

The Trouble with Rational Expectations in Heterogeneous Agent Models: A Challenge for Macroeconomics*

Benjamin Moll[†]

First version: November, 2024

This version: November, 2024

[latest version]

Abstract

The thesis of this essay is that, in heterogeneous agent macroeconomics, the assumption of rational expectations about equilibrium prices is unrealistic, unnecessarily complicates computations, and should be replaced. This is because rational expectations imply that decision makers forecast equilibrium prices like interest rates by forecasting cross-sectional distributions. The result is an extreme version of the curse of dimensionality: dynamic programming problems in which the entire cross-sectional distribution is a state variable (“Master equation” a.k.a. “Monster equation”). This is not only unrealistic but also limits the applicability of the heterogeneous-agent approach to some of the biggest questions in macroeconomics, namely those in which aggregate risk and non-linearities are key, like financial crises. This troublesome feature of the rational expectations assumption poses a challenge: what should replace it? I outline three criteria that alternative approaches should satisfy: (1) simplification of the computational solution, (2) consistency with empirical evidence, and (3) (some) immunity to the Lucas critique. I then discuss some potentially promising directions, including temporary equilibrium approaches, incorporating survey expectations, least-squares learning, and reinforcement learning.

*This paper was originally delivered as the *Economic Journal Lecture* at the Royal Economic Society’s 2024 Annual Conference. I am grateful to the Royal Economic Society for this honor. I thank Klaus Adam, Marios Angeletos, Adrien Auclert, Isaac Baley, Adrien Bilal, Pedro Bordalo, Tobias Broer, Andrew Caplin, Ben Enke, Xavier Gabaix, Joao Guerreiro, Ina Hajdini, Zhen Huo, Greg Kaplan, Alex Kohlhas, Per Krusell, Søren Leth-Petersen, Chen Lian, Guido Lorenzoni, Peter Maxted, Dmitry Mukhin, Galo Nuño, Ricardo Reis, Matt Rognlie, Tom Sargent, Karthik Sastry, Andreas Schaab, Martin Schneider, Chengchun Shi, Alp Simsek, Ben Van Roy, Mirko Wiederholt, and Johannes Wohlfart for helpful discussions and Nicholas Tokay for excellent research assistance. I am also grateful for support from the Leverhulme Trust.

[†]London School of Economics

One of the key developments in macroeconomics research over the last three decades has been the incorporation of explicit heterogeneity into models of the macroeconomy. A benefit of this approach is the empirical discipline from matching macro models to micro data. But some of the biggest questions in macroeconomics have so far been out of reach of this approach. These are questions in which aggregate risk and aggregate non-linearities play a key role, for example the perennial question: why do developed economies experience infrequent but large boom-bust cycles like financial crises?¹

The root cause of these questions being out of reach is one particular assumption: that of rational expectations about equilibrium prices. Under rational expectations, decision makers in economies with heterogeneity forecast cross-sectional distributions in order to forecast prices. This results in Bellman equations in which a (typically) infinite-dimensional cross-sectional distribution is a state variable, an extreme version of the curse of dimensionality (Krusell and Smith, 1998; Den Haan, 1996). This problem arises even though decision makers do not directly “care about” the distribution, i.e. it does not enter their objective functions; instead, as is standard in competitive equilibrium models, they only care about prices. The infinite-dimensional Bellman equation is called “Master equation” in the mathematics literature and has been aptly nicknamed the “Monster equation” due to its complexity.² A recent literature has made impressive advances developing methods for solving this “Monster equation”³ but computations are still extremely costly, severely limiting the reach of heterogeneous-agent macroeconomics.⁴

It is instructive to consider an example: suppose I lived in one of our models and wanted to forecast the evolution of future interest rates, say because I am considering taking out a mortgage to buy a house. According to our theories, I would realize that market-clearing interest rates depend on the entire cross-sectional distribution of different asset holdings in the economy (say the U.S. economy, the Euro area, or indeed the entire world economy). I would therefore forecast interest rates by forecasting this entire cross-sectional distribution.

This paper’s main argument is that we should not make our lives so hard.⁵ It seems self-evident that real-world households and firms do not forecast prices by forecasting distributions and instead solve simpler problems.⁶ Macroeconomists are spending a lot of intellectual and computational horse power solving an unrealistically complex problem. Instead of solving “Monster equations” we should replace the rational expectations assumption and solve

¹For example, as in He and Krishnamurthy (2012, 2013), Brunnermeier and Sannikov (2014), Elenev et al. (2021), Mendoza (2010), and Bianchi (2011). A few recent contributions do take important steps in the direction of solving heterogeneous-agent versions of such models. See the discussion of related work at the end of this introduction.

²To be clear, the Master equation is not specific to continuous-time models. Instead *all* rational-expectations heterogeneous agent models with aggregate risk feature a Master equation. Indeed the Master equation was already there in Krusell and Smith (1998) and Den Haan (1996), though not fully spelled out and not with that name.

³See the end of this introduction as well as Section 2.2.

⁴There is, of course, also the solution method of Krusell and Smith (1998) and Den Haan (1996) in which decision makers forecast prices by forecasting *moments* of cross-sectional distributions. This approach has a bounded-rationality interpretation and considerably simplifies computations relative to solving the full Master equation. I discuss the similarities to and differences from the approach advocated here further below.

⁵To be clear, what I mean is that the *computational* complexity and cost should not so be so high. Better modeling expectations formation may well require more (and not less) mathematics. In Tom Sargent’s words (Rolnick, 2010): “A rule of thumb is that the more dynamic, uncertain and ambiguous is the economic environment that you seek to model, the more you are going to have to roll up your sleeves, and learn and use some math. That’s life.”

⁶On this point, also see the quotes from Manski (2004) and Adam and Marcat (2011) in Section 2.

the simpler equations corresponding to households' and firms' actual price-forecasting behavior. This would open up the door to bringing the empirical and theoretical discipline of heterogeneous-agent macroeconomics to the study of economic booms and busts.

Of course, non-rational optimism and pessimism may also be what drives economic booms and busts in the first place (Minsky, 1977; Kindleberger, 1978, and related literature below). The idea that expectations may themselves be drivers of instability is therefore another argument for departing from rational expectations. Put another way, the promise of alternative approaches is to “kill two birds with one stone”: to simplify the computation of heterogeneous-agent models with aggregate risk while, at the same time, making these models more realistic and more likely to generate interesting macroeconomic phenomena.

But what should replace rational expectations about equilibrium prices? A much more reasonable assumption is that decision makers forecast prices directly rather than indirectly by forecasting cross-sectional distributions. This holds the promise of sidestepping the curse of dimensionality. But how exactly do decision makers form these price expectations? I spell out three criteria for such alternatives to rational expectations and discuss some promising directions for developing such alternative approaches.

Before laying out my main thesis about rational expectations, I revisit the roots of this modeling device in the 1960s and 1970s through the lens of some key early writings like Muth (1961), Lucas and Prescott (1971) and Lucas (1972). Expectations about equilibrium prices played a (if not *the*) key role in the development and popularization of the new hypothesis. These early writings also stressed that a key payoff of adopting rational expectations was to make models of the macroeconomy “operational,” meaning that these models could be readily computed, fit to data, and made clear predictions about the behavior of macroeconomic time series following policy changes – see for example Lucas' FORTRAN analogy quoted below.

In heterogeneous-agent macroeconomics, the rational expectations assumption actually hinders rather than helps this “operability” because it implies the extreme version of the curse of dimensionality discussed above. To explain the problem at a deeper, mathematical level, I use a specific example: the problem of forecasting wages and interest rates in an entirely standard textbook economy, a real business cycle model with household heterogeneity as in Krusell and Smith (1998). This model is a useful laboratory because the same logic applies to *all* heterogeneous-agent models. The state of the economy is the cross-sectional joint distribution of income and wealth. Even though households and firms do *not* directly “care about” this cross-sectional distribution, it becomes a state variable in their dynamic programming problems because they use it to forecast future wages and interest rates. As already noted, the Mean Field Games (MFG) literature has called this equation the “Master equation” (Cardaliaguet et al., 2019) and this name and the associated formalism have also found their way back into economics in recent years (Ahn et al., 2018; Schaab, 2020; Bilal, 2023; Gu et al., 2024).

Next, I briefly review existing solution methods for heterogeneous-agent models with aggregate risk and explain why it may be worth to go back to the drawing board to develop alternative approaches. First, methods that employ “MIT shocks” or linearize with respect to aggregate states are obviously not suitable for studying questions in which aggregate non-

linearities are key. Second, methods that tackle the full rational expectations equilibrium (i.e. the Master equation) are what the criticism of this paper is aimed at: too much intellectual and computational horse power is aimed at solving an unrealistically complex problem. Third, methods that work with moments of the distribution (Krusell and Smith, 1998; Den Haan, 1996) lack realism for similar reasons as Master equation methods: also the idea that decision makers forecast means, variances, etc of distributions in order to forecast prices strains credibility.⁷ One variant of the Krusell-Smith approach is less subject to my criticism: the variant in which the “Krusell-Smith moments” are the prices themselves. While this approach shares some similarities with the approach advocated below, I believe that we can do better, for example by incorporating key empirical evidence.

I then turn to the discussion of what should replace rational expectations about equilibrium prices. In place of a concrete alternative proposal,⁸ I spell out three criteria that such alternatives should satisfy. A common element in this class of alternative approaches is that decision makers form expectations about prices directly rather than indirectly via forecasting distributions. The key question is how to discipline the subjective probability distributions of equilibrium prices, i.e. how to navigate the “wilderness of non-rational expectations” (Sargent, 1993, 1999, 2008).⁹ The goal of my three criteria is to help navigate this wilderness. The three criteria are: (1) simplification of the computational solution, (2) consistency with empirical evidence, and (3) endogeneity of beliefs to model reality (Lucas critique). An implication of Criterion 1 is that it eliminates from the list of candidate alternative approaches any non-rational expectations models that require decision makers to be able to compute the special case of rational expectations. Criteria 2 and 3 further narrow the set of admissible subjective beliefs. Another important requirement is that the candidate alternative approach be applicable to non-linear models and compatible with recursive solution methods for individual optimization problems (Bellman equations).

In the paper’s final section I discuss some potentially promising directions for developing alternative models of price beliefs satisfying these three criteria. First, I discuss the fundamental concept of temporary equilibrium, i.e. competitive equilibrium given specified subjective beliefs, as well as the idea of internal rationality. Second, alternative models of price beliefs should incorporate the large body of empirical work on expectations formation including evidence from survey expectations. I then turn to models of adaptive learning which are natural candidates for delivering on the three criteria. I focus on two classes of learning algorithms: least-squares learning from the economics literature and reinforcement learning from the computer science literature. Reinforcement learning means learning value functions of incompletely-known Markov decision processes and has been at the core of some impressive advances in artificial intelligence. The two forms of learning are linked because both are special cases of a more general class of stochastic approximation algorithms. Finally, I discuss the idea that decision makers may use heuristics and simplified, restricted forecasting models.

⁷See Broer et al. (2022b) for a related critique of the idea that decision makers use moments to forecast prices.

⁸I only know the problem, not the solution!

⁹Sargent attributes the phrase to Sims (1980) but Sims actually writes about the “wilderness of disequilibrium economics”.

Related Literature. A huge literature in macroeconomics analyzes alternatives to the rational expectations paradigm, both empirically and theoretically. I will cover much of this work when I discuss promising directions in Section 4. For now, some useful survey articles and books are: Sargent (1993, 1999, 2008), Woodford (2013), Angeletos et al. (2021), Bordo et al. (2022), Enke (2024), Ortoleva (2024), and Baley and Veldkamp (2025) on the theoretical side and Armantier et al. (2013), Manski (2018), Coibion et al. (2022), D’Acunto and Weber (2024), and Fofana et al. (2024) on the empirical side. Also see the recent Handbook of Economic Expectations (Bachmann et al., eds, 2023) that covers both theory and empirics.

As already noted, I will also review in detail various existing numerical solution methods for heterogeneous-agent models (see Section 2.2). Closest to the approach advocated here is the bounded-rationality approach of Krusell and Smith (1998) and Den Haan (1996), specifically the variant in which the “Krusell-Smith moments” are the prices themselves.

One main motivation of this essay is the pragmatic desire to compute heterogeneous-agent models with aggregate risk and non-linearities so as to be able to apply them to questions like the causes and consequences of financial crises. While these questions have been largely out of reach of the heterogeneous-agent literature, a few recent contributions have taken steps to narrow this gap (e.g. Fernández-Villaverde et al., 2023; Gopalakrishna et al., 2024). However, high computational costs mean that theories with aggregate non-linearities are still highly stylized and cannot yet fully speak to key empirical evidence. Rather than solving “Monster equations”, we should solve simpler, more realistic equations. Rather than “taming the curse of dimensionality” (Fernández-Villaverde et al., 2024a), we should sidestep it.

Finally, as already noted, non-rational expectations may also be the drivers of instability generating booms and busts in the first place. This idea has a long tradition going back to the verbal discussions by Minsky (1977) and Kindleberger (1978). Recent examples of theories in which departures from rational expectations generate boom-bust-cycles include Adam et al. (2017), Bordo et al. (2021), and Maxted (2024).

Roadmap. Section 1 puts the thesis of this essay into perspective by revisiting some key early writings on rational expectations which primarily used it as a device for modeling expectations about endogenous equilibrium prices. Section 2 explains why, in heterogeneous-agent models, the assumption of rational expectations about equilibrium prices is unrealistic, unnecessarily complicates computations, and should be replaced. Section 3 poses the challenge “what should replace rational expectations?” and spells out three criteria that alternative approaches should satisfy. Section 4 discusses some promising directions for developing alternative models of price beliefs. Section 5 concludes.

1 Back to the roots of rational expectations: it was all about equilibrium prices

This section revisits the roots of the rational expectations modeling device in the 1960s and 1970s through the lens of some writings that were key to its development and popularization like Muth (1961), Lucas and Prescott (1971) and Lucas (1972). Expectations about equilibrium prices played a (if not *the*) key role in these developments. This is in contrast to the recent theoretical expectations literature in macroeconomics which mostly focuses on modeling expectations about exogenous variables. The early writings also stressed that a key payoff of adopting rational expectations was to make models of the macroeconomy “operational” with an emphasis on computation.

1.1 Muth (1961) “Rational Expectations and the Theory of Price Movements”

Rational expectations were first proposed by John Muth in this 1961 paper. Already the paper’s title, “*Rational Expectations and the Theory of Price Movements*”, indicates that price expectations play a key role. To “*best explain what the hypothesis is all about*”, Muth first uses a particular example: expectations about equilibrium prices in a single market. Specifically, Muth starts with the following simple demand-supply system (see his Equation 3.1 in Section 3).¹⁰

$$\begin{aligned} D_t &= -\beta p_t && \text{(Demand),} \\ S_t &= \gamma p_t^e + u_t && \text{(Supply),} \\ D_t &= S_t && \text{(Market equilibrium).} \end{aligned} \tag{1}$$

The first equation is the time- t demand curve with demand D_t a decreasing function of the price p_t . The second equation is the supply curve with supply S_t an increasing function of the *expected* price p_t^e and a stochastic, potentially serially correlated supply shock u_t . The third equation is the market clearing condition which equates demand and supply. Production takes place with a one-period lag, i.e. suppliers make their production decision at time $t - 1$, and hence p_t^e is the expectation of p_t based on information available at time $t - 1$.

Muth’s rational expectations hypothesis is a particular assumption about the formation of the price expectation p_t^e . Muth verbally states his hypothesis as follows: “*Expectations of firms (or, more generally, the subjective probability distribution of outcomes) tend to be distributed, for the same information set, about the prediction of the theory (or the ‘objective’ probability distributions of outcomes).*” Applying this idea to the demand-supply system (1), rational expectations is the assumption that the subjective price expectation equals the objective, model-generated price expectation

$$p_t^e = \mathbb{E}_{t-1}[p_t],$$

where the expectations operator \mathbb{E}_{t-1} integrates over the objective, model-generated price distribution. Rational expectations are “model-consistent” in this sense. In his linear model, Muth

¹⁰To make the equations easier to read, I denote demand by D_t instead of C_t and supply by S_t instead of P_t .

solves for the rational price expectation as follows. First, the equilibrium price of the demand-supply system (1) is given by the reduced form

$$p_t = -\frac{\gamma}{\beta}\mathbb{E}_{t-1}[p_t] - \frac{1}{\beta}u_t \quad (2)$$

so that the equilibrium price is low if firms' price expectation is high (because this means that supply is high). Second, taking the time $t - 1$ expectation \mathbb{E}_{t-1} of this equilibrium price yields $\mathbb{E}_{t-1}[p_t] = -\frac{\gamma}{\beta}\mathbb{E}_{t-1}[p_t] - \frac{1}{\beta}\mathbb{E}_{t-1}[u_t]$ or

$$p_t^e = \mathbb{E}_{t-1}[p_t] = -\frac{1}{\beta + \gamma}\mathbb{E}_{t-1}[u_t]$$

which is Muth's equation (3.5). Hence the expected equilibrium price depends only on the expected fundamental, the expectation of the supply shock $\mathbb{E}_{t-1}[u_t]$. A particularly simple case is when u_t is serially uncorrelated and furthermore $\mathbb{E}_{t-1}[u_t] = 0$. In this case the rational price expectation simply equals zero, $p_t^e = \mathbb{E}_{t-1}[p_t] = 0$. From (2) the actual equilibrium price equals $p_t = -u_t/\beta$ which is different from zero and moves around according to the realization of the supply shock u_t . But firms correctly understand that these supply shocks average out. Analogous results obtain when u_t is serially correlated.¹¹

Two appealing features of this approach are as follows. First, to make predictions about equilibrium prices and quantities, no measurement of price expectations is needed because these are instead implied by the theory. Second, because firms understand the determination of equilibrium prices (i.e. they "understand the structure of the economy") and equilibrium prices depend only on the fundamental u_t (and model parameters), modeling firms' price expectations requires only the stochastic process of the fundamental u_t .

1.2 In Bob Lucas' words

Dissatisfaction with prior modeling of price expectations, often via backward-looking "adaptive expectations," was a key impetus for the popularization of the rational expectations hypothesis in the 1960s and 1970s. In his Nobel Prize Lecture, Lucas (1996) describes the situation:

"[1960s-style macroeconomic models] implied behavior of actual equilibrium prices and incomes that bore no relation to, and were in general grossly inconsistent with, the price expectations that the theory imputed to individual agents. [...] This modeling inconsistency became more and more glaring. John Muth's (1961) "Rational Expectations and the Theory of Price Movements" focused on this inconsistency, and showed how it could be removed by taking into account the influences of prices, including future prices, on quantities and simultaneously the effects of quantities on equilibrium prices."

Similarly, Lucas' 1980 methodological review "Methods and Problems in Business Cycle Theory" (Lucas, 1980) stresses the desire to reconcile subjective and objective probability distributions of equilibrium prices:

¹¹Assuming $u_t = \rho u_{t-1} + \varepsilon_t$ for an i.i.d. mean-zero random variable ε_t and correlation coefficient ρ , the rational price expectation is $\mathbb{E}_{t-1}[p_t] = -\frac{1}{\beta+\gamma}\rho u_{t-1}$ and the actual equilibrium price is $p_t = -\frac{1}{\beta+\gamma}\rho u_{t-1} - \frac{1}{\beta}\varepsilon_t$ from (2).

“One needs a principle to reconcile the price distributions implied by the market equilibrium with the distributions used by agents to form their own views of the future. John Muth noted that [...] these distributions could not differ in a systematic way. His term for this latter hypothesis was rational expectations.”

1.3 Lucas and Prescott (1971) “Investment under Uncertainty”

Lucas and Prescott’s important 1971 paper extended rational expectations to dynamic equilibrium models, showed how to formulate a recursive rational expectations equilibrium in a stochastic market economy using dynamic programming techniques, and thus laid the foundation for much of modern macroeconomics up to this day.¹²

The paper models the equilibrium of a single industry in which a large number of competitive firms make production and investment decisions in the face of random demand shocks. Specifically, there is an exogenously-given downward-sloping demand curve for the industry’s output with a demand shifter that follows a Markov process which leads to fluctuations in the output price faced by firms.

The key difficulty then is how to model expectations about future demand and thus equilibrium prices. Lucas and Prescott write:

“From the viewpoint of firms in this industry, forecasting future demand means simply forecasting future output prices. The usual way to formulate this problem is to postulate some forecasting rule for firms, which in turn generates some pattern of investment behavior, which in turn, in conjunction with industry demand, generates an actual price series.

Typically the forecasting rule postulated takes the form of anticipated prices being a fixed function of past prices – ‘adaptive expectations.’ But it is clear that [this implies that] price forecasts and actual prices will have different probability distributions, and this difference will be persistent, costly to forecasters, and readily correctible.

To avoid this difficulty, we shall, in this paper, go to the opposite extreme, assuming that the actual and anticipated prices have the same probability distribution, or that price expectations are rational.”

This last sentence is followed by a footnote stating: *“This term is taken from Muth (1961), who applied it to the case where the expected and actual price (both random variables) have a common mean value. Since Muth’s discussion of this concept applies equally well to our assumption of a common distribution for these random variables, it seems natural to adopt the term here.”* While Muth had already verbally defined rational expectations in this way (see Section 1.1), he had only applied his concept to linear models in which only means matter, and so Lucas and Prescott (1971) first applied rational expectations to entire probability distributions, i.e. in the more general way in which it is commonly used today. Perhaps even more importantly, Lucas and Prescott first showed how to formulate a recursive rational expectations equilibrium in a stochastic market economy using dynamic programming techniques – their key equation (9) is a Bellman equation which indirectly defines the recursive competitive equilibrium by maximizing consumer surplus.

¹²See Sargent (2017) for a useful discussion of Lucas and Prescott (1971).

In summary, also from Lucas and Prescott's pathbreaking paper, it becomes clear that the development of rational expectations was all about improving the modeling of expectations about equilibrium prices.

1.4 Lucas (1972) "Expectations and the Neutrality of Money"

While the economics of Lucas' neutrality paper are, of course, different from the investment paper discussed in the preceding section, the methodology – specifically the construction of a recursive rational expectations equilibrium – is very similar. Lucas is after a Phillips curve, the relation between the change in the price level (inflation) and real output, in an overlapping-generation economy with two physically separated markets (islands). The key variable in his model economy is the equilibrium price level p and the key difficulty is how young individuals form expectations about the future equilibrium price level p' when deciding how much to consume, work, and save for future consumption. This difficulty is exactly analogous to the difficulty of firms predicting future output prices in Lucas and Prescott (1971). So is the solution which Lucas (1972) explains as follows:

"Equilibrium prices and quantities will be characterized mathematically as functions defined on the space of possible states of the economy, which are in turn characterized as finite dimensional vectors. This characterization permits a treatment of the relation of information to expectations which is in some ways much more satisfactory than is possible with conventional adaptive expectations hypotheses."

In a bit more detail, Lucas first writes down the young's optimization problem in terms of an unspecified (subjective) probability distribution F for the future price level p' – see his equation (3.7). He then notes: *"The state of the economy is fully described by the three variables (m, x, θ) [the money supply, the money growth rate, and the fraction of young people on the first island]. [...] If this is so, one can express the equilibrium price as a function $p(m, x, \theta)$ on the space of possible states."*¹³ A recursive rational expectations equilibrium is then a price function $p(m, x, \theta)$ that satisfies the optimality condition of the young's optimization problem (3.7) but with the expectation taken with respect to the joint distribution G of future (m', x', θ') conditional on the current price $p(m, x, \theta)$ in place of the price-level distribution F . Lucas writes: *"We have dispensed with unspecified distribution F , taking the expectation instead with respect to the well-defined distribution G ."*

In a paper written for a conference held at the Federal Reserve Bank of Minneapolis to celebrate the 25th anniversary of the publication of Lucas (1972), Sargent (1996) again emphasizes the importance of rational expectations as a device for modeling expectations about endogenous equilibrium variables: *"The victory of rational expectations owes to its beauty and its utility: the economy with which it eliminates what we had thought were free variables – peoples' expectations about endogenous variables – while adding no free parameters, but bringing instead cross-equation and cross-frequency restrictions."*

¹³Also see useful discussion of Lucas' paper by Chari (1998) who writes *"The key to the technical contribution is that prices are thought of as functions of the state of the economy, where the state is the stock of money and the distribution of young people across islands."*

1.5 The Payoff: “Operational” Macro Theories

The early writings reviewed in this section stress one particular payoff of adopting the rational expectations device: that they make models of the macroeconomy “operational.” By this they mean that models with rational expectations can be readily computed, fit to data, and make clear predictions about policy counterfactuals.

Lucas (1980) writes: *“Our task as I see it [...] is to write a FORTRAN program that will accept specific economic policy rules as ‘input’ and will generate as ‘output’ statistics describing the operating characteristics of time series we care about, which are predicted to result from these policies.”* Similarly, directly after the passage cited in section 1.3, Lucas and Prescott (1971) state: *“[By imposing rational expectations,] we obtain an operational investment theory linking current investment to observable current and past explanatory variables, rather than to “expected” future variables which must, in practice, be replaced by various “proxy variables.”*

This “operability” of rational expectations models, including that they could be computed relatively easily, was a key catalyst of their fast diffusion in macroeconomics. I will return to this section’s themes – expectations about equilibrium prices and operability – later on in the paper.

Finally, while modern macroeconomics owes much to the rational expectations hypothesis, its importance should also not be overstated. Other aspects like micro foundations and general equilibrium are arguably more important for today’s macroeconomics and the label “rational expectations revolution” may be unhelpful. See Werning (2023) for a good discussion.

2 The trouble with rational expectations about equilibrium prices in heterogeneous agent models

With this historical background in mind, I now turn to the main thesis of this essay: in heterogeneous agent macroeconomics, the assumption of rational expectations about equilibrium prices is unrealistic, unnecessarily complicates computations, and should be replaced. This is because rational expectations implies that decision makers (unrealistically) forecast equilibrium prices like interest rates by forecasting cross-sectional distributions, thereby resulting in an extreme version of the curse of dimensionality.

It is instructive to consider the following thought experiment: suppose I lived in one of our models and wanted to forecast the evolution of future interest rates, say because I am considering taking out a mortgage to buy a house. According to our theories, I would realize that market-clearing interest rates depend on the entire cross-sectional distribution of different asset holdings in the economy (say the U.S. economy, the Euro area, or indeed the entire world economy). I would therefore forecast interest rates by forecasting this entire cross-sectional distribution. It seems self-evident that real-world individuals do not do this.

To understand the problem at a deeper, mathematical level, I will use a specific example: the problem of forecasting wages and interest rates in a textbook heterogeneous-agent economy. As a warm-up, I will start with a representative-agent economy (the real business cycle

model). I will then add heterogeneity as in the economy studied by [Krusell and Smith \(1998\)](#).

2.1 Forecasting equilibrium prices in representative agent models

Time is discrete and indexed by $t = 0, \dots, T$ where T may be finite or infinite.¹⁴ The only source of uncertainty is aggregate productivity z_t which follows a Markov process. A representative consumer has preferences over consumption c_t and hours worked n_t given by

$$\mathbb{E}_0 \sum_{t=0}^T \beta^t U(c_t, n_t),$$

where U is a standard utility function and $0 < \beta < 1$ is a discount factor. A representative firm uses capital k_t and labor ℓ_t to produce output according to a technology subject to productivity shocks

$$z_t F(k_t, \ell_t).$$

Capital accumulates according to $k_{t+1} = i_t + (1 - \delta)k_t$ where i_t is investment and $0 < \delta < 1$ the depreciation rate. The economy's resource constraints are:

$$c_t + i_t = z_t F(k_t, \ell_t), \quad \ell_t = n_t, \quad \text{all } t = 0, \dots, T.$$

This completes the description of this simple economy, a real business cycle (RBC) model. The only not-entirely-standard feature is that the time horizon T may be finite (see footnote 14).

The definition of a competitive equilibrium is also standard. To fix ideas, I consider the particular decentralization in which firms own the economy's capital stock which they finance by issuing risk-free bonds to households who use these bonds as their saving vehicle. A competitive equilibrium is then quantities and prices $\{w_t, r_t\}$ such that:

1. Households maximize taking as given $\{w_t, r_t\}$

$$\max_{\{c_t, n_t, a_{t+1}\}} \mathbb{E}_0 \sum_{t=0}^T \beta^t U(c_t, n_t) \quad \text{s.t.} \quad c_t + a_{t+1} = w_t n_t + (1 + r_t) a_t. \quad (3)$$

2. Firms maximize taking as given $\{w_t, r_t\}$

$$\max_{\{x_t, \ell_t, k_{t+1}\}} \mathbb{E}_0 \sum_{t=0}^T R_{0 \rightarrow t}^{-1} (z_t F(k_t, \ell_t) - w_t \ell_t - i_t) \quad \text{s.t.} \quad k_{t+1} = i_t + (1 - \delta)k_t, \quad (4)$$

with $R_{0 \rightarrow t} = \prod_{s=1}^t (1 + r_s)$.

3. Markets clear

$$k_t = a_t, \quad \ell_t = n_t, \quad \text{all } t = 0, \dots, T. \quad (5)$$

The key difficulty, which is the focus of this paper, is that optimizing households and firms

¹⁴The reason for choosing to nest the finite- T case is that the special case of only two time periods $t = 0, 1$ will turn out to be useful below.

need to forecast future wages and interest rates $\{w_t, r_t\}$. To flesh this out, let us focus on wages for now (the logic for interest rates will be exactly analogous). Household optimization (3) yields labor supply $n_t = n(w_t, a_t)$ that depends on the wage and the household's state variables, here wealth a_t . Analogously, firm optimization (4), yields labor demand $\ell_t = \ell(w_t, k_t, z_t)$. From market clearing, therefore, the equilibrium wage is a function of the economy's states, aggregate capital k_t and aggregate productivity z_t :

$$w_t = w^*(k_t, z_t).$$

How then do households and firms forecast future wages? Given rational expectations, they do so by forecasting the economy's state variables (k_t, z_t) . This is because they understand the structure of the economy and, in particular, the mapping from (k_t, z_t) to equilibrium wages w_t . An analogous logic holds for forecasting equilibrium interest rates $r_t = r^*(k_t, z_t)$. Note that this use of "equilibrium price functions" is exactly like in Lucas (1972) – see Section 1.4. It is worth noting that, even in this extremely stylized representative-agent economy, a rational expectations equilibrium is the solution to a relatively complicated fixed-point problem and the informational requirements for decision makers are substantial and arguably unrealistic.

Solution methods for the representative agent case. For questions in which aggregate risk and non-linearities play a key role, one wants a global solution method. There are broadly two types.¹⁵ The first approach is to tackle the competitive equilibrium directly. Even in a simple RBC model like the one I just outlined, global solution methods for the competitive equilibrium are actually surprisingly challenging to implement. A classic method is to use dynamic programming in conjunction with a "Big K , little k " trick (e.g. Prescott and Mehra, 2005; Lucas, 1987; Ljungqvist and Sargent, 2004) but there are various other methods. See, for example, Maliar and Maliar (2014) for an overview and Cao et al. (2023) for one recent and relatively general approach.

The second approach, which most macroeconomists would use to solve the RBC model above, is to appeal to the second welfare theorem and to solve for the equilibrium allocation via a social planner's problem. Of course this only works in special cases. But the key payoff is that the social planner's problem does not feature prices. So this approach completely sidesteps the key difficulty emphasized above, namely expectations about equilibrium prices.

2.2 Forecasting equilibrium prices in heterogeneous agent models

Next consider the exact same economy but with household heterogeneity. Specifically, as in Aiyagari (1994) and Krusell and Smith (1998), assume that households are ex-ante identical but subject to idiosyncratic income risk in the form of labor productivity y_{it} which follows a Markov process. This results in ex-post heterogeneity in household wealth and productivity (a_{it}, y_{it}) and I denote the cross-sectional joint distribution of income and wealth by $G_t(a, y)$.

¹⁵There are, of course, many well-developed local solution methods for representative-agent models with aggregate risk, see e.g. the popular Dynare toolbox.

Analogous to the definition in the representative-agent economy, a competitive equilibrium is then quantities and prices $\{w_t, r_t\}$ such that

1. Households maximize taking as given $\{w_t, r_t\}$

$$\max_{\{c_{it}, n_{it}, a_{it+1}\}} \mathbb{E}_0 \sum_{t=0}^T \beta^t U(c_{it}, n_{it}) \quad \text{s.t.} \quad c_{it} + a_{it+1} = w_t y_{it} n_{it} + (1 + r_t) a_{it}. \quad (6)$$

2. Firms maximize taking as given $\{w_t, r_t\}$ – the same problem as in the representative-agent economy (4) since the economy’s production side is unchanged.

3. Markets clear

$$k_t = \int a dG_t(a, y), \quad \ell_t = \int n_t(a, y) dG_t(a, y), \quad \text{all } t = 0, \dots, T. \quad (7)$$

Importantly, note that in this competitive equilibrium, households and firms do *not* “care about” the cross-sectional distribution G_t , i.e. it does not enter their objective functions. Instead they only care about *prices* $\{w_t, r_t\}$.

As before, the key difficulty is that households and firms need to forecast these future prices. Focusing again on wages, household optimization (6) now yields *household-specific* labor supplies $n_{it} = n(w_t, a_{it}, y_{it})$.¹⁶ Labor demand is identical to the representative-agent case and given by $\ell_t = \ell(w_t, k_t, z_t)$. From market clearing, therefore, the equilibrium wage is a function of the economy’s states, now the *entire cross-sectional distribution* G_t in addition to aggregate productivity z_t :

$$w_t = w^*(G_t(a, y), z_t). \quad (8)$$

A generic feature of heterogeneous-agent models. While this dependence of equilibrium prices on the entire cross-sectional distribution G_t can be avoided in special cases,¹⁷ it is important to note that this distributional dependence is a *generic* feature of heterogeneous-agent models. To this end, denote individual state variables by a vector x_{it} with distribution $G_t(x)$, exogenously evolving aggregate state variables by a vector z_t , and prices by a vector p_t so that the model we considered above is the special case in which $x_{it} = (a_{it}, y_{it})$ and $p_t = (w_t, r_t)$. Then, generically, equilibrium prices satisfy

$$p_t = \mathcal{P}^*(G_t(x), z_t) \quad (9)$$

¹⁶Labor supplies are household-specific except in the knife-edge case without income effects (e.g. GHH utility).

¹⁷For example, in the model above, the equilibrium interest rate r_t depends only on the aggregate capital stock, $r_t = r^*(k_t, z_t)$ with $k_t = \int a dG_t(a, y)$. This follows from the market clearing condition for capital in (7) and the optimality condition of the representative firm. Similarly, if labor supply were perfectly inelastic (as in the benchmark economy of [Krusell and Smith, 1998](#)) or if preferences ruled out income effects, also the equilibrium wage would only depend on the aggregate capital stock k_t . However, even in this case, it would still be true that G becomes a state variable in household’s decision problem. [Rios-Rull \(1997\)](#) explains it nicely (substituting my notation for his): “The pair (z, k) is not, in general, a sufficient statistic for k' : tomorrow’s capital k' depends on the whole distribution of wealth G . Depending on how wealth is distributed, aggregate capital will be different tomorrow, except where individual decision rules are linear in a , which is not the usual case.”

so that the equilibrium price function \mathcal{P}^* again depends on the entire cross-sectional distribution G_t . Solution methods for heterogeneous agent models should be able to handle this generic case. In more general models, other endogenous equilibrium objects besides prices themselves can directly enter decision makers' problems. An example of such "price-like variables" is the job finding rate in search-and-matching models. In equilibrium, these variables will depend on the distribution G_t just like in (9). The vector p_t should therefore be understood as including not only prices themselves but also other price-like variables.

Households and firms with rational expectations forecast prices by forecasting distributions. To clearly see this implication of rational expectations, consider the special case with two time periods $t = 0, 1$.

In this case, the households' problem (6) can be written recursively as follows:

$$\begin{aligned} V_0(a, y, G, z) &= \max_{c, n, a'} U(c, n) + \beta \mathbb{E}[V_1(a', y', G', z') | y, G, z] \quad \text{s.t.} \\ c + a' &= w_0^*(G, z) y n + (1 + r_0^*(G, z)) a \\ V_1(a', y', G', z') &= \max_{c', n'} U(c', n') \quad \text{s.t.} \quad c' = w_1^*(G', z') y' n' + (1 + r_1^*(G', z')) a'. \end{aligned} \quad (10)$$

Here V_0 and V_1 are the value functions at times $t = 0$ and $t = 1$ (which are time-dependent because of the finite horizon). G' is the cross-sectional distribution at time $t = 1$ which is induced by households' optimal saving policy $a' = s_0(a, y, G, z)$ and the transition probabilities for labor income y , starting from the initial distribution G : $G' = \mathcal{T}_{s_0} G$ for some operator \mathcal{T}_{s_0} that depends on the function s_0 .¹⁸ Importantly, the expectation \mathbb{E} in the first line of equation (10) is taken not only over future productivity realizations z' but also over future distributions G' . This is because households understand that equilibrium prices at time $t = 1$ depend on this distribution – see $w_1^*(G', z')$ and $r_1^*(G', z')$ in the third line – and hence so does the future value function $V_1(a', y', G', z')$. Therefore households forecast these prices by forecasting G' in addition to z' .

In terms of computations, the problem is of course the dependence of the value function $V_1(a', y', G', z')$ on the distribution G' . Because this distribution is an infinite-dimensional object (or a very high-dimensional one when it is approximated), this leads to an extreme version of the curse of dimensionality. This problem was already noted by [Krusell and Smith \(1998\)](#) and [Den Haan \(1996\)](#). It is also worth noting that the dependence of the initial prices w_0, r_0 on the distribution is unproblematic: the distribution G is known at $t = 0$ and hence so are w_0, r_0 in the second line; the problem is instead about future distributions G' . Similarly, it is worth noting that the extreme curse of dimensionality arises exclusively because G' enters decision makers' expectations and *not* because of the distribution evolving over time per se. Computing $G_1(a, y)$ given $G_0(a, y)$, a policy function s_0 , and the productivity shock z_t , is computationally straightforward. Heterogeneous-agent models become unrealistically complex only when the distribution enters decision makers' expectations.

¹⁸In the finite-state variant of Section 3.4, the analogue of the operator \mathcal{T}_{s_0} is the big $N \times N$ transition matrix \mathbf{A} .

The firm’s investment decision can be written in a similar fashion and has analogous properties. Just like (10) it features future prices $w_1^*(G', z')$ and $r_1^*(G', z')$, a future value function $J_1(k, G', z')$, and an expectation taken over future distributions G' .

Equation (10) is a two-period version of the “Master equation” from the Mean Field Games literature (Cardaliaguet et al., 2019) which has also found its way into economics in recent years (e.g. Ahn et al., 2018; Schaab, 2020; Bilal, 2023; Gu et al., 2024).¹⁹ To be clear already the work of Krusell and Smith (1998) and Den Haan (1996) effectively featured this Master equation, even though they did not use this name.

Using the general notation introduced above with a distribution $G_t(x)$ over a vector of idiosyncratic states x and a vector of aggregate states z , this Master equation is a Bellman equation for the value function $V(x, G, z)$ which features the infinite-dimensional state variable G . In the MFG literature this Master equation has aptly been nicknamed the “Monster equation,” precisely due it being a Bellman equation with an infinite-dimensional state variable. Not surprisingly, solving this Monster equation is computationally extremely challenging – and often prohibitively so.

In Section 1.4, I cited Lucas (1972) as writing: “Equilibrium prices and quantities will be characterized mathematically as functions defined on the space of possible states of the economy, which are in turn characterized as finite dimensional vectors.” The whole problem is that, in heterogeneous-agent models, these “states of the economy” are extremely high-dimensional: they are either literally infinite-dimensional cross-sectional distributions or, when the state space is discretized, they are extremely high-dimensional histograms, say 100,000-dimensional vectors (Section 3.4 below). To make matters worse, forecasting future distributions G' requires knowledge not only of the current distribution G but also of all *other* decision makers’ optimal choices: the operator \mathcal{T}_{s_0} depends on the saving policy function $s_0(a, y, G, z)$ which summarizes *other* households’ saving decisions. Put differently, using the notation of Section 3.4 below, forecasting future distributions requires knowledge of the huge transition matrix \mathbf{A} for the whole economy that summarizes how all *other* decision makers transition across the economy’s state space.

The unrealism of rational expectations. This paper’s main argument is that we should not make our lives so hard. It seems self-evident that real-world households and firms do not forecast prices by forecasting cross-sectional distributions and instead solve simpler problems. Instead of solving “Monster equations” we should solve the simpler equations corresponding to actual price-forecasting behavior. Below I will spell out some ideas for doing so.

The argument that the rational expectations assumption is often unrealistic is, of course, not new. For example, Manski (2004) writes: “Suppose that the true state of nature actually is the realization of a random variable distributed P . A decision maker attempting to learn P faces the same inferential problems – identification and induction from finite samples – that empirical economists confront in their research. Whoever one is, decision maker or empirical economist, the inferences that

¹⁹As Sargent (2024) points out, the Master equation is an extension of the “Big K , little k ” trick used for solving representative-agent competitive equilibria. Informally, the Master equation is $(k, G(k))$ and (k, K) is the special case in which all individuals are identical (G is a Dirac point mass at K).

one can logically draw are determined by the available data and the assumptions that one brings to bear. Empirical economists seldom are able to completely learn objective probability distributions of interest, and they often cannot learn much at all. It therefore seems hopelessly optimistic to suppose that, as a rule, expectations are either literally or approximately rational.”

Similarly [Adam and Marcet \(2011\)](#) write: “The rational expectations hypothesis (REH) places enormous demands on agents’ knowledge about how the market works. For most models it implies that agents know exactly what market outcome will be associated with any possible contingency that could arise in the future. This appears utterly unrealistic given that state contingent markets that could provide agents with such detailed information often fail to exist.”

The argument here is exactly the same but it applies a fortiori because of the model economy’s high complexity. Also, recall again that decision makers do not even directly care about the cross-sectional distribution G so it seems particularly questionable that they would solve such complicated dynamic programming problems featuring G as a state variable.

Solution methods for the heterogeneous agent case. How then do existing numerical methods go about solving this type of problem? There are broadly three sets of solution methods.

The first set of methods assumes that aggregate uncertainty comes in the form of “MIT shocks” or linearizes with respect to aggregate states. These are the most common methods in the recent heterogeneous-agent literature and, specifically, in most of the HANK literature. For example, [Kaplan et al. \(2018\)](#) use “MIT shocks” and all of [Reiter \(2009\)](#), [Ahn et al. \(2018\)](#), [Auclert et al. \(2021\)](#), [Glawion \(2023\)](#), [Bilal and Goyal \(2024\)](#), and [Bayer et al. \(2024a\)](#) use linearization methods.²⁰ In fact, [Boppart et al. \(2018\)](#) point out a close connection between the two approaches. With MIT shocks, there is certainty about the path of equilibrium prices. To the same effect, linearization implies certainty *equivalence* about prices. Certainty (equivalence) about equilibrium prices means that these methods completely sidestep the key difficulty of price expectations. The downside of this approach is that such methods are not suitable for answering questions in which aggregate risk and non-linearities play a key role, for example understanding financial crises.²¹

The second set of methods tackles the global solution of the full rational expectations equilibrium, i.e. of the Master equation. This is where the recent literature has made the most impressive computational advances ([Schaab, 2020](#); [Maliar et al., 2021](#); [Azinovic et al., 2022](#); [Bilal, 2023](#); [Gu et al., 2024](#); [Gopalakrishna et al., 2024](#); [Han et al., 2021](#); [Lee, 2022](#); [Huang, 2023](#); [Pröhl, 2019](#)). However, despite these advances, solving the Master equation has remained extremely difficult and this set of methods is what this paper’s criticism is aimed at. As stated earlier, in my view, too much intellectual and computational horse power is aimed at solving this unrealistically complex problem.

The third set of methods is that of [Krusell and Smith \(1998\)](#) and [Den Haan \(1996\)](#) and vari-

²⁰ A few authors also use second-order perturbation methods. See for example [Bilal \(2023\)](#), [Bhandari et al. \(2023\)](#) and [Bayer et al. \(2024b\)](#).

²¹ As already noted, what I have in mind are heterogeneous-agent versions of models like [He and Krishnamurthy \(2012, 2013\)](#), [Brunnermeier and Sannikov \(2014\)](#), [Elenev et al. \(2021\)](#), [Mendoza \(2010\)](#), and [Bianchi \(2011\)](#).

ants thereof.²² Compared to solving the full rational expectations equilibrium in which decision makers forecast prices by forecasting distributions, the Krusell-Smith approach instead assumes that they forecast prices by forecasting *moments* of these distributions, say the mean $\bar{a}_t = \int a dG_t(a, y)$ instead of the entire distribution $G_t(a, y)$. Therefore these moments, rather than the entire distribution, become state variables in decision makers' Bellman equations. As [Krusell and Smith \(1998\)](#) discuss, this approach has a bounded rationality interpretation: decision makers assume that distributional moments follow a (typically log-linear) perceived law of motion which differs from the objective law of motion. This approach considerably simplifies the computation of heterogeneous agent models with aggregate risk, including in models with aggregate non-linearities. However, just like the full rational expectations approach, the idea that decision makers forecast means, variances, etc of distributions lacks realism. In terms of the three criteria spelled out below, I conjecture that this approach would fail on Criterion 2, consistency with empirical evidence.²³

There is one variant of the Krusell-Smith approach which seems to me relatively more promising. This is the variant in which the “Krusell-Smith moments” are the prices themselves, i.e. decision makers forecast prices directly rather than indirectly via forecasting other moments like means and variances. See for example [Lee and Wolpin \(2006\)](#), [Storesletten et al. \(2007\)](#), [Gomes and Michaelides \(2008\)](#), [Favilukis et al. \(2017\)](#), [Llull \(2018\)](#), [Kaplan et al. \(2020\)](#), and [Fernández-Villaverde et al. \(2024b\)](#). I will return to this idea in the next section. Still, I believe that it would be beneficial to develop alternatives to rational expectations about equilibrium prices in a more systematic fashion, in particular taking into account empirical evidence about expectations formation.

2.3 Taking stock and what next?

As we saw in Section 1, an explicit goal of Lucas, Prescott, and their contemporaries was to develop “operational” theories of the macroeconomy that can be readily computed, fit to data, and make clear predictions about policy counterfactuals. In representative-agent models, even though they may be unrealistic, rational expectations have achieved exactly this goal.

But in heterogeneous agent macroeconomics, the rational expectations assumption achieves the opposite and severely hampers the models' operability. In particular, the assumption attributes to people an understanding of the price functions that map cross-sectional distributions to equilibrium prices. Concretely, it means that decisions makers forecast prices by forecasting distributions. This implication of rational expectations is not only unrealistic but, at the same time, makes the computational solution of these models challenging and costly.

The obvious way forward is for macroeconomists to go back to the drawing board and to replace rational expectations about equilibrium prices in heterogeneous agent models with alternative assumptions. While there are some attempts at this in the existing literature (most notably the variant of the Krusell-Smith approach I just discussed), it is worth to develop

²²See for example the continuous-time variants by [Fernández-Villaverde et al. \(2023\)](#) and [Lee \(2022\)](#).

²³I am not aware of any empirical evidence that supports the idea that decision makers forecast, say, the mean of the U.S. wealth distribution to forecast interest rates. [Broer et al. \(2022b\)](#) criticize the idea that decision makers use common moments to forecast prices from a theoretical perspective.

such alternatives in a more systematic fashion and to incorporate key empirical evidence. The promise of doing so is to “kill two birds with one stone”: to simplify the computation of heterogeneous-agent models with aggregate risk while, at the same time, making these models more realistic and more likely to generate interesting macroeconomic phenomena.

3 Three criteria for replacing rational expectations in heterogeneous agent models

What should replace the assumption of rational expectations about equilibrium prices? I spell out three criteria that such alternatives should satisfy. A common element in this class of alternative approaches is that decision makers form expectations about prices directly rather than indirectly via forecasting distributions. While this strikes me as the natural solution, any such model of expectations formation necessarily differs from the rational benchmark.²⁴ The key question is how to discipline the subjective probability distributions of equilibrium prices, i.e. how to navigate the “wilderness of non-rational expectations.” The goal of my three criteria is to help navigate this wilderness.

3.1 The natural solution: forecast prices directly

Recall from Section 2 that, with rational expectations, households and firms forecast equilibrium prices by forecasting distributions and solve complex Master equations like (10). The most natural alternative is to instead assume that households and firms forecast prices directly in some way.

To flesh this out in more detail, recall the 2-period Master equation (10) which features an expectation over the cross-sectional distribution G' . The rough idea is that households instead solve

$$\begin{aligned} V_0(a, y, w, r) &= \max_{c, n, a'} U(c, n) + \beta \tilde{\mathbb{E}}[V_1(a', y', w', r') | \cdot] \quad \text{s.t.} \\ c + a' &= w y n + (1 + r) a, \\ V_1(a', y', w', r') &= \max_{c', n'} U(c', n') \quad \text{s.t.} \quad c' = w' y' n' + (1 + r') a', \end{aligned} \tag{11}$$

where the conditional expectation $\tilde{\mathbb{E}}$ is computed using *some* subjective beliefs about future prices (w', r') , i.e. using a conditional probability distribution $\mathbb{P}(w', r' | \cdot)$ that may condition on a number of factors, for example current prices (w, r) or individual factors (so that beliefs are heterogeneous – see below). Having specified such subjective price beliefs, solving the model is not hard: given beliefs solve for actual prices and quantities such that households and firms maximize and markets clear (a *temporary equilibrium* – see Section 4.1 below).

This is true much more generally. Again using the general notation from above with a distribution $G_t(x)$ over a vector of idiosyncratic states x , a vector of aggregate states z , and

²⁴While all alternatives discussed below drop rational expectations about equilibrium prices, for simplicity they all keep rational expectations about non-equilibrium variables, e.g. idiosyncratic y_{it} .

a price vector p , all that is needed is a subjective probability distribution over future prices $\mathbb{P}(p'|\cdot)$. With this probability distribution in hand, one can immediately compute expected values $\tilde{\mathbb{E}}[V(x', z', p')|\cdot]$ and hence solve the Bellman equation for the value function $V(x, z, p)$ as well as solving for the economy's temporary equilibrium given such beliefs. As discussed in Section 2, the vector p may also include other "price-like variables" like the endogenous job finding rate in search-and-matching models. The natural assumption is that decision makers also forecast these other variables directly.

As an aside, it is useful to ask: why can we not just write a Bellman equation (11) with prices p as state variables in the rational expectations case? That is, why can we not just write (11) but with the subjective expectation $\tilde{\mathbb{E}}$ replaced by the rational expectation \mathbb{E} ? The short answer is: because prices p do not follow a Markov process, a point I flesh out in detail in Section 3.4 below.

The focus on a low-dimensional vector of equilibrium variables links the proposed approach to the "sequence space" methods of Auclert et al. (2021). The difference is that my aim is to handle stochastic price sequences outside steady-state neighborhoods, i.e. a global rather than local solution. Handling such sequences in a recursive fashion requires departing from rational expectations precisely because prices do not have the Markov property. It is also worth noting again that the proposed approach in which decision makers forecast prices directly differs not only from the rational-expectations approach (in which decision makers forecast prices indirectly by forecasting cross-sectional distribution) but also from the Krusell-Smith approach (in which decision makers forecast prices indirectly via forecasting moments of the distribution). The one exception is the more promising variant of the Krusell-Smith approach in which these moments are the prices themselves (Section 2.2).

Finally, I wrote the value functions in (11) as featuring only the prices $p = (w, r)$ in addition to the idiosyncratic states $x = (a, y)$. This is more restrictive than necessary. Specifically, the subjective beliefs $\mathbb{P}(p'|\cdot)$ could also condition on (a small number of) other state variables that forecast prices. For example, the percentage of households with mortgages may forecast house prices. The vector of additional states in the value function would then also include these other variables. Decision makers forecasting prices directly means that this vector of forecasting states contains all payoff-relevant prices (but it may contain other variables as well).

3.2 The challenge: disciplining price expectations

The challenge is, of course, how to discipline the subjective probability distributions for future prices $\mathbb{P}(p'|\cdot)$ that are then used to compute price expectations $\tilde{\mathbb{E}}[V(x', z', p')|\cdot]$. This is the challenge of navigating the "wilderness of non-rational expectations." As Sargent (2008) put it: "There is such a bewildering variety of ways to imagine discrepancies between objective and subjective distributions," followed by a pointed footnote "There is an infinite number of ways to be wrong, but only one way to be correct."

To help navigate this wilderness, I next outline three criteria that alternative approaches should fulfil in my view. Sargent's piece advocated "cautious modifications of rational expectations theories [...] to retain the discipline of rational expectations econometrics." The three criteria below

share this goal while, at the same time, aiming to improve the realism and computation of heterogeneous agent models.

3.3 Three criteria for price expectations

The following three criteria are natural requirements on subjective price beliefs \mathbb{P} . I first list the three criteria and then discuss each of them in turn.

1. Simplification of the computational solution
2. Consistency with empirical evidence
3. Endogeneity of beliefs to model reality (Lucas critique)

The challenge posed in this paper is to develop alternatives to rational expectations about equilibrium prices that fulfil these three criteria. The hope is that these criteria *in combination* will provide useful guideposts for navigating the wilderness of non-rational expectations and to considerably narrow down the list of candidate subjective price beliefs \mathbb{P} .

Criterion 1: Simplification of the computational solution. The first criterion is that the candidate alternative approach should actually simplify the computational solution of heterogeneous agent models with aggregate risk (relative to the Master equation approach), i.e. it should make these models “operational.” More precisely, the alternative approach should simplify the problem of individual decision makers and result in Bellman equations with much lower-dimensional state variables.

This criterion actually has quite a bit of bite and rules out (standard versions of) a number of popular expectations models from the literature. In particular it eliminates from the list of candidate subjective price beliefs \mathbb{P} any model that starts from rational expectations as a special case and then obtains the subjective probability distribution by simply “twisting” this objective distribution. This is because the special case with rational expectations causes the extreme curse of dimensionality discussed above and generalizations are necessarily at least as complicated.²⁵

Somewhat more precisely, in the literature, there are a number of models of non-rational expectations in which subjective beliefs are indexed by a parameter θ , i.e. \mathbb{P}^θ , such that $\theta = 0$ is the special case of rational expectations, i.e. $\mathbb{P}^{\theta=0}$ equals the corresponding objective probability distribution \mathbb{P}^{RE} . These include the standard formulations of diagnostic expectations (Bordalo et al., 2018, 2022) and cognitive discounting (Gabaix, 2014, 2020). For example, the standard version of diagnostic expectations would postulate price beliefs \mathbb{P}^θ as

$$\mathbb{P}^\theta(p|D) \propto \mathbb{P}^{RE}(p|D) \left[\frac{\mathbb{P}^{RE}(p|D)}{\mathbb{P}^{RE}(p|-D)} \right]^\theta,$$

²⁵Another class of models that does not satisfy my Criterion 1 for essentially the same reasons is the Limited Information Rational Expectations (LIRE) class, i.e. models that relax only the FI part but not the RE part of Full Information Rational Expectations (FIRE). LIRE includes noisy rational expectations (Lucas, 1972) and rational inattention (Sims, 2003; Maćkowiak et al., 2023).

where D is recently observed data, $-D$ is data in a relevant comparison group, and where $\theta \geq 0$. When $\theta = 0$ the special case of rational expectations obtains. Similarly, the cognitive discounting approach of [Gabaix \(2020\)](#) postulates subjective expectations as $\tilde{\mathbb{E}}[p] = \tilde{\mathbb{E}}[\bar{p} + \hat{p}] = \bar{p} + \bar{m}\mathbb{E}^{RE}[\hat{p}]$, where \bar{p} is a default value (e.g. the steady state), \hat{p} the deviation from this default value, and where $\bar{m} \in [0, 1]$ is a “cognitive discounting” parameter with $\bar{m} = 1$ corresponding to rational expectations (equivalently, $\theta = 1 - \bar{m}$ with rational case $\theta = 0$). A third example is the recursive competitive equilibrium with behavioral agents in Section 5.3 of [Gabaix \(2023\)](#) which requires computation of the rational expectations equilibrium prices whenever attention $m > 0$. These models do not satisfy my first criterion.²⁶

Put differently, for the purposes of this paper, models of non-rational expectations that require decision makers to be able to compute the special case of rational expectations defeat the point of departing from rational expectations in the first place. This is true not just from this paper’s pragmatic, computational viewpoint but arguably also from that of the behavioral macroeconomics literature: while such models have been extremely useful for exploring the implications of deviating from the rational expectations hypothesis (see e.g. the papers above), they are an intermediate step toward more satisfactory theories of bounded rationality that do not require decision makers to access the full rational expectations solution.

Criterion 2: Consistency with empirical evidence. The second criterion is that the candidate alternative approach should be consistent with what is known from the empirical expectations literature. Over the last twenty years, this literature has assembled a large body of evidence, in particular by measuring expectations through surveys (e.g. [Manski, 2004](#); [Mankiw et al., 2004](#); [Armantier et al., 2013](#); [Manski, 2018](#); [Coibion et al., 2022](#); [D’Acunto and Weber, 2024](#)). There are now new surveys like the New York Fed’s Survey of Consumer Expectations (SCE) and even a Handbook of Economic Expectations ([Bachmann et al., eds, 2023](#)).

This literature has documented a number of empirical patterns, some of which I summarize in Section 4.2 below. While survey expectations data need to be interpreted with care (e.g. due to concerns about systematic noise), candidate alternative approaches should take on board this empirical evidence. For example, one important finding in the survey expectations literature is the vast observed heterogeneity of measured beliefs which stands in contrast to the “communism” of rational expectations ([Sargent, 2017](#)), i.e. that all decision makers share the same beliefs. This suggests developing alternatives to rational expectations in which subjective price beliefs differ in the population, that is, modeling subjective probability distributions \mathbb{P}_i where i indexes individual households or firms. It is worth noting that the tools of heterogeneous agent macroeconomics are uniquely suited to modeling belief heterogeneity: after all, modeling and “carrying around” distributions is our bread and butter – so why not add a

²⁶That being said, variants of the standard formulations may well satisfy Criterion 1. For example, both diagnostic expectations and cognitive discounting are specified *relative* to some other beliefs and these other beliefs need not be rational ([Hajdini, 2023](#)). In the equation above \mathbb{P}^{RE} could be replaced by some alternative baseline \mathbb{P} . But this leaves open the question what that alternative baseline \mathbb{P} should be. Similarly, one can imagine a variant of the [Gabaix \(2023\)](#) recursive competitive equilibrium with behavioral agents in which decision makers use a simplified model for equilibrium prices (e.g. an autoregressive process) which endogenously becomes more sophisticated when the stakes increase or the attention cost falls.

distribution of beliefs?²⁷

In economics there is often a trade-off between realism and (computational) simplicity: the real world is a complex place so more of one usually means less of the other. One may have therefore conjectured that Criteria 1 and 2 are in conflict with each other. But this is likely not the case. Instead realism and simplicity may go hand in hand in this case, offering an opportunity to kill two birds with one stone.

There are three other types of empirical evidence that models employing the candidate alternative approach may want to match. First, evidence on the pass-through from subjective beliefs to decision makers' choices (which is what we ultimately care about). Second, evidence on the pass-through from shocks and policies to subjective beliefs (Section 4.2). Third, the time-series properties of observed prices and price-like variables (see the discussion of Criterion 3).

Criterion 3: Endogeneity of beliefs to model reality (Lucas critique). The third criterion is that the candidate alternative approach should imply price beliefs that are endogenous to the behavior of actual equilibrium prices. It is useful to split Criterion 3 into two sub-criteria:

Criterion 3a: In stationary environments with a sufficiently long time series for prices, subjective price beliefs should be approximately consistent with actual (model) equilibrium prices.

Criterion 3b: In all environments, including unfamiliar and non-stationary ones, price beliefs should respond to model reality, specifically to policy changes (Lucas critique).

I now discuss these two sub-criteria in turn. Criterion 3a is concerned with familiar environments in which equilibrium prices follow “regular patterns.” In such situations, price beliefs should not be “too far” from the behavior of actual equilibrium prices. Denoting the subjective probability distribution of prices by $\mathbb{P}(p|\cdot)$ and the objective price distribution by $\mathbb{P}^{obj}(p|\cdot)$, we should have

$$\|\mathbb{P}(p|\cdot) - \mathbb{P}^{obj}(p|\cdot)\| \leq \varepsilon \quad (12)$$

for some distance metric (norm) $\|\cdot\|$ and some “small” scalar $\varepsilon \geq 0$. While rational expectations imposes exact consistency between subjective and objective distributions, $\mathbb{P}^{RE}(p|\cdot) = \mathbb{P}^{obj}(p|\cdot)$ identically, Criterion 3a imposes approximate consistency.

The rationale for Criterion 3a is already in the excerpts from Lucas (1980) and Lucas (1996) in Section 1.2, in particular that “the implied behavior of actual equilibrium prices [should not be] grossly inconsistent with the price expectations that the theory imputes to individual agents.” Without this criterion, alternative approaches would go full circle back to the 1960s. The argument here is that weaker restrictions than rational expectations can also deliver on Lucas' desiderata. This criterion is nothing new and has been advocated by other proponents of (cautious) departures from rational expectations. For example, Woodford (2013) writes: “It makes sense to assume that expectations should not be completely arbitrary and have no relation to the kind of world in which

²⁷Guerreiro (2023) takes a step in this direction but linearizes his model. This simplifies the analysis because he only has to track a distribution of mean beliefs (expectations) across (a relatively small number of) household groups. As noted elsewhere, linearized models are not suitable for modeling phenomena like financial crises. One very tractable but also specific approach to modeling and tracking a distribution of beliefs is the “sticky information” approach of Mankiw and Reis (2002). Also see Carroll (2003), Carroll et al. (2020), and Auclert et al. (2020).

the agents live [...] We should therefore like to replace the RE hypothesis by some weaker restriction that nonetheless implies a substantial degree of conformity between people's beliefs and reality." Also see [Sargent \(1993\)](#), [Sargent \(1999\)](#), and [Ilut and Schneider \(2014\)](#) who impose an approximate consistency criterion similar in spirit to (12) for worst- and best-case means in a model with ambiguity aversion.

The role of Criterion 3b is that Criterion 3a may be too strong in some contexts, specifically in "unfamiliar" situations that have "not been seen before" like the Covid pandemic, the 9/11 terrorist attacks, or the 2000s housing boom. In such contexts, even the best econometricians would struggle to closely approximate the true data-generating process for equilibrium prices. Therefore, price beliefs of decision makers in our models may differ from the true process by a wider margin. This idea is consistent with experimental evidence on reasoning in games that beliefs differ more from rational beliefs in new environments than in familiar ones (e.g. [Nagel, 1995](#)). Similarly, [Angeletos and Lian \(2023\)](#) write: *"The assumption that the subjective statistical model in people's mind is the same as the objective truth, is hard to defend in the context of unusual circumstances, such as the Great Recession and the ZLB on monetary policy."*

Nevertheless, because actual equilibrium prices are endogenous, price beliefs should not be completely exogenous to the model. This is the content of Criterion 3b. For example, when conducting policy counterfactuals, one cannot simply plug empirical survey expectations into the model and compute the corresponding temporary equilibria. Instead, price beliefs need to be modeled as endogenous. The behavior of actual equilibrium prices depends on price beliefs and, conversely, price beliefs depend on the behavior of actual prices – a fixed point problem. Concretely, we need to model how expectations respond to policy changes. If beliefs were completely unresponsive to policy, the model would be subject to the [Lucas \(1976\)](#) critique.

For both Criteria 3a and 3b, an important question is how to evaluate the size of "mistakes", i.e. what distance metric between subjective and objective price distributions to use. A natural approach is to use a welfare-based metric, i.e. to impose that mistakes should not be too costly in terms of welfare. Also this idea is not new. [Lucas and Prescott \(1971\)](#) already criticize models in which the difference between price forecasts and actual prices is *"persistent [and] costly to forecasters"* in the passage quoted in Section 1.3. The idea to judge mistakes in terms of welfare costs similarly lies at the heart of models of both rational and behavioral inattention (but also see the discussion surrounding Criterion 2).

Finally, if we want to impose some consistency between beliefs and model reality (Criterion 3) while, at the same time, imposing consistency with empirical evidence on beliefs about real-world prices (Criterion 2), there should also be some consistency between model-generated and real-world prices. That is, our heterogeneous-agent model should generate empirically reasonable time-series behavior for equilibrium prices.

To describe the philosophy of rational expectations, [Sargent \(2007\)](#) has used Abraham Lincoln's famous phrase: *"You can fool some of the people all of the time, and all of the people some of the time, but you cannot fool all of the people all of the time."* Criterion 3 aims to ensure that this remains true. In fact, Lincoln's quote leaves room for people getting fooled *some* of the time. It is therefore arguably *more* consistent with deviations from rational expectations of the type

advocated here than with rational expectations themselves.

To deliver on both parts of Criterion 3, an obvious approach is to model some form of learning. Models of learning have been a key focus of the existing literature and will occupy a large part of the discussion of promising directions in Section 4 below.

Another requirement: compatibility with non-linear models and recursive methods. As stated above, one of the ultimate goals of developing alternative models of price beliefs is to be able to consider questions in which aggregate risk and aggregate non-linearities play a key role. Therefore, the candidate alternative approach should be applicable in non-linear models. In contrast, much of the theoretical expectations literature has considered linear (or linearized) models, likely because it means that standard techniques like the Kalman filter can be applied. This would need to change.

Similarly, it is crucially important that the alternative approach be compatible with recursive solution methods for household and firm optimization problems (Bellman equations). These lie at the heart of solution methods for current heterogeneous-agent models for two reasons. First, micro non-linearities like borrowing constraints are often important and recursive methods are the natural solution methods for handling these. Second, rather than solving a separate optimization problem for each of a large number of heterogeneous decision makers (i.e. millions of optimization problems), recursive methods mean that we often solve only one single Bellman equation with the resulting policy function summarizing the actions of (ex-post) heterogeneous decision makers at different points of the state space.

Because this is a requirement on the class of models that the approach can be applied to rather than a constraint on subjective price beliefs \mathbb{P} themselves, I did not include it as a fourth criterion in the list above. But, similarly to the three criteria, this requirement should hopefully serve to narrow down the list of applicable alternative approaches.

3.4 With rational expectations, why can we not carry prices as state variables?

Before discussing promising directions, I briefly revisit the question: why can we not just write a Bellman equation (11) with prices p as state variables in the rational expectations case? That is, why can we not just write (11) but with the subjective expectation $\tilde{\mathbb{E}}$ replaced by the rational expectation \mathbb{E} ?

The short answer is that prices do not satisfy the Markov property. As a reminder, a stochastic process has the Markov property if the probability distribution of future states depends only on the present state. For a discrete-time process $\{p_t\}$, the Markov property would mean that one can write the conditional distribution of p_{t+1} as a function of p_t only: $\Pr(p_{t+1}|\cdot) = \Pr(p_{t+1}|p_t)$.

To see that prices do not satisfy this property, it is useful to consider a variant of the general model in Section 2.2 with a distribution over the vector of idiosyncratic states x , aggregate states z , and a price vector p , but with a finite state space for the idiosyncratic states, i.e. x can take only $N < \infty$ possible values, $x \in \{x_1, \dots, x_N\}$, say $N = 100,000$. In this case, the distribution G_t is simply an N -dimensional vector $\mathbf{G}_t = [G_{1,t}, \dots, G_{N,t}]^T$ and it is convenient

to work with the corresponding density $\mathbf{g}_t = [g_{1,t}, \dots, g_{N,t}]^T$ which is simply the “histogram” collecting the fraction of agents at each point of the state space.

As above, the vector of equilibrium prices is a function of the entire distribution:

$$p_t = \mathcal{P}^*(\mathbf{g}_t, z_t), \quad (13)$$

Note that I use regular notation for the vector p_t but bold-face notation for the vector \mathbf{g}_t to signify that p_t is typically much lower-dimensional than \mathbf{g}_t , say $p_t \in \mathbb{R}^2$ versus $\mathbf{g}_t \in \mathbb{R}^{100,000}$.

To see why prices do not satisfy the Markov property, consider the evolution of \mathbf{g}_t . This density evolves according to a Chapman-Kolmogorov equation (the discrete-time analogue of a Kolmogorov-Forward equation):

$$\mathbf{g}_{t+1} = \mathbf{A}(\mathbf{g}_t, z_t, p_t)^T \mathbf{g}_t \quad (14)$$

where $\mathbf{A}(\mathbf{g}_t, z_t, p_t)$ is the transition matrix of the vector of idiosyncratic states x (with x evolving according to the optimal policies from agents’ Bellman equation).²⁸ Note that this is a large matrix of dimension $N \times N$, say with $N = 100,000$, that summarizes how all individual decision makers transition across the economy’s state space. Importantly, after substituting in for p_t from (13), this equation defines a *Markov process for the distribution \mathbf{g}_t and aggregate shock z_t* : we have $\Pr(\mathbf{g}_{t+1}, z_{t+1} | \cdot) = \Pr(\mathbf{g}_{t+1}, z_{t+1} | \mathbf{g}_t, z_t)$.

In summary, prices p_t depend on the distribution and aggregate shock (\mathbf{g}_t, z_t) according to the equilibrium price function \mathcal{P}^* in (13) and (\mathbf{g}_t, z_t) follows a Markov process. This immediately implies that prices p_t themselves do *not* follow a Markov process. Instead, the stochastic process for p_t is extremely complicated. More precisely, the conditional distribution of future prices satisfies $\Pr(p_{t+1} | \cdot) = \Pr(p_{t+1} | \mathbf{g}_t, z_t)$ rather than $\Pr(p_{t+1} | \cdot) = \Pr(p_{t+1} | p_t)$ meaning that the probability distribution of future prices depends on the entire cross-sectional distribution \mathbf{g}_t . Therefore, one cannot simply write a Bellman equation (11) with prices p as state variables. Instead, decision makers with rational expectations forecast the Markov state (\mathbf{g}_t, z_t) in order to forecast the non-Markovian prices p_t .

Another way of looking at things is in analogy with a *Hidden Markov Model* (HMM): the distribution \mathbf{g}_t is a high-dimensional latent state that does satisfy the Markov property and p_t is a low-dimensional observation that is a function of this Markov state. Then the observation p_t violates the Markov property and forecasting it requires forecasting the high-dimensional Markov state.

Knowledge of this aspect of the structure of heterogeneous-agent models should be useful for developing alternatives to the rational expectations approach. For example, assuming that decision makers do not track the high-dimensional cross-sectional distribution and instead only track prices (as proposed in Section 3.1) turns their optimal control problems into *Partially Observable Markov Decision Process* (POMPDs, i.e. controlled HMMs).

²⁸See Achdou et al. (2021) for a continuous-time analogue and an explanation why the Kolmogorov Forward (KF) equation features the transpose of the transition matrix. Intuitively, the transition matrix \mathbf{A} answers the question “where in the state space are people going?” but the KF equation asks “where in the state space are people coming from?” – hence the transpose.

4 Some promising directions

In the final section I discuss some potentially promising directions for developing alternative models of price beliefs satisfying my three criteria. These include temporary equilibrium and internal rationality (section 4.1), incorporating survey expectations (section 4.2), least-squares learning and stochastic approximation (section 4.3), reinforcement learning (section 4.4), and heuristics and simple models (section 4.5). The discussion here is purposely selective. For more complete treatments, see the surveys and books referenced in the introduction.

4.1 Temporary Equilibrium and Internal Rationality

The idea of “temporary equilibrium” is fundamental for developing alternatives to rational expectations about equilibrium prices: all approaches in which decision makers forecast prices directly rather than indirectly via forecasting distributions (Section 3.1) will necessarily involve computing such temporary equilibria as an intermediate step.

Definition: *Temporary equilibrium* at a particular time t is defined as allocations and prices such that (i) households and firms optimize *given expectations of future variables* (including future prices) that are specified in the model but that are *not necessarily rational*, (ii) markets clear at time t .

This idea was originally developed contemporaneously by Hicks (1939) and Lindahl (1939), has been further developed by Grandmont (1977, 1989), and has found a number of applications in the more recent literature.²⁹ Rational expectations equilibrium is the special case of temporary equilibrium in which expectations are rational (model-consistent).³⁰

The approach outlined in Section 3.1 was exactly one of temporary equilibrium: given specified beliefs $\mathbb{P}(p'|\cdot)$, compute value functions $V(x, z, p)$ and associated optimal policies, and solve for equilibrium prices that clear markets. It is important to note that computation of temporary equilibria given price beliefs is relatively straightforward with modest computational costs. In particular, in a temporary equilibrium, there is no fixed point between price beliefs and actual equilibrium prices. Therefore, its computation is only slightly more difficult than the computation of a stationary equilibrium in a heterogeneous-agent model (as in Aiyagari, 1994). In this regard, it is also worth restating that tracking the distribution $G_t(x)$ over time, i.e. computing $G_{t+1}(x)$ given $G_t(x)$, a policy function for the evolution of x , and the aggregate shock z_t is computationally straightforward. For example, in the finite-state-space variant of Section 3.4 tracking the vector \mathbf{g}_t simply means running the difference equation (14) forward in time. Heterogeneous-agent models become unrealistically complex only when the distribution enters decision makers’ expectations.

While computing temporary equilibria is relatively straightforward, to fulfil Criterion 3

²⁹Including Woodford (2013), Piazzesi and Schneider (2016), García-Schmidt and Woodford (2019), Farhi and Werning (2019), and Werning (2022).

³⁰Hicks differentiates “temporary equilibrium” from another notion of equilibrium he calls “equilibrium over time” (which bears some resemblance to a rational-expectations equilibrium) and explains the difference in a dynamic economy with time intervals lasting one week. See Hicks (1939, Chapter X.2): “The wider sense of Equilibrium – Equilibrium over Time, as we may call it, to distinguish it from the Temporary Equilibrium which must rule within any current week – suggests itself when we start to compare the price-situations at any two dates. [...]” Also see Chapter IX.7.

above, this can only be an intermediate step toward solving the bigger fixed point problem in which the behavior of actual equilibrium prices feeds back into endogenous price beliefs.

Internal Rationality. When defining temporary equilibria, it is important to derive individual policy functions from intertemporal decision problems with dynamically consistent subjective beliefs about equilibrium prices. This links temporary equilibrium approaches to the concept of *internal rationality* proposed by Adam and Marcet (2011).

Adam and Marcet distinguish between variables under decision makers' control and those that are beyond their control or "external," with equilibrium prices being the prime example of such external variables. They argue that the correct way of relaxing rational expectations is to start with a well-defined system of (non-rational) subjective beliefs about these external variables and to then derive decision makers' policy functions from their intertemporal problems given such beliefs. They term this approach *internal rationality* to emphasize that decision makers optimize but without being "externally rational," i.e. without knowing the objective probability distributions of external variables. Internal rationality contrasts with an alternative approach that starts from individual optimality conditions derived under rational expectations and then simply replaces the rational expectations operators with some other expectations, an approach that can result in inconsistencies (Preston, 2005).

The temporary equilibria discussed above should thus be "Internally Rational Expectations Equilibria" in Adam and Marcet's language. Starting from Bellman equations with specified subjective beliefs will automatically result in internal rationality. Because this is the natural solution method for heterogeneous agent models, the inconsistencies just discussed should not be a concern in practice.

Self-Confirming Equilibrium and Friends. The economics literature has proposed various concepts of "non-rational expectations equilibrium" or "misspecification equilibrium" in which actual equilibrium outcomes are statistically consistent with decision makers' beliefs (i.e. these beliefs are not disappointed) but with weaker or different consistency requirements than under rational expectations. These alternative equilibrium concepts may be useful building blocks for fulfilling Criterion 3a ("approximate consistency between beliefs and model reality"). In contrast to temporary equilibria in which beliefs are exogenously specified, in these alternative equilibria, beliefs are endogenous and are the fixed point of some mapping from a (potentially misspecified) perceived law of motion to an approximate actual law of motion.

Appendix A summarizes a number of these, specifically: self-confirming equilibrium (SCE), restricted perceptions equilibrium (RPE), and consistent expectations equilibrium (CEE). The appendix also puts these equilibrium concepts in relation to temporary equilibrium (TE), internally rational expectations equilibrium (IREE), and rational expectations equilibrium (REE). As explained there, the relation between the various equilibrium concepts can be summarized as follows:

$$\{REE, RPE, CEE\} \subset SCE \subset IREE \subset TE, \quad (15)$$

where \subset means "is a special case of."

4.2 Survey Expectations and Hypothetical Vignettes

As already discussed, to satisfy Criterion 2 (“consistency with empirical evidence”), alternative approaches should incorporate the findings from the large empirical literature on expectations formation including evidence from survey expectations.

This literature has documented a number of empirical patterns that challenge the rational expectations hypothesis. As already noted, the most obvious of these is the pervasive heterogeneity of subjective beliefs or “belief disagreement.” While some of this disagreement may reflect noise (see below), this finding stands in contrast to the “communism of beliefs” of rational expectations. See for example [Mankiw et al. \(2004\)](#), [Malmendier and Nagel \(2011\)](#), [Malmendier and Nagel \(2015\)](#), [Giglio et al. \(2021\)](#), and [Fofana et al. \(2024\)](#). Similarly, a large literature has documented overreaction of individual-level forecasts to idiosyncratic news and underreaction of average forecasts to aggregate news ([Coibion and Gorodnichenko, 2012, 2015](#); [Bordalo et al., 2022](#)). An example of another interesting finding is “cross-domain extrapolation”.³¹

To empirically discipline subjective beliefs about equilibrium prices, survey evidence on the relevant expectations is needed. These will differ according to the particular model being considered. In the simple heterogeneous-agent model in Section 2 one would need survey expectations of future (real) wages and interest rates. In this regard measures of subjective earnings risk are an important piece of evidence, see for example [Dominitz and Manski \(1997\)](#) and [Caplin et al. \(2023\)](#).

Additionally, one would really need measures of these expectations *conditional* on the state of the aggregate economy (recessions vs booms) and on different policies (policy counterfactuals). Put differently, one would like evidence on the “subjective models” (perceived laws of motion) that decision makers use to form expectations. One way of eliciting such conditional beliefs using surveys is to use hypothetical vignettes ([Haaland et al., 2023](#); [Andre et al., 2022](#)).

An important concern about survey expectations is that a substantial fraction of reported beliefs may reflect “noise”, some of it classical measurement error but some of it also systematic error due to respondents’ cluelessness about the object they are asked to forecast. A common finding in the literature is that reported subjective probabilities are compressed toward 50:50 relative to true objective probabilities, i.e. survey respondents overstate the probability of unlikely events and understate the probability of likely events (e.g. [Viscusi, 1985](#); [Fischhoff and Bruine De Bruin, 1999](#); [Enke and Graeber, 2023](#)). [Enke and Graeber \(2023\)](#) elicit survey respondents’ subjective beliefs and additionally ask respondents how certain they are about these stated beliefs (“cognitive uncertainty”). They find that more cognitively uncertain respondents have more compressed beliefs, i.e. that cognitive uncertainty systematically distorts stated beliefs. Therefore heterogeneity in survey expectations may partly reflect heterogeneity in “cognitive noise” rather than true beliefs.³² This concern about systematic noise in survey

³¹For example, [Cenzon \(2024\)](#) documents that people experiencing credit rejections become pessimistic not only about credit markets but also about inflation, unemployment, and stock prices and [Taubinsky et al. \(2024\)](#) argue that inflation expectations react excessively to household-level income shocks. Also see [Bordalo et al. \(2024\)](#). This evidence, as well as the evidence on subjective earnings risk in the next paragraph, is also of independent interest for heterogeneous-agent modeling, for example it could have important implications for the distribution of wealth.

³²Put differently, the concern about systematic distortions is that cognitively uncertain respondents may appear

expectations is not an argument against the use of such data; rather it calls for interpreting such data with care, particularly in contexts in which cognitive noise may be important, for example when households are asked about objects that are not relevant to them on a daily basis. Another important piece of the puzzle is the pass-through from survey beliefs to actions, see [Giglio et al. \(2021\)](#), [Charles et al. \(2024\)](#), [Yang \(2024\)](#).

While the present section focussed on household expectations, firm expectations about the conditions they operate in are, of course, equally important. See for example [Bachmann et al. \(2013\)](#), [Coibion et al. \(2023\)](#) and [Bloom et al. \(2024\)](#).

Temporary Equilibrium with Measured Expectations. A particularly interesting approach is to compute temporary equilibria in which beliefs are disciplined with survey expectations. This “temporary equilibrium with measured expectations” approach is described in [Piazzesi and Schneider \(2016\)](#) and in Section 2.4 of [Brunnermeier et al. \(2021\)](#), written by Piazzesi). See [Piazzesi et al. \(2015\)](#), [Leombroni et al. \(2020\)](#), [Aron-Dine et al. \(2024\)](#), [Bardóczy and Guerreiro \(2023\)](#), and [Ludwig et al. \(2024\)](#) for recent applications.

[Roth et al. \(2023\)](#) extend this approach to computing policy counterfactuals by means of measured expectations under different counterfactuals, elicited using hypothetical vignettes. While this approach could, in principle, be used to address the critique that beliefs should respond to policy (Lucas critique), a downside is that a new vignette question would need to be fielded for every new policy counterfactual. The more standard alternative approach is to model expectations in a way that incorporates the relevant empirical evidence (e.g. from survey expectations). Vignette proponents argue in favor of their approach by pointing to the falling cost of conducting surveys and economists’ disagreement about modeling expectations.

4.3 Least-squares learning and stochastic approximation

As noted in Section 3.3, an obvious approach to deliver on Criterion 3 is to model some form of learning. One approach that has a long tradition in the economics literature is *least-squares learning*. See for example [Bray \(1982\)](#), [Marcet and Sargent \(1989\)](#), [Sargent \(1993\)](#), chapter 5), [Evans and Honkapohja \(2001\)](#), and [Evans and McGough \(2020\)](#). Least-squares learning is a special case of a more general set of *stochastic approximation methods* (e.g. [Robbins and Monro, 1951](#); [Ljung, 1977](#)) which I also discuss below.³³ Reinforcement learning, which I review in Section 4.4, is another stochastic approximation method.

Least-squares learning. I review least-squares learning by means of a specific example: learning about equilibrium prices in the simple demand-supply system of [Muth \(1961\)](#) which we

either very optimistic or very pessimistic depending on the particular survey question being asked. For example, asking “What is the probability that inflation will exceed 10% over the next year?” (a highly unlikely event), will lead cognitively uncertain households to state a high probability unreasonably close to 50:50 and these households would thus appear very pessimistic about inflation. Conversely, asking “What is the percent chance that inflation will exceed 0%?” (a very likely event), will make them appear very optimistic.

³³[Sargent \(2024\)](#) links stochastic approximation back to [Friedman and Savage \(1947\)](#) and [Hotelling \(1941\)](#).

already encountered in Section 1.1.³⁴ Recall that production takes place with a one-period lag, that supply is an increasing function of the expected (log) price $p_t^e = \mathbb{E}_{t-1}[p_t]$ and of a supply shock u_t , and that equilibrium prices and quantities are determined by the intersection of demand and supply – see (1). As a result, equilibrium prices depend on price beliefs – see (2).

Instead of assuming that producers have rational price expectations about next period's equilibrium price, we now assume that they form it based on past data. That is, producers form estimates $\hat{p}_{t+1}^e = \hat{\mathbb{E}}_t[p_{t+1}]$ from observations of past prices p_1, \dots, p_t and supply shock realizations u_1, \dots, u_t via least squares.

The simplest case is when the supply shock u_t is serially *uncorrelated* so that current and past supply conditions are uninformative about future prices. I cover this simple case first and then return to the more general case momentarily. When supply conditions u_t are serially uncorrelated, producers do not take past u_t 's into account. Instead, they simply form the expectation p_{t+1}^e by taking the sample average of past price observations

$$\hat{p}_{t+1}^e = \frac{1}{t} \sum_{s=1}^t p_s.$$

To link this to least-squares estimation, consider a linear regression model with only an intercept and no independent variable: the sample average is then the least-squares estimate of the intercept parameter. Computation of this sample average can also be implemented recursively as follows:³⁵

$$\hat{p}_{t+1}^e = \hat{p}_t^e + \frac{1}{t} [p_t - \hat{p}_t^e] \quad (16)$$

Notice how this update rule computes the new estimate \hat{p}_{t+1}^e from the previous estimate \hat{p}_t^e and the latest data p_t . This recursive implementation of the sample average is the special case of recursive least squares in a regression model with only an intercept.

More generally, when supply conditions u_t are serially correlated, producers use past observations of u_t to forecast prices. To this end, assume that $u_{t+1} = \rho u_t + \varepsilon_{t+1}$ where the correlation coefficient ρ is unknown to producers and ε_{t+1} is a mean-zero i.i.d. random variable. In this more general case, least-squares learning theory assumes that producers use a log-linear *perceived law of motion*

$$p_{t+1} = \theta_0 + \theta_1 u_t + \varepsilon_{t+1},$$

where θ_0 and θ_1 are unknown parameters. That is, producers conjecture that the current supply shock u_t forecasts next period's u_{t+1} and hence next period's price p_{t+1} but they do not know the strength of these relationships. To form price expectations $\hat{p}_{t+1}^e = \hat{\mathbb{E}}_t[p_{t+1}]$, producers then (i) estimate (θ_0, θ_1) from observations of past prices p_1, \dots, p_t and supply shock realizations u_0, \dots, u_{t-1} via least squares, and (ii) use the time- t estimates $(\hat{\theta}_{0t}, \hat{\theta}_{1t})$ to compute $\hat{p}_{t+1}^e = \hat{\mathbb{E}}_t[p_{t+1}] = \hat{\theta}_{0t} + \hat{\theta}_{1t} u_t$.

³⁴While I use somewhat different notation, this section follows closely the presentation in [Evans and Honkapohja \(2001\)](#) who also start from Muth's example. Also see the presentation of [Bray \(1982\)](#) in chapter 5 of [Sargent \(1993\)](#) which starts directly from the reduced form (2), [Christiano et al. \(2024\)](#), and [Baley and Veldkamp \(2025\)](#).

³⁵The derivation is $\hat{p}_{t+1}^e = \frac{1}{t} \sum_{s=1}^t p_s = \frac{1}{t} \left[p_t + \sum_{s=1}^{t-1} p_s \right] = \frac{1}{t} (p_t + (t-1)\hat{p}_t^e) = \hat{p}_t^e + \frac{1}{t} [p_t - \hat{p}_t^e]$.

Writing the perceived law of motion as $p_t = x_t^T \theta + \varepsilon_t$ where $x_t = [1 \ u_{t-1}]^T$ and $\theta = [\theta_0 \ \theta_1]^T$, the time- t least-squares estimate of the parameter vector θ is given by

$$\hat{\theta}_t = \left[\sum_{s=1}^t x_s x_s^T \right]^{-1} \left[\sum_{s=1}^t x_s p_s \right].$$

Analogously to (16), there is again a recursive implementation that computes the new estimate $\hat{\theta}_{t+1}$ from the previous estimate $\hat{\theta}_t$ and the latest data (p_{t+1}, u_t) – see for example equation (2.9) in [Evans and Honkapohja \(2001\)](#). This recursive implementation is known as *recursive least squares*. In summary, least-squares learning means that decision makers update parameter estimates for a perceived law of motion using recursive least squares.

In heterogeneous-agent models, least-squares learning could be applied to decision makers learning perceived laws of motion of equilibrium prices. A paper by [Nakov and Nuño \(2015\)](#) goes in this direction and models overlapping generations of investors with heterogeneous beliefs learning about equilibrium stock prices using a simple variant of least-squares learning (learning from experience as in [Malmendier and Nagel, 2011, 2015](#)).

Related, there is a link between least-squares learning and Krusell-Smith methods (see Section 2.2).³⁶ In both approaches, decision makers use a (typically log-linear) perceived law of motion and estimate its coefficients via least squares. A difference is that least-squares learning implements this coefficient estimation recursively so that solving for equilibrium and least-squares estimation (belief updating) are done “in one sweep.” This also means that coefficient estimates vary over time (which they do not in Krusell-Smith) and should make least-squares learning more likely to generate interesting phenomena like booms and busts.

Stochastic approximation. The update rule (16) for the sample-average case and the recursive least-squares update rule for $\hat{\theta}_t$ for the more general case take a common form that also arises in a number of other problems (in the next subsection for example). Using the computer science notation to let \leftarrow denote an assignment statement, [Sutton and Barto \(2018, chapter 2.4\)](#) write the general form for this type of update rule as

$$\text{NewEstimate} \leftarrow \text{OldEstimate} + \text{StepSize} [\text{Target} - \text{OldEstimate}]. \quad (17)$$

The “Target” indicates the direction in which the estimate is updated. For example, in the sample-average case (16), the target is the time- t price p_t and the step size is $1/t$ so that updates become smaller and smaller over time. This type of iterative update method is known as a *stochastic approximation method* (e.g. [Robbins and Monro, 1951](#); [Ljung, 1977](#)).

Casting least-squares learning as a stochastic approximation method is useful because convergence results exist for the latter. These results answer questions about the convergence of the least-squares coefficient estimates $\hat{\theta}_t$ as $t \rightarrow \infty$ (as data becomes more and more plentiful). To this end, denote the step size in the update rule (17) by α_t , for example $\alpha_t = 1/t$ in the

³⁶When discussing least-squares learning, [Sargent \(2017\)](#) states: “This is in effect what [Krusell and Smith \(1998\)](#) do, though they do not connect their method to the learning literature.”

sample-average case (16).³⁷ One key result in stochastic approximation theory states that, as $t \rightarrow \infty$, the estimate being updated according to (17) converges with probability 1 under a number of conditions including the following conditions on the step size:³⁸

$$\sum_{t=0}^{\infty} \alpha_t = \infty \quad \text{and} \quad \sum_{t=0}^{\infty} \alpha_t^2 < \infty. \quad (18)$$

Note that both convergence conditions are met for the sample-average case in which $\alpha_t = 1/t$. More specifically, under these conditions, as $t \rightarrow \infty$, the stochastic trajectory of the estimate satisfying the update rule (17) becomes closer and closer to the deterministic trajectory of a particular ordinary differential equation (ODE) derived from this update rule. If this ODE has a locally stable steady state, the estimate converges to a constant that equals this steady state value. See Ljung (1977) as well as the discussions in Marcet and Sargent (1989), Woodford (1990), Evans and Honkapohja (2001), and Sargent (1999).

Using stochastic approximation theory, the literature has shown that, under certain conditions, the equilibrium with least-squares learning converges to the rational expectations equilibrium (Marcet and Sargent, 1989; Woodford, 1990). Evans and Honkapohja (2001) dubbed these conditions “expectational (E) stability.” For example, in the simplest case of the Muth model without serial correlation, it is intuitive that the backward-looking sample average (recursive least-squares expectation) $\hat{\mathbb{E}}_t[p_{t+1}] = \frac{1}{t} \sum_{s=1}^t p_s$ converges to the rational expectation $\mathbb{E}_t[p_{t+1}]$ as $t \rightarrow \infty$. Similar results hold when there is serial correlation or for more general models, though these results require relatively strong assumptions and can be fragile.

As discussed in Section 3.3, to actually simplify computations in economies with heterogeneity (Criterion 1), the perceived law of motion of equilibrium prices would need to be “restricted” to *not* start from rational expectations as a special case. This implies that the least-squares learning process cannot converge to the rational expectations equilibrium. But it may instead converge to a “restricted perceptions equilibrium” (Appendix A) which may therefore be a relevant equilibrium concept in heterogeneous-agent models.

Slow learning. An important caveat is that stochastic approximation convergence theorems only hold asymptotically as $t \rightarrow \infty$ and that learning may be slow in practice. See Christiano et al. (2024) and Sargent (1993) who voices reservations about adaptive algorithms as theories of real-time dynamics but states that he nevertheless “like[s them] as devices for selecting equilibria.”

³⁷Considering more general step sizes links the update rule (16) to adaptive expectations. For a general step size α_t the update rule becomes $\hat{p}_{t+1}^e = \hat{p}_t^e + \alpha_t [p_t - \hat{p}_t^e]$. In the special case of a constant step size $\alpha_t = \alpha$, this becomes the adaptive expectations rule $\hat{p}_{t+1}^e = \hat{p}_t^e + \alpha [p_t - \hat{p}_t^e]$ which implies $\hat{p}_{t+1}^e = (1 - \alpha)^t p_1 + \alpha \sum_{j=0}^t (1 - \alpha)^j p_{t-j}$, a backward-looking weighted average of past prices with exponentially declining weights.

³⁸As Sutton and Barto (2018) explain: “The first condition is required to guarantee that the steps are large enough to eventually overcome any initial conditions or random fluctuations. The second condition guarantees that eventually the steps become small enough to assure convergence.”

4.4 Reinforcement Learning

Reinforcement learning (RL) is learning value functions of incompletely-known Markov decision processes. Decision makers do not know the exact environment they are operating in (which may be extremely complex) and instead learn optimal policies from experience. RL is at the core of some impressive advances in artificial intelligence, e.g. learning to play Go and Atari games better than humans (Mnih et al., 2015; Silver et al., 2016, 2017) and psychologists have argued that RL is responsible for part of human and animal learning (e.g. Niv, 2009; Glimcher, 2011; Caplin and Dean, 2008; Glimcher et al., 2013; Gershman and Daw, 2017).³⁹

RL ideas seem to me a promising direction for developing alternative approaches to rational expectations about equilibrium prices in heterogeneous-agent models: a natural assumption is that agents do not know the true stochastic process for equilibrium prices (which is extremely complex, with prices being driven by an underlying Markov process for an entire cross-sectional distribution – see section 3.4) and instead learn about prices in some way.⁴⁰ RL is linked to least-squares learning because both are stochastic approximation methods (see the previous subsection).⁴¹

Reinforcement learning in a nutshell. I briefly summarize the basics of RL following the brilliant introductory treatment by Sutton and Barto (2018).⁴² RL is concerned with Markov decision processes (MDPs) of the type that are very familiar to economists: an agent maximizes the expected present-discounted value of period rewards

$$\mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t R(x_t, c_t),$$

where $x_t \in \mathcal{X}$ is a state, $c_t \in \mathcal{C}$ is a control (action), R is a reward function, and $0 < \beta < 1$ is a discount factor. The evolution of the state x_t is stochastic and depends on the control c_t according to the transition probabilities $x_{t+1} \sim P(\cdot | x_t, c_t)$. Note that my notation differs from the standard computer science notation (Sutton and Barto, 2018) to make it closer to the notation commonly used in economics. As is standard, the associated Bellman equation for the value function is

$$v(x) = \max_c R(x, c) + \beta \int_{x'} v(x') P(x' | x, c).$$

The difference to standard dynamic programming is that RL is concerned with the case where the agent *does not know the model*, in particular the transition probabilities P . That is, RL is the *optimal control of incompletely-known Markov decision processes*. Somewhat more precisely, agents

³⁹RL ideas have also been applied in game theory (e.g. Erev and Roth, 1995, 1998; Fudenberg and Levine, 2016) but using a different formulation that does not work with value functions and instead directly reinforces the “propensities” of choosing strategies. Also see Barberis and Jin (2023) for a finance application and Fernández-Villaverde et al. (2024a) for references to macroeconomics applications.

⁴⁰A number of papers have applied RL to solving heterogeneous-agent models *without* aggregate risk. See for example Xu et al. (2023) and Laurière et al. (2022). As should hopefully be clear, the present paper is instead about models *with* aggregate risk.

⁴¹I thank Tom Sargent for pointing out this connection.

⁴²Also see Silver (2015) for an equally accessible online course with video recordings and Charpentier et al. (2023) for a summary for economists.

only see the reward R_t they receive after taking action c_t (as well as the current state x_t). Over time, agents learn to take actions that maximize rewards. Rewards are thus like “doggy treats” that reinforce good behavior and this connection to animal psychology is how reinforcement learning got its name. A noteworthy feature of standard RL is that it is “model-free” in that agents learn directly from experience rather than trying to learn the underlying model first (model-based RL). Figure 1 from Sutton and Barto (2018) illustrates the distinction between model-free (or direct) RL and model-based RL.




Figure 1: Model-free vs. model-based reinforcement learning (Sutton and Barto, 2018, ch.8.2).

In mathematical terms, the key idea of RL is to learn value functions from experience (observations of states, actions, and rewards) by means of a stochastic approximation algorithm of the type we encountered in the previous subsection (Jaakkola et al., 1993; Tsitsiklis, 1994). This links RL to least-squares learning and other instances of stochastic approximation algorithms.

To understand how this works, consider first the problem of learning the value function for a fixed deterministic policy π which prescribes a control in feedback form $c = \pi(x)$. Fixing the policy function means that there is no decision and hence this is what the literature calls a Markov reward process (MRP). The corresponding value function is

$$v_\pi(x) = R(x, \pi(x)) + \beta \mathbb{E}_\pi[v_\pi(x_{t+1}) | x_t = x].$$

The standard method for approximating this value function is via *temporal-difference learning* (TD learning). A sketch of the algorithm is as follows:

- Input: a policy $\pi(x)$, a step-size parameter $\alpha > 0$, and a number of time steps T (a large number, say $T = 1000$).
- Initialize $V(x)$ with some initial guess for all $x \in \mathcal{X}$.
- Then, starting from a randomly drawn initial x_0 , for each time step $t = 0, 1, \dots, T$:
 1. given the current state x_t and policy π , compute the next state x_{t+1} using the transition probabilities $P(\cdot | x_t, \pi(x_t))$

2. value prediction: update the guess for the value function according to

$$V(x_t) \leftarrow V(x_t) + \alpha \underbrace{[R_t + \beta V(x_{t+1}) - V(x_t)]}_{\text{Target}}. \quad (19)$$

This algorithm returns a value function $V(x)$ for all $x \in \mathcal{X}$. The recursive updating step (19) takes the form of the general stochastic approximation updating rule (17) and consequently, under certain conditions, $V(x)$ converges to $v_\pi(x)$ as experience $T \rightarrow \infty$ (Jaakkola et al., 1993; Tsitsiklis, 1994; Sutton and Barto, 2018).

The full RL problem is to also solve for the optimal policy π . Because the agent does not know the transition probabilities P , it does not suffice to learn the value function $v_\pi(x)$. Instead, the agent must learn a so-called *action-value function* or *Q function* which equals the value of taking action c in state x . This action-value function is defined for state-action pairs (x, c) as $Q(x, c) := R(x, c) + \beta \int_{x'} v(x') P(x'|x, c)$ and satisfies the Bellman equation

$$Q(x, c) = R(x, c) + \beta \int_{x'} \max_{c'} Q(x', c') P(x'|x, c).$$

Given Q , the optimal policy can then be computed as $\pi(x) = \arg \max_c Q(x, c)$. The agent learns this action-value function using variants of the TD algorithm described above. Two common algorithms are *Sarsa* (which is the most straightforward conceptually) and *Q-learning*.⁴³ RL algorithms for finding optimal policies also involve some variant of “exploration”, i.e. trying out new policies, as opposed to only using policies that maximize the current guess of the action-value function (“exploitation”). For example, so-called ϵ -greedy policies choose actions that are suboptimal given the current action-value function with small probability ϵ .

Big World Hypothesis. Despite exponentially growing computation, applications of RL to real-world learning problems have remained limited. Javed and Sutton (2024) argue that in many learning problems RL agents operate in “big worlds”, i.e. environments that are orders of magnitude more complex than the agent’s information processing capacity. They argue that algorithms should reflect this reality of “small bounded agents learning in large unbounded environments” and propose an approach to algorithm design reminiscent of ideas in the bounded rationality and non-rational expectations literatures in economics.

Heterogeneous-agent economies (and indeed real-world macroeconomies) seem to me good examples of big worlds and so some of these ideas may be applicable to the challenge posed in this paper. The introduction of Javed and Sutton (2024) illustrates the parallels: “*The big world hypothesis says that in many decision-making problems the agent is orders of magnitude smaller than the environment. It can neither fully perceive the state of the world nor can it represent the value or optimal action for every state. Instead, it must learn to make sound decisions using its limited understanding of the environment. The key research challenge for achieving goals in big worlds is to come up with solution methods that efficiently utilize the limited resources of the agent.*”

⁴³The name *Sarsa* comes from the standard computer science notation which denotes states by S_t , actions by A_t and rewards by R_t and that agents update action-value functions from the sequence $(S_t, A_t, R_t, S_{t+1}, A_{t+1})$.

Other contributions have pursued similar ideas. The RL agent of [Dong et al. \(2022\)](#) can operate in complex non-Markovian environments which may be relevant for heterogeneous-agent economies with non-Markovian equilibrium prices (Section 3.4). The agent of [Kumar et al. \(2023\)](#) learns subject to an entropy-based information processing constrained reminiscent of (though different from) the rational inattention literature.

Reinforcement learning about equilibrium prices? In ongoing work I am exploring whether RL ideas can be adapted to the problem of learning about equilibrium prices in heterogeneous-agent models with aggregate risk. One idea is that agents do not know the correct transition probabilities of equilibrium prices (which are extremely complex) and instead learn about the price process from experience using TD learning. The hope is that this type of approach may be able to satisfy Criteria 1 to 3 above. While TD learning about equilibrium prices would not feature an explicit model of price beliefs, the implicit beliefs would still be consistent with model reality (Criterion 3a) in that mistakes that are costly in terms of values would be avoided. Along the lines discussed in Section 3.4, adapting ideas from the theory of RL in non-Markovian environments or POMDPs could be a fruitful avenue.

An interesting question is to what extent delivering on Criterion 2 (consistency with empirical evidence) will require modifications of standard RL ideas. While RL will converge to a different equilibrium than the rational expectations equilibrium, a natural conjecture is that the resulting RL expectations will converge to something “close to” rational. This suggests that modifications of standard RL ideas may be needed to take on board empirical evidence on belief formation.

RL is a machine learning method so it is natural to ask how such an approach relates to a recent literature that has applied such methods to solve heterogeneous agent models with aggregate risk. Typically these papers use deep neural networks to approximate the Master equation, i.e. to approximate value or policy functions on a state space that includes the high-dimensional cross-sectional distribution (e.g. [Han et al., 2021](#); [Azinovic et al., 2022](#); [Gu et al., 2024](#); [Gopalakrishna et al., 2024](#)).⁴⁴ These papers therefore use deep learning to “tame the curse of dimensionality” ([Fernández-Villaverde et al., 2024a](#)). RL about equilibrium prices would instead *sidestep* the curse of dimensionality by modeling decision makers as solving lower-dimensional problems.

Another interesting question is whether standard model-free RL can deliver on Criterion 3 or whether a more model-based approach is needed. An advantage of a model-free approach over typical adaptive learning methods in economics would be not having to specify a perceived law of motion for prices which necessarily involves a degree of arbitrariness. On the other hand, a completely model-free approach in which decision makers do not have a structural understanding of the world may be too simplistic (e.g. [Enke, 2024](#)). Finally, for both model-free and model-based RL methods, learning may be slow in practice. See the discussion

⁴⁴With the exception of [Han et al. \(2021\)](#) these papers also do not use RL ideas and instead use deep neural networks as a powerful function approximation method. Yet another approach is to use a deep neural network to approximate a perceived law of motion for distributional moments (e.g. [Maliar et al., 2021](#); [Fernández-Villaverde et al., 2023, 2024b](#)), i.e. a non-linear version of [Krusell and Smith \(1998\)](#).

on convergence speed in the previous subsection. Because both RL and least-squares learning are special cases of the same general class of stochastic approximation algorithms, this concern applies equally to these two learning methods.

4.5 Heuristics and Simple Models

Like the real world, heterogeneous-agent economies are complex environments that generate complex stochastic processes for equilibrium prices. A natural conjecture is that households and firms operating in this type of complex, stochastic environment employ heuristics and simplified models when making decisions under uncertainty. This idea has a long tradition, see for example [Tversky and Kahneman \(1974\)](#) who argue that *“people rely on a limited number of heuristic principles which reduce the complex tasks of assessing probabilities and predicting values to simpler judgmental operations.”*

In the macroeconomics literature, a number of authors have pursued the related idea that decision makers use simplified, restricted models to forecast economic variables. See for example [Sargent \(1991\)](#), [Brock and Hommes \(1997\)](#), [De Grauwe \(2010\)](#), [Fuster et al. \(2010, 2012\)](#), [Molavi \(2022\)](#), [Hajdini \(2023\)](#), and [Flynn and Sastry \(2024\)](#). Many (but not all) of these authors impose some consistency between subjective beliefs and model reality as in a restricted perceptions equilibrium (Appendix A). This approach – assuming that decision makers form beliefs using “simple models” and finding a restricted perceptions equilibrium – could be applied to price beliefs in heterogeneous-agent models and may be able to satisfy my three criteria. First, the idea that decision makers use simple models and heuristics goes well with Criterion 1 (simplification of the computational solution), in particular the point that models of non-rational expectations should not require decision makers to compute the special case with rational expectations. Second, to make such simple models consistent with key empirical evidence (Criterion 2), one could combine them with popular behavioral biases like diagnostic expectations, cognitive discounting, or models of belief disagreement. Third, the restricted perceptions equilibrium would deliver Criterion 3 (consistency between beliefs and model reality).

The relative sophistication of such simplified forecasting models may differ across decision makers. In particular, financial market participants may be more sophisticated than ordinary households but, over time, the latter may incorporate some of the sophistication of the former ([Caballero and Simsek, 2022](#)). However, even the most sophisticated financial players do not forecast prices by forecasting entire (infinite-dimensional) cross-sectional distributions.

Finally, a recent literature tries to better understand the cognitive foundations of economic decision making in the face of risk and complexity, i.e. to understand “how people think.” See [Enke \(2024\)](#) for a review and [Oprea \(2024\)](#) for an example. [Enke \(2024\)](#) points to evidence that some decision makers fully neglect certain aspects of complex environments (“incomplete representations”), in particular indirect or general equilibrium effects (“system neglect”). Incorporating ideas from this literature could be another way forward.

5 Conclusion

This paper has posed a challenge: to develop alternatives to the assumption of rational expectations about equilibrium prices in heterogeneous agent economies. To this end, I outlined three criteria that alternative approaches should fulfil: (1) simplification of the computational solution, (2) consistency with empirical evidence, and (3) endogeneity of beliefs to model reality (Lucas critique). I then discussed some promising directions, including temporary equilibrium approaches, incorporating survey expectations, least-squares learning, and reinforcement learning.

Developing alternatives to rational expectations holds two main promises. First, to bring the empirical and theoretical discipline of heterogeneous-agent macroeconomics to the study of economic booms and busts and other important questions in which aggregate non-linearities are key. Second, to make these models more realistic, more consistent with empirical evidence on expectations formation, and more likely to generate booms and busts and other interesting phenomena in the first place. While there is often a trade-off between realism and (computational) simplicity, this could be a rare case in which the two go hand in hand.

Macroeconomists specializing in heterogeneous-agent models and those working on the theory and empirics of belief formation should work together and capitalize on the resulting gains from trade. Work combining ideas from these two literatures is starting to emerge. See for example the work on temporary equilibrium with survey expectations discussed above as well as [Guerreiro \(2023\)](#), [Andre et al. \(2024\)](#), [Baley and Turén \(2024\)](#), and [Cai \(2024\)](#). Similarly, [Broer et al. \(2022b\)](#) call for work that considers “more general, dynamic information-choice strategies that can simultaneously match the rich micro-heterogeneity in expectations and explore the subsequent macroeconomic implications.”⁴⁵ I agree.

However, it is also important that researchers not only consider linear or linearized models (as they often do) and instead develop general approaches that are compatible with modeling both micro and macro non-linearities using recursive methods. Similarly, the theoretical expectations literature should allocate more time to modeling expectations about endogenous variables like equilibrium prices (as in [Bastianello and Fontanier, 2024](#), for example). Finally, the literature should focus less on developing models of non-rational expectations that require decision makers to compute the special case of rational expectations. For the purposes of this paper, this defeats the point of departing from rational expectations in the first place.

⁴⁵[Broer et al. \(2022a\)](#) take some steps in this direction but in a way that, counter to my first criterion, makes decision makers’ price forecasting problem more complicated than in the Master equation: the cross-sectional wealth distribution remains a state variable in decision makers’ Bellman equations but, additionally, so is a cross-sectional distribution of beliefs over such wealth distributions.

Appendix

A Various concepts of non-rational expectations equilibrium

The economics literature has proposed various concepts of “non-rational expectations equilibrium” or “misspecification equilibrium.” This appendix briefly summarizes a number of these and puts them in relation to each other. As already stated in Section 4.1, using the acronyms defined there or below, the relation between the various equilibrium concepts is summarized in (15) which I restate here for the reader’s convenience:

$$\{REE, RPE, CEE\} \subset SCE \subset IREE \subset TE,$$

where \subset means “is a special case of.”

Self-confirming equilibrium (SCE). See for example [Sargent \(1999\)](#), [Cho and Sargent \(2016\)](#), and [Fudenberg and Levine \(2016\)](#) in the context of games. In a self-confirming equilibrium, actual equilibrium outcomes are statistically consistent with decision makers’ beliefs, i.e. these beliefs are not disappointed. Self-confirming equilibria can be limiting outcomes of adaptive learning processes of the type described in Section 4.3. A rational expectations equilibrium is a self-confirming equilibrium, but not vice versa. Specifically, beliefs may be incorrect for events that are infrequently observed (e.g. events off the equilibrium parts).

Restricted perceptions equilibrium (RPE). See for example [Sargent \(1991\)](#) and [Branch \(2006\)](#). An RPE is a SCE in which decision makers use restricted forecasting models that do not nest rational expectations, i.e. mis-specified perceived laws of motion. As a result, a RPE is never a REE. Learning processes that update restricted forecasting models may settle down to a RPE (and never a REE). As discussed in Section 3.3, to actually simplify computations in heterogeneous-agent economies (Criterion 1), candidate alternative approaches should *not* start from rational expectations as a special case and so RPE is a natural equilibrium concept. See [Baley and Turén \(2024\)](#) for a recent application of RPE in a model of forecasters’ heterogeneous inflation expectations.

Consistent expectations equilibrium (CEE). See [Hommes and Sorger \(1998\)](#). In a CEE decision makers use autoregressive forecasting models such as an $AR(p)$ process, with a consistency condition for the autocorrelations as for SCE. A CEE is thus a special case of a SCE. In principle, a CEE could coincide with the REE if the RE solution happened to be exactly of the same autoregressive structure as the assumed forecasting model. But RE solutions typically do not take this form (and certainly RE equilibrium prices in heterogeneous-agent models do not) so the CEE forecasting model is typically mis-specified. Thus a CEE is typically a special case of a RPE ([Branch, 2006](#)) so that (15) could also write $CEE \subset RPE$.

Another equilibrium concept not included in (15). *Oblivious equilibrium* is an equilibrium concept for stochastic dynamic games like dynamic industry models of imperfect competition (Weintraub et al., 2005, 2008, 2010). These games suffer from a curse of dimensionality similar to the one in heterogeneous-agent models: in a Markov Perfect Equilibrium, the distribution of other players' states is a state variable in player's dynamic programming problem. Oblivious equilibrium restricts players' policies to be "oblivious" to changes in the distribution of other players' states and instead restricts these policies to be functions of the constant long-run average of this distribution.⁴⁶

References

- Achdou, Yves, Jiequn Han, Jean-Michel Lasry, Pierre-Louis Lions, and Benjamin Moll, "Income and Wealth Distribution in Macroeconomics: A Continuous-Time Approach," *The Review of Economic Studies*, 04 2021, 89 (1), 45–86.
- Adam, Klaus, Albert Marcet, and Johannes Beutel, "Stock Price Booms and Expected Capital Gains," *American Economic Review*, August 2017, 107 (8), 2352–2408.
- and —, "Internal Rationality, Imperfect Market Knowledge and Asset Prices," *Journal of Economic Theory*, May 2011, 146 (3), 1224–1252.
- Ahn, SeHyoun, Greg Kaplan, Benjamin Moll, Thomas Winberry, and Christian Wolf, "When Inequality Matters for Macro and Macro Matters for Inequality," *NBER Macroeconomics Annual*, 2018, 32 (1), 1–75.
- Aiyagari, S. Rao, "Uninsured Idiosyncratic Risk and Aggregate Saving," *The Quarterly Journal of Economics*, August 1994, 109 (3), 659–84.
- Andre, Peter, Carlo Pizzinelli, Christopher Roth, and Johannes Wohlfart, "Subjective Models of the Macroeconomy: Evidence from Experts and Representative Samples," *The Review of Economic Studies*, 2022, 89 (6), 2958–2991.
- , Ingar Haaland, Christopher Roth, Mirko Wiederholt, and Johannes Wohlfart, "Narratives about the Macroeconomy," Working Paper 426, SAFE September 2024.
- Angeletos, George-Marios and Chen Lian, "Chapter 20 - Dampening General Equilibrium: Incomplete Information and Bounded Rationality," in "Handbook of Economic Expectations," Academic Press, 2023, pp. 613–645.
- , Zhen Huo, and Karthik A. Sastry, "Imperfect Macroeconomic Expectations: Evidence and Theory," *NBER Macroeconomics Annual*, 2021, 35 (1), 1–86.
- Armantier, Olivier, Wändi Bruine de Bruin, Simon Potter, Giorgio Topa, Wilbert van der Klaauw, and Basit Zafar, "Measuring Inflation Expectations," *Annual Review of Economics*, May 2013, 5 (1), 273–301.
- Aron-Dine, Shifrah, Johannes Beutel, Monika Piazzesi, and Martin Schneider, "Household Climate Finance: Theory and Survey Data on Safe and Risky Green Assets," NBER Working Papers 32615, National Bureau of Economic Research June 2024.
- Auclert, Adrien, Bence Bardóczy, Matthew Rognlie, and Ludwig Straub, "Using the Sequence-Space Jacobian to Solve and Estimate Heterogeneous-Agent Models," *Econometrica*, September 2021, 89 (5), 2375–2408.

⁴⁶The standard formulation only considers models without aggregate shocks and so is not relevant for our purposes. Weintraub et al. (2010) extend the approach to aggregate shocks by restricting strategies to depend only on a low-dimensional vector that is a function of the history of the aggregate shock (e.g. a truncation of the history).

- , **Matthew Rognlie**, and **Ludwig Straub**, “Micro Jumps, Macro Humps: Monetary Policy and Business Cycles in an Estimated HANK Model,” NBER Working Papers 26647, National Bureau of Economic Research January 2020.
- Azinovic, Marlon, Luca Gaegauf, and Simon Scheidegger**, “Deep Equilibrium Nets,” *International Economic Review*, 2022, 63 (4), 1471–1525.
- Bachmann, Rüdiger, Giorgio Topa, and Wilbert van der Klaauw, eds**, *Handbook of Economic Expectations*, Academic Press, 2023.
- , **Steffen Elstner**, and **Eric R. Sims**, “Uncertainty and Economic Activity: Evidence from Business Survey Data,” *American Economic Journal: Macroeconomics*, April 2013, 5 (2), 217–49.
- Baley, Isaac and Javier Turén**, “Lumpy Forecasts,” *UPF Working Paper Series*, 2024.
- and **Laura L. Veldkamp**, *The Data Economy: Tools and Applications*, Princeton University Press, 2025.
- Barberis, Nicholas and Lawrence Jin**, “Model-free and Model-based Learning as Joint Drivers of Investor Behavior,” NBER Working Papers 31081, National Bureau of Economic Research March 2023.
- Bardóczy, Bence and Joao Guerreiro**, “Unemployment Insurance in Macroeconomic Stabilization with Imperfect Expectations,” *Manuscript*, April, 2023.
- Bastianello, Francesca and Paul Fontanier**, “Expectations and Learning from Prices,” *The Review of Economic Studies*, 10 2024, p. rdae059.
- Bayer, Christian, Benjamin Born, and Ralph Luetticke**, “Shocks, Frictions, and Inequality in US Business Cycles,” *American Economic Review*, May 2024, 114 (5), 1211–47.
- , **Ralph Luetticke, Maximilian Weiss, and Yannick Winkelmann**, “An Endogenous Gridpoint Method for Distributional Dynamics,” CEPR Discussion Paper No. 19067, CEPR, Paris & London 2024.
- Bhandari, Anmol, Thomas Bourany, David Evans, and Mikhail Golosov**, “A Perturbational Approach for Approximating Heterogeneous Agent Models,” Technical Report, National Bureau of Economic Research 2023.
- Bianchi, Javier**, “Overborrowing and Systemic Externalities in the Business Cycle,” *American Economic Review*, 2011, 101 (7), 3400–3426.
- Bilal, Adrien**, “Solving Heterogeneous Agent Models with the Master Equation,” NBER Working Papers 31103, National Bureau of Economic Research 2023.
- and **Shlok Goyal**, “Some Pleasant Sequence-Space Arithmetic in Continuous Time,” Working Paper, Harvard University 2024.
- Bloom, Nick, Mihai Codreanu, and Robert Fletcher**, “Rationalizing Firm Forecasts,” Working Paper, Stanford University 2024.
- Boppart, Timo, Per Krusell, and Kurt Mitman**, “Exploiting MIT Shocks in Heterogeneous-Agent Economies: The Impulse Response as a Numerical Derivative,” *Journal of Economic Dynamics and Control*, 2018, 89, 68–92.
- Bordalo, Pedro, Giovanni Burro, Katherine Coffman, Nicola Gennaioli, and Andrei Shleifer**, “Imagining the Future: Memory, Simulation, and Beliefs,” *The Review of Economic Studies*, 06 2024, p. rdae070.
- , **Nicola Gennaioli, and Andrei Shleifer**, “Diagnostic Expectations and Credit Cycles,” *The Journal of Finance*, 2018, 73 (1), 199–227.
- , —, and —, “Overreaction and Diagnostic Expectations in Macroeconomics,” *Journal of Economic Perspectives*, Summer 2022, 36 (3), 223–244.

- , —, —, and **Stephen J. Terry**, “Real Credit Cycles,” NBER Working Papers 28416, National Bureau of Economic Research December 2021.
- Branch, William A.**, “Restricted Perceptions Equilibria and Learning in Macroeconomics,” *Post Walrasian Macroeconomics: Beyond the Dynamic Stochastic General Equilibrium Model*. Cambridge University Press, New York, 2006, pp. 135–160.
- Bray, Margaret**, “Learning, Estimation, and the Stability of Rational Expectations,” *Journal of Economic Theory*, April 1982, 26 (2), 318–339.
- Brock, William A. and Cars H. Hommes**, “A Rational Route to Randomness,” *Econometrica*, 1997, 65 (5), 1059–1095.
- Broer, Tobias, Alexandre N. Kohlhas, Kurt Mitman, and Kathrin Schlafmann**, “Expectation and Wealth Heterogeneity in the Macroeconomy,” 2022.
- , —, —, and —, “On the Possibility of Krusell-Smith Equilibria,” *Journal of Economic Dynamics and Control*, 2022, 141, 104391.
- Brunnermeier, Markus K. and Yuliy Sannikov**, “A Macroeconomic Model with a Financial Sector,” *American Economic Review*, February 2014, 104 (2), 379–421.
- , **Emmanuel Farhi, Ralph S. J. Koijen, Arvind Krishnamurthy, Sydney C. Ludvigson, Hanno Lustig, Stefan Nagel, and Monika Piazzesi**, “Review Article: Perspectives on the Future of Asset Pricing,” *The Review of Financial Studies*, 2021, 34 (4), 2126–2160.
- Caballero, Ricardo J. and Alp Simsek**, “A Monetary Policy Asset Pricing Model,” NBER Working Papers 30132, National Bureau of Economic Research June 2022.
- Cai, Michael**, “Explaining the Macroeconomic Inertia Puzzle,” Working Paper, Northwestern University 2024.
- Cao, Dan, Wenlan Luo, and Guangyu Nie**, “Global DSGE Models,” *Review of Economic Dynamics*, 2023, 51, 199–225.
- Caplin, Andrew and Mark Dean**, “Dopamine, Reward Prediction Error, and Economics,” *The Quarterly Journal of Economics*, 2008, 123 (2), 663–701.
- , **Victoria Gregory, Eungik Lee, Soren Leth-Petersen, and Johan Sæverud**, “Subjective Earnings Risk,” NBER Working Papers 31019, National Bureau of Economic Research March 2023.
- Cardaliaguet, Pierre, François Delarue, Jean-Michel Lasry, and Pierre-Louis Lions**, *The Master Equation and the Convergence Problem in Mean Field Games*, Vol. 201 of *Annals of Mathematics Studies*, Princeton University Press, 2019.
- Carroll, Christopher D.**, “Macroeconomic Expectations of Households and Professional Forecasters,” *The Quarterly Journal of Economics*, February 2003, 118 (1), 269–298.
- , **Edmund Crawley, Jiri Slacalek, Kiichi Tokuoka, and Matthew N. White**, “Sticky Expectations and Consumption Dynamics,” *American Economic Journal: Macroeconomics*, July 2020, 12 (3), 40–76.
- Cenzon, Josefina**, “Credit Market Experiences and Macroeconomic Expectations: Evidence and Theory,” Working Paper, Barcelona School of Economics 2024.
- Chari, Varadarajan V.**, “Nobel Laureate Robert E. Lucas, Jr.: Architect of Modern Macroeconomics,” *Journal of Economic Perspectives*, 1998, 12 (1), 171–186.
- Charles, Constantin, Cary Frydman, and Mete Kilic**, “Insensitive Investors,” *The Journal of Finance*, 2024, 79 (4), 2473–2503.
- Charpentier, Arthur, Romuald Élie, and Carl Remlinger**, “Reinforcement Learning in Economics and Finance,” *Computational Economics*, June 2023, 62 (1), 425–462.
- Cho, In-Koo and Thomas J. Sargent**, *Self-confirming Equilibria*, Palgrave Macmillan UK, December 2016.

- Christiano, Lawrence J., Martin S. Eichenbaum, and Benjamin K. Johannsen**, “Slow Learning,” NBER Working Papers 32358, National Bureau of Economic Research April 2024.
- Coibion, Olivier and Yuriy Gorodnichenko**, “What Can Survey Forecasts Tell Us about Information Rigidities?,” *Journal of Political Economy*, 2012, 120 (1), 116–159.
- and —, “Information Rigidity and the Expectations Formation Process: A Simple Framework and New Facts,” *American Economic Review*, August 2015, 105 (8), 2644–78.
- , —, and **Bernardo Candia**, “Chapter 11 - The Macroeconomic Expectations of Firms,” in “Handbook of Economic Expectations,” Academic Press, 2023, pp. 321–353.
- , —, **Michael Weber**, and **Francesco D’Acunto**, “The Subjective Inflation Expectations of Households and Firms: Measurement, Determinants, and Implications,” *Journal of Economic Perspectives*, Summer 2022, 36 (3), 157–184.
- D’Acunto, Francesco and Michael Weber**, “Why Survey-Based Subjective Expectations Are Meaningful and Important,” *Annual Review of Economics*, 2024, 16.
- De Grauwe, Paul**, “Top-Down versus Bottom-Up Macroeconomics,” *CESifo Economic Studies*, 2010, 56 (4), 465–497.
- Den Haan, Wouter J.**, “Heterogeneity, Aggregate Uncertainty, and the Short-Term Interest Rate,” *Journal of Business & Economic Statistics*, October 1996, 14 (4), 399–411.
- Dominitz, Jeff and Charles F. Manski**, “Using Expectations Data to Study Subjective Income Expectations,” *Journal of the American Statistical Association*, 1997, 92 (439), 855–867.
- Dong, Shi, Benjamin Van Roy, and Zhengyuan Zhou**, “Simple Agent, Complex Environment: Efficient Reinforcement Learning with Agent States,” *Journal of Machine Learning Research*, 2022, 23 (255), 1–54.
- Elenev, Vadim, Tim Landoigt, and Stijn Van Nieuwerburgh**, “A Macroeconomic Model with Financially Constrained Producers and Intermediaries,” *Econometrica*, 2021, 89 (3), 1361–1418.
- Enke, Benjamin**, “The Cognitive Turn in Behavioral Economics,” Working Paper, Harvard University 2024.
- and **Thomas Graeber**, “Cognitive Uncertainty,” *The Quarterly Journal of Economics*, May 2023, 138 (4), 2021–2067.
- Erev, Ido and Alvin E. Roth**, “Learning in Extensive-Form Games: Experimental Data and Simple Dynamic Models in the Intermediate Term,” *Games and Economic Behavior*, 1995, 8 (1), 164–212.
- and —, “Predicting How People Play Games: Reinforcement Learning in Experimental Games with Unique, Mixed Strategy Equilibria,” *American Economic Review*, 1998, pp. 848–881.
- Evans, George W. and Bruce McGough**, *Adaptive Learning in Macroeconomics*, Oxford University Press, November 2020.
- and **Seppo Honkapohja**, *Learning and Expectations in Macroeconomics*, Princeton University Press, 2001.
- Farhi, Emmanuel and Iván Werning**, “Monetary Policy, Bounded Rationality, and Incomplete Markets,” *American Economic Review*, November 2019, 109 (11), 3887–3928.
- Favilukis, Jack, Sydney C. Ludvigson, and Stijn Van Nieuwerburgh**, “The Macroeconomic Effects of Housing Wealth, Housing Finance, and Limited Risk Sharing in General Equilibrium,” *Journal of Political Economy*, 2017, 125 (1), 140–223.
- Fernández-Villaverde, Jesús, Galo Nuño, and Jesse Perla**, “Taming the Curse of Dimensionality: Quantitative Economics with Deep Learning,” NBER Working Papers 33117, National Bureau of Economic Research November 2024.

- , —, and **Samuel Hurtado**, “Financial Frictions and the Wealth Distribution,” *Econometrica*, May 2023, 91 (3), 869–901.
- , —, **Joël Marbet**, and **Omar Rachedi**, “Inequality and the Zero Lower Bound,” *Journal of Econometrics*, 2024, p. 105819.
- Fischhoff, Baruch and Wändi Bruine De Bruin**, “Fifty–Fifty=50%?,” *Journal of Behavioral Decision Making*, 1999, 12 (2), 149–163.
- Flynn, Joel P. and Karthik Sastry**, “The Macroeconomics of Narratives,” NBER Working Papers 32602, National Bureau of Economic Research June 2024.
- Fofana, Salome, Paula Patzelt, and Ricardo Reis**, “Household Disageement About Expected Inflation,” in Guido Ascari and Riccardo Trezzi, eds., *The Research Handbook of Inflation*, Edward-Elgar, March 2024, chapter 15.
- Friedman, Milton and L. J. Savage**, “Planning Experiments Seeking Maxima,” in Churchill Eisenhart, Millard W. Hartay, and W. Allen Wallis, eds., *Selected Techniques of Statistical Analysis for Scientific and Industrial Research, and Production and Management Engineering*, New York and London: McGraw-Hill, 1947, pp. 363–372.
- Fudenberg, Drew and David K. Levine**, “Whither Game Theory? Towards a Theory of Learning in Games,” *Journal of Economic Perspectives*, 2016, 30 (4), 151–170.
- Fuster, Andreas, Benjamin Hebert, and David Laibson**, “Natural Expectations, Macroeconomic Dynamics, and Asset Pricing,” *NBER Macroeconomics Annual*, 2012, 26 (1), 1–48.
- , **David Laibson, and Brock Mendel**, “Natural Expectations and Macroeconomic Fluctuations,” *Journal of Economic Perspectives*, December 2010, 24 (4), 67–84.
- Gabaix, Xavier**, “A Sparsity-Based Model of Bounded Rationality,” *The Quarterly Journal of Economics*, 2014, 129 (4), 1661–1710.
- , “A Behavioral New Keynesian Model,” *American Economic Review*, 2020, 110 (8), 2271–2327.
- , “Marshall Lecture 2023: Behavioral Macroeconomics via Sparse Dynamic Programming,” *Journal of the European Economic Association*, 2023, 21 (6), 2327–2376.
- García-Schmidt, Mariana and Michael Woodford**, “Are Low Interest Rates Deflationary? A Paradox of Perfect-Foresight Analysis,” *American Economic Review*, January 2019, 109 (1), 86–120.
- Gershman, Samuel J. and Nathaniel D. Daw**, “Reinforcement Learning and Episodic Memory in Humans and Animals: An Integrative Framework,” *Annual Review of Psychology*, 2017, 68 (Volume 68, 2017), 101–128.
- Giglio, Stefano, Matteo Maggiori, Johannes Stroebel, and Stephen Utkus**, “Five Facts about Beliefs and Portfolios,” *American Economic Review*, May 2021, 111 (5), 1481–1522.
- Glawion, René**, “Sequence-Space Jacobians in Continuous Time,” *Available at SSRN: <https://ssrn.com/abstract=4504829>*, 2023.
- Glimcher, Paul W.**, “Understanding Dopamine and Reinforcement Learning: The Dopamine Reward Prediction Error Hypothesis,” *Proceedings of the National Academy of Sciences*, 2011, 108 (supplement.3), 15647–15654.
- , **Colin Camerer, Ernst Fehr, and Russell Poldrack**, *Neuroeconomics: Decision Making and the Brain* Academic Press, Academic Press, 2013.
- Gomes, Francisco and Alexander Michaelides**, “Asset Pricing with Limited Risk Sharing and Heterogeneous Agents,” *Review of Financial Studies*, January 2008, 21 (1), 415–448.
- Gopalakrishna, Goutham, Zhouzhou Gu, and Jonathan Payne**, “Asset Pricing, Participation Constraints, and Inequality,” Working Paper, Princeton Working Paper 2024.

- Grandmont, Jean-Michel**, "Temporary General Equilibrium Theory," *Econometrica*, April 1977, 45 (3), 535–572.
- , "Temporary Equilibrium," in "General Equilibrium," Springer, 1989, pp. 297–304.
- Gu, Zhouzhou, Mathieu Laurière, Sebastian Merkel, and Jonathan Payne**, "Global Solutions to Master Equations for Continuous Time Heterogeneous Agent Macroeconomic Models," *arXiv preprint arXiv:2406.13726*, 2024.
- Guerreiro, Joao**, "Belief Disagreement and Business Cycles," Working Paper, UCLA 2023.
- Haaland, Ingar, Christopher Roth, and Johannes Wohlfart**, "Designing Information Provision Experiments," *Journal of Economic Literature*, March 2023, 61 (1), 3–40.
- Hajdini, Ina**, "Mis-Specified Forecasts and Myopia in an Estimated New Keynesian Model," 2023.
- Han, Jiequn, Yucheng Yang, and Weinan E**, "DeepHAM: A Global Solution Method For Heterogeneous Agent Models With Aggregate Shocks," *arXiv preprint arXiv:2112.14377*, 2021.
- He, Zhiguo and Arvind Krishnamurthy**, "A Model of Capital and Crises," *The Review of Economic Studies*, 2012, 79 (2), 735–777.
- and —, "Intermediary Asset Pricing," *American Economic Review*, April 2013, 103 (2), 732–770.
- Hicks, John R.**, *Value and Capital: An Inquiry into Some Fundamental Principles of Economic Theory*, Clarendon Press. Available at https://benjaminmoll.com/Hicks_Value_and_Capital/, 1939.
- Hommes, Cars and Gerhard Sorger**, "Consistent Expectations Equilibria," *Macroeconomic Dynamics*, 1998, 2 (3), 287–321.
- Hotelling, Harold**, "Experimental Determination of the Maximum of a Function," *The Annals of Mathematical Statistics*, 1941, 12 (1), 20–45.
- Huang, Ji**, "Breaking the Curse of Dimensionality in Heterogeneous-Agent Models: A Deep Learning-Based Probabilistic Approach," Available at SSRN: <https://ssrn.com/abstract=4649043>, 2023.
- Ilut, Cosmin and Martin Schneider**, "Ambiguous Business Cycles," *American Economic Review*, August 2014, 104 (8), 2368–99.
- Jaakkola, Tommi, Michael Jordan, and Satinder Singh**, "Convergence of Stochastic Iterative Dynamic Programming Algorithms," *Advances in Neural Information Processing Systems*, 1993, 6, 703–710.
- Javed, Khurram and Richard S. Sutton**, "The Big World Hypothesis and its Ramifications for Artificial Intelligence," in "Finding the Frame: An RLC Workshop for Examining Conceptual Frameworks" 2024.
- Kaplan, Greg, Benjamin Moll, and Giovanni L. Violante**, "Monetary Policy According to HANK," *American Economic Review*, March 2018, 108 (3), 697–743.
- , **Kurt Mitman, and Giovanni L. Violante**, "The Housing Boom and Bust: Model Meets Evidence," *Journal of Political Economy*, 2020, 128 (9), 3285–3345.
- Kindleberger, Charles**, *Manias, Panics, and Crashes*, Macmillan. Relevant pages available at <https://benjaminmoll.com/kindleberger/>, 1978.
- Krusell, Per and Anthony A. Smith**, "Income and Wealth Heterogeneity in the Macroeconomy," *Journal of Political Economy*, October 1998, 106 (5), 867–896.
- Kumar, Saurabh, Henrik Marklund, Ashish Rao, Yifan Zhu, Hong Jun Jeon, Yueyang Liu, and Benjamin Van Roy**, "Continual Learning as Computationally Constrained Reinforcement Learning," 2023.

- Laurière, Mathieu, Sarah Perrin, Sertan Girgin, Paul Muller, Ayush Jain, Theophile Cabannes, Georgios Piliouras, Julien Perolat, Romuald Elie, Olivier Pietquin, and Matthieu Geist**, “Scalable Deep Reinforcement Learning Algorithms for Mean Field Games,” in Kamalika Chaudhuri, Stefanie Jegelka, Le Song, Csaba Szepesvari, Gang Niu, and Sivan Sabato, eds., *Proceedings of the 39th International Conference on Machine Learning*, Vol. 162 of *Proceedings of Machine Learning Research* PMLR 17–23 Jul 2022, pp. 12078–12095.
- Lee, Donghoon and Kenneth I. Wolpin**, “Intersectoral Labor Mobility and the Growth of the Service Sector,” *Econometrica*, 2006, 74 (1), 1–46.
- Lee, Soyoung**, “The Macroeconomic Effects of Debt Relief Policies during Recessions,” *Available at SSRN 4046177*, 2022.
- Leombroni, Matteo, Monika Piazzesi, Martin Schneider, and Ciaran Rogers**, “Inflation and the Price of Real Assets,” NBER Working Papers 26740, National Bureau of Economic Research February 2020.
- Lindahl, Erik**, *Studies in the Theory of Money and Capital*, Allen and Unwin, 1939.
- Ljung, L.**, “Analysis of Recursive Stochastic Algorithms,” *IEEE Transactions on Automatic Control*, 1977, 22 (4), 551–575.
- Ljungqvist, Lars and Thomas J. Sargent**, *Recursive Macroeconomic Theory*, The MIT Press, 2004.
- Llull, Joan**, “Immigration, Wages, and Education: A Labour Market Equilibrium Structural Model,” *The Review of Economic Studies*, 2018, 85 (3), 1852–1896.
- Lucas, Robert E.**, “Expectations and the Neutrality of Money,” *Journal of Economic Theory*, 1972, 4 (2), 103–124.
- , “Econometric Policy Evaluation: A Critique,” *Carnegie-Rochester Conference Series on Public Policy*, 1976, 1, 19–46.
- , “Methods and Problems in Business Cycle Theory,” *Journal of Money, Credit and Banking*, November 1980, 12 (4), 696–715.
- , *Models of Business Cycles*, Basil Blackwell, Oxford, 1987.
- , “Nobel Lecture: Monetary Neutrality,” *Journal of Political Economy*, August 1996, 104 (4), 661–682.
- and **Edward C. Prescott**, “Investment Under Uncertainty,” *Econometrica*, September 1971, 39 (5), 659–681.
- Ludwig, Alexander, Jochen Mankart, Jorge Quintana, and Mirko Wiederholt**, “Heterogeneity in Expectations and House Price Dynamics,” SAFE Working Papers 432, SAFE 2024.
- Maliar, Lilia and Serguei Maliar**, “Numerical Methods for Large-Scale Dynamic Economic Models,” in “Handbook of Computational Economics,” Vol. 3, Elsevier, 2014, pp. 325–477.
- , —, and **Pablo Winant**, “Deep Learning for Solving Dynamic Economic Models,” *Journal of Monetary Economics*, 2021, 122 (C), 76–101.
- Malmendier, Ulrike and Stefan Nagel**, “Depression Babies: Do Macroeconomic Experiences Affect Risk Taking?,” *The Quarterly Journal of Economics*, 2011, 126 (1), 373–416.
- and —, “Learning from Inflation Experiences,” *The Quarterly Journal of Economics*, October 2015, 131 (1), 53–87.
- Mankiw, N. Gregory and Ricardo Reis**, “Sticky Information Versus Sticky Prices: A Proposal to Replace the New Keynesian Phillips Curve,” *The Quarterly Journal of Economics*, 2002, 117 (4), 1295–1328.
- , —, and **Justin Wolfers**, “Disagreement about Inflation Expectations,” in “NBER Macroeconomics Annual 2003, Volume 18” NBER Chapters, National Bureau of Economic Research, Inc, November 2004, pp. 209–270.

- Manski, Charles F.**, "Measuring Expectations," *Econometrica*, September 2004, 72 (5), 1329–1376.
- , "Survey Measurement of Probabilistic Macroeconomic Expectations: Progress and Promise," *NBER Macroeconomics Annual*, 2018, 32 (1), 411–471.
- Marcet, Albert and Thomas J. Sargent**, "Convergence of Least Squares Learning Mechanisms in Self-referential Linear Stochastic Models," *Journal of Economic Theory*, August 1989, 48 (2), 337–368.
- Maxted, Peter**, "A Macro-Finance Model with Sentiment," *Review of Economic Studies*, 2024, 91 (1), 438–475.
- Maćkowiak, Bartosz, Filip Matějka, and Mirko Wiederholt**, "Rational Inattention: A Review," *Journal of Economic Literature*, March 2023, 61 (1), 226–73.
- Mendoza, Enrique G.**, "Sudden Stops, Financial Crises, and Leverage," *American Economic Review*, 2010, 100 (5), 1941–1966.
- Minsky, Hyman P.**, "The Financial Instability Hypothesis: An Interpretation of Keynes and an Alternative to "Standard" Theory," *Challenge*, 1977, 20 (1), 20–27.
- Mnih, Volodymyr, Koray Kavukcuoglu, David Silver, Andrei A. Rusu, Joel Veness, Marc G. Bellemare, Alex Graves, Martin Riedmiller, Andreas K. Fidjeland, Georg Ostrovski, Stig Petersen, Charles Beattie, Amir Sadik, Ioannis Antonoglou, Helen King, Dharmashan Kumaran, Daan Wierstra, Shane Legg, and Demis Hassabis**, "Human-level control through deep reinforcement learning," *Nature*, Feb 2015, 518 (7540), 529–533.
- Molavi, Pooya**, "Simple Models and Biased Forecasts," *arXiv preprint arXiv:2202.06921*, 2022.
- Muth, John F.**, "Rational Expectations and the Theory of Price Movements," *Econometrica*, July 1961, 29 (3), 315–335.
- Nagel, Rosemarie**, "Unraveling in Guessing Games: An Experimental Study," *The American Economic Review*, 1995, 85 (5), 1313–1326.
- Nakov, Anton and Galo Nuño**, "Learning from Experience in the Stock Market," *Journal of Economic Dynamics and Control*, 2015, 52, 224–239.
- Niv, Yael**, "Reinforcement Learning in the Brain," *Journal of Mathematical Psychology*, 2009, 53 (3), 139–154. Special Issue: Dynamic Decision Making.
- Oprea, Ryan**, "Decisions under Risk Are Decisions under Complexity," *American Economic Review*, December 2024, 114 (12), 3789–3811.
- Ortoleva, Pietro**, "Alternatives to Bayesian Updating," *Annual Review of Economics*, 2024, 16 (Volume 16, 2024), 545–570.
- Piazzesi, Monika and Martin Schneider**, *Housing and Macroeconomics*, Vol. 2 of *Handbook of Macroeconomics*, Elsevier, August 2016.
- , —, and **Tim Landvoigt**, "The Housing Market(s) of San Diego," *American Economic Review*, April 2015, 105 (4), 1371–1407.
- Prescott, Edward C. and Rajnish Mehra**, "Recursive Competitive Equilibrium: The Case of Homogeneous Households," in "Theory Of Valuation," World Scientific, 2005, pp. 357–371.
- Preston, Bruce**, "Learning about Monetary Policy Rules when Long-Horizon Expectations Matter," *International Journal of Central Banking*, September 2005, 1 (2).
- Pröhl, Elisabeth**, "Approximating Equilibria with Ex-Post Heterogeneity and Aggregate Risk," *Swiss Finance Institute Research Paper*, 2019, (17-63).
- Reiter, Michael**, "Solving heterogeneous-agent models by projection and perturbation," *Journal of Economic Dynamics and Control*, March 2009, 33 (3), 649–665.

- Ríos-Rull, José-Víctor**, “Computation of Equilibria in Heterogeneous Agent Models,” Technical Report, Federal Reserve Bank of Minneapolis 1997.
- Robbins, Herbert and Sutton Monro**, “A Stochastic Approximation Method,” *The Annals of Mathematical Statistics*, 1951, 22 (3), 400–407.
- Rolnick, Arthur J.**, “Interview with Thomas Sargent,” *Federal Reserve Bank of Minneapolis*, 2010. Available at: <https://www.minneapolisfed.org/article/2010/interview-with-thomas-sargent>.
- Roth, Christopher, Mirko Wiederholt, and Johannes Wohlfart**, “The Effects of Monetary Policy: Theory with Measured Expectations,” ECONtribute Discussion Papers Series 217, University of Bonn and University of Cologne, Germany January 2023.
- Sargent, Thomas J.**, “Equilibrium with Signal Extraction from Endogenous Variables,” *Journal of Economic Dynamics and Control*, 1991, 15 (2), 245–273.
- , *Bounded Rationality in Macroeconomics* The Arne Ryde Memorial Lectures, Clarendon Press, 1993.
- , “Expectations and the Nonneutrality of Lucas,” *Journal of Monetary Economics*, 1996, 37 (3), 535–548.
- , *The Conquest of American Inflation*, Princeton University Press, 1999.
- , “Rational Expectations,” in “The Concise Encyclopedia of Economics,” Liberty Fund, Indianapolis, available at <https://www.econlib.org/library/Enc/RationalExpectations.html>, 2007.
- , “Evolution and Intelligent Design,” *American Economic Review*, March 2008, 98 (1), 5–37.
- , “Rational Expectations,” in “The New Palgrave Dictionary of Economics,” Palgrave Macmillan UK, 2017, pp. 1–7.
- , “Macroeconomics After Lucas,” 2024. Working Paper.
- Schaab, Andreas**, “Micro and Macro Uncertainty,” Working Paper, UC Berkeley 2020.
- Silver, David**, *Introduction to Reinforcement Learning*, DeepMind x UCL, Video recordings available at <https://www.youtube.com/playlist?list=PLqYmG7hTraZDM-OYHWgPebj2MfCFzF0bQ>, 2015.
- , **Aja Huang, Chris J. Maddison, Arthur Guez, Laurent Sifre, George van den Driessche, Julian Schrittwieser, Ioannis Antonoglou, Veda Panneershelvam, Marc Lanctot, Sander Dieleman, Dominik Grewe, John Nham, Nal Kalchbrenner, Ilya Sutskever, Timothy Lillicrap, Madeleine Leach, Koray Kavukcuoglu, Thore Graepel, and Demis Hassabis**, “Mastering the game of Go with deep neural networks and tree search,” *Nature*, Jan 2016, 529 (7587), 484–489.
- , **Julian Schrittwieser, Karen Simonyan, Ioannis Antonoglou, Aja Huang, Arthur Guez, Thomas Hubert, Lucas Baker, Matthew Lai, Adrian Bolton, Yutian Chen, Timothy Lillicrap, Fan Hui, Laurent Sifre, George van den Driessche, Thore Graepel, and Demis Hassabis**, “Mastering the Game of Go without Human Knowledge,” *Nature*, Oct 2017, 550 (7676), 354–359.
- Sims, Christopher A.**, “Macroeconomics and Reality,” *Econometrica: Journal of the Econometric Society*, 1980, pp. 1–48.
- , “Implications of Rational Inattention,” *Journal of Monetary Economics*, 2003, 50 (3), 665–690.
- Storesletten, Kjetil, Chris Telmer, and Amir Yaron**, “Asset Pricing with Idiosyncratic Risk and Overlapping Generations,” *Review of Economic Dynamics*, October 2007, 10 (4), 519–548.
- Sutton, Richard S. and Andrew G. Barto**, *Reinforcement Learning: An Introduction*, MIT Press, 2018.
- Taubinsky, Dmitry, Luigi Butera, Matteo Saccarola, and Chen Lian**, “Beliefs About the Economy are Excessively Sensitive to Household-Level Shocks: Evidence from Linked Survey and Administrative Data,” NBER Working Papers 32664, National Bureau of Economic Research July 2024.
- Tsitsiklis, John N.**, “Asynchronous Stochastic Approximation and Q-learning,” *Machine Learning*, 1994, 16, 185–202.

- Tversky, Amos and Daniel Kahneman**, “Judgment under Uncertainty: Heuristics and Biases,” *Science*, 1974, 185 (4157), 1124–1131.
- Viscusi, W.Kip**, “A Bayesian Perspective on Biases in Risk Perception,” *Economics Letters*, 1985, 17 (1), 59–62.
- Weintraub, Gabriel Y., C. Lanier Benkard, and Benjamin Van Roy**, “Markov Perfect Industry Dynamics with Many Firms,” *Econometrica*, 2008, 76 (6), 1375–1411.
- , —, and —, “Computational Methods for Oblivious Equilibrium,” *Operations Research*, 2010, 58 (4-part-2), 1247–1265.
- , **Lanier Benkard, and Benjamin Van Roy**, “Oblivious Equilibrium: A Mean Field Approximation for Large-Scale Dynamic Games,” in Y. Weiss, B. Schölkopf, and J. Platt, eds., *Advances in Neural Information Processing Systems*, Vol. 18 MIT Press 2005.
- Werning, Iván**, “Expectations and the Rate of Inflation,” NBER Working Papers 30260, National Bureau of Economic Research July 2022.
- , “Lucas Miracles,” 2023. Available at: <https://economics.mit.edu/sites/default/files/inline-files/Translated%20Lucas%20Miracles%20by%20Ivan%20Werning.pdf>.
- Woodford, Michael**, “Learning to Believe in Sunspots,” *Econometrica*, 1990, 58 (2), 277–307.
- , “Macroeconomic Analysis without the Rational Expectations Hypothesis,” *Annual Review of Economics*, 2013, 5 (1), 303–346.
- Xu, Ruitu, Yifei Min, Tianhao Wang, Michael I. Jordan, Zhaoran Wang, and Zhuoran Yang**, “Finding Regularized Competitive Equilibria of Heterogeneous Agent Macroeconomic Models via Reinforcement Learning,” in Francisco Ruiz, Jennifer Dy, and Jan-Willem van de Meent, eds., *Proceedings of The 26th International Conference on Artificial Intelligence and Statistics*, Vol. 206 of *Proceedings of Machine Learning Research* PMLR 25–27 Apr 2023, pp. 375–407.
- Yang, Jeffrey**, “On the Decision-Relevance of Subjective Beliefs,” Working Paper 2024.