Local rationality*

David Evans
University of Oregon

Jungang Li
University of Oregon

Bruce McGough
University of Oregon

December 2, 2021

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 centered near the economy's restricted perceptions equilibrium. To examine the dynamics implied by our behavioral assumptions, we conduct a calibration exercise. As is well-known, in a calibrated representative-agent RBC model the volatility of consumption is too low relative to the data. Extending the model by either incorporating adaptive learning or heterogeneous agents fails to alter this conclusion. We find that 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

*We thank George Evans and participants at the University of Connecticut Macroeconomics Seminar and the University of Oregon Macro Group for comments.

Bewley models 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 *heterogenous agent* (HA) models had arrived. HA models now figure prominently in the standard cannon 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 also both quantitatively way 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 the-

[1] Stadler (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 inability of the corresponding RA models to match certain aggregate moments.

[2] We know of no 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.

oretical 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) 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 aggregate states evolve and how their 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 this 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 which they update using a *constant gain learning algorithm*. Agents are not required to understand the evolution of the economy's states and they are not required to solve complex, nonlinear dynamic programs.

In a numerical experiment, we evaluate the behavior of a calibrated heterogenous agent economy populated by locally rational agents. To isolate the interaction of bounded rationality and heterogeous agents, we compare outcomes to those of a representative agent model populated with learning agents and a hetereogenous agent model populated by 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 ${}^{\mathrm{std(C)}}/{}_{\mathrm{std(Y)}} = 0.50$ 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 learning[4] the ratio varied from 0.32 to 0.34 while in the HA model 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 gain an understanding of why the interaction of boundedly rational agents and het-

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

[4] We 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.

erogeneity is so potent we extend the concept of *restricted perceptions equilibrium* (RPE) to our heterogeneous agent framework. Start by assuming that all agents have the same beliefs ψ. If these beliefs are held fixed over time then the economy's state vector will converge weakly to an ergodic distribution.[5] We may use this ergodic distribution to compute the linear forecast rule that is optimal in the sense that it minimizes the expected forecast error when averaging over both aggregate and idiosyncratic states. Letting $T(\psi)$ be the coefficients of the optimal forecast rule, we define a RPE as a fixed point: $T(\psi^*) = \psi^*$. In an RPE each agent is using an optimal forecast model in the following restricted sense: if an agent is asked to choose, once and for all, a beliefs vector ψ, and all other agents hold beliefs ψ^*, then the agent will choose $\psi = \psi^*$. RPEs often capture the long run behavior of learning algorithms with small constant gains. We verify, numerically, that this holds in our context: a RPE exists and the locally rational model converges an ergodic distribution centered around the RPE for sufficiently small values of the gain parameter.

A key feature of the RPE is that all agents have the same beliefs when forecasting their future shadow price of savings to make decisions. 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 a TFP shock differs across agents since they differ in both their exposure to and ability to smooth aggregate shocks. As a result, in the locally rational model with a constant gain learning algorithm, rich agents remember what it was like to be poor during a recession and adjust their forecasts appropriately. This results in rich agents, in a recession, being overly pessimistic about the future[6] causing consumption to fall more than in the rational expectations calibration.[7] In the extreme limit when the gain is very small, and the locally rational economy is approximately at the RPE, agents place too much weight on 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.

[5]The distribution is over possible values of the state, which itself includes a distribution as a component.

[6]Pessimism is reflected in higher forecasts of future shadow price of savings relative to fully rational agents.

[7]While the reverse is true that poorer agents are too optimistic, savings and consumption decisions are dominated by that asset rich agents.

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). Guisto 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 5, 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).[8]

[8]See 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.

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.[9] 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 is 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 includes a detailed discussion of restricted perceptions equilibrium. Section 4 provides the calibration details and the methods used for our numerical work. Section 5 presents our computational evidence for existence and stability of the model's RPE, and discusses the results of our calibration exercise. Section 6 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

[9] See 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.

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 more detail than is common.[10]

2.1 The household problem

The household's decision problem is recursive, and under rationally 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 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.[11] 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

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

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

$c_t(a,\varepsilon)$, $l_t(a,\varepsilon)$ and $a_t(a,\varepsilon)$ to satisfy

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

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

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

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

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

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 idioysncratic 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. 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 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.[12] 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 q_t, hires effective labor n_t at real wage w_t, and produces output under perfect competition using CRTS technology $\theta f(k,n)$.

[12]See the Appendix for a formal description of transition dynamics for μ_t.

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 $\{v_t\}$ iid having log-normal distribution ν. There are no capital installation costs. Profit maximizing behavior by the firm implies factors earn their marginal products:

$$\begin{aligned} w_t &= \theta_t f_n(k_t, n_t) \\ q_t &= \theta_t f_k(k_t, n_t) = r_t + \delta, \end{aligned} \tag{6}$$

where δ is the capital depreciation rate.

2.3 Dynamic recursive equilibrium

We define a *dynamic recursive equilibrium* (DRE) 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)

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

  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 $\theta_{t+1} = \nu_{t+1} \theta_t^{\rho}$.

Observe that, given any initial aggregate state (μ_0, θ_0), a DRE, 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 of the definition of the DRE 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

$$u_c(c_t, l_t) \geq \beta \lambda_t^e$$

$$u_l(c_t, l_t) = u_c(c_t, l_t) w_t$$

$$a_t = (1 + r_t) a + w_t \cdot (1 - l_t) - c_t$$

  given forecasts λ_t^e.
- *Agent rationality:* $\lambda_t^e = E_t\left((1 + r_{t+1}) u_c(c_{t+1}, l_{t+1})\right)$.
- *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 $\theta = v = 1$ 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, and noting that, since prices are constant, an agent's behavior depends only on their state (a, ε), we define a *stationary recursive equilibrium* (SRE) as a tuple $\left(\bar{c}, \bar{l}, \bar{a}, \bar{\lambda}^e, \bar{r}, \bar{w}, \bar{\mu}\right)$ 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),

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

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

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

2.6 Looking ahead

To foreshadow what's to come, observe that equations (2), (3) and (5) can be usefully re-interpreted to allow for the inclusion of possibly non-rational forecasts of the shadow price. Using hats to identify boundedly rational decision rules, let $\hat{\lambda}_t^e(a',\varepsilon,\psi)$ be a possibly non-rational agent's forecast of tomorrow's shadow price of claims conditional on their savings choice a', their current labor productivity ε, and finally on some form of beliefs ψ about how today's data inform forecasts of tomorrow's shadow price. For example, ψ could encode the objective conditional distributions of all relevant variables so that $\hat{\lambda}^e$ could align with rational expectations; or, ψ could represent a simple linear forecast model with parameters that are updated over time as new data become available. The shadow-price forecasts $\hat{\lambda}^e$, coupled with the Euler condition (5), can be combined with (2) and (3) to form a system of relations characterizing an agent's contemporaneous decision schedules in terms of prices, observable states, and beliefs. Period t outcomes are then realized via temporary equilibrium.

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 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 smarter than, nor much stupider 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 their beliefs ψ. 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:

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

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

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

Importantly, equations (9) – (11) are taken as *behavioral primitives:* they are imposed assumptions on the behavior the 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 will focus on learning how forecast in the presence of aggregate shocks and, hence, 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, doing so provides a cleaner 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 a given agent's period t beliefs are taken to be a vector $\psi \in \mathbb{R}^n$. This vector is interpreted as a linear functional on the space of aggregate observables $X_t \in \mathbb{R}^n$, and the agent is taken to form expectations as

$$\hat{\lambda}_t^e(a', \varepsilon, \psi) = \bar{\lambda}^e\left(a', \varepsilon\right) \cdot \exp\left(\langle \psi, X_t \rangle\right), \tag{12}$$

where $\bar{\lambda}^e$ is as defined in (8).[13] In this way, the agent's shadow-price forecast is their stationary forecast $\bar{\lambda}^e$ scaled to accommodate aggregate conditions; and the scaling coefficient measures the (exponentiated) action of the linear functional ψ on the observables X.

Equation (12) 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, and indeed it is straightforward to demonstrate that, in a stationary environment with only idiosyncratic risk, if agents are provided forecasting models for $\bar{\lambda}$ then the economy will converge to the associated RPE; and, allowing these forecasting models to depend on higher-order terms will result in an RPE that arbitrarily well approximates the REE. Thus we are, in effect, assuming that agents have *already learned* how to behave optimally in the absence of aggregate uncertainty. 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 γ_t – and may be taken as decreasing or constant over time. We assume that agents use a constant-gain learning (CGL), with gain γ, to implement their estimation procedure.

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

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

[13] An 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.

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

$$\hat{\psi}_t(a,\varepsilon,\psi) = \psi + \gamma \cdot R_{t+1}^{-1} X_{t-1} \left(\log \left(\hat{\lambda}_t(a,\varepsilon,\psi) / \bar{\lambda}(a,\varepsilon) \right) - \langle \psi, X_{t-1} \rangle \right), \tag{13}$$

where

$$R_{t+1} = R_t + \gamma \cdot (X_{t-1} \otimes X_{t-1} - R_t).$$

Note that the term R_t^{-1} depends only on aggregates and so may be taken as common across agents.

Our reliance on CGL is an important feature of local rationality and warrants further comment. Setting $\gamma_t = 1/t$ 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 $\lim_{t\to\infty} \gamma_t \to 0$ at a suitable rate like t^{-1}. In applied work it is common to assumed the gain is a (small) constant: $\gamma_t = \gamma \in (0,1)$. CGL algorithms discount older data at geometric rate $1-\gamma$, and, in stable systems, result in weak convergence to a distribution centered near the RPE or REE.

Several reasons for using constant gain algorithms have been advanced in the literature: see Evans, Evans, and McGough (2021b) for a discussion. Our adoption of 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 date 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

$$k_t = \int a \cdot \mu_t(da, d\varepsilon, d\psi) \text{ and } n_t = \int \left(1 - \hat{l}_t(a, \varepsilon, \psi)\right) \mu_t(da, d\varepsilon, d\psi), \tag{14}$$

and θ_t is the realized TFP shock. The n_t in (14) 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 system on non-linear equations.[14] 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 *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 (14).
3. $R_{t+1} = R_t + \gamma \cdot (X_{t-1} \otimes X_{t-1} - R_t)$
4. $\theta_{t+1} = v_{t+1} \theta_t^{\rho}$
5. μ_{t+1} evolves consistent with $\hat{a}_t$ and $\hat{\psi}_{t+1}$

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 very difficult to compute. Under local rationality, beliefs vary across agents and are updated recursively: less parsimonious but more computationally tractable.

A final observation: just as in the rational case, given any 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 $\{r_t, w_t\}$.

[14] See section B.1 of the appendix for a more detailed formulation.

3.3 Optimal homogeneous beliefs

A key concept to understanding learning behavior in RA models is the notion of a *restricted perceptions equilibrium*. The standard notion of an RPE is a fixed point. Agents choose a forecasting model from a pre-specified class and hold that model fixed over time. An RPE is the model that minimizes the long-run expected forecast error conditional on all agents using that same forecasting model. As noted in the introduction, RPE that are associated with a model's determinate steady state are commonly found to be stable under adaptive learning.[15] In this section we develop the notion of an RPE within the context of our HA model.

In line with the RA case, we assume agents use common beliefs ψ that are held constant over time.[16] Under these beliefs, the distribution of agent states μ_t, the lag of observed aggregates X_{t-1}, and realized shadow prices $\hat{\lambda}_t(a,\varepsilon)$, will converge weakly to an ergodic distribution $\mathscr{M}(\psi)$. We may then ask for the value of ψ that minimizes the expected forecast error, where the expectation is taken over $\mathscr{M}(\psi)$. This value may be computed via projections. Define the operator $T : \mathbb{R}^n \to \mathbb{R}^n$ implicitly as follows:

$$\int \left(\int \log \left({}^{\hat{\lambda}(a,\varepsilon)}/_{\bar{\lambda}(a,\varepsilon)} \right) \mu \left(da, d\varepsilon \right) - \left\langle T \left(\psi \right), X \right\rangle \right) \cdot X \cdot \mathscr{M} \left(\psi \right) \left(d\mu, dX, d\hat{\lambda} \right) = 0. \tag{15}$$

A fixed point ψ^* of this T-map identifies beliefs that are optimal in the sense that no other set of beliefs would lower an agent's unconditional forecast error. We refer to this fixed point as a restricted perceptions equilibrium.

Observe that a fixed point of the T-map is an equilibrium object: the beliefs ψ^* are unconditionally optimal for a given agent only if (almost) all other agents hold them. And while 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, agents' beliefs (locally) converge almost surely to a fixed point of T-map, and that for small constant gains agents' beliefs converge weakly to an ergodic distribution with mean very near a fixed point of the T-map.

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.

[15] See Evans, Evans, and McGough (2021a) for an examination in a non-linear RA models.

[16] The notion could be generalized to allow for heterogenous beliefs held constant over time; however, numerical analysis suggests that even if heterogenous-beliefs RPE exist they are not stable under adaptive learning.

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 (13), just as in the HA case.

4 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 as well as the rational expectations equilibrium.

4.1 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

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

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

We begin by specifying the parameters common to both the rational expectations and boundedly rational model. We assume that the length of 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 $X = \left(1, \log\left(k/\bar{k}\right), \log\left(\theta\right)\right)$. 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 $\gamma = 0.035$ to match our preferred moment, the ratio of consumption to output volatility. Noting that the gain discounts past data at rate $1-\gamma$, 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 γ is consistent with those used in the literature for calibration exercises and applied analysis.[17]

4.2 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).

[17]For 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 $\gamma = 0.05$; 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.

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 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 ratio[18] 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 $z_{\theta,t}\sigma_\nu$ is the response to a one standard deviation shock. The time series of z_t generated by a sequence of shocks ν_t is then constructed by aggregating the effect of all past shocks

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

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[19] and update the aggregate state based on the decision rules $\hat{a}_t$ and $\hat{\psi}_t$.

5 Results

In this section, we study the behavior of an 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 low gain. Turning to simulations, we study the

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

[19]See appendix B.2 for details

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.

5.1 Existence and Stability of Restricted Perceptions Equilibrium

To verify the existence of a restricted perceptions equilibrium we find the fixed point of the finite sample analogue of the T-map (15), 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 $\gamma = 0$, 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 $\psi \to \hat{T}(\psi)$ 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.[20] We numerically verify the existence of an RPE by finding the fixed point of $\hat{T}(\psi)$ when $N = 100{,}000$ and $S = 1{,}000$.[21]

To verify the stability of the RPE under learning we simulate the dynamics of the locally rational economy from two different initial conditions. 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 $\psi = 0$. We plot the path of the average beliefs, across the distribution of agents, over the simulation in figure 1 for a gain $\gamma = 0.001$. The black line represents the path average beliefs initialized at the RPE and the blue line represents the path initialized at beliefs consistent with the SRE, i.e. $\psi = 0$. The black line in figure 1 shows that over the course of a long simulation of 20,000 periods the average beliefs of agents remain consistently around the RPE values. The blue line illustrates that the basis of attraction of the RPE is large with the average beliefs converging to the RPE values by the end of the simulation even if beliefs are initialized far away from the RPE. We should note that the slow convergence of beliefs to the RPE is indicative of the small gain used. We chose this gain to emphasize the local stability of the RPE. Higher gains will have faster convergence but, as we will emphasize in future sections, imply convergence to different long run values due to the non-linearities of the model.

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

[21] Increasing both N and S does not appreciably change the value of the fixed point.

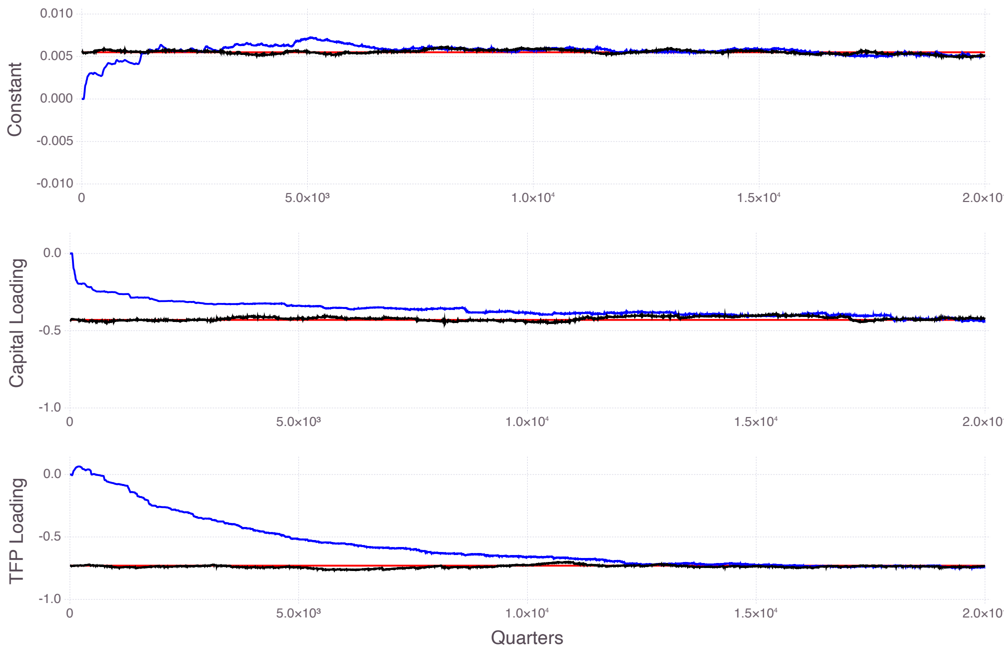

Figure 1: Time path of average beliefs when agents are initialized with the RPE beliefs (black) and $\psi = \mathbf{0}$ (blue). The RPE beliefs are represented with a solid red line.

5.2 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.[22] The same procedure is applied to the U.S. data which runs 240 quarters from 1948Q1 to 2007Q4.

		Representative Agent				Heterogeneous Agent			
	Data	RE	$\gamma = 0.001$	$\gamma = 0.01$	$\gamma = 0.035$	RE	$\gamma = 0.001$	$\gamma = 0.01$	$\gamma = 0.035$
$\frac{\text{std}(C)}{\text{std}(Y)}$	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 1: Business Cycle Statistics

Table 1 reports standard deviations for consumption and investment relative to the stan-

[22]We construct 5000 simulations for each model and average over all simulations.

dard 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 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.

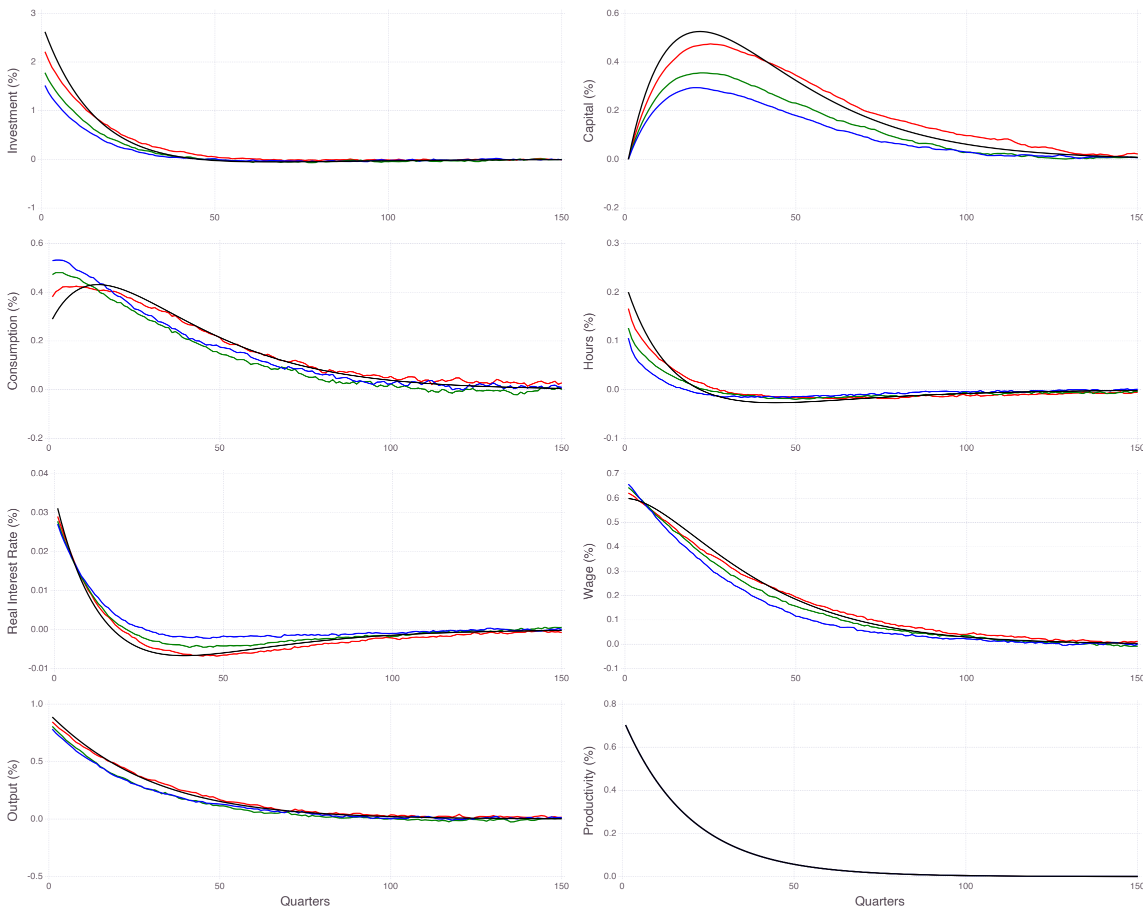

Figure 2: Impulse responses to a one-standard deviation increase in productivity. Black line refers to the linearized rational expectations equilibrium. The blue, green and red lines refer to the median response of the locally rational dynamics with gains equal to 0.001, 0.01, and 0.035 respectively.

Focusing on the last 4 columns of table 1, 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 2. The black line in figure 2 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 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.

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 when the gain is the smallest (0.001) as seen in the blue line which has very muted responses of investment and capital accumulation but amplified responses of consumption. As the gain increases, the response of the locally rational economy converges towards rational expectations with the closest being the red line (gain of 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 Section C of the appendix, 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 gain an understanding of this distribution of beliefs we construct a simple set of summary statistics by running the weighted regression

$$\psi_i^j = \alpha_0 + \alpha_1(a_i - \bar{a}) + \alpha_2(\log(\varepsilon_i) - \overline{\log(\varepsilon)}) + \mu_i, \tag{16}$$

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. Table 2 reports the coefficients of this regression for both the beliefs loading on capital and TFP as well as the corresponding R^2. The terms $\bar{a}$ and $\overline{\log(\varepsilon)}$ represent the average level of wealth and log productivity in the corresponding ergodic distribution and, thus, the intercept terms should be interpreted as the long-run average beliefs of the economy.

There are several pieces of information to be gleaned from examining Table 2. The first is that in all models the average loading on capital and TFP are both negative. The

	Capital Loading				TFP Loading			
	RE	$\gamma = 0.001$	$\gamma = 0.01$	$\gamma = 0.035$	RE	$\gamma = 0.001$	$\gamma = 0.01$	$\gamma = 0.035$
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 2: Coefficients for the regression of beliefs on demeaned wealth and log productivity.

second is that there is a strong correlation of beliefs with individual characteristics under the rational expectations equilibrium capturing a differential exposure of these agents to aggregate shocks. While this relationship is present in the learning models with higher gain, it is absent in the models with the lowest gain. This last fact reflects our previous result that the low gain economy converges to a distribution centered near the RPE, which features all agents having identical beliefs. This convergence is necessary to understand the behavior of the learning economy.

To interpret this heterogeneity in beliefs, consider an agent (labeled i) who, at time t, has wealth $a_{i,t}$, productivity $\varepsilon_{i,t}$, and beliefs $\psi_{i,t}$. Recall that beliefs measure the relative deviation of the expectation of their shadow price of wealth, $\hat{\lambda}^e_{i,t}$, from their steady-state counterpart $\bar{\lambda}^e_i$.[23] For the current information set, X_t, these forecasts are given by

$$\log\left(\frac{\hat{\lambda}^e_{i,t}}{\bar{\lambda}^e_i}\right) = \langle \psi_i, X_t \rangle .$$

To first order, we can decompose an agent's forecast into two components: expectations of the real interest rate, $\hat{r}^e_t$, and expectations concerning their future marginal utility, $\hat{u}^e_{c,i,t}$. Formally this can be expressed as

$$\log\left(\frac{\hat{\lambda}^e_{i,t}}{\bar{\lambda}^e_i}\right) \approx \log\left(\frac{1+\hat{r}^e_t}{1+\bar{r}}\right) + \log\left(\frac{\hat{u}^e_{c,i,t}}{\bar{u}^e_{c,i}}\right).$$

Expectations of interest rates are common across individuals[24] so differences in beliefs are entirely reflected in different forecasts of their marginal utilities. Let agent j and agent i have the same idiosyncratic states but differ in their beliefs. If we subtract the forecasts of agent j from agent i we find

$$\left\langle \psi_i - \psi_j, X_t \right\rangle = \log\left(\hat{u}^e_{c,i,t}\right) - \log\left(\hat{u}^e_{c,j,t}\right).$$

[23] For brevity we subsume the dependence on the idiosyncratic states in the subscript i. $\hat{\lambda}^e_{i,t}$ should be read as $\hat{\lambda}^e_t(a_{i,t}, \varepsilon_{i,t}, \psi_{i,t})$. While we are focusing on boundedly rational agents, this discussion also nests rational agents in the special case when $\psi_{i,t} = \psi^{RE}(a_{i,t}, \varepsilon_{i,t})$.

[24] This is clearly true under rational expectations, but is also true under learning. The log linear structure allows us to interpret the process of learning the shadow price as learning about how the aggregate state affects the interest rate and then, separately, learning about how it affects their own marginal utilities.

When $\langle \psi_i - \psi_j, X_t \rangle$ is positive, we can interpret agent i as being relatively more pessimistic than agent j as they expect their marginal utility to be higher next period. Similarly, when $\langle \psi_i - \psi_j, X_t \rangle$ is negative we can interpret agent i as being relatively optimistic.

Turning now to the beliefs under rational expectations, the TFP column of table 2 has coefficients on wealth and productivity that are both positive.[25] This implies that, in the rational expectations equilibrium, wealthier and more productive agents are generally less optimistic about the future during booms and less pessimistic about the future during recessions than their poorer counterparts. These differences in beliefs reflect the ability of richer agents to use their wealth to buffer themselves against business cycle fluctuations.

The dependence of beliefs on individual characteristics is completely lost in the smaller gain calibrations. The coefficients on wealth and log productivity for the $\gamma = 0.001$ calibration are both orders of magnitude smaller than their rational expectations counterparts. This is the result of the mixing due to idiosyncratic risk being much faster than the discounting of past experiences governed by the gain parameter. Wealthy agents remember (and place near equal weight) 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 roughly the same beliefs and the explanatory power of the belief regressions are nearly zero (R^2s of 0.0004 and 0.0171). This confirms the convergence to the RPE which features uniform beliefs.

The uniformity of beliefs in the RPE holds the key to understanding the locally rational dynamics featured in figure 2. Suppose that the distribution of agents features beliefs that are initially homogeneous at the average REE levels. Rich agents in this economy would be relatively more optimistic in booms and pessimistic in recessions than their counterparts in the rational expectations equilibrium. As a result, richer agents over-consume in booms and under-consume in recessions, which generates the amplified response of consumption observed in table 1 and figure 2.[26] In line with their higher consumption, more productive agents also supply less labor in booms relative to rational expectations, resulting in the smaller increase in hours and interest rates observed in figure 2. Over time, agents internalize the effect of these lower interest rates in their forecasts of the shadow price of savings, which results in average beliefs under learning having a more negative coefficient on TFP than the average beliefs under rational expectations.

As the gain increases, agents place more weight on their current experiences and less weight on the distant pass. This brings the resulting distribution of beliefs more in line with the rational expectations beliefs. We observe this in table 2, as both the mean beliefs and the coefficients on wealth and productivity are generally closer to their rational expectations

[25] To get a feeling for the magnitude of these coefficients note that wealth ranges from 0 to 600 with a standard deviation of 40 while log productivity ranges from -2.5 to 2.5 with a standard deviation of 0.8. A one standard deviation increase in wealth would raise the loading on TFP by 0.06 while a 1 standard deviation increase in productivity would raise the loading on TFP by 0.12.

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

counterparts in models with higher gain. This shift in beliefs is reflected in the figure 2 impulse responses with higher gains being closer to the rational expectations paths. While corresponding moments in table 1 are also closer to rational expectations, they are not identical.[27] The model with a gain of $\gamma = 0.035$ has the best fit, as agents 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).

6 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 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 due to the model's inherent non-linearity, 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

[27]A 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 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.

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

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.

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] Â©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.

Appendices

A Transition Dynamics

A.1 The Transition Dynamics of the Rational Model

We begin with some notation. For complete metric space Y, let $\mathbb{B}(Y)$ be its Borel subsets and $\mathscr{P}(Y)$ be the collection of Borel probability measures on $\mathbb{B}(Y)$. Denote by $\mathscr{A} \subset \mathbb{R}$ the collection of possible claims holdings, by $\mathscr{E} \subset \mathbb{R}$ the state-space of the idiosyncratic efficiency shocks, and by $\mathscr{A} \times \mathscr{E}$ the agent-specific state space. Using these notations we can identify the distribution of agent-states (a, ε) in a given period t with a measure $\mu_t \in \mathscr{P}(\mathscr{A} \times \mathscr{E})$.

Let the period t aggregate state to be $\xi_t = (\mu_t, \theta_t)$. The transition dynamic for μ, denoted H, is constructed assuming that all agents in the economy have the same savings rule a_t. The dynamic H is determined by its values on sets of the form $A \times B$ for $A \in \mathbb{B}(\mathscr{A})$ and $B \in \mathbb{B}(\mathscr{E})$. Letting χ be the Boolean truth operator, the transition H is given by

$$H(\xi_t)(A \times B) = \int_{\mathscr{A} \times \mathscr{E}} \chi(a_t(a, \varepsilon) \in A) \cdot \Pi(B, \varepsilon) \mu_t(da, d\varepsilon). \tag{17}$$

This dynamic has the following interpretation: if, in period t, all agents use the savings rule a_t, if agent-states are distributed over $\mathscr{A} \times \mathscr{E}$ as $\mu_t \in P(\mathscr{A} \times \mathscr{E})$, and if the aggregate state in the current period is $\xi_t = (\mu_t, \theta_t)$ then, in the next period, agent-states are distributed over $\mathscr{A} \times \mathscr{E}$ as $\mu_{t+1} = H(\xi_t)$. Note that the construction of this dynamic does not presuppose a rational expectations equilibrium: it is only necessary that all agents use the same savings rule.

A.2 Transition Dynamics of the LR Economy

To characterize the locally rational model's dynamics it is useful to expand the aggregate state to include the common estimate of the second-moment matrix as well as the previous period's observables: $\xi_t = (\mu_t, \theta_t, R_t, X_{t-1})$. Let $A \in \mathbb{B}(\mathscr{A})$, $B \in \mathbb{B}(\mathbb{R}^n)$, and $C \in \mathbb{B}(\mathscr{E})$. The state dynamics, which condition on contemporaneous prices through the savings behavior of locally rational agents, are given by

$$\hat{H}_t(\mu_t)(A \times B \times C) = \int_{\mathscr{A} \times \mathscr{E} \times \mathbb{R}^n} \chi\Big(\left(\hat{a}_t(a, \varepsilon, \psi), \hat{\psi}_t(a, \varepsilon, \psi)\right) \in A \times B\Big) \cdot \Pi(C, \varepsilon) \cdot \mu_t(da, d\varepsilon, d\psi).$$

In contrast to the rational case, $\hat{H}$ and the agents' savings functions $\hat{a}_t$ are not simultaneously determined. This observation undergirds the computational and cognitive simplicity afforded by local rationality: there is no need for the agent nor the modeler to worry about inter-temporal consistency. The dynamics of the model are causal.

B Temporary Equilibrium

B.1 Theory

Recall that the state of the economy includes lagged observables, which are needed to update beliefs: $\xi_t \equiv (\mu_t, \theta_t, R_t, X_{t-1})$. Expand the definition of $\hat{l}_t$ to include the implicit dependence on current factor prices, i.e. $\hat{l}_t(a, \varepsilon, \psi, r, w)$. Define

$$\mathscr{T}\mathscr{E}_t^r(r,w) = \theta_t f_k\left(\int a \cdot \mu_t(da, dve, d\psi), \int \left(1 - \hat{l}_t(a, \varepsilon, \psi, r, w)\right) \mu_t(da, d\varepsilon, d\psi)\right) - \delta$$

$$\mathscr{T}\mathscr{E}_t^w(r,w) = \theta_t f_n\left(\int a \cdot \mu_t(da, d\varepsilon, d\psi), \int \left(1 - \hat{l}_t(a, \varepsilon, \psi, r, w)\right) \mu_t(da, d\varepsilon, d\psi)\right).$$

As market clearing implies that

$$n_t = \int 1 - \hat{l}_t(a, \varepsilon, \psi, r_t, w_t) \mu_t(da, d\varepsilon, d\psi),$$

the firm optimality implies that the temporary equilibrium factor prices must solve $\mathscr{T}\mathscr{E}_t^r(r_t, w_t) = r_t$ and $\mathscr{T}\mathscr{E}_t^w(r_t, w_t) = w_t$.

B.2 Numerics

Finding the temporary equilibrium can be made more efficient by pre-computing the policy rules for household labor supply. These policy rules, $\hat{l}(a, \varepsilon, \phi, r, w)$, are the choices of an agent, given current prices r and w, with wealth a, labor productivity ε, and beliefs summarized by $\phi = \langle \psi, X \rangle$. We approximate these policy rules using the same basis functions as with the computation of the SRE along the asset dimension, 20th order Chebyshev polynomials along the ϕ dimension, and 10th order Chebyshev polynomials along both the r and w dimension. Aggregate labor supply, given r and w, can then be computed via

$$\hat{N}_t(r,w) = \int 1 - \hat{n}(a, \varepsilon, \phi, r, w) \mu_t(da, d\varepsilon, d\psi).$$

C 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 4.2 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.

D Extension: An Expanded Learning Rule

We modify the agent's expectations function to allow for interactions with idiosyncratic states by modifying the forecasting rule, equation (12), to be

$$\hat{\lambda}_t^e(a', \varepsilon, \psi) = \bar{\lambda}^e(a, \varepsilon) \cdot \exp\left(\langle \psi_t, \mathsf{X}(X_t, a, \varepsilon)\rangle\right). \tag{18}$$

$\mathsf{X}(X, a, \varepsilon)$ is a function to allow for arbitrary interactions of the aggregate observable, X, and the individual states, x . Based on the regressions in table 2 we will study the behavior of learning models when

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

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 (13), is adjusted to include the individual states x_t in a similar manner:

$$\begin{aligned} \hat{R}_t(a, \varepsilon, \psi, R) &= R + \gamma \cdot \left(\mathsf{X}(X_{t-1}, a, \varepsilon) \otimes \mathsf{X}(X_{t-1}, a, \varepsilon) - R\right) \\ \hat{\psi}_t(a, \varepsilon, \psi, R) &= \psi + \gamma \cdot \hat{R}_t(a, \varepsilon, \psi, R)^{-1} \mathsf{X}(X_{t-1}, a, \varepsilon) \left(\log\left(\frac{\hat{\lambda}_t(a, \varepsilon, \psi)}{\bar{\lambda}_t(a, \varepsilon)}\right) - \langle \psi_t, \mathsf{X}(X_{t-1}, a, \varepsilon)\rangle\right). \end{aligned} \tag{19}$$

As each agent will have their own second moment matrix based on their unique experiences, one of the states of the model will be $\mu_t \in \mathscr{P}\left(\mathscr{X} \times \mathbb{R}^n \times (\mathbb{R}^n \times \mathbb{R}^n)\right)$, 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 4. First, the forecasting and learning rules are adjusted according to equations (18) and (19). 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.[28] 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.[29]

Table 3 reports the business cycle statistics for the standard model constructed following the same procedures as in section 5. 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 3 which are almost exactly in line with the rational expectations paths for all of the gains considered.

			Expanded Forecasting Rule		
	Data	RE	$\gamma = 0.001$	$\gamma = 0.005$	$\gamma = 0.01$
$\frac{\text{std}(C)}{\text{std}(Y)}$	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 3: 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 implements, generates stabler paths of beliefs, and produces results that better fit the stylized facts observed in the data.

[28]Agent's beliefs are projected back to the RPE

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

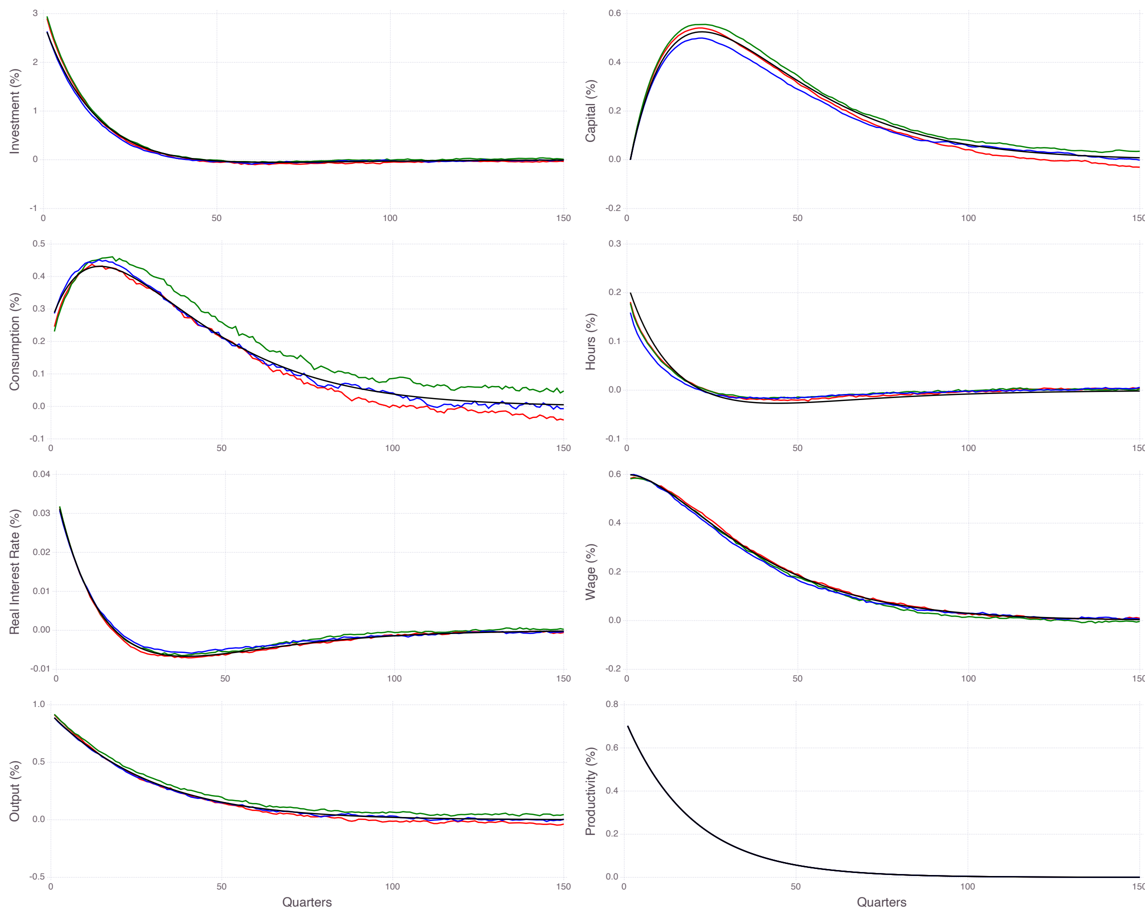

Figure 3: Impulse responses to a one-standard deviation increase in productivity. Black line refers to the rational expectations equilibrium. The blue, green and red lines refer to the median response of the locally rational dynamics with the expanded learning rule and gains equal to 0.001, 0.005, and 0.01 respectively.