

The Economic Journal, 136 (May), 1173–1216 <https://doi.org/10.1093/ej/ueaf104> © The Author(s) 2025. Published by Oxford University Press on behalf of Royal Economic Society. This is an Open Access article distributed under the terms of the Creative Commons Attribution-NonCommercial-NoDerivs licence (<https://creativecommons.org/licenses/by-nc-nd/4.0/>), which permits non-commercial reproduction and distribution of the work, in any medium, provided the original work is not altered or transformed in any way, and that the work is properly cited. For commercial re-use, please contact reprints@oup.com for reprints and translation rights for reprints. All other permissions can be obtained through our RightsLink service via the Permissions link on the article page on our site—for further information please contact journals.permissions@oup.com.

Advance Access Publication Date: 14 November 2025

THE TROUBLE WITH RATIONAL EXPECTATIONS IN HETEROGENEOUS AGENT MODELS: A CHALLENGE FOR MACROECONOMICS*

Benjamin Moll

The thesis of this essay is that, in heterogeneous agent macroeconomics, the assumption of rational expectations about equilibrium prices is unrealistic and should be replaced. Rational expectations imply that decision-makers forecast equilibrium prices like interest rates by forecasting cross-sectional distributions. This leads to an extreme version of the curse of dimensionality: dynamic programming problems in which the entire distribution is a state variable (the ‘Master equation’, also known as the ‘Monster equation’). Frontier computational methods struggle with these infinite-dimensional Bellman equations, making it implausible that real-world agents solve the associated decision problems. These difficulties also limit the applicability of the heterogeneous agent approach to central questions in macroeconomics—those involving aggregate risk and non-linearities such as financial crises. This troublesome feature of the rational expectations assumption poses a challenge: what should replace it? I outline three criteria for alternative approaches: (1) computational tractability, (2) consistency with empirical evidence and (3) (some) immunity to the Lucas critique. I then discuss several promising directions, including temporary equilibrium approaches, incorporating survey expectations, least-squares learning and reinforcement learning.

JEL code: E3

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 remain 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

* Corresponding author: Benjamin Moll, Sir Arthur Lewis Building, LSE Department of Economics, Houghton Street, London WC2A 2AE, United Kingdom. Email: b.moll@lse.ac.uk

This paper was received on 29 November 2024 and accepted on 6 October 2025. The Editor was Franck Portier.

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 honour. I thank Klaus Adam, Marios Angeletos, Adrien Auclert, Isaac Baley, Adrien Bilal, Pedro Bordo, Peter Bossaerts, Jean-Philippe Bouchaud, Tobias Broer, Andrew Caplin, Ben Enke, Xavier Gabaix, Laura Gati, Federico Gavazzoni, Joao Guerreiro, Ina Hajdini, Frank Heinemann, Cars Hommes, Zhen Huo, Alex Imas, Greg Kaplan, Alex Kohlhas, Per Krusell, Søren Leth-Petersen, Chen Lian, Guido Lorenzoni, Charles Manski, Peter Maxted, Dmitry Mukhin, Emi Nakamura, Galo Nuño, Ryan Oprea, Ricardo Reis, Matt Rognlie, Tom Sargent, Karthik Sastry, Andreas Schaab, Martin Schneider, Chengchun Shi, Alp Simsek, Johannes Stroebel, Stéphane Surprenant, Chris Tonetti, Ben Van Roy, Mirko Wiederholt, Johannes Wohlfart, Moto Yogo, Jian-Qiao Zhu, the editor and two anonymous referees for helpful comments. I am also grateful for excellent research assistance from Nicholas Tokay and support from the Leverhulme Trust.

¹ For example, Mendoza (2010), Bianchi (2011), He and Krishnamurthy (2012; 2013), Brunnermeier and Sannikov (2014), Elenev *et al.* (2021) and Krishnamurthy and Li (2025). Some recent work does take steps toward solving heterogeneous agent versions of such models; see the related-literature discussion below.

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 (Den Haan, 1996; Krusell and Smith, 1998). 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 the ‘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 that 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 realise that market-clearing interest rates depend on the entire cross-sectional distribution of different asset holdings in the economy (say the US 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 this forecasting behaviour implied by rational expectations is not just computationally challenging, it is conceptually implausible. If even our most advanced computational tools struggle with the Monster equation, how can we justify the assumption that real-world households and firms solve the associated decision problems?⁵ 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 the simpler equations corresponding to their actual price-forecasting behaviour. In short, I argue against rational expectations on conceptual grounds (in addition to empirical ones): in complex environments with heterogeneous interacting agents, the assumption implies agent forecasting behaviour that is extremely implausible.⁶

Departing from rational expectations would open the door to applying 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 (see Minsky, 1977; Kindleberger, 1978 and the 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

² To 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 Den Haan (1996) and Krusell and Smith (1998), though not fully spelled out and not with that name.

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

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

⁵ For similar points, see the quotes by Morgenstern (1935), Manski (2004) and Adam and Marcet (2011) below.

⁶ A version of this argument arguably also applies in the representative-agent case. Indeed, similar criticisms have a long tradition in the economics literature; see the literature discussion below. However, in my view, the argument is considerably stronger in the heterogeneous agent case because of the model’s complexity, in particular that the economy’s state is an infinite-dimensional distribution.

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 developing my main thesis about rational expectations, I revisit the roots of this modelling 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 popularisation 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 behaviour of macroeconomic time series following policy changes; see, for example, Lucas’ FORTRAN quote 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 going back to the drawing board to develop alternative approaches. First, methods that employ ‘MIT shocks’ or linearise 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 (Den Haan, 1996; Krusell and Smith, 1998) lack realism for similar reasons as Master equation methods: the idea that decision-makers forecast means, variances, etc. of distributions to forecast prices also strains credibility. 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,⁷ I spell out three criteria that such alternatives

⁷ I only know the problem, not the solution!

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).⁸ The goal of my three criteria is to help navigate this wilderness. The three criteria are (1) computational tractability, (2) consistency with empirical evidence and (3) endogeneity of beliefs to model reality (the Lucas critique). An implication of Criterion (1) is that it eliminates from the list of candidate alternative approaches any non-rational expectations models requiring decision-makers to compute the special case of rational expectations. Criteria (2) and (3) further narrow the set of admissible subjective beliefs. The alternative approach should also be compatible with global (rather than local) and recursive solution methods.

In the paper’s final section, I outline some promising directions for developing alternative models of price beliefs that satisfy these three criteria. First, I discuss the fundamental concept of temporary equilibrium, i.e., competitive equilibrium, given specified subjective beliefs, alongside the idea of internal rationality. Second, alternative models of price beliefs should incorporate the substantial body of empirical evidence on expectations formation, including work on survey expectations. I then turn to models of adaptive learning that are natural candidates for meeting 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 driven some impressive advances in artificial intelligence. The two forms of learning are linked because both are special cases of a broader class of stochastic approximation algorithms. Finally, I discuss the idea that decision-makers may use heuristics and simplified, restricted forecasting models. Of course, implementing any of these approaches in practice will entail its own computational, empirical and conceptual challenges, so my case for their promise remains necessarily speculative.

With this caveat, all approaches share a crucial feature: they should enable efficient *global* solutions of heterogeneous agent models with full cross-sectional distributions. Relaxing rational expectations simplifies agents’ decision problems *inside* the model, but does *not* diminish the rich dynamics of the *economy* itself, which still evolves stochastically and non-linearly, driven by the policy functions of forward-looking heterogeneous agents. This is precisely what holds the promise of generating the kind of rich macroeconomic dynamics—such as booms and busts—that rational-expectations heterogeneous agent models have struggled to capture.

Similar critiques of rational expectations have a long tradition in macroeconomics. In fact, they arguably predate the formal development of rational expectations: Morgenstern (1935) criticised the assumption of perfect foresight—the precursor to rational expectations—for demanding that

the forward-looking individual must [...] know not only exactly the influence of his own actions on prices, but also that of all other individuals [and] must grasp all economic interrelations—that is, must master economic theory, [and for] the improbably high demands placed on the intellectual capacity of economic agents. (p. 342 and p. 345)

See [Online Appendix C](#) for further details.

Following this tradition, a huge literature in macroeconomics analyses alternatives to the rational expectations paradigm, both empirically and theoretically. I cover much of this work

⁸ Sargent attributed the phrase to Sims (1980), but Sims actually wrote about the ‘wilderness of disequilibrium economics’.

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), Ortleva (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.* (2025) on the empirical side. See also the recent *Handbook of Economic Expectations* (Bachmann *et al.*, 2023) that covers both theory and empirics.

As already noted, I 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 Den Haan (1996) and Krusell and Smith (1998), 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 study phenomena like financial crises. While largely out of reach for the heterogeneous agent literature, a few recent contributions have taken steps in this direction (e.g., Fernández-Villaverde *et al.*, 2023; Gopalakrishna *et al.*, 2024). Yet high computational costs mean that theories with aggregate non-linearities remain highly stylised and cannot fully speak to key empirical evidence. Rather than solving Monster equations, we should solve simpler, more realistic ones. Rather than 'taming the curse of dimensionality' (Fernández-Villaverde *et al.*, 2024b), 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 Branch and Evans (2011), Adam *et al.* (2017), Barberis (2018), Williams (2018; 2024), Morelli *et al.* (2020), Bordo *et al.* (2021), Maxted (2024) and Krishnamurthy and Li (2025). Agent-based modelling (ABM) is another modelling approach that aims to generate interesting macroeconomic phenomena of this type.⁹ Relative to heterogeneous agent macroeconomics, ABM typically departs from all of the following: (1) rational expectations, (2) optimisation, (3) market clearing. I instead advocate for departing from (1) with regard to equilibrium prices, but retaining (2) and (3).

This paper is organized as follows. Section 1 puts the thesis of this essay into perspective by revisiting some key early writings on rational expectations. Section 2 explains why, in heterogeneous agent models, the assumption of rational expectations about equilibrium prices is unrealistic 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 modelling device in the 1960s and 1970s through the lens of some writings that were key to its development and popularisation, like Muth (1961), Lucas and Prescott (1971) and Lucas (1972). Expectations about equilibrium prices played a (if not *the*) key role in these developments. In contrast, the recent theoretical expecta-

⁹ See, for example, the surveys by Bouchaud (2013) and Axtell and Farmer (2025), the latest volume of the *Handbook of Computational Economics* (Hommes and LeBaron, 2018) or Geanakoplos *et al.* (2012), Gualdi *et al.* (2015) and Wiese *et al.* (2024).

tions literature in macroeconomics mostly focuses on modelling expectations about exogenous variables. The early writings also stressed that a key payoff of adopting rational expectations was to make models ‘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. He 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 that 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 stated 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 expectation operator \mathbb{E}_{t-1} integrates over the objective, model-generated price distribution. Rational expectations are ‘model consistent’ in this sense and this may be a more appropriate terminology (Simon, 1978b). In his linear model, Muth solves for the rational price expectation as follows. First, the equilibrium price of the demand-supply system (1) is given by the reduced form

$$p_t = -\frac{\gamma}{\beta} \mathbb{E}_{t-1}[p_t] - \frac{1}{\beta} u_t, \tag{2}$$

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 $\mathbb{E}_{t-1}[p_t] = -\gamma \mathbb{E}_{t-1}[p_t]/\beta - \mathbb{E}_{t-1}[u_t]/\beta$ or

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

¹⁰ To 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 . Both demand and supply are expressed as deviations from their long-run means, and hence the potentially negative values.

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 realisation of the supply shock u_t . But firms correctly understand that these supply shocks average out. Analogous results are obtained 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 dependence of equilibrium prices on the fundamental u_t , modelling firms' price expectations requires only the stochastic process of the fundamental u_t .

1.2. *In Bob Lucas' Words*

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

[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 modelling 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:

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. (p. 707)

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 that leads to fluctuations in the output price faced by firms.

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

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

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

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. (p. 659 - 660)

This last sentence is followed by a footnote stating that

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. (p. 660)

While Muth had already verbally defined rational expectations in terms of a common distribution (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 that indirectly defines the recursive competitive equilibrium by maximising social surplus.

In summary, also from Lucas and Prescott’s pathbreaking paper, it becomes clear that the development of rational expectations was all about improving the modelling 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) explained 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. (p. 104)

In a bit more detail, Lucas first wrote down the young's optimisation problem in terms of an unspecified (subjective) probability distribution F for the future price level p' ; see his equation (3.7). He then noted that

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. (p. 179)¹³

A recursive rational expectations equilibrium is then a price function $p(m, x, \theta)$ that satisfies the optimality condition of the young's optimisation 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 wrote, '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 twenty-fifth anniversary of the publication of Lucas (1972), Sargent (1996) again emphasised the importance of rational expectations as a device for modelling 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. (p. 542)

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) wrote,

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. (p. 709)

Similarly, directly after the passage cited in Section 1.3, Lucas and Prescott (1971) stated that

[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. (p. 660)

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 return to this section's themes—expectations about equilibrium prices and operability—later in the paper.

Finally, while modern macroeconomics owes much to the rational expectations hypothesis, its importance should 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.

¹³ See also the useful discussion of Lucas' paper by Chari (1998, p. 179), who wrote, '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.'

2. The Trouble with Rational Expectations about Equilibrium Prices in Heterogeneous Agent Models

With this historical background, I now turn to the main thesis of this essay: in heterogeneous agent macroeconomics, the assumption of rational expectations about equilibrium prices is unrealistic 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, producing an extreme version of the curse of dimensionality.

As in the introduction's interest-rate example, rational expectations would have individuals forecast mortgage interest rates by predicting the entire cross-sectional distribution of different asset holdings in the economy—an implausible implication of the theory.

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

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

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 decentralisation in which firms own the economy's capital stock that 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

(I) households maximise, taking $\{w_t, r_t\}$ as given:

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

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

(2) firms maximise, taking $\{w_t, r_t\}$ as given:

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

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

(3) markets clear:

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

The key difficulty, which is the focus of this paper, is that optimising households and firms 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 optimisation (3) yields labour 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 optimisation (4) yields labour 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 $r_t = r^*(k_t, z_t)$. 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 stylised representative-agent economy, a rational expectations equilibrium is the solution to a relatively complicated fixed-point problem and that the informational requirements for decision-makers are substantial and arguably unrealistic.

2.1.1. *Solution methods for the representative-agent case*

For questions where aggregate risk and non-linearities matter, one needs a global solution method. There are two main types.¹⁵ The first tackles the competitive equilibrium directly. Even in a simple RBC model like that I just outlined, global solution methods for the competitive equilibrium are actually surprisingly challenging to implement. A classic method uses dynamic programming with the 'Big K , little k ' trick (e.g., Lucas, 1987; Ljungqvist and Sargent, 2004; Prescott and Mehra, 1980), but there are various other ones; see, for example, Maliar and Maliar (2014) and Cao *et al.* (2023).

The second, more common, approach is to invoke the second welfare theorem and solve for the equilibrium allocation via a social planner's problem. This approach works only in special cases, but the key payoff is that the planner's problem does not feature prices, so it sidesteps the main difficulty emphasised above—expectations about equilibrium prices.

¹⁵ Many well-developed *local* solution methods exist for representative-agent models with aggregate risk; see, e.g., the popular Dynare toolbox.

2.2. Forecasting Equilibrium Prices in Heterogeneous Agent Models

Next, consider the 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 labour productivity y_{it} that 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 $G_t(a, y)$.

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

(1) households maximise, taking $\{w_t, r_t\}$ as given:

$$\max_{\{c_{it}, n_{it}, a_{it+1}\}} \mathbb{E}_0 \sum_{t=0}^T \beta^t U(c_{it}, n_{it}) \quad \text{such that} \quad c_{it} + a_{it+1} = w_t y_{it} n_{it} + (1 + r_t) a_{it}, \quad (5)$$

(2) firms maximise, taking $\{w_t, r_t\}$ as given—the same problem as in the representative-agent economy (4) since the economy's production side is unchanged,

(3) markets clear:

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

Importantly, 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 optimisation (5) now yields *household-specific* labour supplies $n_{it} = n(w_t, a_{it}, y_{it})$.¹⁶ Labour 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 :

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

The same is true for the equilibrium interest $r_t = r^*(G_t(a, y), z_t)$. This follows from the optimality condition of the representative firm $r_t = z_t F_k(k_t, \ell_t) - \delta$ and the fact that $\ell_t = \int n_t(a, y) y dG_t(a, y)$.

2.2.1. 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

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

¹⁷ For example, if labour supply were perfectly inelastic (as in the benchmark economy of Krusell and Smith, 1998) or if preferences ruled out income effects, the equilibrium wage and interest rate would depend only on the aggregate capital stock k_t . However, even in this case, it would still be true that G becomes a state variable in households' decision problem. Ríos-Rull (1997) explained 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.'

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

$$p_t = \mathcal{P}^*(G_t(x), z_t), \tag{6}$$

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 (6). The vector p_t should therefore be understood as including, not only prices themselves, but also other price-like variables.

2.2.2. *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 (5) can be written recursively as

$$V_0(a, y, G, z) = \max_{c, n, a'} U(c, n) + \beta \mathbb{E}[V_1(a', y', G', z') | y, G, z] \tag{7a}$$

$$\text{such that } c + a' = w_0^*(G, z)yn + (1 + r_0^*(G, z))a, \tag{7b}$$

$$V_1(a', y', G', z') = \max_{c', n'} U(c', n') \tag{7c}$$

$$\text{such that } c' = w_1^*(G', z')y'n' + (1 + r_1^*(G', z'))a'. \tag{7d}$$

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 $t = 1$ that is induced by households' optimal saving policy $a' = s_0(a, y, G, z)$ and the transition probabilities for labour income y , starting from the initial distribution G : $G' = \mathcal{T}_{s_0}G$ for some operator \mathcal{T}_{s_0} that depends on function s_0 .¹⁸

Importantly, the expectation \mathbb{E} in the first line of (7) is taken, not only over future productivity realisations 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 fourth 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 Den Haan (1996) and Krusell and Smith (1998). It is also 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.

¹⁸ In the finite-state variant of Section 3.5, the analogue of 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 (7) 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 (7) is a two-period version of the Master equation from the Mean Field Games literature (Cardaliaguet *et al.*, 2019) that 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, the work of Den Haan (1996) and Krusell and Smith (1998) already 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 $V(x, G, z)$ that features the infinite-dimensional state variable G . In the MFG literature this Master equation has aptly been nicknamed the 'Monster equation', precisely due to 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 discretised, they are extremely high-dimensional histograms, say 100,000-dimensional vectors (Section 3.5 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: operator \mathcal{T}_{s_0} depends on the saving policy function $s_0(a, y, G, z)$ that summarises *other* households' saving decisions. Using the notation of Section 3.5 below, forecasting future distributions requires knowledge of the huge transition matrix \mathbf{A} for the whole economy that summarises how all *other* decision-makers transition across the economy's state space. In the words of Morgenstern (1935, p. 342), 'the forward-looking individual must [...] know not only exactly the influence of his own actions on prices, but also that of all other individuals.'

2.2.3. *The unrealism of rational expectations*

This paper's main argument is that this forecasting behaviour implied by rational expectations is extremely implausible. If even our most advanced computational tools struggle with the Monster equation, how can we justify the assumption that real-world households and firms solve the associated decision problems? It seems self-evident that households and firms do not forecast prices by forecasting cross-sectional distributions. Instead of solving Monster equations we should replace the rational expectations assumption and solve the simpler equations corresponding to their actual price-forecasting behaviour. Below I 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) wrote:

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

¹⁹ As Sargent (2024) pointed 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 $(k, G(k))$ with individual state k and distribution $G(k)$ and (k, K) is the special case in which all individuals are identical (G is a Dirac point mass at K).

maker or empirical economist, the inferences that 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. (p. 1335–1336)

Similarly Adam and Marcet (2011) wrote:

‘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. (p. 1225)

See also the excerpts from Morgenstern (1935) in the introduction and [Online Appendix C](#).

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.

2.2.4. *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 linearises 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) used ‘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) used linearisation methods. In fact, Boppart *et al.* (2018) pointed out a close connection between the two approaches. With MIT shocks, there is certainty about the path of equilibrium prices. To the same effect, linearisation 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. The same is true, though to a lesser extent, for other *local* methods like second-order perturbation (Bhandari *et al.*, 2023; Bilal, 2023; Bayer *et al.*, 2024b; Auclert *et al.*, 2025) that may miss important phenomena like non-linear feedback loops and state spaces with crisis regions; see Section 3.4 below.

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 (Pröhl, 2019; Schaab, 2020; Maliar *et al.*, 2021; Azinovic *et al.*, 2022; Han *et al.*, 2022; Bilal, 2023; Huang, 2023; Gopalakrishna *et al.*, 2024; Gu *et al.*, 2024). 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 Den Haan (1996) and Krusell and Smith (1998) and variants thereof.²⁰ Compared to solving the full rational expectations equilibrium in which

²⁰ See, for example, the continuous-time variants by Lee (2022) and Fernández-Villaverde *et al.* (2023).

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 $\bar{a}_t = \int adG_t(a, y)$ 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) discussed, this approach has a bounded rationality interpretation: decision-makers assume that distributional moments follow a (typically log-linear) perceived law of motion that 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 in order to forecast prices 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.²¹

There is one variant of the Krusell–Smith approach that 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), Lull (2018), Kaplan *et al.* (2020) and Fernández-Villaverde *et al.* (2024a). I 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, Lucas, Prescott and their contemporaries aimed to develop operational macroeconomic theories that can be computed, fit to data and that generate clear policy counterfactuals. In representative-agent models, rational expectations—though arguably unrealistic—have achieved exactly this goal.

In heterogeneous agent macroeconomics, however, the rational expectations assumption has the opposite effect, severely hampering the models’ operability. It attributes to people an understanding of the price functions that map cross-sectional distributions to equilibrium prices, implying that decision-makers forecast prices by forecasting distributions. This implication is, not only unrealistic, but also makes such models challenging and costly to solve numerically.

The natural way forward is to replace rational expectations about equilibrium prices in heterogeneous agent models with plausible alternatives. While there are some precedents in the existing literature (most notably the variant of the Krusell–Smith approach discussed above), it is worth developing such alternatives more systematically and to incorporate key empirical evidence. Doing so promises to ‘kill two birds with one stone’: simplifying 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.

²¹ I am not aware of any empirical evidence that supports the idea that decision-makers forecast, say, the mean of the US wealth distribution to forecast interest rates. Broer *et al.* (2022b) criticised the idea that decision-makers use common moments to forecast prices from a theoretical perspective.

3. Three Criteria for Replacing Rational Expectations in Heterogeneous Agent Models

What should replace the assumption of rational expectations about equilibrium prices? I propose 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 by forecasting distributions. While this seems to me the natural solution, any such model of expectations formation necessarily departs from the rational benchmark.²² The central question is how to discipline the subjective probability distributions of equilibrium prices, i.e., how to navigate the ‘wilderness of non-rational expectations’. My three criteria aim 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 (7). The most natural alternative is to instead assume that they forecast prices directly in some way.

To flesh this out in more detail, recall the two-period Master equation (7) that features an expectation over the cross-sectional distribution G' . The rough idea is that households instead solve

$$V_0(a, y, w, r) = \max_{c, n, a'} U(c, n) + \beta \tilde{\mathbb{E}}[V_1(a', y', w', r')|\cdot] \tag{8a}$$

$$\text{such that } c + a' = wyn + (1 + r)a, \tag{8b}$$

$$V_1(a', y', w', r') = \max_{c', n'} U(c', n') \quad \text{such that } c' = w'y'n' + (1 + r')a', \tag{8c}$$

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 maximise 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 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 (8) with prices p as state variables in the rational expectations case? That is, why can we not just write (8) with

²² While all alternatives discussed below drop rational expectations about equilibrium prices, for simplicity, they keep rational expectations about non-equilibrium variables, e.g., idiosyncratic y_{it} .

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 3.5 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 neighbourhoods, 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 (8) 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. Computational tractability,
2. Consistency with empirical evidence,
3. Endogeneity of beliefs to model reality (the 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 (COMPUTATIONAL TRACTABILITY). 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) several popular expectations models. In particular, it eliminates any approach that begins with rational expectations as a special case and then derives the subjective probability distribution \mathbb{P} by merely ‘twisting’ this objective distribution. This is because the rational expectations case already suffers from the extreme curse of dimensionality discussed above, and any generalisation is inevitably at least as complex.²³

More precisely, several models of non-rational expectations in the literature index subjective beliefs by a parameter θ , i.e., \mathbb{P}^θ , with $\theta = 0$ corresponding to rational expectations, so that $\mathbb{P}^{\theta=0}$ equals the objective 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

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

where D is recently observed data, $-D$ is data in a relevant comparison group and $\theta \geq 0$. When $\theta = 0$, the special case of rational expectations is obtained. Similarly, in the cognitive discounting approach of Gabaix (2020), subjective expectations $\tilde{\mathbb{E}}$ are defined from the corresponding rational expectations \mathbb{E}^{RE} 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} is the deviation from this default value, $\bar{m} \in [0, 1]$ is a ‘cognitive discounting’ parameter and $\bar{m} = 1$ corresponds to rational expectations; equivalently, $\theta = 1 - \bar{m}$ with rational case $\theta = 0$. A third example is the recursive competitive equilibrium with behavioural agents in Section 5.3 of Gabaix (2023). All of these models do not satisfy my first criterion.²⁴

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 behavioural 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

²³ Another class of models that does not satisfy 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. LIRE includes noisy rational expectations (Lucas, 1972) and rational inattention (Maćkowiak *et al.*, 2023; Sims, 2003).

²⁴ That 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 behavioural agents in which decision-makers use a simplified model for equilibrium prices (e.g., an autoregressive process) that endogenously becomes more sophisticated when the stakes increase or the attention cost falls.

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 candidate alternative approach should be consistent with established findings 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., Mankiw *et al.*, 2004; Manski, 2004, 2018; Armantier *et al.*, 2013; Coibion *et al.*, 2022; D'Acunto and Weber, 2024). There are now new surveys like the New York Fed's Survey of Consumer Expectations and even a *Handbook of Economic Expectations* (Bachmann *et al.*, 2023).

This literature has documented a number of empirical patterns, some of which I summarise 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 that 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, modelling 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 well suited to modelling belief heterogeneity: after all, modelling and 'carrying around' cross-sectional distributions is our bread and butter—so why not also add a distribution of beliefs to these models?²⁵

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. Second, evidence on the pass-through from shocks and policies to subjective beliefs (Section 4.2 below). 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 (THE LUCAS CRITERION)). The candidate alternative approach should imply price beliefs that are endogenous to the behaviour of actual equilibrium prices. It is useful to split this criterion into two sub-criteria.

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

²⁵ Guerreiro (2023) took a step in this direction, but linearised 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, linearised models are not suitable for modelling phenomena like financial crises. For other dynamic models of belief heterogeneity—though many of them in restrictive linear-Gaussian setups—see, for example, Mankiw and Reis (2002), Carroll (2003), Giannitsarou (2003), Woodford (2003; 2018), Lorenzoni (2009), Angeletos and La'O (2013), Branch and McGough (2018), Carroll *et al.* (2020), Chahrour and Gallo (2020), Bouchaud and Farmer (2023) and Straub and Ulbricht (2023).

(b) In all environments, including unfamiliar and non-stationary ones, price beliefs should respond to model reality, specifically to policy changes (the 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 behaviour 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

$$\|\mathbb{P}(p|\cdot) - \mathbb{P}^{obj}(p|\cdot)\| \leq \varepsilon \quad (9)$$

for some distance metric (norm) $\|\cdot\|$ and some ‘small’ scalar $\varepsilon \geq 0$. While rational expectations impose 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; 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) wrote:

It makes sense to assume that expectations should not be completely arbitrary and have no relation to the kind of world in which 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.

See also Sargent (1993; 1999) and Ilut and Schneider (2014), who imposed an approximate consistency criterion similar in spirit to (9) for worst- and best-case means in a model with ambiguity aversion.²⁶

The role of Criterion 3b is that Criterion 3a may be too strong in many contexts, specifically in ‘unfamiliar’ situations that have ‘not been seen before’. Extreme cases include the COVID pandemic, the 9/11 terrorist attacks or the 2000s housing boom, but many less dramatic real-world events also challenge the assumption of stationarity. 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) wrote:

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. (p. 29)

²⁶ Sargent (2007) invoked Abraham Lincoln’s famous phrase to describe the philosophy behind rational expectations: ‘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 is designed to ensure that this remains true. Note that Lincoln’s quote explicitly allows for people being fooled *some* of the time and therefore arguably aligns more closely with the types of beliefs advocated here than with the strict rational expectations assumption itself.

Bouchaud and Farmer (2023) argued that the rational expectations assumption is difficult to justify even beyond such unusual circumstances because most real-world stochastic processes are neither stationary nor ergodic.

Nevertheless, while such reasoning supports relaxing Criterion 3a, price beliefs should not be completely exogenous to the model. Because equilibrium prices are themselves endogenous, beliefs about them must be modelled as at least partially endogenous as well. 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 modelled as endogenous. The behaviour of actual equilibrium prices depends on price beliefs and, conversely, price beliefs depend on the behaviour 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 criticised 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 behavioural inattention (but 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 behaviour for equilibrium prices.

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.

3.4. *The Need for Global and Recursive Solution Methods*

In addition to these three criteria for the subjective price beliefs \mathbb{P} , an important practical requirement concerns the class of models that the candidate approach can be applied to: non-linear models that are solved using global, recursive solution methods.

3.4.1. *Why global solution 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 fully non-linear models. In contrast, much of the theoretical expectations literature has considered linear (or linearised) models, likely because it means that standard techniques like the Kalman filter can be applied. This would need to change.

As already noted in the introduction, what I have in mind are heterogeneous agent versions of models like Mendoza (2010), Bianchi (2011), He and Krishnamurthy (2012; 2013), Brunnermeier and Sannikov (2014), Elenev *et al.* (2021) and Krishnamurthy and Li (2025) that feature

phenomena like non-linear feedback loops, state spaces with crisis regions or bimodal ergodic distributions for the aggregate economy. Such models need to be solved using global solution methods because local perturbation methods fail to capture their behaviour.

A useful prototype for such models is the scalar ‘double-well’ or ‘bistable’ diffusion process in [Online Appendix A](#) that features a parameter σ that parameterises uncertainty. Without uncertainty, $\sigma = 0$, there are three steady states, two of them stable and one unstable. With uncertainty, $\sigma > 0$, the process has a bimodal stationary distribution. Trajectories spend most time oscillating around the two stable steady states while occasionally hopping between them. Starting from the high steady state, ‘the economy’ stays close to it most of the time, but may be thrown into a ‘crisis’ (low steady state) and get stuck there for a while before ultimately recovering. Perturbation methods around $\sigma = 0$ completely miss this type of behaviour (e.g., Gardiner, 2009, ch.7.2.4). Standard references on perturbation methods also contain other examples of behaviour such methods cannot capture.

Finally, even with less complicated dynamics, perturbation methods may be quantitatively off. See Farmer and Toda (2017) for an example of a representative-agent asset-pricing model in which even a fifth-order perturbation method is highly inaccurate relative to global methods.

3.4.2. Why recursive methods?

Similarly, the alternative approach must be compatible with recursive solution methods for household and firm optimisation problems (Bellman equations). These lie at the heart of solution methods for current heterogeneous agent models for two reasons. First, recursive methods naturally handle micro non-linearities such as borrowing constraints. Second, rather than solving a separate optimisation problem for each of a large number of heterogeneous decision-makers (i.e., millions of optimisation problems), recursive methods often require solving only a single Bellman equation, with the resulting policy function summarising the actions of (ex post) heterogeneous agents across the state space.

3.5. With Rational Expectations, Why Can We Not Carry Prices As State Variables?

Before discussing promising directions, I briefly revisit the question: with rational expectations, why can we not just write a Bellman equation with prices p as state variables? Why can we not write (8) with the subjective expectation $\tilde{\mathbb{E}}$ replaced by the rational expectation \mathbb{E} ?

The short answer is that equilibrium 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, 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 $\mathbf{g}_t = [g_{1,t}, \dots, g_{N,t}]^T$ that 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:

$$p_t = \mathcal{P}^*(\mathbf{g}_t, z_t). \quad (10)$$

I use regular notation for the vector p_t and 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):

$$\mathbf{g}_{t+1} = \mathbf{A}(\mathbf{g}_t, z_t, p_t)^T \mathbf{g}_t. \quad (11)$$

Here $\mathbf{A}(\mathbf{g}_t, z_t, p_t)$ is the transition matrix of the vector of idiosyncratic states x (with x evolving according to the optimal policies from agents' Bellman equation).²⁷ Note that this is a large matrix of dimension $N \times N$, say, with $N = 100,000$, that summarises how all individual decision-makers transition across the economy's state space. Importantly, after substituting in for p_t from (10), this equation defines a *Markov process for the distribution \mathbf{g}_t and aggregate shock z_t* : we have $\Pr(\mathbf{g}_{t+1}, z_{t+1} | \cdot) = \Pr(\mathbf{g}_{t+1}, z_{t+1} | \mathbf{g}_t, z_t)$.

Summarising, prices p_t depend on the distribution and aggregate shock (\mathbf{g}_t, z_t) according to the equilibrium price function \mathcal{P}^* in (10) 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 (8) 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 perspective comes from the analogy with a hidden Markov model (HMM): the distribution \mathbf{g}_t is a high-dimensional latent state satisfying the Markov property, while p_t is a low-dimensional observation derived from it. Since p_t alone is not Markov, forecasting it requires forecasting the high-dimensional Markov state. Recognising this structure in heterogeneous agent models could guide alternatives to rational expectations. For instance, if decision-makers track only prices rather than the full distribution (as in Section 3.1), their optimisation problems become partially observable Markov decision processes (POMDPs), i.e., controlled HMMs.

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.

²⁷ See Achdou *et al.* (2021) for a continuous-time analogue and an explanation of 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.

4.1. *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 optimise, 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.²⁸ Rational expectations equilibrium is the special case of temporary equilibrium in which expectations are rational (model consistent).²⁹

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, then 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 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 3.5 tracking vector \mathbf{g}_t simply means running the difference equation (11) 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, this can only be an intermediate step toward solving the bigger fixed-point problem in which the behaviour of actual equilibrium prices feeds back into endogenous price beliefs.

4.1.1. *Internal rationality*

When defining temporary equilibria, it is important to derive individual policy functions from intertemporal decision problems with dynamically consistent subjective beliefs about equilib-

²⁸ These include Woodford (2013), Asparouhova *et al.* (2016), Piazzesi and Schneider (2016), Farhi and Werning (2019), García-Schmidt and Woodford (2019) and Werning (2022). Although they do not connect to the temporary equilibrium idea, another application is the demand-system asset-pricing approach of Kojien and Yogo (2019), which replaces investor’s beliefs about high-dimensional equilibrium asset returns with a low-dimensional estimated factor model.

²⁹ Hicks (1939, ch.IX.7 and ch.X.2) differentiated ‘temporary equilibrium’ from another notion of equilibrium he calls ‘equilibrium over time’ that bears some resemblance to a rational-expectations equilibrium. He explained the difference in a dynamic economy with time intervals lasting one week: ‘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. [...]’, p. 132. See also Hicks (1979), who defined temporary equilibrium as ‘a momentary market equilibrium in which price-expectations are taken as data’, p. 200.

rium prices. This links temporary equilibrium approaches to the concept of *internal rationality* proposed by Adam and Marcet (2011).

Adam and Marcet distinguished 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 argued 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 termed this approach *internal rationality* to emphasise that decision-makers optimise, 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.

4.1.2. *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.

[Online Appendix B](#) summarises a number of these—specifically, self-confirming equilibrium (SCE), restricted perceptions equilibrium (RPE) and consistent expectations equilibrium (CEE). The [Online 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 summarised as

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

where '⊂' 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; 2015), Giglio *et al.* (2021) and Fofana *et al.* (2025). 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’.³⁰

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), Caplin *et al.* (2023), Mueller and Spinnewijn (2023) and Balleer *et al.* (2024).

Additionally, one would really need measures of these expectations *conditional* on the state of the aggregate economy (recessions versus 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 (Andre *et al.*, 2022; Haaland *et al.*, 2023).

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) elicited survey respondents’ subjective beliefs and additionally asked how certain they were about these stated beliefs (‘cognitive uncertainty’). They found 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.³¹ This concern about systematic noise in survey 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) and 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).

³⁰ For example, Cenyon (2024) documented 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) argued that inflation expectations react excessively to household-level income shocks. See also Bordalo *et al.* (2025). This evidence, as well as the evidence on subjective earnings risk in the next paragraph, is also of independent interest for heterogeneous agent modelling, for example it could have important implications for the distribution of wealth.

³¹ Put differently, the concern about systematic distortions is that cognitively uncertain respondents may appear 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.

4.2.1. *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), Bardóczy and Guerreiro (2023), Aron-Dine *et al.* (2024) and Ludwig *et al.* (2024) for recent applications.

Roth *et al.* (2023) extended 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 (the 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 favour of their approach by pointing to the falling cost of conducting surveys and economists’ disagreement about modelling 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, ch.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) that I also discuss below.³² Reinforcement learning, which I review in Section 4.4 below, is another stochastic approximation method.

4.3.1. *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) that we already encountered in Section 1.1.³³ 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 these expectations from past data. That is, producers form estimates $\widehat{p}_{t+1}^e = \widehat{\mathbb{E}}_t[p_{t+1}]$ from observations of past prices p_1, \dots, p_t and supply shock realisations 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

³² Sargent (2024) linked stochastic approximation back to Hotelling (1941) and Friedman and Savage (1947).

³³ While I use somewhat different notation, this section follows closely the presentation of Evans and Honkapohja (2001), who also started from Muth’s example. See also the presentation of Bray (1982) in chapter 5 of Sargent (1993) that starts directly from the reduced form (2), Christiano *et al.* (2024) and Baley and Veldkamp (2025).

taking the sample average of past price observations

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

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 as³⁴

$$\widehat{p}_{t+1}^e = \widehat{p}_t^e + \frac{1}{t} [p_t - \widehat{p}_t^e]. \quad (12)$$

Note how this update rule computes the new estimate \widehat{p}_{t+1}^e from the previous estimate \widehat{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 independent and identically distributed (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 $\widehat{p}_{t+1}^e = \widehat{\mathbb{E}}_t[p_{t+1}]$, producers then (i) estimate (θ_0, θ_1) from observations of past prices p_1, \dots, p_t and supply shock realisations u_0, \dots, u_{t-1} via least squares, and (ii) use the time- t estimates $(\widehat{\theta}_0, \widehat{\theta}_1)$ to compute $\widehat{p}_{t+1}^e = \widehat{\mathbb{E}}_t[p_{t+1}] = \widehat{\theta}_0 + \widehat{\theta}_1 u_t$.

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

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

Analogously to (12), there is again a recursive implementation that computes the new estimate $\widehat{\theta}_{t+1}$ from the previous estimate $\widehat{\theta}_t$ and the latest data (p_{t+1}, u_t) ; see, for example, equation (2.9) of 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 can be applied to decision-makers learning perceived laws of motion of equilibrium prices. Jacobson (2025) implemented this approach in a housing model, where households learn about house prices using a simple variant of least-squares learning (stochastic gradient learning as in Evans and Honkapohja, 1998) and used it to study the US housing boom of the 2000s. Nakov and Nuño (2015) modelled overlapping generations of investors learning about equilibrium stock prices from experience (as in Malmendier and Nagel, 2011; 2015).

³⁴ The derivation is $\widehat{p}_{t+1}^e = (1/t) \sum_{s=1}^t p_s = (1/t)[p_t + \sum_{s=1}^{t-1} p_s] = (p_t + (t-1)\widehat{p}_t^e)/t = \widehat{p}_t^e + [p_t - \widehat{p}_t^e]/t$.

Related, there is a link between least-squares learning and Krusell–Smith methods (see Section 2.2).³⁵ 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’ (Moll and Ryzhik, 2025). This also means that coefficient estimates vary over time (which they do not in the Krusell–Smith model) and makes least-squares learning more likely to generate interesting phenomena, such as booms and busts (Jacobson, 2025).

4.3.2. Stochastic approximation

The update rule (12) 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 (for example, in the next subsection). Using the computer science notation to let ‘←’ denote an assignment statement, Sutton and Barto (2018, chapter 2.4) wrote the general form for this type of update rule as

$$\text{NewEstimate} \leftarrow \text{OldEstimate} + \text{StepSize}[\text{Target} - \text{OldEstimate}]. \quad (13)$$

The ‘Target’ is the target value that the algorithm attempts to drive the estimate to, i.e., it indicates the direction in which the estimate is updated, while the step size controls how large this update is. For example, in the sample-average case (12), the target is the time- t price p_t , meaning that the new estimate \hat{p}_{t+1}^e is updated in the direction of 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). See Zhao (2025, ch.6) for a good textbook discussion.

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 become more and more plentiful). To this end, denote the step size in the update rule (13) by α_t , for example $\alpha_t = 1/t$ in the sample-average case (12).³⁶ One key result in stochastic approximation theory states that, as $t \rightarrow \infty$, the estimate being updated according to (13) converges with probability 1 under a number of conditions, including the following conditions on the step size:³⁷

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

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 (13) becomes closer and closer to the deterministic trajectory of a particular

³⁵ When discussing least-squares learning, Sargent (2017, p. 879) stated that ‘This is in effect what Krusell and Smith (1998) do, though they do not connect their method to the learning literature.’ The connection is explicit in Giusto (2014), who studied least-squares learning about the aggregate capital stock (rather than prices) in the Krusell–Smith model.

³⁶ Considering more general step sizes links the update rule (12) 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 that $\hat{p}_{t+1}^e = (1 - \alpha)^{t+1}\hat{p}_0^e + \alpha \sum_{j=0}^t (1 - \alpha)^j p_{t-j}$, a backward-looking weighted average of past prices with exponentially declining weights.

³⁷ As Sutton and Barto (2018, p. 39) explained: ‘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.’

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), Sargent (1999) and Evans and Honkapohja (2001).

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) $\widehat{\mathbb{E}}_t[p_{t+1}] = (1/t) \sum_{s=1}^t p_s$ 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’ (Online Appendix B) that may therefore be a relevant equilibrium concept in heterogeneous agent models.

4.3.3. *Slow or non-convergent learning*

An important caveat is that stochastic approximation convergence theorems hold only asymptotically as $t \rightarrow \infty$ and that learning may be slow in practice. See Christiano *et al.* (2024) as well as Sargent (1993), who voiced reservations about adaptive algorithms as theories of real-time dynamics, but stated that he nevertheless ‘like[s them] as devices for selecting equilibria.’ If the environment is sufficiently complex and non-stationary, it is also possible that learning never converges (e.g., Bouchaud and Farmer, 2023; Garnier-Brun *et al.*, 2024).

4.3.4. *Learning as a process*

In contrast to rational expectations, learning models explicitly specify the *process* by which individuals form expectations. This approach echoes Simon’s (1978b) distinction between ‘procedural’ and ‘substantive’ rationality. In a similar vein, Simon (1978a) wrote,

The theory of the business cycle is another important candidate area, for a procedural theory of the forming of expectations and the making of decisions. [...] One direction of progress is to erect theories that postulate [...] how expectations for the future are in fact formed by economic actors, and how those expectations enter into the calculations of actions. A realistic procedural theory would almost certainly have to include learning mechanisms. (p. 505)

As an aside, Simon (1978a) also featured a fascinating discussion of artificial intelligence from the perspective of its time.

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., Caplin and Dean, 2008; Niv, 2009; Glimcher, 2011; Glimcher *et al.*, 2013; Gershman and Daw, 2017).³⁸

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 3.5) and instead learn about prices in some way.³⁹ RL is linked to least-squares learning because both are stochastic approximation methods (see the previous subsection).

4.4.1. Reinforcement learning in a nutshell

I briefly summarise the basics of RL following the brilliant introductory treatments by Sutton and Barto (2018) and Zhao (2025).⁴⁰ RL is concerned with Markov decision processes of the type that are very familiar to economists: an agent maximises 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 state x_t is stochastic and depends on control c_t according to the transition probabilities $x_{t+1} \sim P(\cdot | x_t, c_t)$. My notation differs from the standard computer science notation (Sutton and Barto, 2018; Zhao, 2025) 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 from 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 only see the reward R_t they receive after taking action c_t (as well as the current state x_t and successor state x_{t+1}). Over time, agents learn to take actions that maximise rewards. Rewards are thus like ‘doggy treats’ that reinforce good behaviour, 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.

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

³⁸ RL 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. See also Barberis and Jin (2023), Chen *et al.* (2023), de la Barrera and de Silva (2024), Duarte *et al.* (2024a,b) and the references in Fernández-Villaverde *et al.* (2024b) for applications in finance and macroeconomics.

³⁹ A number of papers have applied RL to solving heterogeneous agent models *without* aggregate risk. See, for example, Laurière *et al.* (2022) and Xu *et al.* (2023). As should hopefully be clear, the present paper is instead about models *with* aggregate risk.

⁴⁰ See also Silver (2015), Murphy (2025) for an online course with video recordings, and Charpentier *et al.* (2023).




Fig. 1. *Model-Free versus Model-Based Reinforcement Learning* (Sutton and Barto, 2018, ch.8.2).

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 that 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. 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 $T = 1,000$).
- Initialise $V(x)$ 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

$$V(x_t) \leftarrow V(x_t) + \alpha \underbrace{[R_t + \beta V(x_{t+1}) - V(x_t)]}_{\text{Target}}. \quad (14)$$

This algorithm returns a value function $V(x)$ for all $x \in \mathcal{X}$. The recursive updating step (14) takes the form of the general stochastic approximation updating rule (13) 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; Zhao, 2025).

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* that 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*.⁴¹ 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 maximise 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 .

4.4.2. *Big world hypothesis*

Despite exponentially growing computational power, applications of RL to real-world learning problems have remained limited. Javed and Sutton (2024) argued 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 argued that algorithms should reflect this reality of ‘small bounded agents learning in large unbounded environments’ and proposed 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) illustrated 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. (p. 1)

4.4.3. *Reinforcement learning about equilibrium prices?*

In ongoing work with co-authors, 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. 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 3.5, adapting ideas from the theory of RL in non-Markovian environments or POMDPs could be a fruitful avenue.

⁴¹ The name *Sarsa* comes from the standard computer science notation that denotes states by S_t , actions by A_t and rewards by R_t and that agents update action-value functions from the sequence $(S_t, A_t, R_t, S_{t+1}, A_{t+1})$.

An interesting question is whether delivering on Criterion 2 (consistency with empirical evidence) will require modifying standard RL ideas. While RL will converge to an equilibrium different from rational expectations, a natural conjecture is that the resulting RL expectations will be ‘close to’ rational. This suggests that modifications of standard RL ideas may be needed to take on board empirical evidence on belief formation.

Another interesting question is whether standard model-free RL can satisfy Criterion 3 or whether a more model-based approach is needed. An advantage of model-free RL over typical adaptive learning methods in economics is that it avoids specifying a perceived law of motion for prices, which necessarily involves some arbitrariness. However, a completely model-free approach, in which decision-makers lack any structural understanding of the world, may be too simplistic (e.g., Enke, 2024). For both model-free and model-based RL methods, learning may also be slow in practice; see the discussion on convergence speed above. Since RL and least-squares learning are both special cases of stochastic approximation algorithms, this concern applies equally to both.

RL is a machine learning method, so it is natural to ask how such an approach relates to recent work applying deep learning methods to solve heterogeneous agent models with aggregate risk (e.g., Azinovic *et al.*, 2022; Han *et al.*, 2022; Gopalakrishna *et al.*, 2024; Gu *et al.*, 2024).⁴² These papers typically 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. Such work uses deep learning to ‘tame the curse of dimensionality’ (Fernández-Villaverde *et al.*, 2024b), whereas RL about equilibrium prices would instead *sidestep* this curse by modelling decision-makers as solving lower-dimensional problems.

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, p. 1), who argued 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’, and Gigerenzer *et al.* (2000).

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.* (2012; 2010), Hajdini (2023), Flynn and Sastry (2024) and Molavi (2024). Many (but not all) of these authors impose some consistency between subjective beliefs and model reality, as in a restricted perceptions equilibrium (Online Appendix B). 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 (computational tractability), in particular the point that models of non-rational expectations

⁴² With the exception of Han *et al.* (2022), these papers 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; 2024a), i.e., a non-linear version of Krusell and Smith (1998).

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 use popular behavioural biases like diagnostic expectations or cognitive discounting to ‘distort’ the simplified forecasts (not the rational ones). Third, the restricted perceptions equilibrium would deliver Criterion 3 (endogeneity of beliefs to 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) pointed 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) computational tractability, (2) consistency with empirical evidence and (3) endogeneity of beliefs to model reality (the 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 specialising in heterogeneous agent models and those working on the theory and empirics of belief formation should work together and capitalise 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, p. 13) called 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.’⁴³ I agree. However, it is also important that researchers, not only consider deviations from rational expectations in linear or linearised models (as they often do).

⁴³ Broer *et al.* (2022a) took 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.

In a similar vein, we need to develop efficient *global* solution methods for heterogeneous agent models. By focussing on linear or other local solution methods, the heterogeneous agent literature risks repeating the mistakes of the pre-financial crisis representative-agent literature. See, for example, Blanchard (2014), who wrote:

The reason for this assumption, called linearity, was technical: models with nonlinearities [...] were difficult, if not impossible, to solve under rational expectations. Thinking about macroeconomics was largely shaped by those assumptions. We in the field did think of the economy as roughly linear, constantly subject to different shocks, constantly fluctuating, but naturally returning to its steady state over time. [...] The problem is that we came to believe that this was indeed the way the world worked.

A weaker form of Blanchard's criticism applies more generally to methods that solve models only locally around a steady state and which may therefore miss important global dynamics; see Section 3.4. Such local methods remind me a bit of the old joke about the drunk who is looking under a lamppost for a key that he has lost on the other side of the street because 'this is where the light is'.

The theoretical expectations literature should also allocate more time to modelling expectations about endogenous variables like equilibrium prices (as in Bastianello and Fontanier, 2025 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. In the context of this paper, this defeats the very purpose of departing from rational expectations.

London School of Economics and Political Science, UK

Additional Supporting Information may be found in the online version of this article:

Online Appendix

References

- Achdou, Yves, Jiequn Han, Jean-Michel Lasry, Pierre-Louis Lions, and Benjamin Moll. (2021). 'Income and wealth distribution in macroeconomics: A continuous-time approach', *The Review of Economic Studies*, vol. 89(1), pp. 45–86.
- Adam, Klaus, and Albert Marcet. (2011). 'Internal rationality, imperfect market knowledge and asset prices', *Journal of Economic Theory*, vol. 146(3), pp. 1224–52.
- Adam, Klaus, Albert Marcet, and Johannes Beutel. (2017). 'Stock price booms and expected capital gains', *American Economic Review*, vol. 107(8), pp. 2352–408.
- Ahn, SeHyoun, Greg Kaplan, Benjamin Moll, Thomas Winberry, and Christian Wolf. (2018). 'When inequality matters for macro and macro matters for inequality', *NBER Macroeconomics Annual*, vol. 32(1), pp. 1–75.
- Aiyagari, S. Rao. (1994). 'Uninsured idiosyncratic risk and aggregate saving', *The Quarterly Journal of Economics*, vol. 109(3), pp. 659–84.
- Andre, Peter, Ingar Haaland, Christopher Roth, Mirko Wiederholt, and Johannes Wohlfart. (2024). 'Narratives about the macro-economy', Working paper, SAFE.
- Andre, Peter, Carlo Pizzinelli, Christopher Roth, and Johannes Wohlfart. (2022). 'Subjective models of the macroeconomy: Evidence from experts and representative samples', *The Review of Economic Studies*, vol. 89(6), pp. 2958–91.
- Angeletos, George-Marios, Zhen Huo, and Karthik A. Sastry. (2021). 'Imperfect macroeconomic expectations: Evidence and theory', *NBER Macroeconomics Annual*, vol. 35(1), pp. 1–86.
- Angeletos, George-Marios, and Jennifer La'O. (2013). 'Sentiments', *Econometrica*, vol. 81(2), pp. 739–79.
- Angeletos, George-Marios, and Chen Lian. (2023). 'Dampening general equilibrium: Incomplete information and bounded rationality', in (Rüdiger Bachmann, Giorgio Topa and Wilbert van der Klaauw, eds.), *Handbook of Economic Expectations*, pp. 613–45, Academic Press.

- Armantier, Olivier, Wändi Bruine de Bruin, Simon Potter, Giorgio Topa, Wilbert van der Klaauw, and Basit Zafar. (2013). 'Measuring inflation expectations', *Annual Review of Economics*, vol. 5(1), pp. 273–301.
- Aron-Dine, Shifrah, Johannes Beutel, Monika Piazzesi, and Martin Schneider. (2024). 'Household climate finance: Theory and survey data on safe and risky green assets', Working paper, National Bureau of Economic Research.
- Asparouhova, Elena, Peter Bossaerts, Nilanjan Roy, and William Zame. (2016). '“Lucas” in the laboratory', *The Journal of Finance*, vol. 71(6), pp. 2727–80.
- Auclert, Adrien, Bence Bardóczy, Matthew Rognlie, and Ludwig Straub. (2021). 'Using the sequence-space Jacobian to solve and estimate heterogeneous agent models', *Econometrica*, vol. 89(5), pp. 2375–408.
- Auclert, Adrien, Rodolfo Rigato, Matthew Rognlie, and Ludwig Straub. (2025). 'Second order solutions in the sequence space', Work in Progress.
- Axtell, Robert L., and J. Doyne Farmer. (2025). 'Agent-based modelling in economics and finance: Past, present, and future', *Journal of Economic Literature*, vol. 63, pp. 197–287.
- Azinovic, Marlon, Luca Gaegauf, and Simon Scheidegger. (2022). 'Deep equilibrium nets', *International Economic Review*, vol. 63(4), pp. 1471–525.
- Bachmann, Rüdiger, Steffen Elstner, and Eric R. Sims (2013). 'Uncertainty and economic activity: Evidence from business survey data', *American Economic Journal: Macroeconomics*, vol. 5(2), pp. 217–49.
- Bachmann Rüdiger, Topa Giorgio and van der Klaauw Wilbert (eds.). (2023). *Handbook of Economic Expectations*, Academic Press.
- Baley, Isaac, and Javier Turén. (2024). 'Lumpy forecasts', Working paper, Universitat Pompeu Fabra.
- Baley, Isaac, and Laura L. Veldkamp (2025). *The Data Economy: Tools and Applications*, Princeton, NJ: Princeton University Press.
- Balleer, Almut, Georg Duernecker, Susanne K. Forstner, and Johannes Goensch. (2024). 'The effects of biased labor market expectations on consumption, wealth inequality, and welfare', Working paper.
- Barberis, Nicholas. (2018). 'Psychology-based models of asset prices and trading volume', in (B. Douglas Bernheim, Stefano DellaVigna and David Laibson, eds.), *Handbook of Behavioral Economics: Applications and Foundations*, vol. 1, pp. 79–175, North-Holland.
- Barberis, Nicholas, and Lawrence Jin. (2023). 'Model-free and model-based learning as joint drivers of investor behavior', Working Paper 31081, National Bureau of Economic Research.
- Bardóczy, Bence, and Joao Guerreiro. (2023). 'Unemployment insurance in macroeconomic stabilization with imperfect expectations', Working paper.
- de la Barrera, Marc, and Tim de Silva. (2024). 'Model-agnostic dynamic programming', Working paper, Stanford University.
- Bastianello, Francesca, and Paul Fontanier. (2025). 'Expectations and learning from prices', *The Review of Economic Studies*, vol. 92, pp. 1341–74.
- Bayer, Christian, Benjamin Born, and Ralph Luetticke. (2024a). 'Shocks, frictions, and inequality in US business cycles', *American Economic Review*, vol. 114(5), pp. 1211–47.
- Bayer, Christian, Ralph Luetticke, Maximilian Weiss, and Yannick Winkelmann. (2024b). 'An endogenous gridpoint method for distributional dynamics', Discussion Paper 19067, Centre for Economic Policy Research.
- Bhandari, Anmol, Thomas Bourany, David Evans, and Mikhail Golosov. (2023). 'A perturbational approach for approximating heterogeneous agent models', Working Paper 31744, National Bureau of Economic Research.
- Bianchi, Javier. (2011). 'Overborrowing and systemic externalities in the business cycle', *American Economic Review*, vol. 101(7), pp. 3400–26.
- Bilal, Adrien. (2023). 'Solving heterogeneous agent models with the master equation', Working Paper 31103, National Bureau of Economic Research.
- Bilal, Adrien, and Shlok Goyal. (2024). 'Some pleasant sequence-space arithmetic in continuous time', Working paper, Harvard University.
- Blanchard, Olivier. (2014). 'Where danger lurks', *Finance and Development*, vol. 51(3), pp. 22–3.
- Bloom, Nick, Mihai Codreanu, and Robert Fletcher. (2024). 'Rationalizing firm forecasts', Working paper, Stanford University.
- Boppart, Timo, Per Krusell, and Kurt Mitman. (2018). 'Exploiting MIT shocks in heterogeneous agent economies: The impulse response as a numerical derivative', *Journal of Economic Dynamics and Control*, vol. 89, pp. 68–92.
- Bordalo, Pedro, Giovanni Burro, Katherine Coffman, Nicola Gennaioli, and Andrei Shleifer. (2025). 'Imagining the future: Memory, simulation, and beliefs', *The Review of Economic Studies*, vol. 92, pp. 1532–63.
- Bordalo, Pedro, Nicola Gennaioli, and Andrei Shleifer. (2018). 'Diagnostic expectations and credit cycles', *The Journal of Finance*, vol. 73(1), pp. 199–227.
- Bordalo, Pedro, Nicola Gennaioli, and Andrei Shleifer. (2022). 'Overreaction and diagnostic expectations in macroeconomics', *Journal of Economic Perspectives*, vol. 36(3), pp. 223–44.
- Bordalo, Pedro, Nicola Gennaioli, Andrei Shleifer, and Stephen J. Terry (2021). 'Real credit cycles', Working Paper 28416, National Bureau of Economic Research.
- Bouchaud, Jean-Philippe. (2013). 'Crises and collective socio-economic phenomena: Simple models and challenges', *Journal of Statistical Physics*, vol. 151, pp. 567–606.
- Bouchaud, Jean-Philippe, and Roger E.A. Farmer (2023). 'Self-fulfilling prophecies, quasi nonergodicity, and wealth inequality', *Journal of Political Economy*, vol. 131(4), pp. 947–93.

- Branch, William A., and George W. Evans (2011). 'Learning about risk and return: A simple model of bubbles and crashes', *American Economic Journal: Macroeconomics*, vol. 3(3), pp. 159–91.
- Branch, William A., and Bruce McGough. (2018). 'Heterogeneous expectations and micro-foundations in macroeconomics', in (Cars Hommes and Blake LeBaron, eds.), *Handbook of Computational Economics*, pp. 3–62, Cambridge, MA: Elsevier.
- Bray, Margaret. (1982). 'Learning, estimation, and the stability of rational expectations', *Journal of Economic Theory*, vol. 26(2), pp. 318–39.
- Brock, William A., and Cars H. Hommes (1997). 'A rational route to randomness', *Econometrica*, vol. 65(5), pp. 1059–95.
- Broer, Tobias, Alexandre N. Kohlhas, Kurt Mitman, and Kathrin Schlafmann. (2022a). 'Expectation and wealth heterogeneity in the macroeconomy'.
- Broer, Tobias, Alexandre N. Kohlhas, Kurt Mitman, and Kathrin Schlafmann. (2022b). 'On the possibility of Krusell–Smith equilibria', *Journal of Economic Dynamics and Control*, vol. 141, p. 104391.
- Brunnermeier, Markus K., Emmanuel Farhi, Ralph S.J. Koijen, Arvind Krishnamurthy, Sydney C. Ludvigsson, Hanno Lustig, Stefan Nagel, and Monika Piazzesi. (2021). 'Review article: Perspectives on the future of asset pricing', *The Review of Financial Studies*, vol. 34(4), pp. 2126–60.
- Brunnermeier, Markus K., and Yuliy Sannikov. (2014). 'A macroeconomic model with a financial sector', *American Economic Review*, vol. 104(2), pp. 379–421.
- Caballero, Ricardo J., and Alp Simsek. (2022). 'A monetary policy asset pricing model', Working Paper 30132, National Bureau of Economic Research.
- Cao, Dan, Wenlan Luo, and Guangyu Nie. (2023). 'Global DSGE models', *Review of Economic Dynamics*, vol. 51, pp. 199–225.
- Cai, Michael. (2024). 'Explaining the macroeconomic inertia puzzle', Working paper, Northwestern University.
- Caplin, Andrew, and Mark Dean. (2008). 'Dopamine, reward prediction error, and economics', *The Quarterly Journal of Economics*, vol. 123(2), pp. 663–701.
- Caplin, Andrew, Victoria Gregory, Eungik Lee, Soren Leth-Petersen, and Johan Sæverud. (2023). 'Subjective earnings risk', Working Paper 31019, National Bureau of Economic Research.
- Cardaliaguet, Pierre, François Delarue, Jean-Michel Lasry, and Pierre-Louis Lions. (2019). *The Master Equation and the Convergence Problem in Mean Field Games*, Princeton, NJ: Princeton University Press.
- Carroll, Christopher D. (2003). 'Macroeconomic expectations of households and professional forecasters', *The Quarterly Journal of Economics*, vol. 118(1), pp. 269–98.
- Carroll, Christopher D., Edmund Crawley, Jiri Slacalek, Kiichi Tokuoka, and Matthew N. White (2020). 'Sticky expectations and consumption dynamics', *American Economic Journal: Macroeconomics*, vol. 12(3), pp. 40–76.
- Cenzon, Josefina. (2024). 'Credit market experiences and macroeconomic expectations: Evidence and theory', Working paper, Barcelona School of Economics.
- Chahrour, Ryan, and Gaetano Gaballo. (2020). 'Learning from house prices: Amplification and business fluctuations', *The Review of Economic Studies*, vol. 88(4), pp. 1720–59.
- Chari, Varadarajan V. (1998). 'Nobel Laureate Robert E. Lucas, Jr: Architect of modern macroeconomics', *Journal of Economic Perspectives*, vol. 12(1), pp. 171–86.
- Charles, Constantin, Cary Frydman, and Mete Kilic. (2024). 'Insensitive investors', *The Journal of Finance*, vol. 79(4), pp. 2473–503.
- Charpentier, Arthur, Romuald Élie, and Carl Remlinger. (2023). 'Reinforcement learning in economics and finance', *Computational Economics*, vol. 62(1), pp. 425–62.
- Chen, Mingli, Andreas Joseph, Michael Kumhof, Xinlei Pan, and Xuan Zhou. (2023). 'Deep reinforcement learning in a monetary model', Preprint, <https://arxiv.org/abs/2104.09368v2>.
- Christiano, Lawrence J., Martin S. Eichenbaum, and Benjamin K. Johansson (2024). 'Slow learning', Working Paper 32358, National Bureau of Economic Research.
- Coibion, Olivier, and Yuriy Gorodnichenko. (2012). 'What can survey forecasts tell us about information rigidities?', *Journal of Political Economy*, vol. 120(1), pp. 116–59.
- Coibion, Olivier, and Yuriy Gorodnichenko. (2015). 'Information rigidity and the expectations formation process: A simple framework and new facts', *American Economic Review*, vol. 105(8), pp. 2644–78.
- Coibion, Olivier, Yuriy Gorodnichenko, and Bernardo Candia. (2023). 'The macroeconomic expectations of firms', in (Rüdiger Bachmann, Giorgio Topa and Wilbert van der Klaauw, eds.), *Handbook of Economic Expectations*, pp. 321–53, Cambridge, MA: Academic Press.
- Coibion, Olivier, Yuriy Gorodnichenko, Michael Weber, and Francesco D'Acunto. (2022). 'The subjective inflation expectations of households and firms: Measurement, determinants, and implications', *Journal of Economic Perspectives*, vol. 36(3), pp. 157–84.
- D'Acunto, Francesco, and Michael Weber. (2024). 'Why survey-based subjective expectations are meaningful and important', *Annu. Rev. Econ.*, vol. 16, pp. 329–57.
- De Grauwe, Paul. (2010). 'Top-down versus bottom-up macroeconomics', *CESifo Economic Studies*, vol. 56(4), pp. 465–97.
- Den Haan, Wouter J. (1996). 'Heterogeneity, aggregate uncertainty, and the short-term interest rate', *Journal of Business and Economic Statistics*, vol. 14(4), pp. 399–411.

- Dominitz, Jeff, and Charles F. Manski (1997). 'Using expectations data to study subjective income expectations', *Journal of the American Statistical Association*, vol. 92(439), pp. 855–67.
- Duarte, Victor, Diogo Duarte, and Dejanir H. Silva (2024a). 'Machine learning for continuous-time finance', *The Review of Financial Studies*, vol. 37(11), pp. 3217–71.
- Duarte, Victor, Julia Fonseca, Aaron S. Goodman, and Jonathan A. Parker (2024b). 'Simple allocation rules and optimal portfolio choice over the lifecycle', Working paper, MIT.
- Elenev, Vadim, Tim Landvoigt, and Stijn Van Nieuwerburgh. (2021). 'A macroeconomic model with financially constrained producers and intermediaries', *Econometrica*, vol. 89(3), pp. 1361–418.
- Enke, Benjamin. (2024). 'The cognitive turn in behavioral economics', Working paper, Harvard University.
- Enke, Benjamin, and Thomas Graeber. (2023). 'Cognitive uncertainty', *The Quarterly Journal of Economics*, vol. 138(4), pp. 2021–67.
- Erev, Ido, and Alvin E. Roth (1995). 'Learning in extensive-form games: Experimental data and simple dynamic models in the intermediate term', *Games and Economic Behavior*, vol. 8(1), pp. 164–212.
- Erev, Ido, and Alvin E. Roth (1998). 'Predicting how people play games: Reinforcement learning in experimental games with unique, mixed strategy equilibria', *American Economic Association*, vol. 88 (4), 848–81.
- Evans, George W., and S. Honkapohja (1998). 'Stochastic gradient learning in the cobweb model', *Economics Letters*, vol. 61(3), pp. 333–7.
- Evans, George W., and Seppo Honkapohja. (2001). *Learning and Expectations in Macroeconomics*, Princeton, NJ: Princeton University Press.
- Evans, George W., and Bruce McGough. (2020). *Adaptive Learning in Macroeconomics*, New York, NY: Oxford University Press.
- Farhi, Emmanuel, and Iván Werning. (2019). 'Monetary policy, bounded rationality, and incomplete markets', *American Economic Review*, vol. 109(11), pp. 3887–928.
- Farmer, Leland E., and Alexis A. Toda. (2017). 'Discretizing nonlinear, non-Gaussian Markov processes with exact conditional moments', *Quantitative Economics*, vol. 8(2), pp. 651–83.
- Favilukis, Jack, Sydney C. Ludvigson, and Stijn Van Nieuwerburgh. (2017). 'The macroeconomic effects of housing wealth, housing finance, and limited risk sharing in general equilibrium', *Journal of Political Economy*, vol. 125(1), pp. 140–223.
- Fernández-Villaverde, Jesús, Galo Nuño, and Samuel Hurtado. (2023). 'Financial frictions and the wealth distribution', *Econometrica*, vol. 91(3), pp. 869–901.
- Fernández-Villaverde, Jesús, Galo Nuño, Joël Marbet, and Omar Rachedi. (2024a). 'Inequality and the zero lower bound', *Journal of Econometrics*, vol. 249, 105819.
- Fernández-Villaverde, Jesús, Galo Nuño, and Jesse Perla. (2024b). 'Taming the curse of dimensionality: Quantitative economics with deep learning', Working Paper 33117, National Bureau of Economic Research.
- Fischhoff, Baruch, and Wändi Bruine De Bruin. (1999). 'Fifty-fifty=50%?', *Journal of Behavioral Decision Making*, vol. 12(2), pp. 149–63.
- Flynn, Joel P., and Karthik Sastry. (2024). 'The macroeconomics of narratives', Working Paper 32602, National Bureau of Economic Research.
- Fofana, Salome, Paula Patzelt, and Ricardo Reis. (2025). 'Household disagreement about expected inflation', in (Guido Ascari and Riccardo Trezzi, eds.), *The Research Handbook of Inflation*, pp. 335–57, Cheltenham: Edward Elgar.
- Friedman, Benjamin M., and L.J. Savage (1947). '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*, pp. 363–72, New York: McGraw-Hill.
- Fudenberg, Drew, and David K. Levine (2016). 'Whither game theory? Towards a theory of learning in games', *Journal of Economic Perspectives*, vol. 30(4), pp. 151–70.
- Fuster, Andreas, Benjamin Hebert, and David Laibson. (2012). 'Natural expectations, macroeconomic dynamics, and asset pricing', *NBER Macroeconomics Annual*, vol. 26(1), pp. 1–48.
- Fuster, Andreas, David Laibson, and Brock Mendel. (2010). 'Natural expectations and macroeconomic fluctuations', *Journal of Economic Perspectives*, vol. 24(4), pp. 67–84.
- Gabaix, Xavier. (2014). 'A sparsity-based model of bounded rationality', *The Quarterly Journal of Economics*, vol. 129(4), pp. 1661–710.
- Gabaix, Xavier. (2020). 'A behavioral New Keynesian model', *American Economic Review*, vol. 110(8), pp. 2271–327.
- Gabaix, Xavier. (2023). 'Marshall lecture 2023: Behavioral macroeconomics via sparse dynamic programming', *Journal of the European Economic Association*, vol. 21(6), pp. 2327–76.
- García-Schmidt, Mariana, and Michael Woodford. (2019). 'Are low interest rates deflationary? A paradox of perfect-foresight analysis', *American Economic Review*, vol. 109(1), pp. 86–120.
- Gardiner, Crispin. (2009). *Stochastic Methods: A Handbook for the Natural and Social Sciences*, 4th edn., Berlin: Springer.
- Garnier-Brun, Jerome, Michael Benzaquen, and Jean-Philippe Bouchaud. (2024). 'Unlearnable games and "satisficing" decisions: A simple model for a complex world', *Physical Review X*, vol. 14(2), 021039.
- Geanakoplos, John, Robert Axtell, Doyne J. Farmer, Peter Howitt, Benjamin Conlee, Jonathan Goldstein, Matthew Hendrey, Nathan M. Palmer, and Chun-Yi Yang. (2012). 'Getting at systemic risk via an agent-based model of the housing market', *The American Economic Review*, vol. 102(3), pp. 53–8.

- Gershman, Samuel J., and Nathaniel D. Daw (2017). 'Reinforcement learning and episodic memory in humans and animals: An integrative framework', *Annual Review of Psychology*, vol. 68(1), pp. 101–28.
- Giannitsarou, Chryssi. (2003). 'Heterogeneous learning', *Review of Economic Dynamics*, vol. 6(4), pp. 885–906.
- Gigerenzer, Gerd, and Peter M. Todd and ABC Research Group. (2000). *Simple Heuristics that Make Us Smart*, New York: Oxford University Press.
- Giglio, Stefano, Matteo Maggiori, Johannes Stroebel, and Stephen Utkus. (2021). 'Five facts about beliefs and portfolios', *American Economic Review*, vol. 111(5), pp. 1481–522.
- Giusto, Andrea. (2014). 'Adaptive learning and distributional dynamics in an incomplete markets model', *Journal of Economic Dynamics and Control*, vol. 40, pp. 317–33.
- Glawion, René. (2023). 'Sequence-space jacobians in continuous time', Preprint, <https://ssrn.com/abstract=4504829>.
- Glimcher, Paul W. (2011). 'Understanding dopamine and reinforcement learning: The dopamine reward prediction error hypothesis', *Proceedings of the National Academy of Sciences*, vol. 108(supplement_3), pp. 15647–54.
- Glimcher, Paul W., Colin Camerer, Ernst Fehr, and Russell Poldrack. (2013). *Neuroeconomics: Decision Making and the Brain*, New York: Academic Press.
- Gomes, Francisco, and Alexander Michaelides. (2008). 'Asset pricing with limited risk sharing and heterogeneous agents', *Review of Financial Studies*, vol. 21(1), pp. 415–48.
- Gopalakrishna, Goutham, Zhouzhou Gu, and Jonathan Payne. (2024). 'Asset pricing, participation constraints, and inequality', Preprint, <http://dx.doi.org/10.2139/ssrn.5030405>.
- Grandmont, Jean-Michel. (1977). 'Temporary general equilibrium theory', *Econometrica*, vol. 45(3), pp. 535–72.
- Grandmont, Jean-Michel. (1989). 'Temporary equilibrium', in (John Eatwell, Murray Milgate and Peter Newman, eds.), *General Equilibrium*, pp. 297–304, London: Palgrave Macmillan.
- Gu, Zhouzhou, Mathieu Laurière, Sebastian Merkel, and Jonathan Payne. (2024). 'Global solutions to master equations for continuous time heterogeneous agent macroeconomic models', Preprint, <https://arxiv.org/abs/2406.13726v1>.
- Gualdi, Stanislaw, Marco Tarzia, Francesco Zamponi, and Jean-Philippe Bouchaud. (2015). 'Tipping points in macroeconomic agent-based models', *Journal of Economic Dynamics and Control*, vol. 50, pp. 29–61.
- Guerreiro, Joao. (2023). 'Belief disagreement and business cycles', Working paper, UCLA.
- Haaland, Ingar, Christopher Roth, and Johannes Wohlfart. (2023). 'Designing information provision experiments', *Journal of Economic Literature*, vol. 61(1), pp. 3–40.
- Hajdini, Ina. (2023). 'Mis-specified forecasts and myopia in an estimated New Keynesian model', *American Economic Journal: Macroeconomics (Forthcoming)*.
- Han, Jiequn, Yucheng Yang, and Weinan E. (2022). 'DeepHAM: A global solution method for heterogeneous agent models with aggregate shocks', Preprint, <https://arxiv.org/abs/2112.14377v2>.
- He, Zhiguo, and Arvind Krishnamurthy. (2012). 'A model of capital and crises', *The Review of Economic Studies*, vol. 79(2), pp. 735–77.
- He, Zhiguo, and Arvind Krishnamurthy. (2013). 'Intermediary asset pricing', *American Economic Review*, vol. 103(2), pp. 732–70.
- Hicks, John R. (1939). *Value and Capital: An Inquiry into Some Fundamental Principles of Economic Theory*, Oxford: Clarendon Press.
- Hicks, John R. (1979). 'The formation of an economist', *PSL Quarterly Review*, doi: 10.13133/2037-3643/11143.
- Hommes, Cars, and Blake LeBaron (eds.) (2018). *Handbook of Computational Economics*, Amsterdam: North-Holland: Elsevier.
- Hotelling, Harold. (1941). 'Experimental determination of the maximum of a function', *The Annals of Mathematical Statistics*, vol. 12(1), pp. 20–45.
- Huang, Ji. (2023). 'Breaking the curse of dimensionality in heterogeneous agent models: A deep learning-based probabilistic approach', Preprint, <http://dx.doi.org/10.2139/ssrn.4649043>.
- Ilut, Cosmin, and Martin Schneider. (2014). 'Ambiguous business cycles', *American Economic Review*, vol. 104(8), pp. 2368–99.
- Jaakkola, Tommi, Michael Jordan, and Satinder Singh. (1993). 'Convergence of stochastic iterative dynamic programming algorithms', *Advances in Neural Information Processing Systems*, vol. 6, pp. 703–10.
- Jacobson, Margaret M. (2025). 'Beliefs, aggregate risk, and the U.S. housing boom', Finance and Economics Discussion Series 2022-061r1, Board of Governors of the Federal Reserve System (U.S.).
- Javed, Khurram, and Richard S. Sutton (2024). 'The big world hypothesis and its ramifications for artificial intelligence', in *Finding the Frame: An RLC Workshop for Examining Conceptual Frameworks*, <https://openreview.net/forum?id=Sv7DazuCn8>. (Last accessed: 27 October 2025).
- Kaplan, Greg, Kurt Mitman, and Giovanni L. Violante (2020). 'The housing boom and bust: Model meets evidence', *Journal of Political Economy*, vol. 128(9), pp. 3285–345.
- Kaplan, Greg, Benjamin Moll, and Giovanni L. Violante (2018). 'Monetary policy according to HANK', *American Economic Review*, vol. 108(3), pp. 697–743.
- Kindleberger, Charles. (1978). *Manias, Panics, and Crashes*, Cham: Springer.
- Koijen, Ralph S.J., and Motohiro Yogo. (2019). 'A demand system approach to asset pricing', *Journal of Political Economy*, vol. 127(4), pp. 1475–515.
- Krishnamurthy, Arvind, and Wenhao Li. (2025). 'Dissecting mechanisms of financial crises: Intermediation and sentiment', *Journal of Political Economy*, vol. 133(3), pp. 935–85.

- Krusell, Per, and Anthony A. Smith (1998). 'Income and wealth heterogeneity in the macroeconomy', *Journal of Political Economy*, vol. 106(5), pp. 867–96.
- Laurière, Mathieu, Sarah Perrin, Sertan Girgin, Paul Muller, Ayush Jain, Theophile Cabannes, Georgios Piliouras, Julien Perolat, Romuald Elie, Olivier Pietquin, and Matthieu Geist. (2022). '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, pp. 12078–95, PMLR.
- Lee, Donghoon, and Kenneth I. Wolpin (2006). 'Intersectoral labor mobility and the growth of the service sector', *Econometrica*, vol. 74(1), pp. 1–46.
- Lee, Soyoung. (2022). 'The macroeconomic effects of debt relief policies during recessions', Preprint, <http://dx.doi.org/10.2139/ssrn.4046177>.
- Leombroni, Matteo, Monika Piazzesi, Martin Schneider, and Ciaran Rogers. (2020). 'Inflation and the price of real assets', Working Paper 26740, National Bureau of Economic Research.
- Lindahl, Erik. (1939). *Studies in the Theory of Money and Capital*, London: Allen and Unwin.
- Ljung, L. (1977). 'Analysis of recursive stochastic algorithms', *IEEE Transactions on Automatic Control*, vol. 22(4), pp. 551–75.
- Ljungqvist, Lars, and Thomas J. Sargent (2004). *Recursive Macroeconomic Theory*, Cambridge, MA: MIT Press.
- Lull, Joan. (2018). 'Immigration, wages, and education: A labour market equilibrium structural model', *The Review of Economic Studies*, vol. 85(3), pp. 1852–96.
- Lorenzoni, Guido. (2009). 'A theory of demand shocks', *American Economic Review*, vol. 99(5), pp. 2050–84.
- Lucas, Robert E. (1972). 'Expectations and the neutrality of money', *Journal of Economic Theory*, vol. 4(2), pp. 103–24.
- Lucas, Robert E. (1976). 'Econometric policy evaluation: A critique', *Carnegie-Rochester Conference Series on Public Policy*, vol. 1, pp. 19–46.
- Lucas, Robert E. (1980). 'Methods and problems in business cycle theory', *Journal of Money, Credit and Banking*, vol. 12(4), pp. 696–715.
- Lucas, Robert E. (1987). *Models of Business Cycles*, Oxford: Basil Blackwell.
- Lucas, Robert E. (1996). 'Nobel lecture: Monetary neutrality', *Journal of Political Economy*, vol. 104(4), pp. 661–82.
- Lucas, Robert E., and Edward C. Prescott (1971). 'Investment under uncertainty', *Econometrica*, vol. 39(5), pp. 659–81.
- Ludwig, Alexander, Jochen Mankart, Jorge Quintana, and Mirko Wiederholt. (2024). 'Heterogeneity in expectations and house price dynamics', Working paper, SAFE.
- Maćkowiak, Bartosz, Filip Matějka, and Mirko Wiederholt. (2023). 'Rational inattention: A review', *Journal of Economic Literature*, vol. 61(1), pp. 226–73.
- Maliar, Lilia, and Serguei Maliar. (2014). 'Numerical methods for large-scale dynamic economic models', in (Karl Schmedders and Kenneth L. Judd, eds.), *Handbook of Computational Economics*, vol. 3, pp. 325–477, Amsterdam: Elsevier.
- Maliar, Lilia, Serguei Maliar, and Pablo Winant. (2021). 'Deep learning for solving dynamic economic models', *Journal of Monetary Economics*, vol. 122(C), pp. 76–101.
- Malmendier, Ulrike, and Stefan Nagel. (2011). 'Depression babies: Do macroeconomic experiences affect risk taking?', *The Quarterly Journal of Economics*, vol. 126(1), pp. 373–416.
- Malmendier, Ulrike, and Stefan Nagel. (2015). 'Learning from inflation experiences', *The Quarterly Journal of Economics*, vol. 131(1), pp. 53–87.
- Mankiw, N. Gregory, and Ricardo Reis. (2002). 'Sticky information versus sticky prices: A proposal to replace the New Keynesian Phillips curve', *The Quarterly Journal of Economics*, vol. 117(4), pp. 1295–328.
- Mankiw, N. Gregory, Ricardo Reis, and Justin Wolfers. (2004). 'Disagreement about inflation expectations', *NBER Macroeconomics Annual 2003*, vol. 18, pp. 209–70.
- Manski, Charles F. (2004). 'Measuring expectations', *Econometrica*, vol. 72(5), pp. 1329–76.
- Manski, Charles F. (2018). 'Survey measurement of probabilistic macroeconomic expectations: Progress and promise', *NBER Macroeconomics Annual*, vol. 32(1), pp. 411–71.
- Marcet, Albert, and Thomas J. Sargent (1989). 'Convergence of least squares learning mechanisms in self-referential linear stochastic models', *Journal of Economic Theory*, vol. 48(2), pp. 337–68.
- Maxted, Peter. (2024). 'A macro-finance model with sentiment', *Review of Economic Studies*, vol. 91(1), pp. 438–75.
- Mendoza, Enrique G. (2010). 'Sudden stops, financial crises, and leverage', *American Economic Review*, vol. 100(5), pp. 1941–66.
- Minsky, Hyman P. (1977). 'The financial instability hypothesis: An interpretation of Keynes and an alternative to "standard" theory', *Challenge*, vol. 20(1), pp. 20–7.
- 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, Dhharshan Kumaran, Daan Wierstra, Shane Legg, and Demis Hassabis. (2015). 'Human-level control through deep reinforcement learning', *Nature*, vol. 518(7540), pp. 529–33.
- Molavi, Pooya. (2024). 'Simple models and biased forecasts', Preprint, <https://arxiv.org/abs/2202.06921v5>.
- Moll, Benjamin, and Lenya Ryzhik. (2025). 'Mean field games without rational expectations', Working paper.

- Morelli, Federico Guglielmo, Michael Benzaquen, Marco Tarzia, and Jean-Philippe Bouchaud. (2020). 'Confidence collapse in a multihousehold, self-reflexive DSGE model', *Proceedings of the National Academy of Sciences*, vol. 117(17), pp. 9244–9.
- Morgenstern, Oskar. (1935). 'Vollkommene voraussicht und wirtschaftliches gleichgewicht', *Zeitschrift für Nationalökonomie/Journal of Economics*, vol. 6, pp. 337–57.
- Mueller, Andreas I., and Johannes Spinnewijn. (2023). 'Expectations data, labor market, and job search', in (Rüdiger Bachmann, Giorgio Topa and Wilbert van der Klaauw, eds.), *Handbook of Economic Expectations*, pp. 677–713, Cambridge, MA: Academic Press.
- Murphy, Kevin. (2025). 'Reinforcement learning: An overview', Preprint, <https://arxiv.org/abs/2412.05265v4>.
- Muth, John F. (1961). 'Rational expectations and the theory of price movements', *Econometrica*, vol. 29(3), pp. 315–35.
- Nagel, Rosemarie. (1995). 'Unraveling in guessing games: An experimental study', *The American Economic Review*, vol. 85(5), pp. 1313–26.
- Nakov, Anton, and Galo Nuño. (2015). 'Learning from experience in the stock market', *Journal of Economic Dynamics and Control*, vol. 52, pp. 224–39.
- Niv, Yael. (2009). 'Reinforcement learning in the brain', *Journal of Mathematical Psychology*, vol. 53(3), pp. 139–54.
- Oprea, Ryan. (2024). 'Decisions under risk are decisions under complexity', *American Economic Review*, vol. 114(12), pp. 3789–811.
- Ortoleva, Pietro. (2024). 'Alternatives to Bayesian updating', *Annual Review of Economics*, vol. 16, pp. 545–70.
- Piazzesi, Monika, Martin Schneider, and Tim Landvoigt. (2015). 'The housing market(s) of San Diego', *American Economic Review*, vol. 105(4), pp. 1371–407.
- Piazzesi, Monika, and Martin Schneider. (2016). 'Housing and Macroeconomics', in (John B. Taylor and Harald Uhlig, eds.), *Handbook of Macroeconomics*, vol. 2, pp. 1547–640, Amsterdam: Elsevier.
- Prescott, Edward C., and Rajnish Mehra. (1980). 'Recursive competitive equilibrium: The case of homogeneous households', in *Econometrica*, pp. 1365–79.
- Preston, Bruce. (2005). 'Learning about monetary policy rules when long-horizon expectations matter', *International Journal of Central Banking*, vol. 1(2), 81–126.
- Pröhl, Elisabeth. (2019). 'Approximating equilibria with ex-post heterogeneity and aggregate risk', Research Paper 17-63, Swiss Finance Institute.
- Reiter, Michael. (2009). 'Solving heterogeneous agent models by projection and perturbation', *Journal of Economic Dynamics and Control*, vol. 33(3), pp. 649–65.
- Ríos-Rull, José-Víctor. (1997). 'Computation of equilibria in heterogeneous agent models', Technical Report, Federal Reserve Bank of Minneapolis.
- Robbins, Herbert, and Sutton Monro. (1951). 'A stochastic approximation method', *The Annals of Mathematical Statistics*, vol. 22(3), pp. 400–7.
- Roth, Christopher, Mirko Wiederholt, and Johannes Wohlfart. (2023). 'The effects of monetary policy: Theory with measured expectations', ECONtribute Discussion Paper 217, University of Bonn and University of Cologne.
- Sargent, Thomas J. (1991). 'Equilibrium with signal extraction from endogenous variables', *Journal of Economic Dynamics and Control*, vol. 15(2), pp. 245–73.
- Sargent, Thomas J. (1993). *Bounded Rationality in Macroeconomics: The Arne Ryde Memorial Lectures*, Oxford: Clarendon Press.
- Sargent, Thomas J. (1996). 'Expectations and the nonneutrality of lucas', *Journal of Monetary Economics*, vol. 37(3), pp. 535–48.
- Sargent, Thomas J. (1999). *The Conquest of American Inflation*, Princeton, NJ: Princeton University Press.
- Sargent, Thomas J. (2007). 'Rational expectations', in *The Concise Encyclopedia of Economics*, Indianapolis, IN: Liberty Fund, <https://www.econlib.org/library/Enc/RationalExpectations.html>. (Last accessed: 27 October 2025).
- Sargent, Thomas J. (2008). 'Evolution and intelligent design', *American Economic Review*, vol. 98(1), pp. 5–37.
- Sargent, Thomas J. (2017). 'Rational expectations', in (Steven N. Durlauf and Lawrence E. Blume, eds.), *The New Palgrave Dictionary of Economics*, pp. 5389–5394, London: Palgrave Macmillan.
- Sargent, Thomas J. (2024). 'Macroeconomics after lucas', Working paper.
- Schaab, Andreas. (2020). 'Micro and macro uncertainty', Working paper, UC Berkeley.
- Silver, David. (2015). 'Introduction to reinforcement learning', DeepMind x UCL, Video recordings available at <https://www.youtube.com/playlist?list=PLqYmG7hTraZDM-OYHWgPebj2MfCFzFObQ>. (Last accessed: 27 October 2025).
- Silver, David, 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. (2016). 'Mastering the game of go with deep neural networks and tree search', *Nature*, vol. 529(7587), pp. 484–9.
- Silver, David, 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. (2017). 'Mastering the game of go without human knowledge', *Nature*, vol. 550(7676), pp. 354–9.
- Simon, Herbert A. (1978a). 'On how to decide what to do', *The Bell Journal of Economics*, vol. 9(2), pp. 494–507.

- Simon, Herbert A. (1978b). 'Rationality as process and as product of thought', *The American Economic Review*, vol. 68(2), pp. 1–16.
- Sims, Christopher A. (1980). 'Macroeconomics and reality', *Econometrica: Journal of the Econometric Society*, vol. 48, 1–48.
- Sims, Christopher A. (2003). 'Implications of rational inattention', *Journal of Monetary Economics*, vol. 50(3), pp. 665–90.
- Storesletten, Kjetil, Chris Telmer, and Amir Yaron. (2007). 'Asset pricing with idiosyncratic risk and overlapping generations', *Review of Economic Dynamics*, vol. 10(4), pp. 519–48.
- Straub, Ludwig, and Robert Ulbricht. (2023). 'Endogenous uncertainty and credit crunches', *The Review of Economic Studies*, vol. 91(5), pp. 3085–115.
- Sutton, Richard S., and Andrew G. Barto (2018). *Reinforcement Learning: An Introduction*, Cambridge, MA: MIT Press.
- Taubinsky, Dmitry, Luigi Butera, Matteo Saccarola, and Chen Lian. (2024). 'Beliefs about the economy are excessively sensitive to household-level shocks: evidence from linked survey and administrative data', Working Paper 32664, National Bureau of Economic Research.
- Tsitsiklis, John N. (1994). 'Asynchronous stochastic approximation and Q-learning', *Machine Learning*, vol. 16, pp. 185–202.
- Tversky, Amos, and Daniel Kahneman. (1974). 'Judgment under uncertainty: Heuristics and biases', *Science*, vol. 185(4157), pp. 1124–31.
- Viscusi, W. Kip. (1985). 'A Bayesian perspective on biases in risk perception', *Economics Letters*, vol. 17(1-2), pp. 59–62.
- Werning, Iván. (2022). 'Expectations and the rate of inflation', Working Paper 30260, National Bureau of Economic Research.
- Werning, Iván. (2023). 'Lucas miracles', <https://economics.mit.edu/sites/default/files/inline-files/Translated%20Lucas%20Miracles%20by%20Ivan%20Werning.pdf>. (Last accessed: 27 October 2025).
- Wiese, Samuel, Jagoda Kaszowska-Mojsa, Joel Dyer, Jose Moran, Marco Pangallo, Francois Lafond, John Muellbauer, Anisoara Calinescu, and J. Doynne Farmer. (2024). 'Forecasting macroeconomic dynamics using a calibrated data-driven agent-based model', Preprint, <https://arxiv.org/abs/2409.18760v1>.
- Williams, Noah. (2018). 'Escape dynamics in learning models', *The Review of Economic Studies*, vol. 86(2), pp. 882–912.
- Williams, Noah. (2024). 'A learning model of financial instability', Working paper, University of Miami.
- Woodford, Michael. (1990). 'Learning to believe in sunspots', *Econometrica*, vol. 58(2), pp. 277–307.
- Woodford, Michael. (2003). 'Imperfect common knowledge and the effects of monetary policy', in (Philippe Aghion, Roman Frydman, Joseph Stiglitz and Michael Woodford, eds.), *Expectations in Modern Macroeconomics: In Honor of Edmund S. Phelps*, pp. 25–58, Princeton, NJ: Princeton University Press.
- Woodford, Michael. (2013). 'Macroeconomic analysis without the rational expectations hypothesis', *Annual Review of Economics*, vol. 5(1), pp. 303–46.
- Woodford, Michael. (2018). 'Monetary policy analysis when planning horizons are finite', *NBER Macroeconomics Annual*, vol. 33, pp. 1–50.
- Xu, Ruitu, Yifei Min, Tianhao Wang, Michael I. Jordan, Zhaoran Wang, and Zhuoran Yang. (2023). '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, pp. 375–407, PMLR.
- Yang, Jeffrey. (2024). 'On the decision-relevance of subjective beliefs', SSRN Working Paper 4425080.
- Zhao, Shiyu. (2025). *Mathematical Foundations of Reinforcement Learning*, Singapore: Springer, <https://github.com/MathFoundationRL/Book-Mathematical-Foundation-of-Reinforcement-Learning>. (Last accessed: 27 October 2025).