Local rationality*

David Evans
University of Oregon

Jungang Li
University of Oregon

Bruce McGough
University of Oregon

July 6, 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, and 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, the natural equilibrium concept for this environment. 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 Bewley models to serve as laboratories for the study of incomplete markets in general

*We thank George Evans and the University of Oregon Macro Group for comments.

equilibrium environments. By developing the needed technical machinery to incorporate aggregate risk into Aiyagari's model, Krusell and Smith (1998) realized this illustrated potential, and the era of *heterogenous agent* (HA) macro 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. 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 her dynamic program and thereby make fully rational choices, an agent must understand the transition dynamics of her 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?

The criticism that REE, as an equilibrium concept, is too demanding of a model's agents is often countered by claims that REE are most appropriately viewed as emergent outcomes of learned behaviors, in the same way that the child learns to catch a ball not by studying physics but instead by practice. In macroeconomic environments, adaptive learning has become the dominant mechanism by which agents are modeled as learning to make forecasts and take decisions. Whereas under RE agents are assumed to use the endogenous conditional distributions of the relevant variables when forming forecasts and making decisions, under adaptive learning agents are assumed to estimate these distributions using standard econometric methods. This paradigm adheres to the *cognitive consistency principle* that a model's agents should be neither much smarter than, nor much stupider than the agents' modeler.

If an economy populated with learning agents converges in an appropriate sense to an

[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.

REE then it is natural to view that equilibrium as an emergent outcome of learned behavior; and the associated literature has met with success in providing conditions that guarantee stability under adaptive learning in linearized models. The early, ground-breaking work of Bray and Savin (1986) demonstrated the stability under least-squares learning of the cobweb model's REE. Their results were established from first principles, making generalizations challenging; this impediment was circumvented by Marcet and Sargent (1989) who brought to bear techniques from the engineering literature, thus allowing for the examination of stability under adaptive learning of REE in many linear models. See Evans and Honkapohja (2001) for detailed discussion of adaptive learning in macroeconomics.

The stability of REE in many models were subsequently studied by many authors, and from these studies a central theme emerges: neither existence nor uniqueness of an REE establishes it as an emergent outcome arising from learned behavior. For example, McCallum (2007) establishes the equivalence of stability and determinacy in some linear models, while in other classes of linear models, including models of applied importance, the unique REE may be unstable: see Bullard and Mitra (2002) for an example of a linearize new Keynesian model with a determinate steady state and an unlearnable equilibrium.[3]

In case of equilibrium multiplicity the outcomes are even more nuanced. Woodford (1990) showed sunspot equilibria in simple over-lapping generations model are stable under adaptive learning. However, Evans and Honkapohja (2001) found that sunspot equilibria in non-convex real economies like those studied by Farmer and Guo (1994) and Benhabib and Farmer (1996) are not stable, thus suggesting that the sunspot equilibria in these models are not naturally viewed as emergent outcomes and thus are unlikely to be of applied importance. On the other hand, McGough, Meng, and Xue (2013) provide an example of a non-convex real economy for which the sunspot equilibria are stable under learning. The lesson from the learning literature is clear: an REE is not necessarily an emergent outcome arising from learned behavior.

Adaptive learning has also been used to help dynamic stochastic general equilibrium (DSGE) models meet the data. In these types of examinations, agents are typically assumed use *constant gain learning* (CGL) when updating their forecast rules. CGL is a form of weighed least-squares with higher weight placed on more recent data, and has the potential to introduce endogenous persistence into environments like the RBC model that lack strong intrinsic propagation mechanisms. For example, Milani (2007) found that a benchmark new Keynesian model did not need the inertia induced by habit persistence or inflation indexation to match the data, provided that the economy's agents were modeled as adaptive learners using CGL. Eusepi and Preston (2011) showed, in an RBC environment, that adaptive learning reduces the volatility of the exogenous technology shocks needed to match the standard deviation of output, and that the learning model's persistence properties are much closer than REE to the data.[4]

[3]See Honkapohja and Mitra (2020) for a more recent application of adaptive learning to questions of monetary policy in the form of price-level targeting.

[4]Adaptive learning has also provided useful explanations of asset price movements. Branch and Evans

Even with learning agents, the benchmark RA RBC model struggles to match some moments, consumption volatility relative to output volatility being a particular difficult statistic to match. 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.[5]

In this paper we develop a new behavioral concept – *local rationality* – that is intended to address the criticism that RE demands far too much of agents in HA environments, and to improve the ability of the benchmark RBC model to match moments. Informally, a locally rational agent knows how to respond optimally to idiosyncratic (local) shocks but must learn how to account for aggregate (global) shocks, the idea being that its 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 decisions.

Operationally, locally rational agents know how to behave optimally in the Aiyagari model, but not the Krusell-Smith model. We use the *shadow price approach*, developed in Evans and McGough (2020), to model how agents respond to variation in aggregates.[6] The advantage of this approach is the simplicity it affords the agents: they make decisions based on perceived trade offs which they measure using shadow prices, and form expectations of future shadow prices using simple linear forecast rules. In particular, agents are not required to understand the evolution of the economy's states and they are not required to solve complex, nonlinear dynamic programs.[7]

The simplicity afforded the model's agents extends also to the agents' modelers. To

(2010) develop an asset-price model with learning agents and find that it reproduces regime-switching returns and volatilities – see also Branch and Evans (2011). Adam, Marcet, and Nicolini (2016) use a consumption-based asset pricing model with learning agents to match the volatility and persistence of the price/dividend ratio.

[5]The aforementioned 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.

[6]In a complete markets model, Honkapohja and Mitra (2006) consider the impact of expectations heterogeneity on equilibrium stability, and apply their results to policy considerations in new Keynesian model. Giusto (2014) opened the door to using adaptive learning in HA models. He demonstrates that if agents follow Krusell and Smith (1998) and use average capital ownership as a sufficient statistic for the wealth distribution, and if they use recursive least squares estimate a model for forecasting average capital ownership, then the equilibrium in the Krusell-Smith model is stable under learning. Guisto also finds that an implementation of CLG improves the model's ability to match the wealth distribution found in the US data.

[7]The shadow-price approach 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 (2020) and Gabaix (2017). 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.

compute the REE of an HA model, the modeler must characterize the transition dynamic of the economy's state by solving an infinite-dimensional functional equation. This is no small feat, and the computational technology needed to approximate solutions is still in the developmental stages. In contrast, an HA model with locally rational agents can be simulated recursively: each period local and global shocks are realized; agents' forecast rules are updated and their associated decision schedules are thereby determined; the schedules are then coordinated in temporary equilibrium to realize individual and aggregate variables; and then the process repeats.

A given agents' beliefs, denoted by ψ and interpreted as the coefficients of her forecast rule, feature prominently in her decision making, and thus the determination of beliefs must be carefully disciplined. If all agents use forecast rules with the same regressors and with the same beliefs coefficients ψ, and if these beliefs are held fixed over time (as opposed to being updated using adaptive learning) then, under certain regularity conditions, the economy's state vector will converge weakly to an ergodic distribution.[8] We may then use this ergodic distribution to compute the linear forecast rule, conditioning on the same set of common regressors, that is optimal in the sense that it minimize the expected forecast error. Letting $T(\psi)$ be the coefficients of the optimal forecast rule, we define a *restricted perceptions equilibrium* (RPE) as a fixed point: $T(\psi^*) = \psi^*$. This equilibrium concept has a venerated history that goes back to Sargent (1993); see Branch (2006) for a survey.

In an RPE, agents will, on average, have no incentive to adjust their beliefs. In this way, an RPE in a natural counterpart to an REE, and an appropriate discipline benchmark. However, conditional on her individual state, the optimal forecast rule for a given agent may differ considerably from the RPE rule. To account for this, locally rational agents use a constant gain learning algorithm, as discussed above, which improves forecasts and decisions as agent-states evolve.

To examine the potential for local rationality to improve moment matching, we use a standard heterogeneous agent environment along the lines of Krusell and Smith (1998). We first develop the model under rationality to provide a benchmark, and then introduce local rationality as an alternate behavioral primitive. As is standard in the HA literature, we use a computational approach to evidence the existence of an RPE, and we further find that it is stable under adaptive learning. Agent's beliefs converge weakly to an ergodic distribution, and, for small gain, mean beliefs approximate the RPE. In this way, local rationality generates an ergodic distribution of beliefs that is an emergent outcome of a natural learning process.

We then turn to our calibration exercise, which is based on parameter values used by Boppart, Krusell, and Mitman (2018). As noted above, simple HA models of the type under examination here are unable to match the volatility of consumption relative to output found in the data when they are endowed with rational agents; and the analogous represen-

[8] This distribution will be over possible values of the state vector, which itself includes a distribution as a component.

tative agent models can't match these moments either, regardless of whether the agents are taken as rational or as adaptive learners. Our principal interest here is to assess whether local rationality, i.e. the combination of heterogeneity and adaptive learning, can improve matters. To this end, we first compute various moments from de-trended US data, and then, using simulations, we compute the corresponding moments in the HA and RA models, both under rationality and local rationality. Our findings confirm the literature: in the data $\text{std(C)}/\text{std(Y)} = 0.50$, whereas across all RA treatments and for the HA model with rational agents, this ratio is between 0.32 and 0.36. With heterogeneity and learning combined, and a gain consistent with similar models, the resulting ratio exactly matches the data.

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 considers alternative forecast models for the agents. Section 7 concludes.

2 The rational model

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

Heterogeneous agent models with rational agents are, by now, so commonplace in the literature that their presentation is often high-level and brief, with emphasis placed only on the novelty under examination. The reader typically is assumed sufficiently familiar with the many technical details that they can either proceed with confidence of the model's internal consistency or they can work through the analysis themselves. As our work here re-imagines the agents' behavioral primitives, it is ground-level and necessarily detailed. To motivate our modified primitives and to facilitate comparison to the benchmark case, we develop the well-known rational model in more detail than is common.[9]

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

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 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 her flow utility as a function of period t consumption c_t and leisure l_t. Different agents have different efficiency units of labor per hour worked. In return to supplying labor, each agent receives a wage that can be separated into two parts: an aggregate component w_t that is the same across all agents; and an idiosyncratic efficiency component ε_t that is independent and identically distributed across all agents. We assume that $\{\varepsilon_t\}$ is a Markov process with time-invariant transition function Π. An agent cannot fully insure against idiosyncratic risk, but in each period t she can trade one-period claims to capital up to an exogenously given borrowing constraint $\underline{a}$, for net return r_t. We denote by a_t the number of claims held from period t to period $t+1$. Goods and factor markets are assumed competitive.

We imagine the following narrative for the economy's evolution: in period t a given agent finds herself holding claims a_{t-1}, experiencing idiosyncratic efficiency ε_t, and facing prices r_t and w_t. Additionally, the agent has at her disposal a host of additional data and information useful for forming forecasts and making decisions.[10] She proceeds to make period t decisions by choosing values for c_t, l_t and a_t to satisfy

$$u_c(c_t, l_t) \geq \beta E_t \lambda_{t+1} \text{ and } a_t \geq \underline{a}, \;\; \textit{with c.s.} \tag{1}$$

$$u_l(c_t, l_t) = u_c(c_t, l_t)\, w_t \tag{2}$$

$$a_t = (1+r_t)\, a_{t-1} + w_t \cdot \varepsilon_t \cdot (1-l_t) - c_t \tag{3}$$

$$\lambda_t = (1+r_t)\, u_c(c_t, l_t) \tag{4}$$

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 saving. Equation (2) balances her intra-temporal trade-off between labor and leisure.

To emphasize per-period decision making, which will be useful when connecting the rational case to our locally rational implementation below, let λ_t^e represent the agent's period t forecast of her period $t+1$ shadow price. In the rational case under examination here, i.e. $\lambda_t^e = E_t \lambda_{t+1}$, the data used to make this forecast include contemporaneous states and decisions: $\lambda_t^e = \tilde{\lambda}_t^e(a_t, \varepsilon_t)$, where we use tildes to identify functional dependence in the rational model with aggregate shocks (we will have a distinct notation for the model

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

without aggregate risk).[11] Inserting $\tilde{\lambda}_t^e(a_t, \varepsilon_t)$ into the Euler condition, we see that (1) – (4) can be used to determine decision schedules as functions of states: $c_t = \tilde{c}_t(a_{t-1}, \varepsilon_t)$, $l_t = \tilde{l}_t(a_{t-1}, \varepsilon_t)$, and $a_t = \tilde{a}_t(a_{t-1}, \varepsilon_t)$.

A note on notation is warranted. The reader no-doubt knows that, for example, the symbol c_t in equation (1) can be interpreted in a variety of ways, depending on context: c_t may be viewed as the quantity consumed by a given agent in period t; as part of a stochastic process whose period t realization coincides with the quantity consumed by the agent in period t; or as a policy function that conditions on the agent's state and determines the quantity consumed by the agent in period t. This latter interpretation is the most precise, but also leads to overly-cumbersome notation when all the associated dependencies are tracked. For our work here, it will be helpful to, in effect, split the notational difference.

We will distinguish between realized values and functions by decorating the later with a tilde in the rational case with aggregate risk (as mentioned above), with a bar in the rational case without aggregate risk, and with a hat in the locally rational case. Also, when using functions we will often suppress some of the arguments in order to keep notation manageable. Thus, for example, the rational agent's decision schedule above is written $c_t = \tilde{c}_t(a_{t-1}, \varepsilon_t)$, which we take to mean that the value c_t of her consumption in period t is given by the function $\tilde{c}_t$, which takes as arguments (a_{t-1}, ε_t). A common approach is to have the rational agents' decision depend on the the distribution of agent-states and the aggregate productivity level – here this dependency has been subsumed into the subscript t. One final comment: we will often only display the dependencies salient to a particular discussion, and so the displayed dependencies will differ across sections of the paper; further it may be useful to suppress all dependencies, especially when referencing values that require forecasting, e.g. $E_t\tilde{c}_{t+1}$ is the rational forecast of next period's consumption.

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 $\theta f(k,n)$. 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:

$$\begin{aligned} w_t &= \theta_t f_n(k_t, n_t) \\ q_t &= \theta_t f_k(k_t, n_t) = r_t + \delta, \end{aligned} \tag{5}$$

where δ is the capital depreciation rate.

[11] The data also include, for example, prices, which we do not explicitly reference as conditioning arguments in order to thin the notation; instead, we subsume some arguments into the time subscript – a notational convenience that is discussed in more detail below.

2.3 The transition dynamics

The transition function of the wealth distribution, H, is an endogenous object of the economy, and its characterization is central both to the analysis of equilibrium when agents are rational and to the development below of local rationality. We therefore take time to develop the characterization of H with care.

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 $\mathscr{X}$ the agent-specific state space, i.e. $\mathscr{X} = \mathscr{A} \times \mathscr{E}$. Using these notations we can identify the distribution of agent-states $x_t = (a_{t-1}, \varepsilon_t)$ in a given period t with a measure $\mu_t \in \mathscr{P}(\mathscr{X})$.

Rational forecasts in heterogenous agent models like the one under examination here require knowledge of the agent-state distribution as well as the aggregate shocks, thus we take 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 $\tilde{a}_t$. 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

$$H(\xi_t)(A \times B) = \int_{\mathscr{X}} \chi(\tilde{a}_t(x) \in A) \cdot \Pi(\varepsilon, B)\, \mu_t(dx). \tag{6}$$

This dynamic has the following interpretation: if, in period t, all agents use the savings rule $\tilde{a}_t$, if agent-states are distributed over $\mathscr{X}$ as $\mu_t \in P(\mathscr{X})$, 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 $\mathscr{X}$ 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.

2.4 Dynamic recursive equilibrium

Rather than defining an equilibrium as a collection of stochastic processes satisfying certain restrictions, we characterize our various equilibrium concepts in terms of the endogenous behavioral rules and transition dynamics implied by the adopted behavioral primitives and exogenous drivers. A *dynamic recursive equilibrium* (DRE) is a collection of behavioral rules $\left(\tilde{c}_t, \tilde{l}_t, \tilde{a}_t, \tilde{\lambda}_t^e\right)$ and pricing functions (r_t, w_t) satisfying

- *Agent optimality:* For all $x_t = (a_{t-1}, \varepsilon_t) \in \mathscr{X}$, the choices $c_t = \tilde{c}_t(x_t)$, $l_t = \tilde{l}_t(x_t)$, and $a_t = \tilde{a}_t(x_t)$ satisfy (1) – (3)
- *Agent rationality:* For all $a_t \in \mathscr{A}$, $\varepsilon_t \in \mathscr{E}$, and $\theta_t > 0$,

$$\tilde{\lambda}_t^e(a_t, \varepsilon_t) = E_t\left[\tilde{\lambda}_{t+1} \middle| a_t, \varepsilon_t\right] \tag{7}$$

where $\tilde{\lambda}_t = (1+r_t)u_c\left(\tilde{c}_t, \tilde{l}_t\right)$ is the period t shadow price.

- *Market clearing:* Markets clear:

$$k_t = \int_{\mathscr{X}} a \cdot \mu_t(dx) \text{ and } n_t = \int_{\mathscr{X}} \left(1 - \tilde{l}_t(x, r_t, w_t)\right) \cdot \mu_t(dx)$$

- *Firm optimality:* Prices r_t and w_t satisfy (5)
- *State dynamics:* The state evolves as $\mu_{t+1} = H(\xi_t)$ and $\theta_{t+1} = v_{t+1}\theta_t^{\rho}$.

Observe that, given any initial aggregate state $\xi_0 = (\mu_0, \theta_0)$, a DRE, 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 $\{r_t, w_t\}$.

2.5 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. Equations (1) – (4) and (5) still hold, and by identifying agent-specific variables with corresponding aggregates, the model's dynamics are quite simple to characterize:

- *Agent optimality:* Aggregate capital, consumption, and leisure satisfy (1) – (3)
- *Agent rationality:* $\lambda_t^e = E_t(1+r_{t+1})u_c(c_{t+1}, l_{t+1})$
- *Market clearing:* Markets clear: $k_t = a_{t-1}$ and $n_t = 1 - l_t$
- *Firm optimality:* Prices r_t and w_t satisfy (5)
- *State dynamics:* Capital evolves as $k_{t+1} = \theta_t f(k_t, n_t) + (1-\delta)k_t - c_t$

2.6 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. Replacing tildes with over-bars to distinguish this special case, and noting that, since prices are constant, an agent's behavior depends only on her state $x \in \mathscr{X}$, 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 $x = (a, \varepsilon) \in \mathscr{X}$, the choices $c = \bar{c}(x)$, $l = \bar{l}(x)$, and $a = \bar{a}(x)$ satisfy (1) – (3) when $r = \bar{r}$ and $w = \bar{w}$
- *Agent rationality:* For all $a \in \mathscr{A}$ and $\varepsilon \in \mathscr{E}$,
$$\bar{\lambda}^e(a, \varepsilon) = \int_{\mathscr{E}} \bar{\lambda}\left(a, \varepsilon'\right) \Pi(\varepsilon, d\varepsilon'), \tag{8}$$
where $\bar{\lambda} = (1 + \bar{r}) u_c\left(\bar{c}, \bar{l}\right)$.
- *Market clearing:* Markets clear when $r = \bar{r}$ and $w = \bar{w}$:
$$k = \int_{\mathscr{X}} \bar{a}(x) \cdot \bar{\mu}(dx) \quad \text{and} \quad n = \int_{\mathscr{X}} \left(1 - \bar{l}(x)\right) \cdot \bar{\mu}(dx)$$
- *Firm optimality:* Prices $\bar{r}$ and $\bar{w}$ satisfy (5)
- *State dynamics:* $\bar{\mu} = H_{\bar{a}}(\bar{\mu})$

The asymmetry in the definitions of DRE and SRE, i.e. that an SRE includes an additional equilibrium object, reflects the focus in the latter case on the steady state distribution of agent-states: in effect, the definition of an SRE includes a particular initial condition for the aggregate state dynamics.

2.7 Looking ahead

To foreshadow what's to come, observe that equations (1) - (3) 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
$$\hat{\lambda}_t^e = \hat{\lambda}_t^e\left(a_t, \varepsilon_t, X_t, \mathscr{B}_t\right) \tag{9}$$
be a possibly non-rational agent's forecast of tomorrow's shadow price of claims conditional on her state ε_t, her savings choice a_t, some collection of observable aggregates X_t, and finally on some form of beliefs $\mathscr{B}_t$ about how today's data inform forecasts of tomorrow's shadow price. For example, $\mathscr{B}$ could encode the objective conditional distributions of all relevant variables so that λ^e could align with rational expectations; or, $\mathscr{B}$ could represent a simple linear forecast model with parameters that are updated over time as new data become available. The shadow-price forecasts $\hat{\lambda}_t^e$, coupled with the Euler condition (1), 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*, emphasized by Evans and Honkapohja (2013), which asserts that the model's agents should be taken as neither much smarter, nor much stupider than the 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

In an REE, the model's agents know not only the current distribution of agent-states but also its law of motion and its 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 (2020), 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. A novel aspect of our approach here is that we assume agents are adept at responding to their idiosyncratic circumstances – they are *locally rational* – which we operationalize by assuming they know how to behave optimally in the absence of aggregate risk and are only learning how aggregate shocks should affect their decisions.

In period t, a given agent is identified by her state x_t and her beliefs $\mathscr{B}_t$. She, together with all other agents, is assumed to observe some common vector of aggregates $X_t \in \mathbb{R}^n$, and she conditions her forecasts, $\hat{\lambda}_t^e$, on these aggregates. She then uses this forecast rule to determine her period t decisions.

It remains to specify how the expectation $\hat{\lambda}_t^e$ is formed. Our *local rationality* assumption is that, absent aggregate risk, an agent knows how to form forecasts optimality; we now assume that in the presence of aggregate risk, the agent forms expectations *relative* to the rational forecasts she would have made in a stationary environment. Additionally, consistent with the agent-level approach that emphasizes the natural simplicity of linear forecast models, we assume that the agent adjusts her stationary expectations by a scaling factor that depends linearly on observables and that is updated over time using a simple estimation procedure.

Operationally, a given agent's period t beliefs $\mathscr{B}_t$ are taken to include a vector $\psi_t \in \mathbb{R}^n$. This vector is interpreted as a linear functional on the space of aggregate observables $\mathbb{R}^n$, and the agent is taken to form expectations as

$$\hat{\lambda}_t^e = \bar{\lambda}^e\left(a_t, \varepsilon_t\right) \cdot \exp\left(\langle \psi_t, X_t \rangle\right), \tag{10}$$

where $\bar{\lambda}^e$ is as defined in (8).[12] In this way, the agent's shadow-price forecast is her 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.

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 convex 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*, and may be taken as decreasing or constant over time. We assume that agents use a constant-gain learning (CGL) to implement their estimation procedure. Our reliance on CGL is an important feature of local rationality, and its implications are discussed in detail below.

To update their beliefs agents regress log deviations of the realized shadow price $\hat{\lambda}_t$ from its stationary counterpart $\bar{\lambda}_t$ on to the previous period's observables X_{t-1}. Letting R measure the agent's estimate of the second-moments of X, and setting $\mathscr{B}_t = (\psi_t, R_t)$, the recursive formulation of this updating rule is given by

$$\begin{aligned} R_{t+1} &= R_t + \gamma \cdot \left(X_{t-1} \otimes X_{t-1} - R_t\right) \equiv \mathscr{R}\left(R_t, X_{t-1}\right) \\ \psi_{t+1} &= \psi_t + \gamma \cdot \mathscr{R}\left(R_t, X_{t-1}\right)^{-1} X_{t-1} \left(\log\left(\frac{\hat{\lambda}_t\left(x_t\right)}{\bar{\lambda}_t\left(x_t\right)}\right) - \langle \psi_t, X_{t-1} \rangle\right) \equiv \Psi_t\left(x_t, \psi_t, R_t, X_{t-1}\right), \end{aligned} \tag{11}$$

where the last equivalence defines notation that will be useful when characterizing equilibrium dynamics. Here $\gamma \in (0,1)$ is the gain, and is usually taken to be small, e.g. $\gamma \leq 0.1$,

[12] 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.

though there is nothing in principle that precludes higher values. Also, note that the term R_t^{-1}, which captures the needed forecast-error adjustment mentioned above, depends only on aggregates and so may be taken as common across agents.

We may now characterize locally rational behavior. Given state (a_{t-1}, ε_t), observables X_t, beliefs ψ_t, and prices (r_t, w_t), the agent then chooses c_t, l_t, and a_t to satisfy

$$u_c(c_t, l_t) \geq \beta \left(\int_{\mathscr{E}} \bar{\lambda}(a_t, \varepsilon') \Pi(\varepsilon_t, d\varepsilon') \right) \cdot \exp(\langle \psi_t, X_t \rangle), \quad a_t \geq \underline{a}, \quad \textit{with c.s.} \tag{12}$$

$$u_l(c_t, l_t) = u_c(c_t, l_t) w_t \tag{13}$$

$$a_t = (1 + r_t) a_{t-1} + w_t \cdot \varepsilon_t \cdot (1 - l_t) - c_t. \tag{14}$$

Importantly, equations (12) – (14) are taken as *behavioral primitives:* they are imposed assumptions on the behavior the the households. Equation (12) balances the agent's intertemporal consumption/savings trade off, and equation (13) balances her intra-temporal labor/leisure trade off. In analog to the rational case, these equations determine decision schedules as functions of states and prices *and beliefs*: $c_t = \hat{c}_t(x_t, X_t, r_t, w_t, \mathscr{B}_t)$, $l_t = \hat{l}_t(\cdot)$, and $a_t = \hat{a}_t(\cdot)$.

Before turning to equilibrium dynamics it is worth reflecting on the simple nature of our agent's behavior. She enters a period with beliefs $\mathscr{B}_t$, holding claims a_{t-1} and having idiosyncratic productivity ε_t; she observes prices r_t and w_t as well as other aggregate information X_t; she makes decisions by balancing perceived trade-offs between consumption and savings (using her expected shadow price) and between labor and leisure; and finally she updates her perceptions by measuring her realized shadow price of claims relative to its known stationary value, and then regressing (the log of) this measure on aggregate information.

3.2 Locally rational dynamics

Given SRE behavior $\bar{\lambda}$, agent-specific states and beliefs $(x_t, \mathscr{B}_t)$, prices (r_t, w_t) and observable aggregates X_t, the conditions (12) - (14) determine agents' decision schedules. 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. Because R_t is common across agents, the distribution of beliefs need only account for variation in ψ_t.[13]

Let $\mu_t \in \mathscr{P}(\mathscr{X} \times \mathbb{R}^n)$ be the contemporaneous distribution of agent-states and ψ-beliefs. Then temporary equilibrium imposes that $r_t = \theta_t f_k(k_t, n_t) - \delta$ and $w_t = \theta_t f_n(k_t, n_t)$,

[13]It would perhaps be more accurate to say that the agent-specific state space must be expanded to include beliefs, but for notational consistency we will continue to refer to $x_t = (a_{t-1}, \varepsilon_t)$ as a representative agent state, and to ψ_t as a representative agent-specific belief.

where

$$k_t = \int_{\mathscr{X} \times \mathbb{R}^n} a \cdot \mu_t(dx, d\psi) \text{ and } n_t = \int_{\mathscr{X} \times \mathbb{R}^n} \left(1 - \hat{l}_t(x, \psi)\right) \mu_t(dx, d\psi) \tag{15}$$

and θ_t is the realized TFP shock. Note that, just as in the RE model, prices are determined by the distribution μ_t and the productivity shock θ_t, i.e. $r_t = \hat{r}(\xi_t), w = \hat{w}(\xi_t)$, with $\xi_t = (\mu_t, \theta_t)$ as before. The difference is that the distribution μ_t is over states x *and* beliefs ψ.

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

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

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.

To elaborate on this point, first consider the agent's view: she enters the period with individual states $x_t = (a_{t-1}, \varepsilon_t)$ and individual beliefs ψ_t and R_t (though R is common across agents). She observes the aggregate X_t and then forms her schedules. The schedules of all agents are coordinated in temporary equilibrium, resulting in realized prices, which in turn results in realized choices c_t, l_t, and a_t for her. She goes to work, gets her wage, goes to her broker to trade claims, and stops by the store on the way home to collect her consumables. Finally, she measures her realized shadow price, λ_t, and updates her beliefs. Now she's ready to relax, no more decisions or actions being needed until tomorrow.

Now consider the modeler's view. The state of the economy is $\xi_t = (\mu_t, \theta_t, R_t, X_{t-1})$, and the functions $\left(\hat{c}_t, \hat{l}_t, \hat{a}_t\right)$ are known. Define

$$\mathscr{T}\mathscr{E}_t^r(r_t, w_t) = \theta_t f_k\left(\int_{\mathscr{X} \times \mathbb{R}^n} a \cdot \mu_t(dx, d\psi), \int_{\mathscr{X} \times \mathbb{R}^n} \left(1 - \hat{l}(x, \psi, r_t, w_t)\right) \mu_t(dx, d\psi)\right) - \delta$$
$$\mathscr{T}\mathscr{E}_t^w(r_t, w_t) = \theta_t f_n\left(\int_{\mathscr{X} \times \mathbb{R}^n} a \cdot \mu_t(dx, d\psi), \int_{\mathscr{X} \times \mathbb{R}^n} \left(1 - \hat{l}(x, \psi, r_t, w_t)\right) \mu_t(dx, d\psi)\right)$$

Denote by Γ the map that takes the full aggregate state ξ_t to the observed aggregate state 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: for the aggregate state $\xi_t = (\mu_t, \theta_t, R_{t-1}, X_{t-1})$,

1. $X_t = \Gamma(\xi_t)$
2. (r_t, w_t) solve $w_t = \mathscr{T}\mathscr{E}_t^w(r_t, w_t)$
   $r_t = \mathscr{T}\mathscr{E}_t^r(r_t, w_t)$

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. $\mu_{t+1} = \hat{H}_t (\mu_t, r_t, w_t)$

The recursive causality of this dynamic system simplifies the computational burden faced by the modeler: it is no longer necessary to search for (an approximation of) a distributional transition dynamic that is consistent with rational expectations on the part of agents. This simplification comes at a cost: resolution of the temporary equilibrium (item 2) and approximation of the distributional dynamics (item 5) require analysis of an agent-specific state-space that has been expanded to include beliefs. Under rationality, beliefs are homogenous 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 $\xi_0 = (\mu_0, \theta_0, R_0)$, the LR-dynamics, together with a sequence of innovation draws $\{v_t\}$, uniquely determines a time path of aggregate states $\{\xi_t\}$, and thus of agent-state distributions $\{\mu_t\}$, as well as a time path of prices $\{r_t, w_t\}$.

3.3 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. RPE that are associated with a model's determinate steady state are commonly found to be stable under adaptive learning: see Evans, Evans, and McGough (2021) for an examination in a non-linear RBC model. 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.[14] Under these beliefs, the distribution of agent states μ_t, observed aggregates and their lags (X_t, X_{t-1}), and prices (r_t, w_t), will converge weakly to an ergodic distribution $\mathscr{M}(\psi)$ over $\mathscr{P}(\mathscr{X}) \times \mathbb{R}^{2n} \times \mathbb{R}^2_+$, where we recall that $\mathscr{P}(\mathscr{X})$ is the space of Borel measures on $\mathscr{X}$.

We may now ask for the values of ψ that minimize the expected forecast error, where the expectation is taken over the endogenously determined ergodic distribution $\mathscr{M}(\psi)$ implied by agents' beliefs ψ. This value may be computed via projections. Define the

[14]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.

operator $T:\mathbb{R}^n\to\mathbb{R}^n$ implicitly as follows:

$$\Lambda(X_1,r,w,\mu)\equiv\int_{\mathscr{X}}\log\left(\frac{\hat{\lambda}(x,X_1,r,w)}{\bar{\lambda}(x)}\right)\mu(dx), \tag{16}$$
$$\int_{\mathscr{P}(\mathscr{X})\times\mathbb{R}^{2n}\times\mathbb{R}^2_+}\Big(\Lambda(X_1,r,w,\mu)-\langle T(\psi),X_2\rangle\Big)\cdot X_2\cdot\mathscr{M}(\psi)(d\mu,dX_1\otimes dX_2,dr,dw)=0.$$

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 her 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.

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 (2020).

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 (12) – (14) 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 (11), 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 is 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 $X = \left(1, \log\left(k/\bar{k}\right), \log(\theta)\right)$. 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.[15]

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). 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. 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[16] that

[15]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 (2006) 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.

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

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 and update the aggregate state according to the algorithm presented in section 3.2.

Finding the temporary equilibrium can made more efficient by pre-computing the policy rules for household labor supply. These policy rules, $\hat{n}(a,\varepsilon,\phi,r,w)$, 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

$$\hat{N}(r,w) = \int_{\mathscr{X}\times\mathbb{R}^n} \hat{n}(a,\varepsilon,\phi,r,w)\,\mu(da,d\varepsilon,d\psi).$$

5 Numerical 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 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 (16) , 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 $\lambda_{i,t}(\psi)$ and $\bar{\lambda}_{i,t}(\psi) = (1+\bar{r})\bar{\lambda}(a_{i,t}(\psi), \varepsilon_{i,t})$ 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

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

This map converges point-wise to T as N and S approach infinity.[17] We numerically verify the existence of an RPE by finding the fixed point of $\hat{T}(\psi)$ when $N = 100,000$ and $S = 1,000$.[18]

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 values due to the non-linearities of the model.

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

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

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

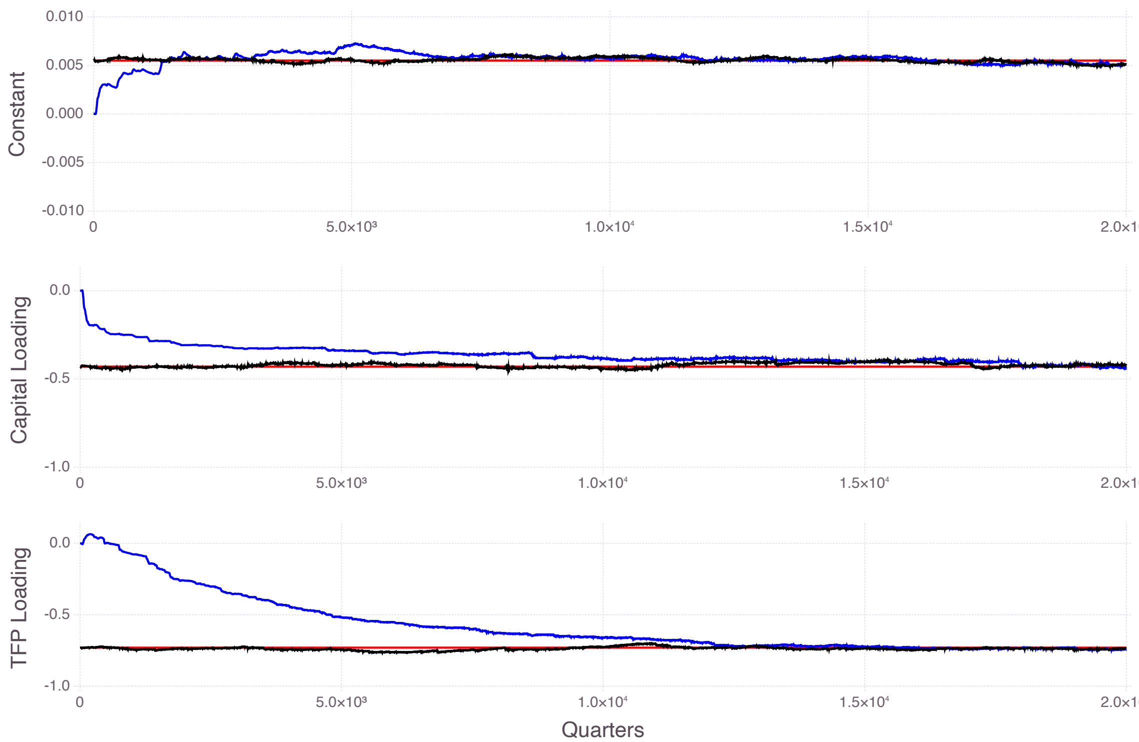

Figure 1: Time path of average beliefs when agents are initialized with the RPE beliefs (black) and $\psi = \mathbf{0}$ (blue). The RPE beliefs are represented with a solid red line.

HP-filtering the log of all relevant variables.[19] 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 standard deviation of output for all models and the data. As has been well documented in the literature (see Romer (2012)), the benchmark real business cycle model both overstates the variation of investment and, correspondingly, understates the variation of consumption relative to the data. Neither the introduction of bounded rationality through shadow price learning nor the introduction of heterogeneous agents is able to significantly change any of these moments. However, the interaction of bounded rationality and agent heterogeneity leads to substantially different second moments, bringing them closer to the data by increasing the standard deviation of consumption while decreasing the standard deviation of investment.

Focusing on the last 4 columns of table 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

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

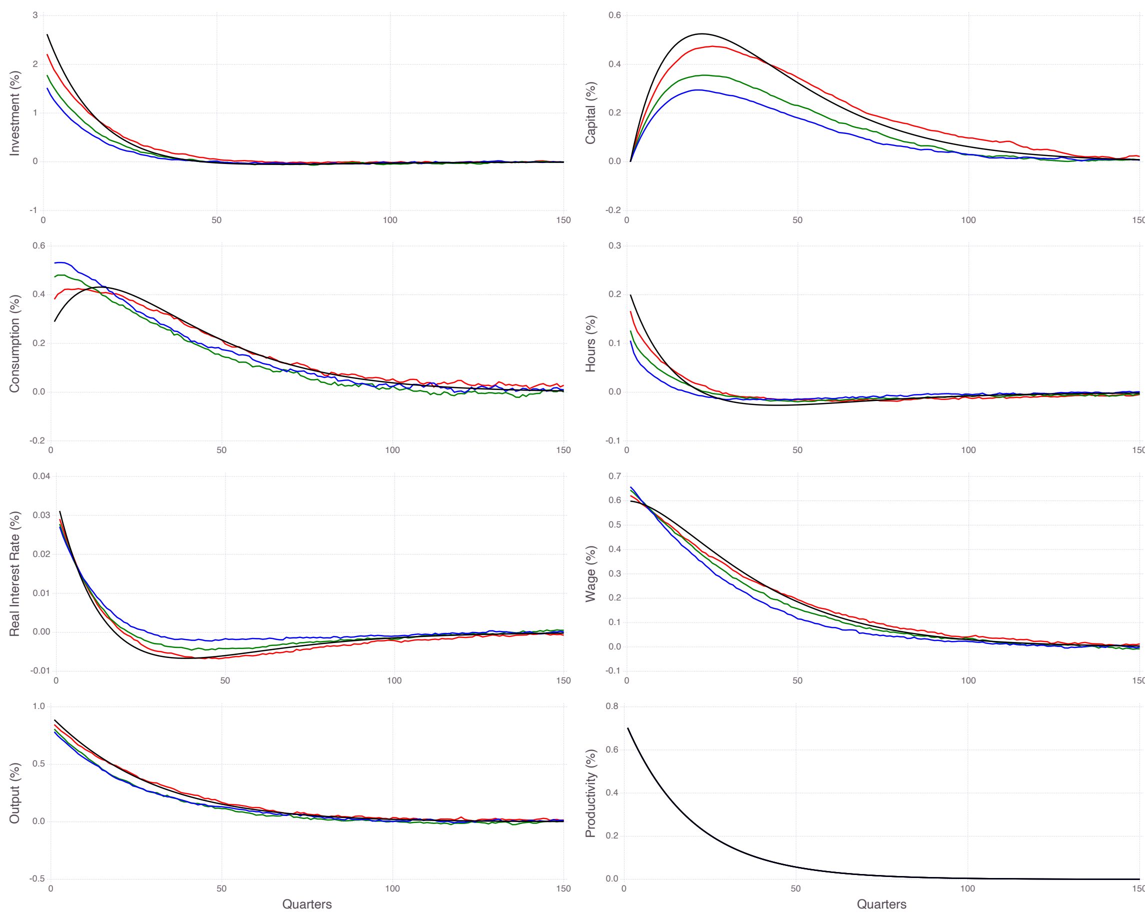

Figure 2: Impulse responses to a one-standard deviation increase in productivity. Black line refers to the linearized rational expectations equilibrium. The blue, green and red lines refer to the median response of the locally rational dynamics with gains equal to 0.001, 0.01, and 0.035 respectively.

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 the beliefs arising from 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 and, for each point in the state space (a,ε), the log deviation of the expected shadow price of wealth from its steady state counterpart:

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

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. Note, that we allow the rational expectations beliefs to vary based on individual states. This codifies the fact that different agents have different experiences in recessions depending on their current situation. These different experiences are what given rise to an endogenous distribution of beliefs present in the locally rational model.

	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.

To gain an understanding of this distribution of beliefs we construct a simple set of summary statistics by running the weighted regression

$$\psi_i^j = \alpha_0 + \alpha_1(a_i - \bar{a}) + \alpha_2(\log(\varepsilon_i) - \overline{\log(\varepsilon)}) + \mu_i, \tag{17}$$

which regresses beliefs of agents on their state variables. The term ψ_i^j represents the jth 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 average loading on capital and TFP are both negative. The 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 that low gain economy converged to the restricted perceptions equilibrium which features all agents having identical beliefs. These features are necessary to understand the behavior of the learning economy.

To understand how to interpret this heterogeneity in beliefs, recall that beliefs measure the relative deviation of the expectation of shadow price of wealth, $\lambda^e_{i,t}$, from their steady-state counterpart $\bar{\lambda}^e_i$.[20] For a given information set, X_t, these forecasts are given by

$$\log\left(\frac{\lambda^e_{i,t}}{\bar{\lambda}^e_i}\right) = \psi_i' X_t.$$

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

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

Expectations of interest rates are common across individuals[21] 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

$$(\psi_i - \psi_j)' X_t = \log\left(u^e_{c,i,t}\right) - \log\left(u^e_{c,j,t}\right).$$

When $(\psi_i - \psi_j)' X_t$ 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 $(\psi_i - \psi_j)' X_t$ 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.[22] 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

[20] For brevity we subsuming the dependence on the idiosyncratic states in the subscript i. $\lambda^e_{i,t}$ should be read as $\lambda^e_t(a_i, \varepsilon_i, \psi_i)$

[21] 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.

[22] 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.

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.

A feature of the RPE beliefs is that richer agents will be relatively more optimistic in booms and pessimistic in recessions than their counterparts in the rational expectations equilibrium. This would be true even if beliefs were homogeneous and held at the average REE levels. 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.[23] 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 on their shadow price of wealth which results in average beliefs under learning differing from 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 counterparts in the 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 impulse responses. While corresponding moments in table 1 are also closer to the rational expectation, they are not identical. The model with a gain of $\gamma = 0.035$ has the best fit, as agents who 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).

[23]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.

6 Extension: An Expanded Learning Rule

One feature of table 2 is the R^2 for the rational expectation regression is higher for both coefficients. The linear model explains 73% of the loading on capital and 52% of the loading on TFP. This suggests that including an interaction term of X_t with idiosyncratic states may bring the learning models closer in line rational expectations. We explore this idea in this section.

We modify the agent's expectations function to allow for these interaction by modifying the forecasting rule, equation (10), to be

$$\hat{\lambda}_t^e = \bar{\lambda}^e(a_t, \varepsilon_t) \cdot \exp\left(\langle \psi_t, \mathsf{X}(X_t, x_t)\rangle\right). \tag{18}$$

$\mathsf{X}(X,x)$ 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

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

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_t will differ across agents. The recursive formulation of the updating rule, equation (11), is adjusted to include the individual states x_t in a similar manner:

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

As each agent will have their own second moment matrix based on their unique experiences, one of the states of the model will be $\mu_t \in \mathscr{P}(\mathscr{X} \times \mathbb{R}^n \times (\mathbb{R}^n \times \mathbb{R}^n))$, i.e. the contemporaneous distribution of agent-states x, 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.[24] 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.[25]

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.

7 Conclusion

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

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

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

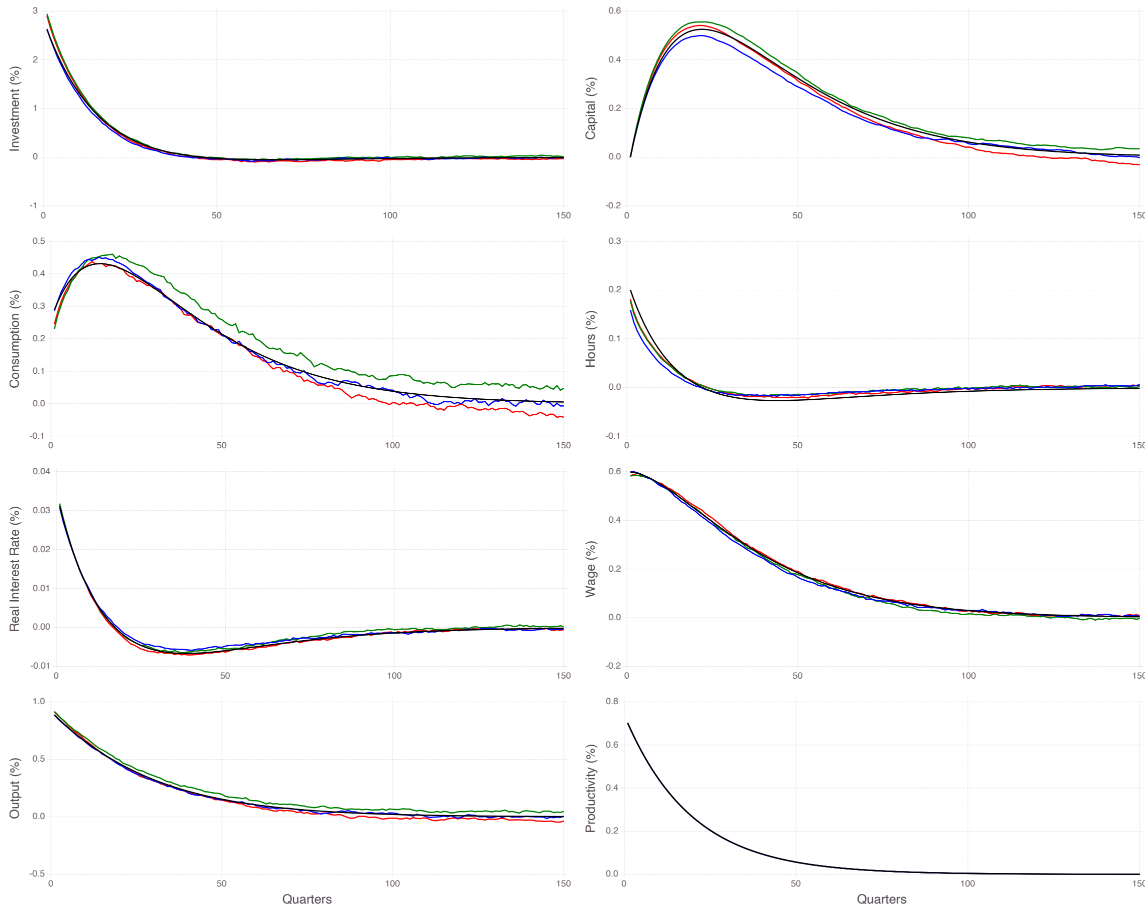

Figure 3: Impulse responses to a one-standard deviation increase in productivity. Black line refers to the rational expectations equilibrium. The blue, green and red lines refer to the median response of the locally rational dynamics with the expanded learning rule and gains equal to 0.001, 0.005, and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

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 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., A. MARCET, AND J. P. NICOLINI (2016): "Stock Market Volatility and Learning," *The Journal of Finance*, 71(1), 33–82.

AIYAGARI, S. R. (1994): "Uninsured idiosyncratic risk and aggregate saving," *The Quarterly Journal of Economics*, 109(3), 659–684.

BENHABIB, J., AND R. E. FARMER (1996): "Indeterminacy and sector-specific externalities," *Journal of Monetary Economics*, 37(3), 421–443.

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 Colander (2006), pp. 135–160.

BRANCH, W. A., AND G. W. EVANS (2006): "A simple recursive forecasting model," *Economics Letters*, 91(2), 158–166.

——— (2010): "Asset Return Dynamics and Learning," *The Review of Financial Studies*, 23(4), 1651–1680.

——— (2011): "Learning about Risk and Return: A Simple Model of Bubbles and Crashes," *American Economic Journal: Macroeconomics*, 3(3), 159–91.

BRANCH, W. A., G. W. EVANS, AND B. MCGOUGH (2013): "Finite Horizon Learning," in Sargent and Vilmunen (2013), chap. 9.

BRANCH, W. A., AND B. MCGOUGH (2011): "Business cycle amplification with heterogeneous expectations," *Economic Theory*, 47(2/3), 395–421.

BULLARD, J., AND K. MITRA (2002): "Learning about monetary policy rules," *Journal of Monetary Economics*, 49(6), 1105–1129.

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.

COLANDER, D. (2006): *Post Walrasian Macroeconomics*. Cambridge, Cambridge, U.K.

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 (2021): "THe RPEs of RBCs," 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.

——— (2013): "Learning as a Rational Foundation for Macroeconomics and Finance," in Frydman and Phielps (2013).

EVANS, G. W., AND B. MCGOUGH (2020): "Learning to optimize," Discussion paper.

FARMER, R. E., AND J.-T. GUO (1994): "Real Business Cycles and the Animal Spirits Hypothesis," *Journal of Economic Theory*, 63(1), 42–72.

FRYDMAN, R., AND E. PHIELPS (2013): *Rethinking Expectations: The Way Forward for Macroeconomics*. Princeton University Press.

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.

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.

HONKAPOHJA, S., AND K. MITRA (2006): "Learning Stability in Economies with Heterogeneous Agents," *Review of Economic Dynamics*, 9(2), 284–309.

——— (2020): "Price level targeting with evolving credibility," *Journal of Monetary Economics*, 116, 88–103.

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.

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 mechanisms in self-referential linear stochastic models," *Journal of Economic Theory*, 48(2), 337–368.

MCCALLUM, B. T. (2007): "E-stability vis-a-vis determinacy results for a broad class of linear rational expectations models," *Journal of Economic Dynamics and Control*, 31(4), 1376–1391.

MCGOUGH, B., Q. MENG, AND J. XUE (2013): "Expectational stability of sunspot equilibria in non-convex economies," *Journal of Economic Dynamics and Control*, 37(6), 1126–1141.

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. (1993): *Bounded Rationality in Macroeconomics*. Oxford University Press, Oxford.

SARGENT, T. J., AND J. VILMUNEN (eds.) (2013): *Macroeconomics at the Service of Public Policy*. Oxford University Press.

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.

WOODFORD, M. (1990): "Learning to Believe in Sunspots," *Econometrica*, 58(2), 277–307.

——— (2018): "Monetary Policy Analysis when Planning Horizons are Finite," Working Paper 24692, National Bureau of Economic Research.