117, 603–636.
- AIYAGARI, S. R. (1994): “Uninsured idiosyncratic risk and aggregate saving,” *The Quarterly Journal of Economics*, 109(3), 659–684.
- ANUFRIEV, M., T. ASSENZA, C. HOMMES, AND D. MASSARO (2013): “INTEREST RATE RULES AND MACROECONOMIC STABILITY UNDER HETEROGENEOUS EXPECTATIONS,” *Macroeconomic Dynamics*, 17(8), 1574–1604.
- BOPPART, T., P. KRUSELL, AND K. MITMAN (2018): “Exploiting MIT shocks in heterogeneous-agent economies: the impulse response as a numerical derivative,” *Journal of Economic Dynamics and Control*, 89, 68–92.
- BRANCH, W. (2006): “Restricted Perceptions Equilibria and Learning in Macroeconomics,” in *Post Walrasian Macroeconomics*, ed. by D. Colander, pp. 135–160. Cambridge University Press, Cambridge, UK.
- BRANCH, W., AND B. MCGOUGH (2009): “A New Keynesian model with heterogeneous expectations,” *Journal of Economic Dynamics and Control*, 33(5), 1036–1051.
- BRANCH, W., B. MCGOUGH, AND M. ZHU (2021): “Statistical Sunspots,” *Theoretical Economics*.
- BRANCH, W. A., AND G. W. EVANS (2006a): “Intrinsic Heterogeneity in Expectation Formation,” *Journal of Economic Theory*, 127, 264–295.
- (2006b): “A simple recursive forecasting model,” *Economics Letters*, 91(2), 158–166.
- (2007): “Model Uncertainty and Endogenous Volatility,” *Review of Economic Dynamics*, 10, 207–237.
- BRANCH, W. A., G. W. EVANS, AND B. MCGOUGH (2013): “Finite Horizon Learning,” in *Macroeconomics at the Service of Public Policy*, ed. by T. J. Sargent, and J. Vilmunen, chap. 9. Oxford University Press.
- BRANCH, W. A., AND B. MCGOUGH (2011): “Business cycle amplification with heterogeneous expectations,” *Economic Theory*, 47(2/3), 395–421.
- (2018): “Chapter 1 - Heterogeneous Expectations and Micro-Foundations in Macroeconomics,” in *Handbook of Computational Economics*, ed. by C. Hommes, and B. LeBaron, vol. 4, pp. 3–62. Elsevier.

- BULLARD, J., G. W. EVANS, AND S. HONKAPOHJA (2008): “Monetary Policy, Judgment and Near-Rational Exuberance,” *American Economic Review*, 98, 1163–1177.
- CAO, D. (2020): “Recursive equilibrium in Krusell and Smith (1998),” *Journal of Economic Theory*, 186, 104978.
- CARROLL, C. D. (2006): “The method of endogenous gridpoints for solving dynamic stochastic optimization problems,” *Economics letters*, 91(3), 312–320.
- DEN HAAN, W. J., K. L. JUDD, AND M. JUILLARD (2010): “Computational suite of models with heterogeneous agents: Incomplete markets and aggregate uncertainty,” *Journal of Economic Dynamics and Control*, 34(1), 1–3.
- EUSEPI, S., AND B. PRESTON (2011): “Expectations, learning, and business cycle fluctuations,” *American Economic Review*, 101(6), 2844–72.
- EVANS, D., G. EVANS, AND B. MCGOUGH (2021a): “The RPEs of RBCs and other DSGEs,” Discussion paper.
- EVANS, D., G. W. EVANS, AND B. MCGOUGH (2021b): “The RPEs of RBCS and other DSGEs,” Discussion paper.
- (2022): “The RPEs of RBCs and other DSGEs,” *Journal of Economic Dynamics and Control*, 143, 104492.
- EVANS, G. W., AND S. HONKAPOHJA (2001): *Learning and Expectations in Macroeconomics*. Princeton University Press, Princeton, New Jersey.
- (2006): “Monetary Policy, Expectations and Commitment*,” *The Scandinavian Journal of Economics*, 108(1), 15–38.
- EVANS, G. W., S. HONKAPOHJA, AND T. J. SARGENT (1993): “On the Preservation of Deterministic Cycles When Some Agents Perceive Them to be Random Fluctuations,” *Journal of Economic Dynamics and Control*, 17, 705–721.
- EVANS, G. W., AND B. MCGOUGH (2018): “Equilibrium Selection, Observability and Backward-Stable Solutions,” *Journal of Monetary Economics*, 98, 1–10.
- (2020a): “Equilibrium Stability in a Nonlinear Cobweb Model,” *Economics Letters*, 193, 109130. <https://doi.org/10.1016/j.econlet.2020.109130>.
- EVANS, G. W., AND B. MCGOUGH (2020b): “Learning to optimize,” Discussion paper.
- EVANS, G. W., AND B. MCGOUGH (2020c): “Stable Near-rational Sunspot Equilibria,” *Journal of Economic Theory*, 186, 104987. <https://doi.org/10.1016/j.jet.2019.104987>.
- GABAIX, X. (2017): “Behavioral Macroeconomics Via Sparse Dynamic Programming [Working Paper],” Discussion paper.

- GABAIX, X. (2020): “A Behavioral New Keynesian Model,” *American Economic Review*, 110(8), 2271–2327.
- GASTEIGER, E. (2018): “DO HETEROGENEOUS EXPECTATIONS CONSTITUTE A CHALLENGE FOR POLICY INTERACTION?,” *Macroeconomic Dynamics*, 22(8), 2107–2140.
- GIUSTO, A. (2014): “Adaptive learning and distributional dynamics in an incomplete markets model,” *Journal of Economic Dynamics and Control*, 40(C), 317–333.
- HOMMES, C. (2013): *Behavioral Rationality and Heterogeneous Expectations in Complex Economic Systems*. Cambridge University Press.
- HOMMES, C. H., AND M. ZHU (2014): “Behavioral Learning equilibria,” *Journal of Economic Theory*, 150, 778–814.
- HONKAPOHJA, S., AND K. MITRA (2006): “Learning Stability in Economies with Heterogeneous Agents,” *Review of Economic Dynamics*, 9(2), 284–309.
- KOPECKY, K. A., AND R. M. SUEN (2010): “Finite state Markov-chain approximations to highly persistent processes,” *Review of Economic Dynamics*, 13(3), 701–714.
- KREPS, D. M. (1998): “Anticipated Utility and Dynamic Choice,” pp. 242–274. Cambridge University Press, Cambridge, Cambridge.
- KRUEGER, D., K. MITMAN, AND F. PERRI (2016): “Macroeconomics and household heterogeneity,” in *Handbook of Macroeconomics*, vol. 2, pp. 843–921. Elsevier.
- KRUSELL, P., AND A. A. SMITH, JR (1998): “Income and wealth heterogeneity in the macroeconomy,” *Journal of political Economy*, 106(5), 867–896.
- MALMENDIER, U., AND S. NAGEL (2011): “Depression Babies: Do Macroeconomic Experiences Affect Risk Taking?,” *The Quarterly Journal of Economics*, 126(1), 373–416.
- MARCET, A., AND T. J. SARGENT (1989): “Convergence of Least-Squares Learning in Environments with Hidden State Variables and Private Information,” *Journal of Political Economy*, 97, 1306–1322.
- MILANI, F. (2007): “Expectations, Learning and Macroeconomic Persistence,” *Journal of Monetary Economics*, 54, 2065–2082.
- ORPHANIDES, A., AND J. C. WILLIAMS (2003): “Imperfect Knowledge, Inflation Expectations, and Monetary Policy,” Working Paper 9884, National Bureau of Economic Research.
- PRESTON, B. (2005): “Learning about Monetary Policy Rules when Long-Horizon Expectations Matter,” *International Journal of Central Banking*, 1, 81–126.
- REITER, M. (2009): “Solving heterogeneous-agent models by projection and perturbation,” *Journal of Economic Dynamics and Control*, 33(3), 649–665.

- ROMER, D. (2012): *Advanced macroeconomics*. Fourth edition. New York : McGraw-Hill/Irwin, [2012] ©2012.
- SARGENT, T. J. (1991): “Equilibrium with Signal Extraction from Endogenous Variables,” *Journal of Economic Dynamics and Control*, 15, 245–273.
- SLOBODYAN, S., AND R. WOUTERS (2012): “Estimating a Medium-Scale DSGE Model with Expectations Based on Small Forecasting Models,” *American Economic Journal: Macroeconomics*, 4, 65–101.
- STADLER, G. W. (1994): “Real Business Cycles,” *Journal of Economic Literature*, 32(4), 1750–1783.
- WILLIAMS, N. (2003): “Adaptive Learning and Business Cycles,” Discussion paper.
- (2018): “Escape Dynamics in Learning Models,” *The Review of Economic Studies*, 86(2), 882–912.
- (2019): “Escape Dynamics in Learning Models,” *Review of Economic Studies*, 86, 882–912.
- WOODFORD, M. (2018): “Monetary Policy Analysis when Planning Horizons are Finite,” Working Paper 24692, National Bureau of Economic Research.
- YOUNG, E. R. (2010): “Solving the incomplete markets model with aggregate uncertainty using the Krusell–Smith algorithm and non-stochastic simulations,” *Journal of Economic Dynamics and Control*, 34(1), 36–41.

Appendix: On-line Supplementary Materials

A Appendices for Section 3

The locally rational model can be written as

$$\begin{aligned}
 u_c(\hat{c}_t(a, \varepsilon, \psi), \hat{l}_t(a, \varepsilon, \psi)) - \beta \bar{\lambda}^e(\hat{a}_t(a, \varepsilon, \psi), \varepsilon) \exp(\langle \psi, X_t \rangle) &\geq 0 \\
 u_l(\hat{c}_t(a, \varepsilon, \psi), \hat{l}_t(a, \varepsilon, \psi)) - u_c(\hat{c}_t(a, \varepsilon, \psi), \hat{l}_t(a, \varepsilon, \psi)) w_t &= 0 \\
 \hat{a}_t(a, \varepsilon, \psi) + \hat{c}_t(a, \varepsilon, \psi) - (1 + r_t)a - w_t \varepsilon (1 - \hat{l}_t(a, \varepsilon, \psi)) &= 0
 \end{aligned} \tag{A.1}$$

with the first equation holding with strict inequality only if $\hat{a}_t(a, \varepsilon, \psi) = \underline{a}$. Define $Y_t = (r_t, w_t)$ as the vector of aggregate variables relevant to the consumer. For a given history of individual shocks, ε^t , the time path of individual choices is given recursively by

$$\begin{aligned}
 \hat{a}_t(\varepsilon^t, \psi) &= \hat{a}_t(\hat{a}_{t-1}(\varepsilon^{t-1}, \psi), \varepsilon_t, \psi) \\
 \hat{c}_t(\varepsilon^t, \psi) &= \hat{c}_t(\hat{a}_{t-1}(\varepsilon^{t-1}, \psi), \varepsilon_t, \psi) \\
 \hat{l}_t(\varepsilon^t, \psi) &= \hat{l}_t(\hat{a}_{t-1}(\varepsilon^{t-1}, \psi), \varepsilon_t, \psi).
 \end{aligned}$$

Similarly one can define the time path of the shadow price of wealth as

$$\hat{\lambda}_t(\varepsilon^t, \psi) = (1 + r_t) u_c(\hat{c}_t(\varepsilon^t, \psi), \hat{l}_t(\varepsilon^t, \psi)).$$

Taking the path of aggregates as given, a fixed point of the T-map ψ^* solves the following system of equations:

$$\lim_{t \rightarrow \infty} \mathbb{E}_0 \left[\left(\log \left(\hat{\lambda}_t(\varepsilon^t, \psi) / \bar{\lambda}(\hat{a}_{t-1}(\varepsilon^{t-1}, \psi), \varepsilon_t) \right) - \langle \psi, X_{t-1} \rangle \right) X_{t-1} \right] = 0. \tag{A.2}$$

This can be written more succinctly as

$$\lim_{t \rightarrow \infty} \mathbb{E}_0 \left[\left(d \log \hat{\lambda}_t(\hat{a}_{t-1}(\varepsilon^{t-1}, \psi^*), \varepsilon_t) - \langle \psi^*, X_{t-1} \rangle \right) X_{t-1} \right] = 0 \tag{A.3}$$

where $d\star_t = \star_t - \bar{\star}$ represents a deviation from the SRE values.

A.1 Proof of Proposition 1

We're focusing on RPE which are local to the SRE, which implies that $dY_t = Y_t - \bar{Y} = \mathcal{O}(\sigma_v)$ and $\mathbb{E}_0[dY_t] = \mathcal{O}(\sigma_v^2)$, where σ_v is the standard deviation of the innovations to TFP. Without loss of generality we will assume that $X_t = \begin{pmatrix} 1 \\ d\hat{X}_t \end{pmatrix}$ where \hat{X}_t is a vector of observables.²⁷ As we are

²⁷Our benchmark model is in exactly this form. More generally one can always demean all variables.

studying an RPE local to the SRE we can assume that $d\hat{X}_t = \mathcal{O}(\sigma)$ and $\mathbb{E}_0 [d\hat{X}_t] = \mathcal{O}(\sigma_v^2)$. It will also prove convenient to partition $\psi = \begin{pmatrix} \psi_0 \\ \psi_1 \end{pmatrix}$, so that $\langle \psi, X_t \rangle = \psi_0 + \langle \psi_1, d\hat{X}_t \rangle$. Under these assumptions, equation (A.3) can be written as two equations

$$\lim_{t \rightarrow \infty} \mathbb{E}_0 \left[d \log \hat{\lambda}_t(\hat{a}_{t-1}(\varepsilon^{t-1}, \psi^*), \varepsilon_t) - \langle \psi^*, X_{t-1} \rangle \right] = 0 \quad (\text{A.4})$$

$$\lim_{t \rightarrow \infty} \mathbb{E}_0 \left[\left(d \log \hat{\lambda}_t(\hat{a}_{t-1}(\varepsilon^{t-1}, \psi^*), \varepsilon_t) - \langle \psi^*, X_{t-1} \rangle \right) d\hat{X}_{t-1} \right] = 0. \quad (\text{A.5})$$

An object of particular interest for us is the variable $\phi_t^* = \langle \psi^*, X_t \rangle$ capturing the deviations in forecasts from SRE values. If we define $\bar{\phi}^* = \langle \psi^*, \bar{X} \rangle = \psi_0^*$ then by construction $d\phi_t^* = \phi_t^* - \bar{\phi}^*$ satisfies $d\phi_t^* = \mathcal{O}(\sigma_v)$ and $\mathbb{E}_0 d\phi_t^* = \mathcal{O}(\sigma_v^2)$. Finally, let $\bar{c}(a, \varepsilon, \phi)$, $\bar{a}(a, \varepsilon, \phi)$, and $\bar{l}(a, \varepsilon, \phi)$ be defined by solving the following non-linear equations

$$\begin{aligned} u_c(\bar{c}(a, \varepsilon, \phi), \bar{l}(a, \varepsilon, \phi)) - \beta \bar{\lambda}^e(\bar{a}(a, \varepsilon, \phi), \varepsilon) \exp(\phi) &\geq 0 \\ u_l(\bar{c}(a, \varepsilon, \phi), \bar{l}(a, \varepsilon, \phi)) - u_c(\bar{c}(a, \varepsilon, \phi), \bar{l}(a, \varepsilon, \phi)) \bar{w} &= 0 \\ \bar{a}(a, \varepsilon, \phi) + \bar{c}(a, \varepsilon, \phi) - (1 + \bar{r})a - \bar{w}\varepsilon(1 - \bar{l}(a, \varepsilon, \phi)) &= 0, \end{aligned}$$

where the first equation holds with strict inequality only if $\bar{a}(a, \varepsilon, \phi) = \underline{a}$. We can use these objects to construct a first order approximation of the policy rules in the following Lemma:

Lemma 1. *There exists functions*

$$c_Y(a, \varepsilon, \bar{\phi}^*), c_\phi(a, \varepsilon, \bar{\phi}^*), a_Y(a, \varepsilon, \bar{\phi}^*), a_\phi(a, \varepsilon, \bar{\phi}^*), l_Y(a, \varepsilon, \bar{\phi}^*), \text{ and } l_\phi(a, \varepsilon, \bar{\phi}^*)$$

such, that almost everywhere,

$$\begin{aligned} d\hat{c}_t(a, \varepsilon, \psi^*) &= \left(\bar{c}(a, \varepsilon, \bar{\phi}^*) - \bar{c}(a, \varepsilon) \right) + c_Y(a, \varepsilon, \bar{\phi}^*) dY_t + c_\phi(a, \varepsilon, \bar{\phi}^*) d\phi_t^* + \mathcal{O}(\sigma_v^2) \\ d\hat{a}_t(a, \varepsilon, \psi^*) &= \left(\bar{a}(a, \varepsilon, \bar{\phi}^*) - \bar{a}(a, \varepsilon) \right) + a_Y(a, \varepsilon, \bar{\phi}^*) dY_t + a_\phi(a, \varepsilon, \bar{\phi}^*) d\phi_t^* + \mathcal{O}(\sigma_v^2) \\ d\hat{l}_t(a, \varepsilon, \psi^*) &= \left(\bar{l}(a, \varepsilon, \bar{\phi}^*) - \bar{l}(a, \varepsilon) \right) + l_Y(a, \varepsilon, \bar{\phi}^*) dY_t + l_\phi(a, \varepsilon, \bar{\phi}^*) d\phi_t^* + \mathcal{O}(\sigma_v^2) \end{aligned}$$

Proof. When $\sigma_v = 0$, equations (A.1) satisfy

$$\begin{aligned} u_c(\hat{c}_t(a, \varepsilon, \psi^*), \hat{l}_t(a, \varepsilon, \psi^*)) - \beta \bar{\lambda}^e(\hat{a}_t(a, \varepsilon, \psi^*), \varepsilon) \exp(\langle \psi^*, \bar{X} \rangle) + \mathcal{O}(\sigma_v) &\geq 0 \\ u_l(\hat{c}_t(a, \varepsilon, \psi^*), \hat{l}_t(a, \varepsilon, \psi^*)) - u_c(\hat{c}_t(a, \varepsilon, \psi^*), \hat{l}_t(a, \varepsilon, \psi^*)) \bar{w} + \mathcal{O}(\sigma_v) &= 0 \quad (\text{A.6}) \\ \hat{a}_t(a, \varepsilon, \psi^*) + \hat{c}_t(a, \varepsilon, \psi^*) - (1 + \bar{r})a - \bar{w}\varepsilon(1 - \hat{l}_t(a, \varepsilon, \psi^*)) + \mathcal{O}(\sigma) &= 0, \end{aligned}$$

which implies that, when $\sigma_v = 0$, $\hat{c}_t(a, \varepsilon, \psi^*) = \bar{c}(a, \varepsilon, \bar{\phi}^*)$, and similarly for the other policy rules. Expanding (A.1), for points of the state space (a, ε) where the borrowing constraint does not bind ($\bar{a}(a, \varepsilon) > \underline{a}$), around this $\sigma_v = 0$ limit implies yields

$$A^+(a, \varepsilon, \bar{\phi}^*) \begin{pmatrix} \hat{c}_t(a, \varepsilon, \psi^*) - \bar{c}(a, \varepsilon, \bar{\phi}^*) \\ \hat{l}_t(a, \varepsilon, \psi^*) - \bar{l}(a, \varepsilon, \bar{\phi}^*) \\ \hat{a}_t(a, \varepsilon, \psi^*) - \bar{a}(a, \varepsilon, \bar{\phi}^*) \end{pmatrix} + B^+(a, \varepsilon, \bar{\phi}^*) dY_t + C^+(a, \varepsilon, \bar{\phi}^*) d\phi_t^* + \mathcal{O}(\sigma_v^2) = 0,$$

where

$$A^+(a, \varepsilon, \bar{\phi}^*) = \begin{pmatrix} u_{cc}(a, \varepsilon, \bar{\phi}^*) & u_{cl}(a, \varepsilon, \bar{\phi}^*) & -\beta \bar{\lambda}_a^e(\bar{a}(a, \varepsilon, \bar{\phi}^*), \varepsilon) \exp(\bar{\phi}^*) \\ u_{cl}(a, \varepsilon, \bar{\phi}^*) - u_{cc}(a, \varepsilon, \bar{\phi}^*) \bar{w} & u_{ll}(a, \varepsilon, \bar{\phi}^*) - u_{cl}(a, \varepsilon, \bar{\phi}^*) \bar{w} & 0 \\ 1 & -\bar{w} \varepsilon & 1 \end{pmatrix}$$

$$B^+(a, \varepsilon, \bar{\phi}^*) = \begin{pmatrix} 0 & 0 \\ 0 & -u_l(a, \varepsilon, \bar{\phi}^*) \\ -a & \varepsilon(1 - \bar{l}(a, \varepsilon, \bar{\phi}^*)) \end{pmatrix} \text{ and } C^+(a, \varepsilon, \bar{\phi}^*) = \begin{pmatrix} -\beta \bar{\lambda}^e(\bar{a}(a, \varepsilon, \bar{\phi}^*), \varepsilon) \exp(\bar{\phi}^*) \\ 0 \\ 0 \end{pmatrix},$$

with $u_{cc}(a, \varepsilon, \bar{\phi}^*) = u_{cc}(\bar{c}(a, \varepsilon, \bar{\phi}^*), \bar{l}(a, \varepsilon, \bar{\phi}^*))$, etc. Similarly, the first order expansion for the points for the points where the borrowing constraint binds ($\bar{a}(a, \varepsilon, \bar{\phi}^*) = 0$) implies

$$A^-(a, \varepsilon, \bar{\phi}^*) \begin{pmatrix} \hat{c}_t(a, \varepsilon, \psi^*) - \bar{c}(a, \varepsilon, \bar{\phi}^*) \\ \hat{l}_t(a, \varepsilon, \psi^*) - \bar{l}(a, \varepsilon, \bar{\phi}^*) \\ \hat{a}_t(a, \varepsilon, \psi^*) - \bar{a}(a, \varepsilon, \bar{\phi}^*) \end{pmatrix} + B^-(a, \varepsilon, \bar{\phi}^*) dY_t + C^-(a, \varepsilon, \bar{\phi}^*) d\phi_t^* + \mathcal{O}(\sigma_v^2) = 0,$$

where

$$A^-(a, \varepsilon, \bar{\phi}^*) = \begin{pmatrix} 0 & 0 & 1 \\ u_{cl}(a, \varepsilon, \bar{\phi}^*) - u_{cc}(a, \varepsilon, \bar{\phi}^*) \bar{w} & u_{ll}(a, \varepsilon, \bar{\phi}^*) - u_{cl}(a, \varepsilon, \bar{\phi}^*) \bar{w} & 0 \\ 1 & -\bar{w} \varepsilon & 1 \end{pmatrix}$$

$$B^-(a, \varepsilon, \bar{\phi}^*) = \begin{pmatrix} 0 & 0 \\ 0 & -u_l(a, \varepsilon, \bar{\phi}^*) \\ -a & \varepsilon(1 - \bar{l}(a, \varepsilon, \bar{\phi}^*)) \end{pmatrix}$$

$$C^-(a, \varepsilon, \bar{\phi}^*) = \begin{pmatrix} 0 \\ 0 \\ 0 \end{pmatrix}.$$

Combining these two facts together implies that

$$\begin{pmatrix} d\hat{c}_t(a, \varepsilon, \psi^*) \\ d\hat{l}_t(a, \varepsilon, \psi^*) \\ d\hat{a}_t(a, \varepsilon, \psi^*) \end{pmatrix} = \begin{pmatrix} \bar{c}(a, \varepsilon, \bar{\phi}^*) - \bar{c}(a, \varepsilon) \\ \bar{l}(a, \varepsilon, \bar{\phi}^*) - \bar{l}(a, \varepsilon) \\ \bar{a}(a, \varepsilon, \bar{\phi}^*) - \bar{a}(a, \varepsilon) \end{pmatrix} + \begin{cases} A^+(a, \varepsilon, \bar{\phi}^*)^{-1} B^+(a, \varepsilon, \bar{\phi}^*) dY_t & \text{if } \bar{a}(a, \varepsilon, \bar{\phi}^*) > 0 \\ A^-(a, \varepsilon, \bar{\phi}^*)^{-1} B^-(a, \varepsilon, \bar{\phi}^*) dY_t & \text{if } \bar{a}(a, \varepsilon, \bar{\phi}^*) = 0 \end{cases}$$

$$+ \begin{cases} A^+(a, \varepsilon, \bar{\phi}^*)^{-1} C^+(a, \varepsilon, \bar{\phi}^*) dY_t & \text{if } \bar{a}(a, \varepsilon, \bar{\phi}^*) > 0 \\ A^-(a, \varepsilon, \bar{\phi}^*)^{-1} C^-(a, \varepsilon, \bar{\phi}^*) dY_t & \text{if } \bar{a}(a, \varepsilon, \bar{\phi}^*) = 0 \end{cases} + \mathcal{O}(\sigma_v^2),$$

which completes the proof. \square

With these policy rules in hand we can directly construct λ_Y and λ_ϕ such that

$$d \log \hat{\lambda}_t(a, \varepsilon, \psi^*) = \left(\log \bar{\lambda}(a, \varepsilon, \bar{\phi}^*) - \log \bar{\lambda}(a, \varepsilon) \right) + \lambda_Y(a, \varepsilon, \bar{\phi}^*) dY_t + \lambda_\phi(a, \varepsilon, \bar{\phi}^*) d\phi_t^* + \mathcal{O}(\sigma_v^2) \quad (\text{A.7})$$

We note that if $\bar{\phi}^* \neq 0$ then, to zeroth order, the policies of the locally rational agent will not align with the rational agent. A direct corollary of the expansion above shows that if ψ^* is chosen optimally this cannot be the case.

Corollary 1. $\bar{\phi}^* = \langle \psi^*, \bar{X} \rangle = \psi_0^* = 0 + \mathcal{O}(\sigma_V^2)$.

Proof. We have that $d \log \hat{\lambda}_t(a, \varepsilon) = dr_t - \gamma d\hat{c}_t(a, \varepsilon, \psi^*)$, thus a first order expansion of (A.4) yields

$$\lim_{t \rightarrow \infty} \mathbb{E}_0 \left[\left(d \log \hat{\lambda}_t(\bar{a}(\varepsilon^t, \bar{\phi}^*), \varepsilon_t, \psi^*) - \langle \psi^*, \bar{X} + dX_{t-1} \rangle \right) + \mathcal{O}(\sigma_V^2) \right] = 0,$$

where $\bar{a}(\varepsilon^t, \bar{\phi}^*)$ represents the zeroth order path of assets defined recursively via $\bar{a}(\varepsilon^t, \bar{\phi}^*) = \bar{a}(\bar{a}(\varepsilon^{t-1}, \bar{\phi}^*), \varepsilon_t, \bar{\phi}^*)$. Applying equation (A.7) along with $\mathbb{E}_0[dY_t] = \mathcal{O}(\sigma_V^2) = \mathbb{E}_0[dX_t]$ implies

$$\lim_{t \rightarrow \infty} \mathbb{E}_0 \left[\log \bar{\lambda}(\bar{a}(\varepsilon^{t-1}, \bar{\phi}^*), \varepsilon_t, \bar{\phi}^*) - \log \bar{\lambda}(\bar{a}(\varepsilon^{t-1}, \bar{\phi}^*), \varepsilon_t) - \bar{\phi}^* \right] + \mathcal{O}(\sigma_V^2) = 0. \quad (\text{A.8})$$

We now compare $\bar{\lambda}(a, \varepsilon, \bar{\phi}^*)$ to $\bar{\lambda}(a, \varepsilon)$. First, we note that by definition $\bar{\lambda}(a, \varepsilon, 0) = \bar{\lambda}(a, \varepsilon)$. Next we consider $\bar{\lambda}(a, \varepsilon, \bar{\phi}^*)$ for $\bar{\phi}^* > 0$. Increasing $\bar{\phi}^*$ raises the shadow price of savings, therefore we can conclude that $\bar{a}(a, \varepsilon, \bar{\phi}^*)$ is increasing in $\bar{\phi}^*$ and hence

$$\log \bar{\lambda}^e(\bar{a}(a, \varepsilon, \bar{\phi}^*), \varepsilon) < \bar{\lambda}^e(a, \varepsilon).$$

Thus

$$\begin{aligned} \bar{\lambda}(a, \varepsilon, \bar{\phi}^*) &= (1 + \bar{r}) u_c(\bar{c}(a, \varepsilon, \bar{\phi}^*), \bar{l}(a, \varepsilon, \bar{\phi}^*)) \\ &= (1 + \bar{r}) \beta \bar{\lambda}^e(\bar{a}(a, \varepsilon, \bar{\phi}^*), \varepsilon) \exp(\bar{\phi}^*) < \bar{\lambda}(a, \varepsilon) \exp(\bar{\phi}^*). \end{aligned}$$

For constrained agents we have that $\bar{\lambda}(a, \varepsilon, \bar{\phi}^*) = \bar{\lambda}(a, \varepsilon)$ for all $\bar{\phi}^* > 0$, and thus we have

$$\log \left(\bar{\lambda}(a, \varepsilon, \bar{\phi}^*) / \bar{\lambda}(a, \varepsilon) \right) < \bar{\phi}^*$$

for all (a, ε) when $\bar{\phi}^* > 0$. We conclude that

$$\lim_{t \rightarrow \infty} \mathbb{E}_0 \left[\log \bar{\lambda}(\bar{a}(\varepsilon^{t-1}, \bar{\phi}^*), \varepsilon_t, \bar{\phi}^*) - \log \bar{\lambda}(\bar{a}(\varepsilon^{t-1}, \bar{\phi}^*), \varepsilon_t) - \bar{\phi}^* \right] + \mathcal{O}(\sigma_V^2) < 0$$

for all $\bar{\phi}^* > 0$. Similar arguments show that

$$\lim_{t \rightarrow \infty} \mathbb{E}_0 \left[\log \bar{\lambda}(\bar{a}(\varepsilon^{t-1}, \bar{\phi}^*), \varepsilon_t, \bar{\phi}^*) - \log \bar{\lambda}(\bar{a}(\varepsilon^{t-1}, \bar{\phi}^*), \varepsilon_t) - \bar{\phi}^* \right] + \mathcal{O}(\sigma_V^2) > 0$$

for all $\bar{\phi}^* < 0$. We thus conclude, via equation (A.8), that $\bar{\phi}^* = 0 + \mathcal{O}(\sigma_V^2)$. \square

We can now turn to proving the main result. As $d\phi_t^* = d\hat{X}_t^T \psi_1^*$, this corollary implies that

$$d \log \hat{\lambda}_t(a, \varepsilon, \psi^*) = \lambda_Y(a, \varepsilon) dY_t + \lambda_\phi(a, \varepsilon) d\hat{X}_t^T \psi_1^* + \mathcal{O}(\sigma_V^2),$$

where $\lambda_Y(a, \varepsilon) = \lambda_Y(a, \varepsilon, 0)$ and $\lambda_\phi(a, \varepsilon) = \lambda_\phi(a, \varepsilon, \bar{\phi}^*)$. A second order expansion of (A.5) then yields

$$\lim_{t \rightarrow \infty} \mathbb{E}_0 \left[d\hat{X}_{t-1} \left(dY_t^T \lambda_Y(\bar{a}(\varepsilon^{t-1}), \varepsilon_t)^T + \lambda_\phi(\bar{a}(\varepsilon^{t-1}), \varepsilon_t) d\hat{X}_t^T \psi_1^* - d\hat{X}_{t-1}^T \psi_1^* \right) \right] + \mathcal{O}(\sigma_v^3) = 0,$$

We note that $(\bar{a}(\varepsilon^{t-1}), \varepsilon_t)$ is independent of dY_t and $d\hat{X}_t$, and converges in probability to the stationary joint distribution of assets and productivities Ω^* . If we define $\Sigma_{XY} = \lim_{t \rightarrow \infty} \mathbb{E}_0 [d\hat{X}_{t-1} dY_t^T]$, $\Sigma_{XX} = \lim_{t \rightarrow \infty} \mathbb{E}_0 [d\hat{X}_{t-1} d\hat{X}_{t-1}^T]$, and $\Sigma_{XX+} = \lim_{t \rightarrow \infty} \mathbb{E}_0 [d\hat{X}_{t-1} d\hat{X}_t^T]$, then we have

$$\Sigma_{XY} D_Y + D_\phi \Sigma_{XX+} \psi_1^* - \Sigma_{XX} \psi_1^* + \mathcal{O}(\sigma_v^3) = 0,$$

where

$$D_Y = \int \lambda_Y(a, \varepsilon)^T d\Omega^*(a, \varepsilon)$$

$$D_\phi = \int \lambda_\phi(a, \varepsilon) d\Omega^*(a, \varepsilon).$$

The same arguments as the proof of Corollary 1 imply that $|\lambda_\phi(a, \varepsilon)| < 1$ and hence $|D_\phi| < 1$. This, along with stationarity of $d\hat{X}_t$, implies that $\Sigma_{XX} - D_\phi \Sigma_{XX+}$ invertible and hence

$$\psi_1^* = (\Sigma_{XX} - D_\phi \Sigma_{XX+})^{-1} \Sigma_{XY} D_Y + \mathcal{O}(\sigma_v)$$

is uniquely determined for small enough σ_v .²⁸ Thus for small enough σ_v the RPE must be unique.

B Appendices for Section 4

We begin by studying the CARA economy under the assumption of perfect foresight. As noted by Acharya and Dogra (2020) and Boppart, Krusell, and Mitman (2018) linearizing the the model with perfect foresight is equivalent to linearizing the full stochastic model.

Let $\bar{c}_t(a, \varepsilon)$ represent the individual consumption policy rules when $\sigma_\eta = 0$ for a given r_0 . Guess the following functional form for the consumption policy rule for the rational agent

$$\bar{c}_t(a, \varepsilon) = \bar{\mathcal{C}}_t + \bar{\mu}_t((1+r_t)a + \varepsilon).$$

This implies the following policy rule for individual wealth

$$\bar{a}_t(a, \varepsilon) = (1 - \bar{\mu}_t)((1+r_t)a + \varepsilon) - \bar{\mathcal{C}}_t,$$

which implies that

$$\bar{c}_{t+1}(\bar{a}_t(a, \varepsilon), \varepsilon') = \bar{\mathcal{C}}_{t+1} + \bar{\mu}_{t+1}((1+r_{t+1})((1 - \bar{\mu}_t)((1+r_t)a + \varepsilon) - \bar{\mathcal{C}}_t) + \varepsilon').$$

²⁸Note that $\Sigma_{XX} - D_\phi \Sigma_{XX+} = \mathcal{O}(\sigma_v)^2$

Therefore, conditional on (a, ε) , $-\gamma \bar{c}_{t+1}(\bar{a}_t(a, \varepsilon), \varepsilon')$ has mean

$$-\gamma(\bar{\mathcal{C}}_{t+1} + \bar{\mu}_{t+1}((1+r_{t+1})((1-\bar{\mu}_t)((1+r_t)a + \varepsilon) - \bar{\mathcal{C}}_t) + 1))$$

and variance

$$\gamma^2 \bar{\mu}_{t+1}^2 \sigma_\varepsilon^2$$

Exponentiating $-\gamma \bar{c}_{t+1}(\bar{a}_t(a, \varepsilon), \varepsilon')$ and taking expectations yields

$$\begin{aligned} \log \mathbb{E}_t \exp(-\gamma \bar{c}_{t+1}(\bar{a}_t(a, \varepsilon), \varepsilon')) &= -\gamma(\bar{\mathcal{C}}_{t+1} + \bar{\mu}_{t+1}(1+r_{t+1})((1-\bar{\mu}_t)((1+r_t)a + \varepsilon) - \bar{\mathcal{C}}_t)) \\ &\quad + \frac{1}{2} \gamma^2 \bar{\mu}_{t+1}^2 \sigma_\varepsilon^2 \\ &= -\gamma \bar{\mu}_{t+1}(1+r_{t+1})(1-\bar{\mu}_t)((1+r_t)a + \varepsilon) - \gamma \bar{\mu}_{t+1} \\ &\quad - \gamma(\bar{\mathcal{C}}_{t+1} - \bar{\mu}_{t+1}(1+r_{t+1})\bar{\mathcal{C}}_t) + \frac{1}{2} \gamma^2 \bar{\mu}_{t+1}^2 \sigma_\varepsilon^2 \end{aligned}$$

which can be plugged into the Euler equation to get

$$\begin{aligned} -\gamma(\bar{\mathcal{C}}_t + \bar{\mu}_t((1+r_t)a + \varepsilon)) &= \log(\beta \phi(1+r_{t+1})) - \gamma \bar{\mu}_{t+1}(1+r_{t+1})(1-\bar{\mu}_t)((1+r_t)a + \varepsilon) \\ &\quad - \gamma \bar{\mu}_{t+1} - \gamma(\bar{\mathcal{C}}_{t+1} - \bar{\mu}_{t+1}(1+r_{t+1})\bar{\mathcal{C}}_t) + \frac{1}{2} \gamma^2 \bar{\mu}_{t+1}^2 \sigma_\varepsilon^2. \end{aligned}$$

This equation must hold for all (a, ε) which implies

$$\bar{\mu}_t = \bar{\mu}_{t+1}(1+r_{t+1})(1-\bar{\mu}_t) \tag{B.9}$$

$$\bar{\mathcal{C}}_t = -\frac{1}{\gamma} \log(\beta \phi(1+r_{t+1})) + \bar{\mathcal{C}}_{t+1} - \bar{\mu}_{t+1}((1+r_{t+1})\bar{\mathcal{C}}_t + 1) - \frac{1}{2} \gamma \bar{\mu}_{t+1}^2 \sigma_\varepsilon^2 \tag{B.10}$$

B.1 Proof of Proposition 2

Assuming $r_0 = \bar{r}$ we have $r_t = \bar{r}$ which implies

$$\bar{c}_t(a, \varepsilon) = \bar{c}(a, \varepsilon) = \bar{\mathcal{C}} + \bar{\mu}((1+\bar{r})a + \varepsilon).$$

Equations (B.9) and (B.10) imply that $\bar{\mu}$ and $\bar{\mathcal{C}}$ must satisfy

$$\bar{\mu} = \bar{\mu}(1+\bar{r})(1-\bar{\mu})$$

and

$$\bar{\mathcal{C}} = -\frac{1}{\gamma} \log(\beta \phi(1+\bar{r})) + \bar{\mathcal{C}} - \bar{\mu}(1+\bar{r})\bar{\mathcal{C}} - \bar{\mu} - \frac{1}{2} \gamma \bar{\mu}^2 \sigma_\varepsilon^2$$

Simplifying these two equations implies

$$\bar{\mu} = \frac{\bar{r}}{(1+\bar{r})}$$

and

$$\bar{\mathcal{C}} = \frac{-\frac{1}{\gamma} \log(\beta \phi(1+\bar{r})) - \bar{\mu} - \frac{1}{2} \gamma \bar{\mu}^2 \sigma_\varepsilon^2}{\bar{r}}.$$

B.2 Proof of Proposition 3

Define $d\star_t \equiv \star_t - \bar{\star}$ represent the deviation of an object from its steady-state value. Differentiating equation (B.9) implies

$$d\bar{\mu}_t = d\bar{\mu}_{t+1} - \bar{\mu}(1 + \bar{r})d\bar{\mu}_t + \bar{\mu}(1 - \bar{\mu})dr_{t+1} + \mathcal{O}(dr_t^2).$$

We require $\bar{\mu}_r$, which is the change in $\bar{\mu}_t$ given a change in r_t , accounting for the impact of the implied changes in future values of r on current $\bar{\mu}$. Differentiating (17) of the main text implies $dr_{t+1} = \rho dr_t + \mathcal{O}(dr_t^2)$, so that, to first order,

$$d\bar{\mu}_t = d\bar{\mu}_{t+1} - \bar{\mu}(1 + \bar{r})d\bar{\mu}_t + \bar{\mu}(1 - \bar{\mu})(1 + \bar{r})\rho dr_t.$$

Since $\bar{\mu}(1 - \bar{\mu})(1 + \bar{r}) = \mu$ and $1 + \bar{\mu}(1 + \bar{r}) = 1 + \bar{r}$, we have

$$d\bar{\mu}_t = (1 + \bar{r})^{-1} (d\bar{\mu}_{t+1} + \bar{\mu}\rho dr_t).$$

Forward iterating, we may finally conclude that

$$\bar{\mu}_r = \left(\frac{\bar{\mu}\rho}{1 + \bar{r}} \right) \sum_{n \geq 0} \left(\frac{\rho}{1 + \bar{r}} \right)^n dr_t = \frac{\bar{\mu}\rho}{(1 + \bar{r}) - \rho}. \quad (\text{B.11})$$

Similarly, differentiating (B.10) implies

$$d\bar{\mathcal{C}}_t = -\frac{1}{\gamma}dr_{t+1} + d\bar{\mathcal{C}}_{t+1} - (1 + \bar{r})\bar{\mathcal{C}}d\bar{\mu}_{t+1} - \bar{\mu}(1 + \bar{r})d\bar{\mathcal{C}}_t - \bar{\mu}\bar{\mathcal{C}}dr_{t+1} - d\bar{\mu}_{t+1} - \gamma\bar{\mu}\sigma_\varepsilon^2 d\bar{\mu}_{t+1} + \mathcal{O}(dr_t^2),$$

which we may write as $d\bar{\mathcal{C}}_t = (1 + \bar{r})^{-1} (d\bar{\mathcal{C}}_{t+1} + \nabla dr_t)$, where

$$\nabla = -\rho (\gamma^{-1} + \bar{\mu}\bar{\mathcal{C}}(1 + \bar{r}) + \bar{\mu}_r(1 + \bar{\mathcal{C}}(1 + r)) + \gamma\bar{\mu}\sigma_\varepsilon^2 \bar{\mu}_r\rho).$$

Forward iterating as above, we conclude

$$\bar{\mathcal{C}}_r = \frac{\nabla}{(1 + \bar{r}) - \rho}. \quad (\text{B.12})$$

As $\bar{c}_t(a, \varepsilon) = \bar{\mathcal{C}}_t + \bar{\mu}_t((1 + r_t)a + \varepsilon)$ we have

$$\begin{aligned} d\bar{c}_t(a, \varepsilon) &= d\bar{\mathcal{C}}_t + d\bar{\mu}_t((1 + \bar{r})a + \varepsilon) + \bar{\mu}(1 + \bar{r})adr_t \\ &= (\bar{\mathcal{C}}_r + \bar{\mu}_r((1 + \bar{r})a + \varepsilon) + \bar{\mu}(1 + \bar{r})a) dr_t \end{aligned}$$

and

$$d\bar{a}_t(a, \varepsilon) = a(1 + \bar{r})dr_t - d\bar{c}_t(a, \varepsilon) = (a - \bar{\mathcal{C}}_r - \bar{\mu}_r((1 + \bar{r})a + \varepsilon)) dr_t.$$

Next, we consider the locally rational decisions rules that solve

$$\begin{aligned} -\gamma \hat{c}_t(a, \varepsilon, \psi) &= \log(\beta \phi) + \log\left(\bar{\lambda}^e(\hat{a}_t(a, \varepsilon, \psi), \varepsilon)\right) + \psi \log\left(\frac{(1+r_t)}{(1+\bar{r})}\right) \\ \hat{a}_t(a, \varepsilon, \psi) &= (1+r_t)a + y - \hat{c}_t(a, \varepsilon, \psi). \end{aligned}$$

When $r_t = \bar{r}$ this is solved by $\hat{c}_t(a, \varepsilon, \psi) = \bar{c}(a, \varepsilon)$ and $\hat{a}_t(a, \varepsilon, \psi) = \bar{a}(a, \varepsilon)$ so for small deviations dr_t we have

$$\begin{aligned} -\gamma d\hat{c}_t(a, \varepsilon, \psi) &= \frac{\bar{\lambda}_a^e(\bar{a}(a, \varepsilon), \varepsilon)}{\bar{\lambda}^e(\bar{a}(a, \varepsilon), \varepsilon)} d\hat{a}_t(a, \varepsilon, \psi) + \psi dr_t + \mathcal{O}(dr_t^2) \\ d\hat{a}_t(a, \varepsilon, \psi) &= (1+\bar{r})adr_t - d\hat{c}_t(a, \varepsilon, \psi) + \mathcal{O}(dr_t^2). \end{aligned}$$

To simplify this expression note that

$$\begin{aligned} \bar{\lambda}_a^e(a', \varepsilon) &= \int \bar{\lambda}_a(a', \varepsilon') d\Pr(\varepsilon') \\ &= -\gamma \int (1+\bar{r}) \exp(-\gamma \bar{c}(a', \varepsilon')) \underbrace{\bar{c}_a(a', \varepsilon')}_{\bar{\mu}(1+\bar{r})} d\Pr(\varepsilon') \\ &= -\gamma \bar{\mu}(1+\bar{r}) \int (1+\bar{r}) \exp(-\gamma \bar{c}(a', \varepsilon')) d\Pr(\varepsilon') \\ &= -\gamma \bar{\mu}(1+\bar{r}) \bar{\lambda}^e(a', \varepsilon), \end{aligned}$$

which then implies that $d\hat{c}_t$ and $d\hat{a}_t$ satisfy

$$\begin{aligned} d\hat{c}_t(a, \varepsilon, \psi) &= \bar{\mu}(1+\bar{r})d\hat{a}_t(a, \varepsilon, \psi) - \frac{\psi}{\gamma} dr_t + \mathcal{O}(dr_t^2) \\ d\hat{a}_t(a, \varepsilon, \psi) &= (1+\bar{r})adr_t - d\hat{c}_t(a, \varepsilon, \psi) + \mathcal{O}(dr_t^2). \end{aligned}$$

Solving for $d\hat{a}_t(a, \varepsilon, \psi)$ and $d\hat{c}_t(a, \varepsilon, \psi)$ then yields

$$d\hat{c}_t(a, \varepsilon, \psi) = \left(\bar{\mu}(1+\bar{r})a - \frac{\psi}{\gamma(1+\bar{r})} \right) dr_t + \mathcal{O}(dr_t^2) \quad (\text{B.13})$$

$$d\hat{a}_t(a, \varepsilon, \psi) = \left(a + \frac{\psi}{\gamma(1+\bar{r})} \right) dr_t + \mathcal{O}(dr_t^2). \quad (\text{B.14})$$

We conclude by noting that under rational expectations the perfect foresight and stochastic economies are equivalent to first order and so

$$dc_t(a, \varepsilon) = (\bar{\mathcal{C}}_r + \bar{\mu}_r((1+\bar{r})a + \varepsilon) + \bar{\mu}(1+\bar{r})a) dr_t + \mathcal{O}(\sigma_\eta^2) \quad (\text{B.15})$$

$$da_t(a, \varepsilon) = (a - \bar{\mathcal{C}}_r - \bar{\mu}_r((1+\bar{r})a + \varepsilon)) dr_t + \mathcal{O}(\sigma_\eta^2). \quad (\text{B.16})$$

As $\mathcal{O}(dr_t^2) = \mathcal{O}(\sigma_\eta^2)$, by comparing equations (B.14) and (B.16) we conclude that

$$a_t(a, \varepsilon) = \hat{a}(a, \varepsilon, \psi^{RE}(a, \varepsilon)) + \mathcal{O}(\sigma_\eta^2)$$

and, hence,

$$c_t(a, \varepsilon) = \hat{c}(a, \varepsilon, \psi^{RE}(a, \varepsilon)) + \mathcal{O}(\sigma_\eta^2)$$

for

$$\psi^{RE}(a, \varepsilon) = \underbrace{-\gamma(1+\bar{r})(\bar{\mathcal{C}}_r + \bar{\mu}_r)}_{\psi_0} + \underbrace{-\gamma(1+\bar{r})^2 \bar{\mu}_r a}_{\psi_a} + \underbrace{-\gamma(1+\bar{r}) \bar{\mu}_r (\varepsilon - 1)}_{\psi_\varepsilon}.$$

Equation (B.11) implies that $\psi_a < 0$ while adding together (B.11) and (B.12) implies²⁹

$$\bar{\mathcal{C}}_r + \bar{\mu}_r = \frac{-\frac{1}{\gamma}\rho - 2\bar{\mu}_r\rho - \gamma\bar{\mu}\sigma_\varepsilon^2\bar{\mu}_r\rho}{(1+\bar{r}) - \rho},$$

which guarantees $\psi_0 > 0$ as desired.

B.3 Proof of Theorem 1

Let ι_t the shock that captures whether the agent dies at the end of period t . Define $s_t = (\varepsilon_t, \iota_t, \eta_t)$ as the vector of time t shocks. For a given belief ψ , the agent's stochastic path for wealth is given by

$$a_t(s^t, \psi) = (1 - \iota_t) \hat{a}_t(a_{t-1}(s^{t-1}, \psi), \varepsilon_t, \psi)$$

with initial conditions $a_{-1} = 0$. Define

$$\bar{a}_t(s^t) = (1 - \iota_t) \bar{a}(\bar{a}_{t-1}(s^{t-1}), \varepsilon_t) = (1 - \iota_t) (\bar{a}_{t-1}(s^{t-1}) + (1 - \bar{\mu}) \varepsilon_t)$$

as the stochastic wealth position of the agent when $\sigma_\eta = 0$ with initial conditions $\bar{a}_{-1} = 0$. Finally, define

$$a_t^{(1)}(s^t, \psi) = (1 - \iota_t) \left(a_{t-1}^{(1)}(s^{t-1}, \psi) + \left(\bar{a}_{t-1}(s^{t-1}) + \frac{\psi}{\gamma(1+\bar{r})} \right) dr_t \right)$$

with the initial condition $a_{-1}^{(1)} = 0$. By construction $a_t^{(1)}(s^t, \psi) = \mathcal{O}(\sigma_\eta)$, we show in the following claim that that it captures the first order approximation to $a_t(s^t, \psi)$ relative to $\bar{a}_t(s^t)$.

Claim 1. $a_t(s^t, \psi) = \bar{a}_t(s^t) + a_t^{(1)}(s^t, \psi) + \mathcal{O}(\sigma_\eta^2)$

²⁹Recall we are focusing on the case where average assets are zero so $\bar{\mathcal{C}}(1+\bar{r}) = 1$

Proof. We proceed by induction, as the claim holds for $t = -1$. Assume true for $t - 1$ then

$$\begin{aligned}
a_t(s^t, \psi) &= (1 - \iota_t) \hat{a}_t(a_{t-1}(s^{t-1}, \psi), \varepsilon_t, \psi) \\
&= (1 - \iota_t) \bar{a}(a_{t-1}(s^{t-1}, \psi), \varepsilon_t) + (1 - \iota_t) \left(a_{t-1}(s^{t-1}, \psi) + \frac{\psi}{\gamma(1 + \bar{r})} \right) dr_t + \mathcal{O}(\sigma_\eta^2) \\
&= (1 - \iota_t) \left(\bar{a}_{t-1}(s^{t-1}) + a_{t-1}^{(1)}(s^{t-1}, \psi) + \mathcal{O}(\sigma_\eta^2) + (1 - \bar{\mu}) \varepsilon_t \right) \\
&\quad + (1 - \iota_t) \left(\bar{a}_{t-1}(s^{t-1}) + a_{t-1}^{(1)}(s^{t-1}, \psi) + \mathcal{O}(\sigma_\eta^2) + \frac{\psi}{\gamma(1 + \bar{r})} \right) dr_t + \mathcal{O}(\sigma_\eta^2) \\
&= (1 - \iota_t) \left(\bar{a}_{t-1}(s^{t-1}) + (1 - \bar{\mu}) \varepsilon_t \right) \\
&\quad + (1 - \iota_t) \left(a_{t-1}^{(1)}(s^{t-1}, \psi) + \left(\bar{a}_{t-1}(s^{t-1}) + \frac{\psi}{\gamma(1 + \bar{r})} \right) dr_t \right) + \mathcal{O}(\sigma_\eta^2) \\
&= \bar{a}_t(s^t) + a_t^{(1)}(s^t, \psi) + \mathcal{O}(\sigma_\eta^2).
\end{aligned}$$

The third equality used our knowledge that $a_{t-1}^{(1)}(s^{t-1}, \psi) = \mathcal{O}(\sigma_\eta)$ so that

$$a_{t-1}^{(1)}(s^{t-1}, \psi) dr_t = \mathcal{O}(\sigma_\eta^2).$$

□

We can then directly use the methodology of Evans, Evans, and McGough (2022) to show that the process for $a_t(s^t, \psi)$ is locally a contraction and thus there must exist and ergodic distribution. Our next step is to characterize an RPE in this environment. The log shadow price of wealth, at history s^t , for the individual is given by

$$\log \lambda_t(s^t, \psi) = \log \hat{\lambda}(a_{t-1}(s^{t-1}), \varepsilon_t, \psi) = \log(1 + r_t) - \gamma \hat{c}_t(a_{t-1}(s^{t-1}), \varepsilon_t, \psi).$$

In the definition of an RPE, the agent chooses ψ optimally once and for all in order to best forecast deviations in the shadow price from the steady state according to the PLM in equation (18). That optimal choice is characterized by the following optimality condition

$$0 = \lim_{t \rightarrow \infty} \mathbb{E}_0 \left[\log \left(\frac{1 + r_{t-1}}{1 + \bar{r}} \right) \left(\log \left(\lambda_t(s^t, \psi) / \bar{\lambda}(a_{t-1}(s^{t-1}), \varepsilon_t) \right) - \psi \log \left(\frac{1 + r_{t-1}}{1 + \bar{r}} \right) \right) \right].$$

A first order expansion of $\log((1 + r_{t-1})/(1 + \bar{r}))$ implies

$$\log((1 + r_{t-1})/(1 + \bar{r})) = dr_{t-1} + \mathcal{O}(\sigma_\eta^2).$$

As $\log \bar{\lambda}(a, \varepsilon) = \log((1 + \bar{r})) - \gamma \bar{c}(a, \varepsilon)$, we compute the following first order expansion:

$$\begin{aligned} \log \left(\frac{\lambda_t(s^t, \psi)}{\bar{\lambda}(a_{t-1}(s^{t-1}), \varepsilon_t)} \right) &= dr_t - \gamma d\hat{c}_t(a_{t-1}(s^{t-1}), \varepsilon_t, \psi) + \mathcal{O}(\sigma_\eta^2) \\ &= dr_t - \gamma \left[\bar{\mu}(1 + \bar{r}) \left(\bar{a}_{t-1}(s^{t-1}) + a_{t-1}^{(1)}(s^{t-1}, \psi) + \mathcal{O}(\sigma_\eta^2) \right) \right. \\ &\quad \left. - \frac{\psi}{\gamma(1 + \bar{r})} \right] dr_t + \mathcal{O}(\sigma_\eta^2) \\ &= dr_t - \gamma \left(\bar{\mu}(1 + \bar{r}) \bar{a}_{t-1}(s^{t-1}) - \frac{\psi}{\gamma(1 + \bar{r})} \right) dr_t + \mathcal{O}(\sigma_\eta^2). \end{aligned}$$

A second order expansion of the optimality condition for the RPE implies that ψ^{RPE} must satisfy

$$0 = \lim_{t \rightarrow \infty} \mathbb{E}_0 \left[dr_{t-1} \left(dr_t - \gamma \left(\bar{\mu}(1 + \bar{r}) \bar{a}_{t-1}(s^{t-1}) - \frac{\psi^{RPE}}{\gamma(1 + \bar{r})} \right) dr_t - \psi^{RPE} dr_{t-1} \right) \right] + \mathcal{O}(\sigma_\eta^3).$$

Finally, we note that $\bar{a}_t(s^t)$ and dr_t are independent with our model set such that $\mathbb{E}_0[\bar{a}_t(s^t)] = 0$ and also that $\lim_{t \rightarrow \infty} \mathbb{E}_0[dr_t dr_{t-1}] = \rho \lim_{t \rightarrow \infty} \mathbb{E}_0[dr_{t-1} dr_{t-1}] = \rho \frac{\sigma_\eta^2}{1 - \rho^2}$, which implies that ψ^{RPE} must satisfy

$$0 = \rho \frac{\sigma_\eta^2}{1 - \rho^2} + \frac{\psi^{RPE}}{(1 + \bar{r})} \rho \frac{\sigma_\eta^2}{1 - \rho^2} - \psi^{RPE} \frac{\sigma_\eta^2}{1 - \rho^2} + \mathcal{O}(\sigma_\eta^3).$$

which implies that

$$\psi^{RPE} = \frac{\rho}{1 - \frac{\rho}{(1 + \bar{r})}} + \mathcal{O}(\sigma_\eta)$$

is uniquely determined for small enough σ_η .

B.4 Decreasing Gain Convergence

Here we verify that the decreasing gain algorithm converges to the the RPE. In Figure 7 we plot the time paths of beliefs, ψ , of an agent learning with a decreasing gain learning algorithm, $g_t = 0.5t^{-0.65}$ for 1,000,000 periods. The black line represents the time path of beliefs when ψ is initialized at 0, while the blue line represents the time path of beliefs when ψ is initialized at the RPE values. As can be clearly seen both lines converge to the RPE value for beliefs (red line).

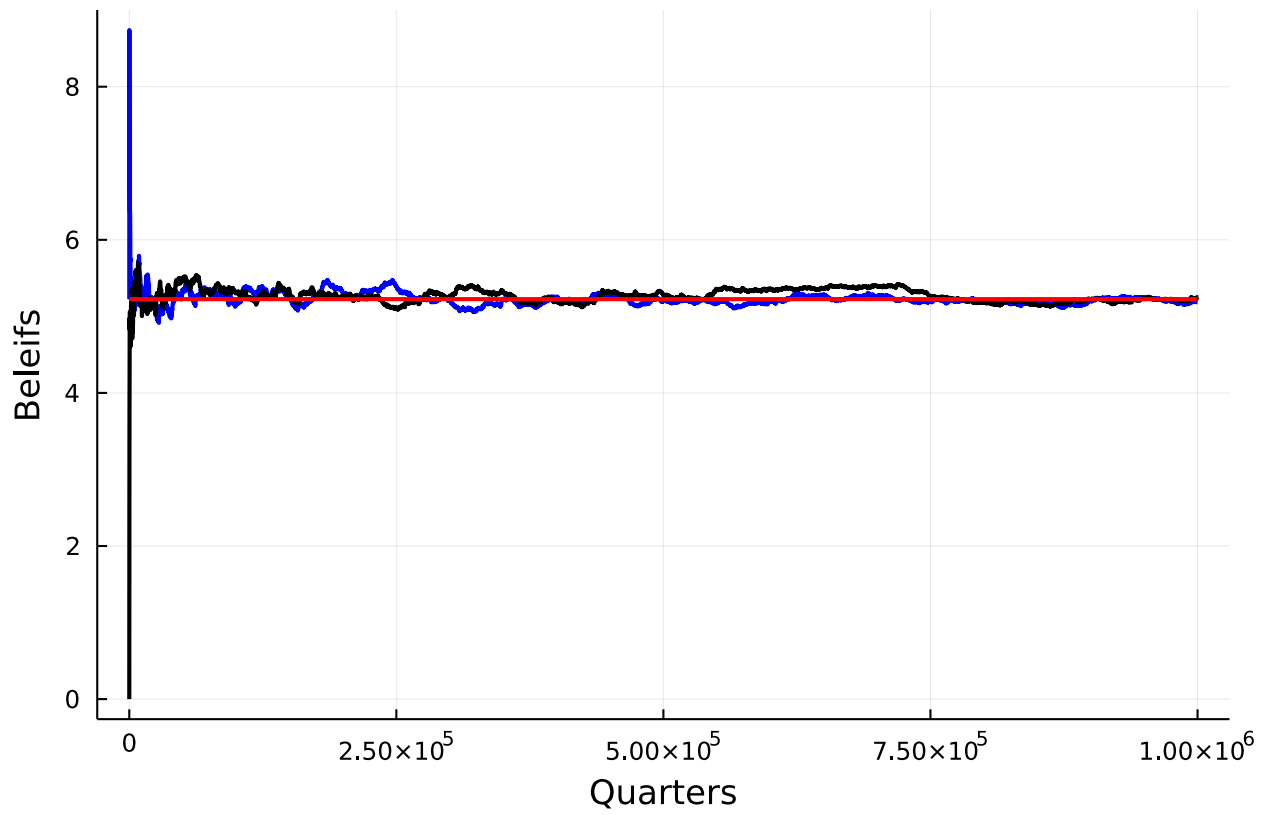


Figure 7: Time path of beliefs generated by an agent using a decreasing gain learning algorithm. The black line represents the time path of beliefs with initial beliefs are zero. The red line represents the time path of beliefs when beliefs are initialized at the RPE. The red line represents the RPE.

C Appendices for Section 5

Here we present additional numerical details of the algorithm used to simulate the locally rational economy. We will take as given the ability to solve for a Bewely-Aiyagari steady state, and instead focus on the steps necessary to simulate the behavior of the locally rational agents. We first describe the initialization process, which includes precomputing objects that will be required for later simulating. We then turn to the simulation itself, which is we detail as a step-by-step procedure for going from the beginning of time t to the beginning of time $t + 1$.

C.1 Initialization

Initialization involves three steps, as identified in Algorithm 1.

Algorithm 1 Initialization

- 1: Approximate SRE objects $\{\bar{\mu}, \bar{r}, \bar{w}, \bar{c}(a, \varepsilon), \bar{a}(a, \varepsilon), \bar{l}(a, \varepsilon), \bar{n}(a, \varepsilon)\}$
 - 2: Approximate LR policy rules $\hat{c}(a, \varepsilon, \phi, r, w)$, $\hat{l}(a, \varepsilon, \phi, r, w)$ and $\hat{a}(a, \varepsilon, \phi, r, w)$
 - 3: Construct initial state from $\bar{\mu}$ and initial ψ
-

Step 1 is a standard procedure, and so we will only report the general parameters used for approximating the objects. The individual policy functions $(\bar{c}, \bar{a}, \bar{l}, \bar{n})$ were approximated using cubic spline approximations along the asset dimension with $N_a = 150$ non-equally spaced gridpoints given by

$$a_i = \underline{a} + (\bar{a} - \underline{a}) \left(\frac{i-1}{N_a - 1} \right)^{1.7}.$$

We set lower bound on assets to $\bar{a} = 0$ while the upper bound is set to $\bar{a} = 600$. We use a finite state approximation to the AR(1) process using the Rouwenhsorst method with 11 grid points provided by the QuantEcon.jl package. For a given r and w the optimal individual decisions are found using the endogenous grid method of Carroll (2006) with a terminal tolerance of 10^{-5} in the consumption policies under the sup norm. Finally, the distribution is approximated using 5000 equally spaced points along the asset dimension using the algorithm of Young (2010).

Next, we turn to step 2, which pre-computes the LR policy rules necessary to simulate the LR model. Define the policy functions $\hat{c}(a, \varepsilon, \phi, r, w)$, $\hat{l}(a, \varepsilon, \phi, r, w)$ and $\hat{a}(a, \varepsilon, \phi, r, w)$ that solve the following system of equations

$$u_c(\hat{c}(a, \varepsilon, \phi, r, w), \hat{l}(a, \varepsilon, \phi, r, w)) \geq \beta \bar{\lambda}^e (\hat{a}(a, \varepsilon, \phi, r, w), \varepsilon) \exp(\phi) \quad (\text{C.17})$$

$$u_l(\hat{c}(a, \varepsilon, \phi, r, w), \hat{l}(a, \varepsilon, \phi, r, w)) = u_c(\hat{c}(a, \varepsilon, \phi, r, w), \hat{l}(a, \varepsilon, \phi, r, w)) w \quad (\text{C.18})$$

$$\hat{a}(a, \varepsilon, \phi, r, w) + \hat{c}(a, \varepsilon, \phi, r, w) = (1+r)a + w \cdot \varepsilon \cdot (1 - \hat{l}(a, \varepsilon, \phi, r, w)), \quad (\text{C.19})$$

where equation (C.17) holds with strict inequality only if $\hat{a}(a, \varepsilon, \phi, r, w) = \underline{a}$. Here ϕ represents the log deviation of the expected shadow price of savings from their SRE values. It is straightforward to show that $\hat{c}_t(a, \varepsilon, \psi) = \hat{c}(a, \varepsilon, \langle \psi, X_t \rangle, r_t, w_t)$ and similarly for \hat{l}_t and \hat{a}_t . Precomputing $(\hat{c}, \hat{l}, \hat{a})$ can therefore allow for faster simulation of the economy. To do this we maintain the same gridpoints along the a and ε dimensions used to solve for the SRE. We use 20^{th} order Chebyshev polynomials along the ϕ dimension and 10^{th} order Chebyshev polynomials along the r and w dimension, using the BasisMatrices package provided by QuantEcon.jl. We then solve equations (C.17)-(C.19) at each gridpoint to construct approximations to the policy rules.³⁰

Finally, in step 3 we construct the initial state for simulation. As presented in Section 3.2, the full aggregate state for the economy is $\xi_t \equiv (\mu_t, \theta_t, R_t, X_{t-1})$, where μ_t is the joint distribution over (a, ε, ψ) . To construction the initial conditions ξ_0 we set $\theta_0 = 1$ and set $X_{-1} = [1, 0, 0]$ to correspond to being initially at the SRE.³¹ We set R_0 to the second-moment matrix computed from simulating the REE of our economy.³² μ_t is approximated with $N = 100,000$ agents. As such, we draw N agents randomly from the distribution $\bar{\mu}$ and then endow all agents with some initial beliefs ψ_0 .

C.2 Simulation

We now present the algorithm used for simulating the dynamics of the LR economy by detailing the steps for going from ξ_t to ξ_{t+1} . Assume we enter with $\xi_t \equiv (\mu_t, \theta_t, R_t, X_{t-1})$, where μ_t is approximated by N agents with states $(a_{i,t}, \varepsilon_{i,t}, \psi_{i,t})$.

Algorithm 2 Time t Simulation

- 1: Compute $X_t = (1, \log(k_t/\bar{k}), \log(\theta_t))$ where $k_t = \frac{1}{N} \sum_i a_{i,t}$
 - 2: Solve for the TE prices (r_t, w_t)
 - 3: Update $R_{t+1} = R_t + g_t(X_{t-1}X_{t-1}^T - R_t)$
 - 4: Draw v_{t+1} and update $\theta_{t+1} = v_{t+1}\theta_t^\rho$
 - 5: Update μ_{t+1}
-

To solve for the temporary equilibrium we first construct the labor demand curve from the FONC of the representative firm: $w = \theta_t f_n(k_t, n_t^d(w))$. The calibration presented in Section 5 implies that labor demand is given by

$$n_t^d(w) = \left(\frac{\theta_t}{w} \right)^{1/\alpha} k_t.$$

³⁰Recall $\bar{\lambda}^\varepsilon(a', \varepsilon)$ can be easily constructed from the stationary policy rules $\bar{c}(a, \varepsilon)$.

³¹Recall $X_t = (1, \log(k_t/\bar{k}), \log(\theta_t))$.

³²One could also set R_0 from the ergodic second-moment matrix of the RPE. As we focus on the ergodic distribution of the learning economy our results are not sensitive to this choice.

At this labor demand, the interest rate that clears the asset market is given by $r_t(w) = \theta_t f_k(k_t, n_t^d(w)) - \delta$. We then construct aggregate labor supply from the individual policy rules $\hat{l}(a, \varepsilon, \phi, r, w)$ computed in the initialization. First, compute each agent's forecasted log deviation in the expected shadow prices from the SRE value: $\phi_{i,t} = \langle \psi_{i,t}, X_t \rangle$. Next, we aggregate the individual labor supply to compute aggregate labor supply as a function of the wage

$$n_t^s(w) = \frac{1}{N} \sum_i 1 - \hat{l}(a_{i,t}, \varepsilon_{i,t}, \phi_{i,t}, r_t(w), w).$$

We then use the Brent root-solving algorithm to find w_t that solves $n_t^s(w_t) = n_t^d(w_t)$, and set $r_t = r_t(w_t)$.

Finally, we turn to updating the distribution μ_{t+1} which consists of

$$\{a_{i,t+1}, \varepsilon_{i,t+1}, \psi_{i,t+1}\}_{i=1}^N.$$

For each agent i , their asset next period can be computed from the pre-stored policy rules

$$a_{i,t+1} = \hat{a}(a_{i,t}, \varepsilon_{i,t}, \phi_{i,t}, r_t, w_t).$$

The next period's idiosyncratic productivity, $\varepsilon_{i,t+1}$, is drawn randomly from the Rouwenhorst process conditional on $\varepsilon_{i,t}$. Finally, we can construct the agent's current shadow price of savings as

$$\lambda_{i,t} = (1 + r_t) u_c(\hat{c}(a_{i,t}, \varepsilon_{i,t}, \phi_{i,t}, r_t, w_t), \hat{l}(a_{i,t}, \varepsilon_{i,t}, \phi_{i,t}, r_t, w_t)),$$

and use it to update beliefs via the recursive learning algorithm (15):

$$\psi_{i,t+1} = \psi_{i,t} + g_t R_{t+1}^{-1} X_{t-1} (\log(\lambda_{i,t} / \bar{\lambda}(a_{i,t}, \varepsilon_{i,t})) - \langle \psi_{i,t}, X_{t-1} \rangle).$$

D Appendices for Section 6

In this section of the Appendix we explore two additional aspects of locally rational dynamics that are commonly addressed in analyses that involve agents modeled as adaptive learners. First, our formal results are local in nature, as are as most theoretical results from the learning literature; so, in the first subsection below, we consider the robustness of our findings to perturbations in beliefs by examining the size the basin of attraction associated with our RPE. Second, our results are predicated upon the assumption that all agents use a particular (and, in our view, particularly natural) PLM. In the second subsection below, we consider outcomes of simulations that allow for subpopulations of agents to use distinct PLMs, thus generating RPE with beliefs heterogeneity.

D.1 Stability Analysis

To examine the extent of our equilibrium's basin of attraction, we simulated a range of economies, initializing agents with the same beliefs chosen randomly from a region around the RPE. Agents'

beliefs are randomly chosen from $\psi \in [\psi_1^{RPE} - 0.05, \psi_1^{RPE} + 0.05] \times [\psi_2^{RPE} - 0.75, \psi_2^{RPE} - 0.75] \times [\psi_3^{RPE} - 1.25, \psi_3^{RPE} - 1.25]$, and they are assumed to use a decreasing gain algorithm. Figure 9 (which, due to spacing issues, follows a figure discussed later in this section) plots the time paths of the cross-sectional average of beliefs for each of 60 simulations. To give scale, the horizontal grey bars represent the support of the ergodic distribution of average beliefs under the constant gain algorithm with gain 0.035.

Initially, we see a wide discrepancy in mean beliefs across simulations, reflecting the effect of the randomly drawn initial conditions. Beliefs far from RPE are quickly tamed, as discrepancy quickly shrinks and agents learn to coordinate on the RPE; however, due to our use of decreasing gain, which puts very small weight on new information as the simulations evolve, the variance in beliefs remains visible even after 60,000 periods (of course, asymptotically, it will vanish). The main take-away: even for a wide range of initial beliefs, long run beliefs consistently converge to the RPE.

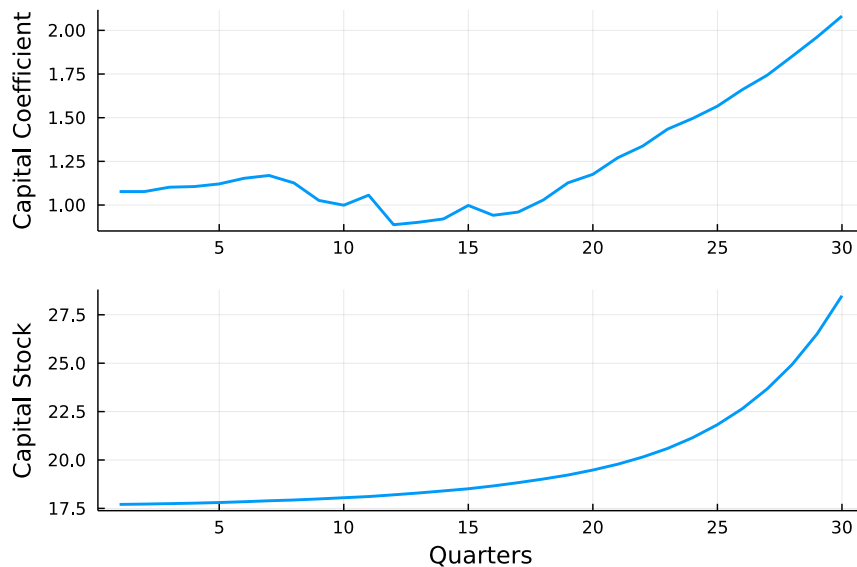


Figure 8: Unstable initial beliefs

As noted, stability results in the learning literature are local in nature: some beliefs may induce behavior that destabilizes the economy while at the same time pushing beliefs even farther from the basin of attract. Consider, for example, an economy in which the beliefs coefficient of capital is modestly above one. An increase in capital leads to an increase in the forecasted shadow-price of savings, thus inducing more savings and lower consumption. When all agents hold these beliefs their savings behavior increases aggregate capital while at the same time raising their realized shadow prices, thus perpetuating their beliefs.

To illustrate this behavior, we consider a simulation in which we initialize all agents with the same beliefs $\psi = [0., 1.07, 1.0]$, and then plot the time paths of the average coefficient on capital and the capital stock in Figure 8. We can see that both feature explosive paths.




Figure 9: Basin of attraction: stability of the RPE is robust to varied initial beliefs

D.2 Alternative Forecasting Rules

In this section we document the behavior of our learning agents with alternative forecasting rules. We begin by studying the stability and dynamics of the economy when populated by different types of agents with different forecasting rules. We then document how the economy behaves when agents are endowed with a forecasting rule which includes interactions with the individual states.

D.2.1 Heterogeneous Forecasting Rules

The model presented in the main text assumed all agents uses the same PLM; further, under this assumption we established that, for small enough aggregate shocks, agents held common beliefs in the RPE. It is straightforward to extend the environment to allow for multiple types of forecasting rules and hence heterogeneous beliefs in the RPE. In subsection we specify four agent types, distinguished by the function form of the PLM they use, and analyze the associated RPE. The following table identifies our agent types.

Table 3: Agent types

Agent type	PLM
Type 0	$\log(\hat{\lambda}_t/\bar{\lambda}_t) = \psi_0 + \psi_k \log(k_{t-1}/\bar{k}) + \psi_\theta \log \theta_{t-1}$
Type 1	$\log(\hat{\lambda}_t/\bar{\lambda}_t) = \psi_0$
Type 2	$\log(\hat{\lambda}_t/\bar{\lambda}_t) = \psi_0 + \psi_k \log(k_{t-1}/\bar{k})$
Type 3	$\log(\hat{\lambda}_t/\bar{\lambda}_t) = \psi_0 + \psi_\theta \log \theta_{t-1}$

The benchmark model obtains when all the agents in the economy are type 0. Here we consider the behavior economies populated with two agent types, each comprising 50% of the population. We begin by studying the long run behavior under a decreasing gain learning algorithm. For each of the six economies³³ and simulate the dynamics for 80,000, periods starting with all agents having SRE beliefs, i.e. all relevant belief parameters set to zero. We plot the type paths of average beliefs for the six economies in Figure 10. As can be readily observed from the figure, the average beliefs of all agents settle down and eventually converge to their RPE levels. However, as the agents have different forecasting rules, their long run belief coefficients can end up being different from each other; also, RPE beliefs associated with a given agent type depend on the other type of agent populating the economy.

³³We consider economies populated with agents of all possible pairings of types (0,1), (0,2), (0,3), (1,2), (1,3), and (2,3)

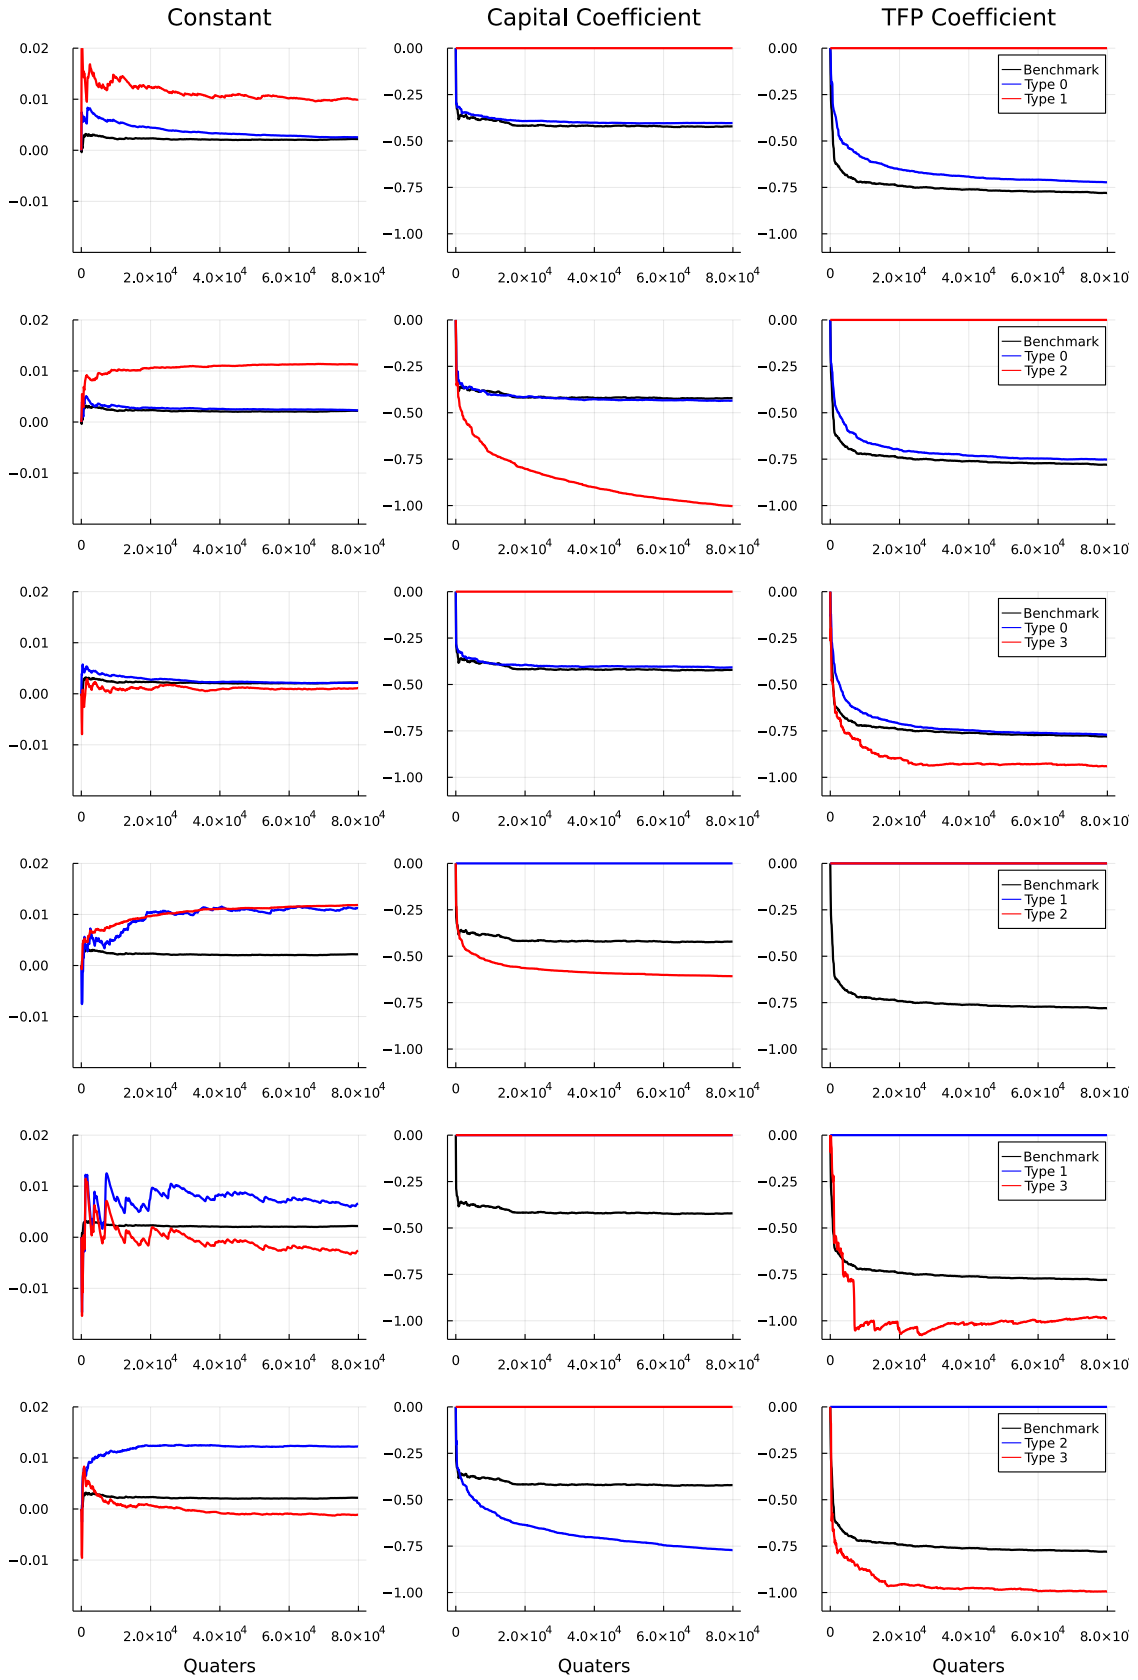


Figure 10: Time paths of average beliefs for the economies (from top to bottom) populated with types (0, 1), (0, 2), (0, 3), (1, 2), (1, 3), and (2, 3).

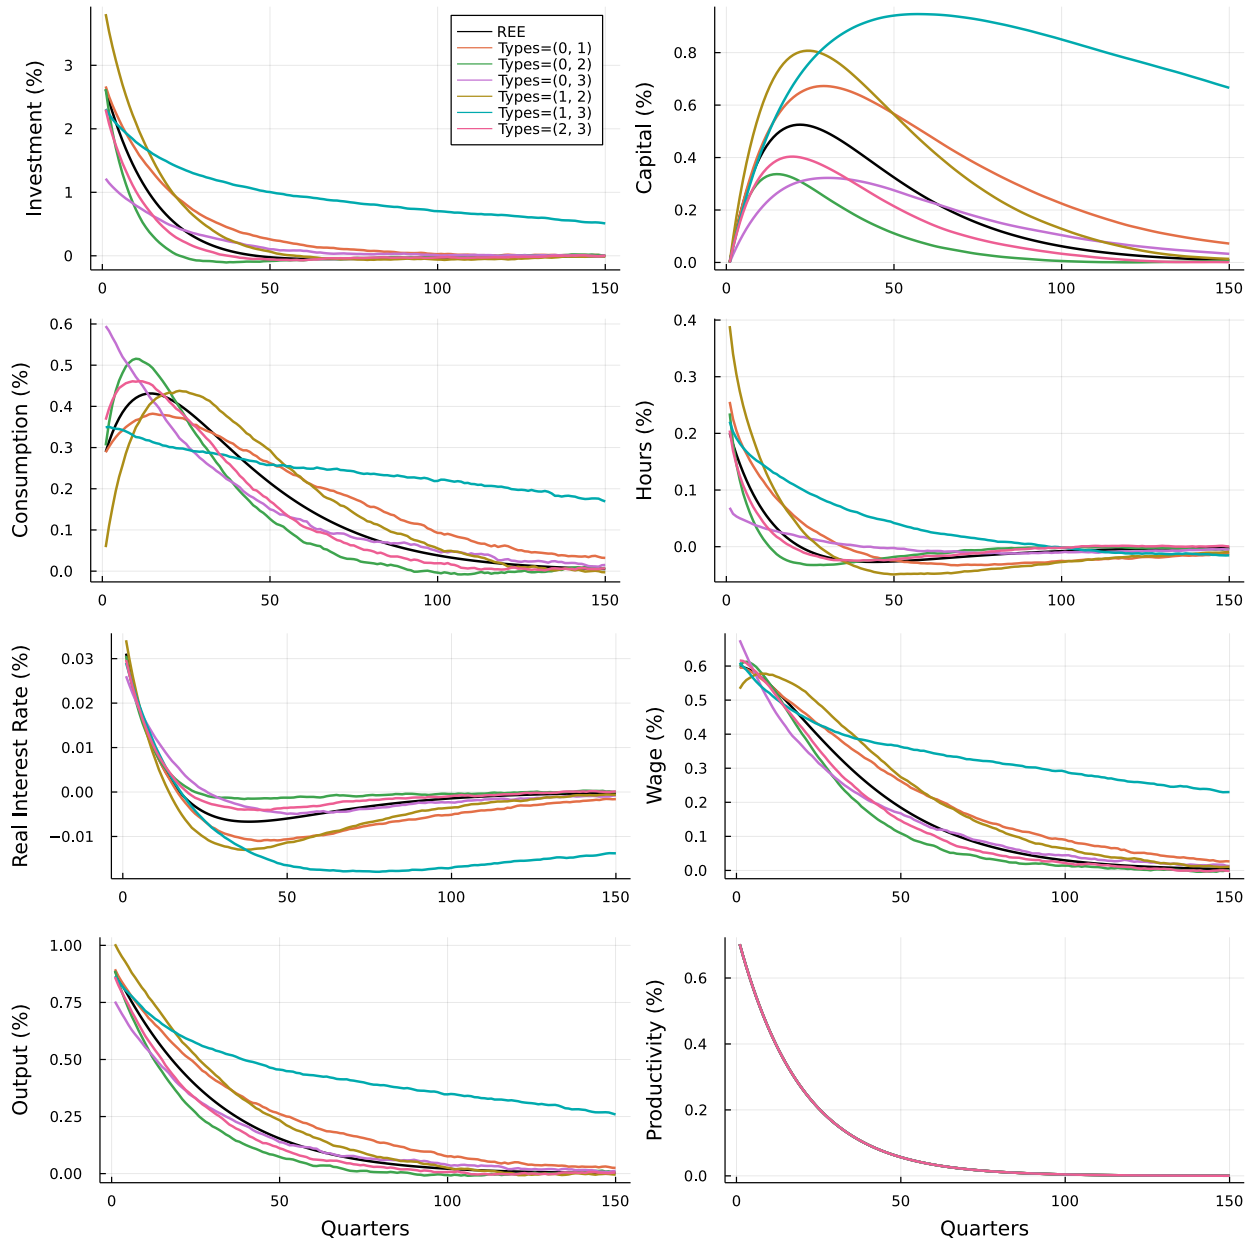


Figure 11: Impulse responses to a one-standard deviation increase in productivity. Black line refers to the linearized rational expectations equilibrium. The other lines represent the mean impulse responses for the six economies populated by agents with different types of forecasting rules. They are assumed to learn with a constant gain learning algorithm with gain 0.001.

The RPEs identified in Figure 11 are necessarily distinct: after all, no two economies have nested PLMs; however, the distinction does not only reflect agents' use of different forecast models: aggregate equilibrium dynamics are different as well.

To illustrate this point, in Figure 11 we plot the impulse responses of various aggregates to a one-standard deviation productivity shock. The black line represents the same REE line as in Figure 4 of the main text. The remaining lines represent the impulse responses of the six locally rational economies populated by heterogeneous types. The agents are assumed to be learning with a constant gain learning algorithm with gain $g = 0.001$.

The impulse responses are constructed in the same manner as those in the main text, i.e. by repeatedly drawing an initial distribution of assets, productivities and beliefs from the ergodic distribution generated by long simulation. We then record the impulse responses to a one-standard-deviation productivity shock from those initial starting points and plot the median response of all variables as a percentage deviation from the path which would prevail in absence of a shock.

As can be readily observed from the figure, the forecasting rules used can have a significant impact on the aggregate responses. It is of particular interest that the introduction of heterogeneous PLMs can serve to amplify the propagation of shocks, a feature of the data that benchmark rational models struggle to replicate. These observations suggest that expectations heterogeneity in the form of forecast model selection in HA environments may be a fruitful area of future research.

D.2.2 Expanded Learning Rule

We modify the agent's expectations function to allow for interactions with idiosyncratic states by modifying the forecasting rule, equation (13), to be

$$\hat{\lambda}_t^e(a', \varepsilon, \psi) = \bar{\lambda}^e(a, \varepsilon) \cdot \exp(\langle \psi_t, X(X_t, a, \varepsilon) \rangle). \quad (\text{D.20})$$

$X(X, a, \varepsilon)$ is a function to allow for arbitrary interactions of the aggregate observable, X , and the individual states, x . Based on our analysis in Section 6 we will study the behavior of learning models when

$$X(X, a, \varepsilon) = \begin{pmatrix} 1 \\ \log(k/\bar{k}) \\ \log(\theta) \\ \log(k/\bar{k})(a - \bar{a}) \\ \log(k/\bar{k})(\log(\varepsilon) - \overline{\log(\varepsilon)}) \\ \log(\theta)(a - \bar{a}) \\ \log(\theta)(\log(\varepsilon) - \overline{\log(\varepsilon)}) \end{pmatrix}$$

where \bar{a} and $\overline{\log(\varepsilon)}$ are the average levels of wealth and log productivity in the stationary recursive competitive equilibrium.

In addition to changing the agent's forecasting rule, the agent's learning behavior must be adjusted slightly as the second moment matrix R will differ across agents. The recursive formulation of the updating rule, equation (15), is adjusted to include the individual states x_t in a similar manner:

$$\begin{aligned} \hat{R}_t(a, \varepsilon, \psi, R) &= R + g_t \cdot (X(X_{t-1}, a, \varepsilon) \otimes X(X_{t-1}, a, \varepsilon) - R) \\ \hat{\psi}_t(a, \varepsilon, \psi, R) &= \psi + g_t \cdot \hat{R}_t(a, \varepsilon, \psi, R)^{-1} X(X_{t-1}, a, \varepsilon) \left(\log \left(\frac{\hat{\lambda}_t(a, \varepsilon, \psi)}{\bar{\lambda}_t(a, \varepsilon)} \right) - \langle \psi_t, X(X_{t-1}, a, \varepsilon) \rangle \right). \end{aligned} \quad (\text{D.21})$$

As each agent will have their own second moment matrix based on their unique experiences, one of the states of the model will be μ_t , i.e. the contemporaneous distribution of agent-states (a, ε) , beliefs ψ , and second moment matrices R .

We explore the behavior of this model through simulation. Numerically, there are two changes required relative to the procedure described in Section 5. First, the forecasting and learning rules are adjusted according to equations (D.20) and (D.21). This requires tracking the individual specific second moment matrix along with individual beliefs and states. Second, the persistence of the individual states can lead to a collinearity of the regressors which results in unstable paths of beliefs not present in the more parsimonious learning model. This is particularly problematic for higher gains since agents will put most weight on recent periods when idiosyncratic states will be most similar. We resolve this problem by employing a projection facility when beliefs become too extreme.³⁴ As such, we will only report results for models with gains of 0.001, 0.005, and 0.01 when the projection facility is rarely implemented.³⁵

³⁴Agent's beliefs are projected back to the RPE

³⁵For the gain of 0.01 the projection facility is active for 0.06% of agents every period.

Table 4 reports the business cycle statistics for the standard model constructed following the same procedures as in Section 6. We see that at the lowest gain the moments are nearly identical to the rational expectations equilibrium and changing the gain has little effect on the moments. These results are mirrored in the impulse response plotted in Figure 12 which are almost exactly in line with the rational expectations paths for all of the gains considered.

	Data	RE	Expanded Forecasting Rule		
			$g = 0.001$	$g = 0.005$	$g = 0.01$
$\frac{\text{std}(C)}{\text{std}(Y)}$	0.50	0.36	0.37	0.33	0.35
$\frac{\text{std}(I)}{\text{std}(Y)}$	2.73	2.91	2.93	3.11	3.04

Table 4: Business Cycle Statistics for Expanded Model

As anticipated, the expanded learning rule allows the model dynamics to converge to those that closely match the rational expectations equilibrium. We favor the parsimonious learning rule for its simplicity and tractability. The parsimonious rule is easier for agents to implement, generates stabler paths of beliefs, and produces results that better fit the stylized facts observed in the data.

D.3 Rational Expectations Beliefs

Here we construct the beliefs of agents in the rational expectations equilibrium. To generate these beliefs we simulate the linearized rational expectations equilibrium as described in 5.3 and record both the path of observables X_t . For each point in the state space (a, ε) , we compute the log deviation of the expected shadow price of wealth from its steady state counterpart:

$$\hat{\lambda}_t^{RE}(a, \varepsilon) \equiv \mathbb{E}_t \left[\log \left(\frac{\lambda_{t+1}}{\bar{\lambda}_{t+1}} \right) \middle| a, \varepsilon \right].$$

We then, for each (a, ε) , project $\hat{\lambda}_t^{RE}$ on X_t to construct the beliefs, $\psi^{RE}(a, \varepsilon)$, that rationalize the rational expectations equilibrium.

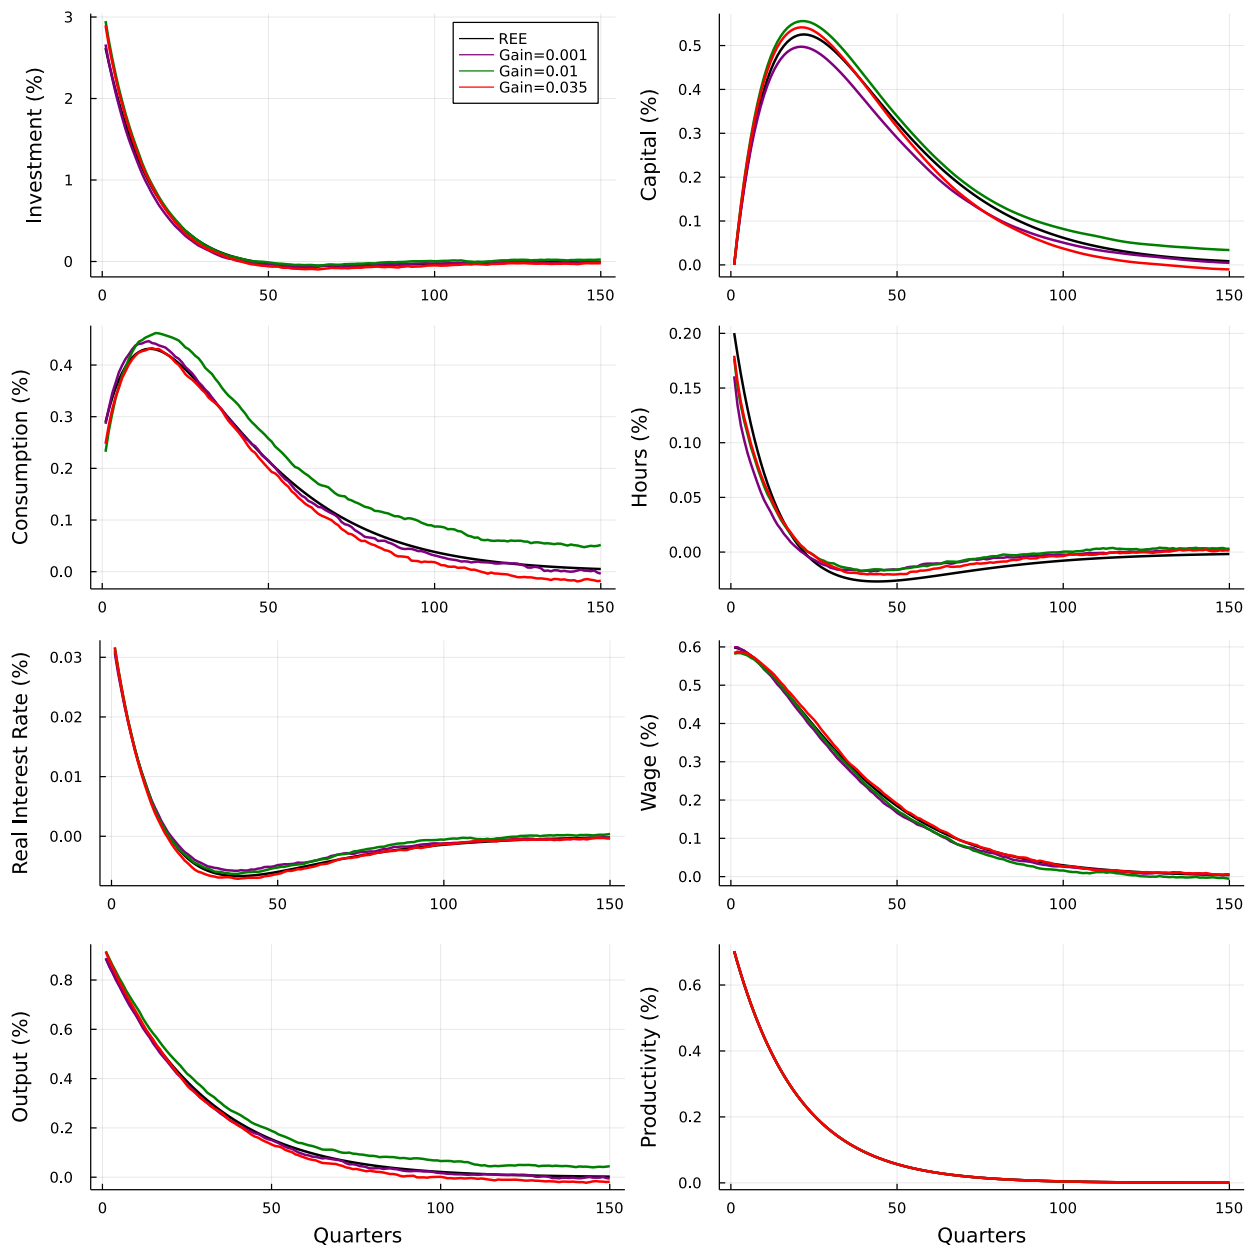


Figure 12: Impulse responses to a one-standard deviation increase in productivity. Black line refers to the rational expectations equilibrium. The blue, green and red lines refer to the mean response of the locally rational dynamics with the expanded learning rule and gains equal to 0.001, 0.005, and 0.01 respectively.

D.4 Belief Regression Table

To construct the best linear predictors used in Section 6.2 we run the weighted regression

$$\psi_i^j = \alpha_0 + \alpha_1(a_i - \bar{a}) + \alpha_2(\log(\varepsilon_i) - \overline{\log(\varepsilon)}) + \mu_i, \quad (\text{D.22})$$

which regresses beliefs of agents on their state variables. The term ψ_i^j represents the j^{th} component of the belief vector for agent i with wealth a_i and productivity ε_i . The regression is weighted by the fraction of agents with states $(a_i, \varepsilon_i, \psi_i)$ in the ergodic distribution constructed through simulation. The estimated coefficients from (D.22) are reported in table 5 below. As can be seen in the coefficients, increasing the gain of the constant gain learning algorithm brings beliefs more in line with rational expectations both in slope and in level.

	Capital Coefficient				TFP Coefficient			
	REE	$g = 0.001$	$g = 0.01$	$g = 0.035$	REE	$g = 0.001$	$g = 0.01$	$g = 0.035$
Intercept	-0.4303	-0.4204	-0.3735	-0.3537	-0.5197	-0.7403	-0.6802	-0.6392
Wealth	0.0015	-0.0000	0.0001	0.0002	0.0015	0.0001	0.0005	0.0012
LogWage	-0.0394	-0.0006	-0.0025	0.0054	0.1593	0.0041	0.0273	0.0336
R^2	0.7315	0.0004	0.0024	0.0032	0.5246	0.0171	0.0809	0.0663

Table 5: Coefficients for the regression of beliefs on demeaned wealth and log productivity.