

Roles of Rational Expectations

Consistency, Inference, and Policy in Economic Models

Thomas J. Sargent*

October 16, 2026

Abstract

Rational expectations is a property of a model's equilibrium, not of real people. A rational expectations equilibrium requires the subjective statistical models of agents inside a model to equal the objective statistical model that their decisions generate. I describe that requirement as a fixed point of a mapping from subjective to objective models and trace its roles. The resulting communism of models disciplines beliefs, supplies cross-equation restrictions that help identify structural parameters, turns models into machines that compute how probability distributions of outcomes respond to policies, and settles whose probabilities welfare comparisons should use. Often under other names, the restriction pervades economic theory and econometrics. Agents need to forecast only the variables that enter their own decision problems. Computing the equilibrium is the model builder's task. The restriction is least plausible for rare and off-path events. Alternatives to rational expectations put three roles most at risk: discipline on beliefs, identification, and welfare comparisons.

Keywords: Rational expectations; communism of models; cross-equation restrictions; identification; policy evaluation; determinacy; global games; heterogeneous agent models; Kullback–Leibler divergence; robustness; restricted perceptions; learning.

*University of Pennsylvania, Wharton School. Email: sargentt@upenn.edu. I am especially grateful to Fernando Alvarez and Francesco Lippi, who read drafts of this paper and made many useful comments. I thank In-Koo Cho, Lars Peter Hansen, Kenneth Kasa, Eric Leeper, Ben Moll, and Todd Walker for discussions of the topic of this paper.

1 Introduction

A rational expectations equilibrium restriction requires that the conditional probability distributions that agents use to forecast equal the conditional probability distributions generated by their decisions and the model's other assumptions. This paper is about that restriction, which Muth (1961) proposed and Lucas (1987) called a consistency axiom for economic models. Over six decades, the restriction has become the default across economics (section 7.2) and in the econometric methods that estimate economic models. Experimental economists test predictions of rational expectations equilibria in laboratory markets.

I organize the paper around a mapping T from the subjective statistical models of agents inside an equilibrium Markov model to the objective statistical model that their decisions generate. A rational expectations equilibrium is a fixed point of T at which subjective and objective models coincide. I call that a communism of models. Agents share a joint probability distribution but need not share information (section 5). Each agent needs to forecast only the variables that enter his or her own decision problem (section 11.2). The communism of models disciplines beliefs by eliminating free parameters that would otherwise describe them. It supplies cross-equation restrictions that help identify structural parameters. It lets a model builder compute how people inside a model respond to proposed policies, and it lets economists analyze time inconsistency, credibility, reputations, and information aggregation. It settles whose probabilities welfare comparisons should use.

Survey data on expectations display disagreement and predictable forecast errors (Coibion and Gorodnichenko, 2015; Bordalo et al., 2020). In heterogeneous agent models with aggregate risk, a rational expectations equilibrium requires agents to forecast prices that depend on cross-section distributions, and Moll (2026) suggests not imposing the restriction in some such models. Evaluating that suggestion and others like it requires an appreciation of what the restriction does and an accounting of what abandoning it would lose. I intend the paper as a reference for students and researchers. It draws on Sargent (2024) and Sargent (2025).

The paper has four parts. Part I (sections 2–5) describes what rational expectations is: words, statistical models, equilibrium Markov models, and the restriction as a fixed point of T . Part II (sections 6–10) describes what the restriction does in econometrics, across economics, and in policy analysis, how economists respond when models that impose it fail, and equilibria in which agents share but distrust an approximating model. Part III (sections 11–13) describes what the restriction asks of people, four defenses of asking it, and where it strains, and closes by collecting the restriction's roles. Part IV (section 14) describes disciplined alternatives and what they keep and lose. Section 15 concludes.

Part I. What rational expectations is

2 Words

Economics is what economists say about the economy. Economists use ordinary words in peculiar ways. Peculiar does not necessarily mean strange or incorrect, but it can. Lucas (1987, p. 13) complained that the phrase “rational expectations” is often misused:

The term ‘rational expectations’, as Muth used it, refers to a consistency axiom for economic models, so it can be given precise meaning only in the context of specific models. I think this is why attempts to define rational expectations in a model-free way tend to come out either vacuous (‘People do the best they can with the information they have’) or silly (‘People know the true structure of the world they live in’).

The acronym FIRE—“full-information rational expectations”—is another peculiar use of words. It contains two misleading modifiers and omits a noun. The misleading modifiers are “full” and “information.” Some rational expectations equilibrium models contain agents who have different information sets, i.e., different filtrations (Townsend, 1983; Huo and Pedroni, 2020; Hellwig and Venkateswaran, 2025). Their predictive densities differ. Suppose that z equals 0 or 1 with equal probabilities. An agent who observes z forecasts z ; an agent who does not forecasts 1/2. Both forecasts are correct. In such contexts, I do not know what “full” means. “Full information” also conflates information with knowledge.¹

The missing noun is “equilibrium” or “model.” A *model* either does or does not impose the rational expectations equilibrium restriction (section 5). Real people do not “have” or “not have” rational expectations. Neither, in isolation, do the artificial people who live inside a model: whether one agent’s forecasts are correct depends on outcomes that all agents’ decisions produce (section 5.2). Rational expectations is a property of a model’s equilibrium: a relation between agents’ beliefs and the outcomes that their decisions produce.

Optimization and correct forecasts are different restrictions: an agent can optimize under a misspecified model. Correct forecasts and full information are different too: an agent can forecast correctly while conditioning on little information. Internal consistency and empirical accuracy are different again: an internally consistent model can fit data poorly.

¹Throughout, I use *information* to mean data, i.e., the history or filtration on which a forecast conditions, and *knowledge* to mean knowing the structural parameters θ_s of a model (section 3). A rational expectations econometrician and the agents inside the model that the econometrician estimates can have the same information but different knowledge: the agents know θ_s and the econometrician does not (section 6.2).

3 Statistical models

I want to use words unambiguously, so I start with a dictionary. A *statistical model* is a joint probability distribution $f(Y | \theta)$ for a vector stochastic process $Y = \{y_t\}$, indexed by a vector of parameters θ . A *manifold of statistical models* $\{f(Y | \theta)\}_{\theta \in \Theta}$ emerges as θ sweeps through a set Θ of possible parameter vectors. I take Lucas’s “economic model” to mean a statistical model whose parameters describe objective functions and constraint sets.

Lucas was using a long-standing distinction between two types of statistical models. In a *descriptive* model, the parameters and the functional form of the probability distribution are data compression devices. Vector autoregressions and dynamic factor models are descriptive models. Christopher Sims called them “economically uninterpreted” models.² In a *structural* model, an economic theorist can interpret the parameters. They describe preferences, technologies, information sets, and government strategies. Sims called such models “fully interpreted.”

Koopmans (1947) viewed descriptive models as necessary inputs into estimating the parameters of structural models, not as ends in themselves. He illustrated distinct roles of descriptive and structural models with examples from the history of physics: Tycho Brahe and Kepler produced descriptive models, the “Kepler stage”; Newton used them, together with Galileo’s experiments, to construct a structural model, the “Newton stage.”

I denote parameters of a structural model by θ_s and parameters of a descriptive model by θ_d . An equilibrium Markov model with structural parameters θ_s induces, given agents’ subjective models, an objective statistical model $f(Y | \theta)$ (section 4). Without a subscript, θ denotes the parameters of that objective law of motion. *Objective* means generated by the model. Whether the model’s law also generates actual data is an empirical question.

Three kinds of people appear in this paper (cf. Hansen, 2014). *Agents* live inside a model. The *model builder* designs the model and computes its equilibria. The *econometrician* lives outside the model and uses it to interpret data. One person often plays the last two roles.

4 Equilibrium Markov models

Lucas and Prescott (1971) and Prescott and Mehra (1980) showed how to impose the rational expectations restriction on models with a Markov structure that Lucas (1987, sec. I) later outlined informally.

A Markov decision problem (MDP) consists of a set of states X , a set of actions U , a

²Chris Sims used both phrases in a conversation with me in a hallway at the University of Minnesota in the late 1970s.

controlled transition law $P(x' | x, u)$ that gives the probability that next period's state is x' when the current state is x and the action is u , a reward function $r(x, u, x')$, and a discount factor $\beta \in [0, 1)$. Its Bellman equation is

$$V(x) = \max_{u \in U} E[r(x, u, x') + \beta V(x') | x, u], \quad (1)$$

where the expectation is taken with respect to $P(x' | x, u)$.

Definition 4.1. *An equilibrium Markov model contains (1) a collection of decision makers, (2) associated MDPs defined over a common state space, and (3) budget and resource constraints that bind decision makers' MDPs together (Lucas, 1987, sec. I; Sargent, 2025). An equilibrium Markov process is the stochastic process of outcomes that such a model generates.*

Each agent's controlled transition law is a *subjective* probability distribution. An equilibrium concept organizes the agents' MDPs into a coherent set. It need not impose a rational expectations restriction. Each agent $i = 1, \dots, I$ inside the model, or each type of agent when there is a continuum of them, has a subjective model $f_i(Y | \theta_i)$.³ Agents' subjective distributions describe two kinds of outcomes: outcomes determined *outside* the model, i.e., exogenous shock processes; and outcomes determined *inside* the model, i.e., prices and quantities.

Each MDP contains its own theory of cause and effect: the decision maker regards the controlled transition law P as “causal” in the sense that it is invariant across admissible choices of his or her actions. An equilibrium Markov model thus contains as many assumptions about causality as there are decision makers (Sargent, 2025).

Prescott and Mehra (1980) formulated an equilibrium Markov model as a recursive competitive equilibrium. They replaced a search for equilibrium sequences of history-contingent functions with a search for equilibrium decision rules. A decision rule expresses a current action as a time-invariant function of the current state. In their model, households form a continuum of measure one. The state of household i is a triple (k_i, K, ζ) : the household's own capital k_i ; aggregate capital $K = \int k_i di$; and a Markov shock ζ . In a symmetric equilibrium, $k_i = K$ for every i . To pose its decision problem, each household takes the law of motion of K as given and must be free to choose $k_i \neq K$. This is the “Big K , little k ” device of Lucas and Prescott (1971). Prescott and Mehra described what a household needs to know:

The consumer does not have a well defined decision problem, until the equilibrium law of motion, $K' = H(K, \zeta)$, and the pricing function, $p(K, \zeta)$, are specified.

³ Y is common across agents' models, but the functional forms f_i and parameters θ_i can differ across agents.

Knowledge of these elements along with $F(\zeta' \mid \zeta)$ is sufficient for forecasting future prices (actually forming predictive probability distributions of future prices) and selection of optimal current actions. (Prescott and Mehra, 1980, p. 1371)⁴

The functions H and p are components of a household's subjective model.

5 The rational expectations restriction

Muth (1961, p. 316) stated his hypothesis this way:

... 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).

Muth stated his hypothesis for probability distributions but applied it to conditional means in linear models (Lucas and Prescott, 1971, p. 660, n. 4). Lucas and Prescott (1971, p. 664) applied it to entire distributions (section 11.1).

Let $\theta_i \in \Theta_i$ be parameters of agent i 's subjective model and $\theta \in \Theta$ be parameters of the objective statistical model that the equilibrium Markov model induces. Agents' decisions depend on their subjective distributions, and the equilibrium joint distribution depends on their decisions. Equilibrium conditions thereby define a mapping

$$\theta = \mathsf{T}(\theta_1, \theta_2, \dots, \theta_I; \theta_s) \quad (2)$$

from agents' subjective joint distributions to an objective joint distribution. The mapping T encodes preferences, technologies, and information, which the structural parameters θ_s describe, together with market clearing and aggregation. It also depends on the functional forms of agents' families $\mathcal{F}_i = \{f_i(\cdot \mid \theta_i)\}_{\theta_i \in \Theta_i}$. A more complete notation is $\mathsf{T}_{\mathcal{F}}$, where $\mathcal{F} = (\mathcal{F}_1, \dots, \mathcal{F}_I)$. I suppress θ_s and the subscript $\mathcal{F}$ when that causes no confusion.

I write $f(Y \mid \theta)$ and $f_{\theta}(Y)$ interchangeably for the same joint density. The first notation reads it as a likelihood, a function of θ for given data. The second treats it as a point in a space of probability distributions. That is how I use it when I compare distributions with Kullback–Leibler (KL) divergences and information projections (below and section 6). Similarly, agent i 's subjective model $f_i(Y \mid \theta_i)$ is also written $f_{i,\theta_i}(Y)$. When Y is an infinite sequence, KL divergences between two processes are to be read as per-period rates, or as

⁴I have replaced Prescott and Mehra's symbols for aggregate capital, its law of motion, and the shock with K , H , and ζ .

divergences between one-step-ahead conditional densities such as $f_\theta(x_{t+1} \mid x_t)$, either state by state or averaged over the stationary distribution of the Markov state x_t .⁵

The KL divergence, or relative entropy, of the objective probability measure relative to agent i 's subjective measure is

$$KL(f_\theta, f_{i,\theta_i}) = \int \log \left(\frac{f_\theta(Y)}{f_{i,\theta_i}(Y)} \right) f_\theta(Y) dY \geq 0. \quad (3)$$

Its first argument is the measure under which the expectation is taken. KL divergences are positive whenever subjective and objective models differ on events to which the objective model assigns positive probability. For example, the worst-case models of the robust decision makers of Hansen and Sargent (2001, 2008) have positive KL divergences from the objective model.

Muth (1961) proposed building models in which subjective and objective distributions coincide. When they coincide, the θ_i 's on the right side of (2) become endogenous. A rational expectations equilibrium is a fixed point

$$\theta = T(\theta, \theta, \dots, \theta) \quad (4)$$

of the mapping from subjective distributions to an objective distribution. At the fixed point, all KL divergences (3) equal zero. When the fixed point is unique, it expresses θ as a function $\theta = \theta(\theta_s)$ of the structural parameters. The restrictions that $\theta(\theta_s)$ imposes on θ are the cross-equation restrictions of rational expectations. When the fixed point is not unique, the model builder (section 8.1) or the data must select an equilibrium. In 1993 I put it this way: “When an equilibrium is unique, ‘beliefs’ contribute no free parameters to the model. When equilibria are not unique, ... some of the free parameters index beliefs” (Sargent, 1993, p. 13, n. 7). Those free parameters can be estimated along with the others, and the model can remain overidentified. Sargent and Wallace (1987) showed this for a model of hyperinflation with a four-dimensional continuum of equilibria, some driven by sunspots, and İmrohoroglu (1993) estimated that model for the German hyperinflation by maximum likelihood, letting the data choose among the equilibria.

Rational expectations restricts functional forms as well as parameters. It requires each agent's subjective model to equal the objective model, $f_i(\cdot \mid \theta_i) = f(\cdot \mid \theta)$. Without loss of generality, each agent's family can then be taken to be the objective family, $\mathcal{F}_i = \{f(\cdot \mid \theta)\}_{\theta \in \Theta}$,

⁵These readings have limits. A zero per-period rate does not imply equal initial distributions, and an average over a stationary distribution tests one-step conditional densities only where that distribution puts weight. For those reasons, equation (5) below states the rational expectations restriction as an equality of conditional distributions given each agent's information.

with $\theta_i = \theta$. Restriction (4) asserts a communism of statistical models: everyone inside the model shares a joint probability distribution with everyone else and with the model builder, but not necessarily the same information set. Agents who observe different signals draw different conditional distributions from the same joint distribution, and each forecast is correct given what its maker observes. Restriction (4) eliminates as free objects all of the f_i 's and θ_i 's.

Stated conditionally, the restriction requires

$$f_i(z_{t+1} \mid \mathcal{I}_{it}; \theta_i) = f(z_{t+1} \mid \mathcal{I}_{it}; \theta) \quad (5)$$

for every agent i and every information set $\mathcal{I}_{it}$, including those reached only after deviations. Here z_{t+1} is a future variable that agent i forecasts and $\mathcal{I}_{it}$ is agent i 's information at t . The example in section 2 illustrates (5). When Y includes every variable that agents observe, equal joint distributions of Y imply (5) almost surely under $f(Y \mid \theta)$, that is, along the equilibrium path. They say nothing at information sets that the equilibrium never generates, so by themselves they impose only the self-confirming restriction of section 14.1. To restrict beliefs after deviations, I read θ in (4) more broadly than in section 3, as indexing decision rules and controlled transition laws at every state (section 4); the right side of (5) is then computed from them. When such an information set does not reveal the state, beliefs about the state must also be specified. As in a sequential equilibrium, Bayes' rule does not pin them down. They can index multiple equilibria, which equilibrium-path data can distinguish only through their implications for outcomes on that path.

Prescott and Mehra (1980, pp. 1370–71) completed their model with rational expectations: “the prices and price distributions on which the economic agents base their consumption-investment-production decisions are exactly the same as those that result as a consequence of their decisions through market clearing.” That sentence describes fixed point (4). The functions H and p of section 4 are *outputs* of the equilibrium; households use them as *inputs* to their MDPs.

The fixed point answers a question that troubled economists before Muth. Morgenstern (1928) argued that economic forecasting is impossible because people react to published forecasts in ways that can falsify them. Grunberg and Modigliani (1954) showed that when outcomes depend continuously on a published forecast and both lie in the same closed, bounded interval, there exists a forecast that remains correct after people react to it. A rational expectations equilibrium extends that fixed point from a single published forecast to the entire subjective probability distributions of all agents inside a model.

Arrow had already imposed fixed point (4) on the price block of a general equilibrium

model before Grunberg and Modigliani (1954) proved the single-forecast version and eight years before Muth (1961) named the hypothesis. In a lecture that he read in Paris in May 1952 and published in French in 1953, Arrow (1964) studied a two-period economy with S states and C goods and compared two ways of allocating risk: a market at time 0 in SC claims to particular goods in particular states, and a market in S securities payable in money followed, after the state is realized, by C ordinary spot markets. The second arrangement supports the same allocation as the first if consumers correctly forecast the spot prices that will prevail in each state.⁶

The step from one published forecast to a whole general equilibrium model is not automatic. Grunberg and Modigliani's existence argument assumes that the outcome depends *continuously* on the forecast. Hart (1975) exhibited a general equilibrium model in which the counterpart of that continuity fails and in which no self-fulfilling forecast exists. His securities pay goods rather than money, so the dollar payoff of a security is the forecast spot price vector times its dividend bundle, and forecasts of spot prices determine which risks the securities span. Arrow's construction avoids that failure because his securities pay a unit of account, one security for each state, which fixes the span of portfolio payoffs before any forecast is made.

5.1 An example of T

A stochastic version of the Cagan (1956) model studied by Sargent and Wallace (1973) and Sargent (2024) exhibits T. Let m_t be the log of the money supply, p_t the log of the price level, $\mu_t = m_{t+1} - m_t$ the rate of money creation, $\pi_t = p_{t+1} - p_t$ the rate of inflation, and π_t^* the public's expected rate of inflation between t and $t + 1$. The demand for real balances is

$$m_t - p_t = -\alpha\pi_t^*, \quad \alpha > 0. \quad (6)$$

Money creation is governed by

$$x_{t+1} = Ax_t + C\epsilon_{t+1}, \quad \mu_t = Gx_t, \quad (7)$$

where x_t is an $n \times 1$ vector that carries information about future rates of money creation and ϵ_{t+1} is an i.i.d. shock. To exhibit T, suppose that the public forecasts inflation with a linear

⁶I owe this reading of Arrow's paper to Kihlstrom (2019), who notes that the dynamic programming arguments that make the equivalence of Arrow's two market arrangements transparent, namely those that decompose a once-and-for-all choice of state-contingent claims into a portfolio choice followed by spot purchases, were being created by Bellman while Arrow was writing. Arrow (1964, p. 91, n. 1) records the 1952 lecture and the 1953 French publication.

rule $\pi_t^* = \tilde{h}x_t$, where the $1 \times n$ vector $\tilde{h}$ parameterizes its subjective model. Equation (6) then implies $p_t = m_t + \alpha\tilde{h}x_t$, so actual inflation is $\pi_t = \mu_t + \alpha\tilde{h}(x_{t+1} - x_t)$. Its conditional expectation under the objective model is

$$E_t\pi_t = \mathsf{T}(\tilde{h}; \theta_s) x_t, \quad \mathsf{T}(\tilde{h}; \theta_s) = G + \alpha\tilde{h}(A - I), \quad (8)$$

where $\theta_s = (\alpha, A, G, C)$. Here T maps the public's subjective forecasting rule $\tilde{h}$ into the objective forecasting rule, a conditional-mean coefficient that summarizes the objective model. Rational expectations imposes the fixed point $h = \mathsf{T}(h)$, whose solution is

$$h : \quad \pi_t^* = hx_t, \quad h(\theta_s) = \frac{1}{1 + \alpha} G \left[I - \left(\frac{\alpha}{1 + \alpha} \right) A \right]^{-1}. \quad (9)$$

When $(1 + \alpha)I - \alpha A$ is invertible, the fixed point is unique among linear rules. Restricting forecasts to rules that are linear in x_t rules out bubbles when x_t has no component that grows at the rate $(1 + \alpha)/\alpha$ at which a bubble grows. When the eigenvalues of $\alpha A/(1 + \alpha)$ lie inside the unit circle, solving (6) forward and ruling out bubbles gives the same answer,

$$\pi_t^* = \frac{1}{1 + \alpha} \sum_{j=0}^{\infty} \left(\frac{\alpha}{1 + \alpha} \right)^j E_t \mu_{t+j}.$$

Equation (9) exhibits the hallmark cross-equation restrictions of rational expectations: the parameters (A, G) of the money creation process (7) appear in the rule h that describes expectations. A change in (A, G) changes h . With a scalar state and $A = \rho$, expected inflation is $\pi_t^* = \mu_t/(1 + \alpha - \alpha\rho)$, so more persistent money creation makes expected inflation respond more to current money creation. Lucas (1976) and Lucas and Sargent (1981) argued that econometric practice should impose these restrictions.⁷

5.2 Several agents

The Cagan example has a single representative forecaster. The communism of models matters most when there are several. Consider a version of Muth's (1961) market with two groups of suppliers, with all variables measured as deviations from their means. Demand is $D_t = -ap_t$ and supply is $S_t = b[\omega p_{1t}^e + (1 - \omega)p_{2t}^e] + u_t$, where p_t is the price, p_{it}^e is the price that group i expects, at $t - 1$, to prevail at t , $\omega \in (0, 1)$ is the share of group 1, $a, b > 0$, and the supply shock obeys $u_t = \rho u_{t-1} + \varepsilon_t$ with ε_t i.i.d. Suppose that group i forecasts with a

⁷The shock loading C does not appear in h because the conditional expectations $E_t x_{t+j} = A^j x_t$ do not depend on C . That fact also underlies the Simon (1956)–Theil (1957) certainty equivalence property of linear-quadratic control problems.

linear rule $p_{it}^e = h_i u_{t-1}$. Market clearing, $D_t = S_t$, then implies that the objective conditional expectation of the price is $E_{t-1} p_t = T(h_1, h_2; \theta_s) u_{t-1}$, where

$$T(h_1, h_2; \theta_s) = -\frac{b[\omega h_1 + (1 - \omega)h_2] + \rho}{a}, \quad \theta_s = (a, b, \omega, \rho). \quad (10)$$

Rational expectations imposes $h_1 = h_2 = h = T(h, h; \theta_s)$, whose solution $h(\theta_s) = -\rho/(a + b)$ is Muth's. The share ω drops out because, under the communism of models, there is only one forecasting rule. Suppose instead that group 2 uses a fixed rule h_2 that differs from $h(\theta_s)$, while group 1's forecasts equal the objective conditional expectation given group 2's rule. Then group 1's rule solves $h_1 = T(h_1, h_2; \theta_s)$, so that $h_1 = -[\rho + b(1 - \omega)h_2]/(a + b\omega)$. Now both ω and h_2 matter. The free parameters that the communism of models eliminated return as soon as one group leaves it.

5.3 Inference from prices

In the Cagan and cobweb examples, agents forecast prices. When some agents have private information, prices can also reveal that information, and an agent who interprets a price needs a model of how it was generated. A one-auction version of Kyle (1985) exhibits T for that inference problem. An asset pays v . An informed trader observes v and submits the order νv . Noise traders submit the order w . The payoff v and the noise order w are independent, mean-zero normal random variables with standard deviations σ_v and σ_w . A competitive, risk-neutral market maker observes total order flow $q = \nu v + w$ and sets the price equal to the conditional expectation of the payoff. If the market maker conjectures that the informed trader's coefficient is $\tilde{\nu}$, the price is

$$p = \tilde{E}[v | q] = \kappa(\tilde{\nu}) q, \quad \kappa(\tilde{\nu}) = \frac{\tilde{\nu} \sigma_v^2}{\tilde{\nu}^2 \sigma_v^2 + \sigma_w^2}, \quad (11)$$

where $\tilde{E}$ is an expectation under the conjecture $\tilde{\nu}$. Facing that pricing rule, the informed trader chooses an order d to maximize expected profit $d[v - \kappa(\tilde{\nu})d]$, so the trader's actual coefficient is

$$\nu = T(\tilde{\nu}) = \frac{1}{2\kappa(\tilde{\nu})} = \frac{\tilde{\nu}^2 \sigma_v^2 + \sigma_w^2}{2\tilde{\nu} \sigma_v^2}. \quad (12)$$

Rational expectations imposes $\nu = T(\nu)$, whose positive solution is $\nu = \sigma_w/\sigma_v$, so that $\kappa = \sigma_v/(2\sigma_w)$. The fixed point determines the market maker's inference rule and the informed trader's decision rule together. How much a price reveals is an output of the fixed point (Radner, 1979; Grossman and Stiglitz, 1980).

5.4 Two ways to build a rational expectations model

Beware of theorists bearing free parameters. (Arthur Goldberger; R. E. Lucas, Jr.; see Sargent, 1993, ch. 7, p. 168)

Lucas and Sargent (1981, pp. xv–xvi) described what is wrong with the alternative that preceded Muth’s idea:

What is “wrong” with [Milton Friedman’s adaptive expectations] is not [expressing] forecasts of future variables as distributed lags of current and lagged variables. The future must be forecast on the basis of the past, and it is surely acceptable to simplify things by modeling agents as using linear forecasting rules. . . . The difficulty lies not in postulating forecasts which are linear functions of history but rather in introducing the coefficients in these linear functions as so many additional “free parameters,” unrestricted by theory.

A model builder can equate subjective distributions to an objective distribution in two ways. One way is to compute a fixed point of (4), as in sections 5.1, 5.2, and 5.3. Another way is to start with a subjective forecasting rule and reverse engineer an environment in which it equals the objective conditional expectation. Muth (1961) did it the first way. A year earlier, Muth (1960) had done it the second way.⁸

Muth (1960) started with Friedman’s (1957) adaptive expectations, i.e., geometric distributed lag, model of how a consumer forecasts permanent income. He reverse engineered a *univariate* stochastic process for income for which Friedman’s formula for a consumer’s “permanent income” coincides with linear least squares forecasts. Let income be

$$y_t = y_t^p + \sigma_\eta \eta_t, \quad y_t^p = y_{t-1}^p + \sigma_\varepsilon \varepsilon_t,$$

where $\{\eta_t\}$ and $\{\varepsilon_t\}$ are mutually orthogonal white noises with means of 0 and standard deviations of 1. Then the linear least squares forecast is

$$\hat{E} \left[y_{t+k} \mid y^t \right] = (1 - \lambda) \sum_{j=0}^{\infty} \lambda^j y_{t-j}, \quad \forall k \geq 1,$$

⁸Constantinides and Duffie (1996) did it the second way for asset pricing. They started with a stochastic discount factor process that prices an arbitrary list of observed returns and reverse engineered an incomplete markets rational expectations economy whose equilibrium supports it. The key ingredient is a cross-section variance of idiosyncratic income growth that varies with aggregate consumption growth. Bhandari (2015) reverse engineered a different economy that supports the same stochastic discount factor. That economy has complete markets and heterogeneous beliefs, represented by a panel of likelihood ratio processes. Because agents’ subjective beliefs differ from the objective distribution, that economy does not impose rational expectations. Bhandari (2015) disciplines that departure by mapping beliefs into the cross-section dispersion of return forecasts.

where $\lambda \in (0, 1)$ solves $\sigma_\varepsilon^2/\sigma_\eta^2 = (1 - \lambda)^2/\lambda$, so that λ falls as the signal-to-noise ratio $\sigma_\varepsilon^2/\sigma_\eta^2$ rises. With Gaussian shocks, $\hat{E}$ is also the conditional expectation. Furthermore $y_t - y_{t-1} = e_t - \lambda e_{t-1}$ where e_t is the innovation in y_t relative to the history y^{t-1} . Muth's rationalization restricts the stochastic process for income. Before Lucas (1976), economists had not digested the implications of that finding for econometric practice (Sargent, 2024).⁹

Sargent (1977) performed another reverse engineering exercise in the spirit of Muth (1960), using the Cagan model that section 5.1 solved as a fixed point. He started with Cagan's (1956) adaptive expectations model of how people forecast inflation,

$$\pi_t^* = (1 - \lambda) \sum_{j=0}^{\infty} \lambda^j \pi_{t-1-j}.$$

Because $\pi_t = p_{t+1} - p_t$ is not known at t , the forecast uses inflation rates through $t - 1$. He reverse engineered a *bivariate* stochastic process for money creation and inflation for which Cagan's forecasts coincide with linear least squares forecasts of inflation conditioned on histories of money creation *and* inflation. That puts a restriction on the Granger–Sims causality structure of the joint (μ_t, π_t) process: money creation fails to Granger-cause inflation, while inflation can Granger-cause money creation. That restriction fit the 1922–1923 German hyperinflation well.¹⁰

5.5 Computing a rational expectations equilibrium

Computing a rational expectations equilibrium means finding a fixed point θ of (4). Write $\mathsf{T}(\theta)$ for $\mathsf{T}(\theta, \dots, \theta)$. Because T need not be a contraction, iterating $\theta^{k+1} = \mathsf{T}(\theta^k)$ can fail to converge. A relaxation algorithm can converge where that iteration does not:

$$\theta^{k+1} = \theta^k + \gamma_k [\mathsf{T}(\theta^k) - \theta^k], \quad \gamma_k \in (0, 1]. \quad (13)$$

Iterations (13) are Euler's method for solving the ordinary differential equation (ODE)

$$\frac{d\theta}{d\tau} = \mathsf{T}(\theta) - \theta. \quad (14)$$

Rest points of (14) are fixed points of (4). Stable rest points are the ones that relaxation (13) with small enough γ_k can find. A rest point is stable when the eigenvalues of the derivative of T at it have real parts less than one. Relaxation cannot find other fixed points. For example,

⁹However, Sargent (1971) and Lucas (1972a) had incorporated Muth's cross-equation restrictions into econometric tests of the natural rate hypothesis.

¹⁰Lucas and Sargent (1981) reprint that analysis. The joint distribution that emerges from the equilibrium is not invariant to monetary-fiscal interventions that alter the money creation process.

$T(\theta) = 2\theta + 1$ has the fixed point $\theta = -1$, but each step of (13) multiplies the distance from it by $1 + \gamma_k$. Adding noise to (13) leads to stochastic approximation and quasi-stochastic approximation algorithms for finding a θ that solves (4) (Meyn, 2022).

Four questions differ. Does an equilibrium exist? Is it unique? Can an algorithm compute it? Can a learning process find it? The last three presuppose the first but are otherwise distinct: a determinate equilibrium (section 7.1) can be unstable under a particular learning rule.

Part II. What rational expectations does

6 Rational expectations econometrics

Lucas (1987, p. 5) wrote that econometric implications and applications “may well be the most lasting contribution of John F. Muth’s idea of rational expectations.” Rational expectations econometrics involves two problems (Sargent, 2025). A *direct problem* takes a vector of known structural parameters θ_s and computes an equilibrium Markov process. A solution of the direct problem lets us draw random samples from $f(Y \mid \theta(\theta_s))$. An *inverse problem* takes an observed data set as known and infers unknown structural parameters θ_s .

Rational expectations econometrics assigns different knowledge to two types of people:

- agents *inside* the model know the model, i.e., they know θ_s , which equilibrium is selected, and hence $\theta = \theta(\theta_s)$;
- an econometrician *outside* the model knows the manifold $\{f(Y \mid \theta(\theta_s))\}_{\theta_s \in \Theta_s}$ but not θ_s .

These objects are in play:

- structural parameters $\theta_s \in \Theta_s$;
- hypothetical data $\{y_t\}_{t=0}^N$ generated by sampling from $f(Y \mid \theta(\theta_s))$;
- an estimator $\hat{\theta}_{s,t} = \hat{\theta}_s(y^t)$, where y^t is the history $y_t, y_{t-1}, \dots, y_0$;
- regularity conditions under which $\hat{\theta}_{s,t} \rightarrow \theta_s$ in probability or almost surely;
- a parameter-estimation algorithm, for example, a recursion

$$\begin{bmatrix} \hat{\psi}_{t+1} \\ R_{t+1} \end{bmatrix} = \mathcal{G} \left(\begin{bmatrix} \hat{\psi}_t \\ R_t \end{bmatrix}, z_t \right), \quad (15)$$

where $\hat{\psi}_t$ is an estimate of a parameter vector ψ , R_t is an empirical second-moment matrix that is updated jointly with $\hat{\psi}_t$, and z_t is a vector of observations. The econometrician

applies such a recursion with $\psi = \theta_s$. Learning agents inside a model apply it to the parameters of their own subjective models (section 14.2).

To solve the inverse problem, an econometrician computes, i.e., numerically approximates, an equilibrium for each candidate θ_s and assumes that nature drew the sample from it. Estimators include maximum likelihood (Hansen and Sargent, 1980), indirect inference (Gouriéroux et al., 1993), the efficient method of moments (Gallant and Tauchen, 1996), and the generalized method of moments (Hansen, 1982), which section 6.3 shows can also be applied without computing an equilibrium. An econometrician can pretend to be a frequentist and use maximum likelihood. An econometrician can instead pretend to be a Bayesian, put a prior over θ_s , and approximate a posterior.

Maximum likelihood minimizes a sample counterpart of (3). For a family of models $\{g_\delta\}_{\delta \in \Delta}$, $KL(f_\theta, g_\delta) = E \log f_\theta(Y) - E \log g_\delta(Y)$, where E is an expectation under the objective model f_θ , the first argument of (3). Minimizing over δ therefore maximizes the expected log likelihood $E \log g_\delta(Y)$. Maximum likelihood replaces that expectation with a sample average. Under regularity conditions, when the family contains f_θ , estimates converge to the δ at which $g_\delta = f_\theta$. When it does not, they converge to a minimizer of (3), an information projection (section 14.1; White, 1982).¹¹ Bayesian posteriors concentrate on the same minimizers (Berk, 1966).

6.1 Identification

The Cowles Commission constructed a special descriptive model called the reduced form of a structural model. Population moments implied by the structural model determine θ_d through a function

$$\theta_d = \varphi(\theta_s). \tag{16}$$

Structural parameters θ_s are (globally) identified when φ has a well defined inverse $\theta_s = \varphi^{-1}(\theta_d)$ (Koopmans, 1950; Hood and Koopmans, 1953).¹² Sims (1980) argued that rational expectations makes many Cowles Commission exclusion restrictions incredible. A rational expectations equilibrium restriction replaces them with cross-equation restrictions like (9). Under rational expectations, φ is a composition of two maps: the map $\theta_s \mapsto \theta(\theta_s)$ defined by fixed point (4), followed by a map from the objective model to the descriptive model. That second map delivers the parameters of the Cowles Commission reduced form, or those of an

¹¹Information theorists call such a minimizer a reverse information projection; their information projection minimizes over the first argument of (3) (Csiszár and Matúš, 2003).

¹²The classic Cowles Commission reduced form differs from the contemporary debased “reduced form” that denotes a descriptive model not tied explicitly to a structural model (Sargent, 2024).

information projection when the descriptive model is misspecified (sections 8.4 and 14.1). Cross-equation restrictions can make φ invertible (Sargent, 2025). In the market of section 5.2, the forecast coefficient $h = -\rho/(a + b)$ identifies only the ratio $\rho/(a + b)$. Separating a from b requires other observations.

6.2 Who is learning?

Statistical learning is always *about* a parameter vector. A rational expectations econometrician outside the model is learning about θ_s , for example via recursion (15). Agents inside the model are not learning about θ_s . They know it. In rational expectations models whose agents are Bayesian learners, agents instead know a prior over θ_s and the model that generates data from it; the prior becomes an additional parameter of a bigger rational expectations model, provided that nature draws θ_s from it (Bray and Kreps, 1987).

That asymmetry makes predictive densities differ. An agent inside the model who has observed y^t forecasts with $f(y_{t+1} | y^t; \theta)$. A Bayesian econometrician who has observed the same y^t forecasts with

$$\int f(y_{t+1} | y^t; \theta(\tilde{\theta}_s)) p(\tilde{\theta}_s | y^t) d\tilde{\theta}_s,$$

where $p(\tilde{\theta}_s | y^t)$ is a posterior. Agents inside the model have the same *information* as the econometrician, namely, y^t , but more *knowledge*: they know θ_s . Their predictive density is the objective one. The expected KL divergence (3) of the objective predictive density relative to the econometrician's is generally positive for $t < +\infty$ and, under regularity conditions, vanishes as $t \rightarrow +\infty$. A “behavioral bias” is a difference between two predictive distributions, so its size and sign depend on the benchmark.

6.3 Imposing the restriction partially

A widespread empirical use of the rational expectations restriction imposes it only partially. Hansen's (1982) generalized method of moments estimates parameters that appear in Euler equations without specifying complete laws of motion for agents' information variables. An agent who optimizes under subjective probabilities Q satisfies the Euler equation

$$E_t^Q [m_{t+1}(\gamma)R_{t+1} - 1] = 0, \quad m_{t+1}(\gamma) = \beta \frac{u'_\gamma(c_{t+1})}{u'_\gamma(c_t)}, \quad (17)$$

where R_{t+1} is a gross return and γ indexes preferences. Optimization alone implies (17). Rational expectations replaces the subjective expectation E_t^Q with the objective expectation E_t . The pricing error $m_{t+1}(\gamma)R_{t+1} - 1$ is then orthogonal to every variable z_t in agents'

information set:

$$E[z_t(m_{t+1}(\gamma)R_{t+1} - 1)] = 0. \quad (18)$$

Hansen and Singleton (1982) used such conditions to estimate preference parameters without solving for an equilibrium. Semiparametric efficiency bounds (Hansen and Singleton, 1991) measure how precisely parameters of interest can be estimated when nuisance parameters are left unrestricted. Hansen and Sargent (1993) and Sims (1993) constructed examples of misspecified models in which seasonal adjustment improves estimates of preference and technology parameters, the parameters of interest, while degrading estimates of nuisance parameters that govern exogenous processes that agents forecast but cannot control (Sargent, 2025). Imposing the restriction partially is a good response to an econometrician’s fear of model misspecification.

7 What the restriction buys and where it is used

7.1 What the restriction buys

The rational expectations restriction

- eliminates free parameters that describe subjective distributions. It converts agents’ subjective conditional distributions of prices and quantities from *inputs* into *outputs* of theorizing;
- sets KL divergences (3) between subjective distributions and the objective distribution to zero. That renders the model immune to Blume and Easley’s (2006) question “If you’re so smart, why aren’t you rich?” because, when beliefs are correct, market selection has no differences in beliefs to act on, although it still acts on differences in discount factors (section 11.4);
- respects and strengthens Harsanyi’s (1967) common prior doctrine. The fixed point (4) implies a common prior. It also requires that common prior to equal the objective distribution. A common prior can be commonly wrong.

The communism of models makes agents’ conditional distributions mutually consistent. That consistency lets model builders use these tools:

- recursive representations of equilibria, including nested Bellman equations for Ramsey and credibility problems (section 8.2);
- applications of Wiener–Kolmogorov optimal prediction theory to construct equilibria;¹³

¹³See Whiteman (1983), Kasa (2000), Walker (2007), Nimark (2017), Huo and Pedroni (2020), and Huo and Takayama (2025).

- analysis of counterfactuals;
- explicit models of “expectations management” through policy choices;
- models of runs, e.g., Diamond and Dybvig (1983), and the global games that select among their equilibria (section 7.3);
- statistical inference about parameters. In a time-invariant environment, a recursive equilibrium has a time-invariant law of motion, so that, when the resulting process is stationary and ergodic (after suitable scaling), laws of large numbers apply; cross-equation restrictions like (9) help identify structural parameters, even when there is a continuum of equilibria (section 5).

In a rational expectations equilibrium, nobody misunderstands the consequences of his or her choices.

Policy evaluation. Rational expectations made macroeconomic models usable for evaluating policy rules. Lucas (1980, pp. 709–10) described the task this way:

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.

Rational expectations lets the program compute agents’ beliefs from the policy rule instead of asking its user to supply them.

Welfare comparisons. The restriction settles whose probabilities welfare comparisons should use. Let E^{Q_i} denote expectation under agent i ’s subjective model. Under rational expectations, agent i ’s utility $U_i(c)$ from an allocation c satisfies

$$E^{Q_i} [U_i(c) \mid \mathcal{I}_{it}] = E [U_i(c) \mid \mathcal{I}_{it}],$$

so a model builder can rank policies by agents’ own criteria. The equality holds at a given date and information set. Welfare weights and the choice between ex ante and interim comparisons remain for the model builder to specify. When beliefs differ from the objective model, a welfare criterion must take a stand on whose probabilities to use. Brunnermeier et al. (2014) propose judging an allocation inefficient only if it is inefficient under every convex combination of agents’ beliefs.

Determinacy. The restriction makes determinacy of equilibrium a criterion for policy design. Blanchard and Kahn (1980) gave conditions under which a linear rational expectations

model has a unique stable solution. Monetary economists judge interest rate rules partly by whether they deliver a unique, or determinate, rational expectations equilibrium; in the basic New Keynesian model, the Taylor principle, which requires the nominal interest rate to move more than one for one with inflation, is a condition of that kind (Woodford, 2003).

Benchmark. Rational expectations equilibria serve as benchmarks for models that do not impose the restriction. Optimism and pessimism are optimism and pessimism *relative to* something. Rational expectations equilibria are the benchmarks against which “behavioral biases” are defined (section 6.2). They are the approximating models that robust agents, model builders, and econometricians doubt (section 10). Algorithms that compute them can be converted into models of non-rational learning (sections 5.5 and 14.2).

7.2 The restriction at work across economics

The restriction is not confined to macroeconomics. It often enters as part of the definition of equilibrium, under other names: equilibrium price expectations, a common prior from which nature draws types, correct beliefs about other players’ strategies, or knowledge of a transition law. Two distinctions organize its uses.¹⁴

The first distinction is between rational expectations about *exogenous* processes and rational expectations about *equilibrium* objects. A model in which an agent knows the law of motion of an exogenous shock, such as the distribution from which a job searcher draws wage offers, imposes rational expectations about something that nobody’s beliefs can affect, which requires no fixed point. Prices, job-finding rates, and other agents’ strategies are different. Their objective laws of motion depend on beliefs through the mapping T , so rational expectations about them requires a fixed point.

In some models the communism of models decomposes into the two blocks that the first distinction separates, an exogenous block and an endogenous one, and not every use of it needs both. Arrow’s (1964) welfare theorem for securities markets needs communism about the *endogenous* block only. His consumers’ subjective probabilities of states may differ from one another and from nature’s, so they need not share a common prior. Yet if consumers share and are right about state-contingent spot prices, his two market arrangements support the same allocation, which is Pareto optimal by consumers’ own subjective probabilities. That is a fixed point in the price block of T without $f_i = f$. The split keeps the welfare theorem

¹⁴Executable versions of several of the models described in this section appear in Sargent and Stachurski (2026b), including lectures titled “Rational Expectations Equilibrium,” “Investment Under Uncertainty,” “Markov Perfect Equilibrium,” “Job Search I: The McCall Search Model,” “Estimating Euler Equations by Generalized Method of Moments,” “Competitive Equilibria with Arrow Securities,” and “Heterogeneous Beliefs and Bubbles.”

but not the settlement of whose probabilities count (section 7.1). Moll (2026) proposes the reverse: imposing the restriction on exogenous processes but not on equilibrium prices (section 14.3). What lets Arrow's split support his welfare theorem is that his securities pay a unit of account and complete the market (section 5). When securities pay goods, as in Hart (1975), the fixed point in the price block can fail to exist or can select which markets function, whether or not consumers agree about probabilities (section 12.3). What a split delivers therefore depends on the model to which it is applied.

The second distinction is between rational expectations *on* the equilibrium path and *off* it. Some equilibrium concepts require beliefs to be correct only about events that occur with positive probability in equilibrium. Others also require correct beliefs about what would happen after deviations that never occur. Game theory draws that distinction most sharply.

General equilibrium. Radner (1972) defined an equilibrium of plans, prices, and price expectations in a sequence of markets and proved that one exists, under a bound on forward transactions that Hart (1975) showed is necessary: without it equilibrium can fail to exist, and with it the equilibrium in general depends on the bound chosen, which is arbitrary. Radner's equilibrium requires that the prices that agents expect in each future event be the prices that clear markets in that event. Agents need not agree about the probabilities of events, but they must agree about, and be right about, the price that will prevail in each event. Radner's sequence of markets extends the requirement to many dates and events; in Arrow's two periods, correct state-by-state spot prices amount to perfect foresight once the state is drawn. Radner (1979) studied rational expectations equilibria in which prices also transmit private information, and Grossman and Stiglitz (1980) showed that when information is costly to acquire, prices cannot transmit it fully.

Macroeconomics. Lucas (1972b) characterized equilibrium prices as functions of the state of the economy and required agents' forecasts to use those functions. The real business cycle models that followed Kydland and Prescott (1982), the New Keynesian models of Woodford (2003), and estimated models such as Smets and Wouters (2007) all impose the restriction. So, exactly or approximately, do heterogeneous agent models. In Aiyagari's (1994) stationary equilibrium, households know the constant interest rate that clears the market. In Krusell and Smith (1998), households facing aggregate risk forecast prices from a few moments of the cross-section distribution of wealth, in practice its mean. Their algorithm fits the coefficients of those forecasting rules by least squares to data that the model generates, so it computes what section 14.1 calls a restricted perceptions equilibrium. Krusell and Smith found that those forecasting rules predict prices almost perfectly, a property they called

approximate aggregation. Heterogeneous agent models such as Achdou et al. (2022) and Kaplan et al. (2018) descend from those models. The Federal Reserve Board’s FRB/US model can be simulated under either VAR-based or model-consistent, that is, rational, expectations (Brayton et al., 1997).

Consumption and asset pricing. Hall (1978) showed that under rational expectations and the life cycle–permanent income hypothesis, marginal utility is a martingale when the discount factor times the gross interest rate equals one: no variable in consumers’ information set other than current marginal utility helps to predict future marginal utility. With quadratic utility, the same statement holds for consumption itself. Lucas (1978) made asset prices functions of the state and priced assets with a stochastic discount factor built from a representative consumer’s marginal rates of substitution and the objective law of motion of endowments. Hansen and Singleton (1982) estimated the preference parameters in that stochastic discount factor from orthogonality conditions that rational expectations implies (section 6.3). The expectations theory of the term structure combines rational expectations with the assumption that term premia are constant. It makes long rates averages of expected future short rates plus those premia. Sargent (1979) estimated that theory by maximum likelihood, imposing its cross-equation restrictions. Campbell and Shiller (1991) found that yield spreads forecast long-run movements in short rates in the direction that the theory predicts, but forecast short-run movements in long rates with the wrong sign. The equity premium and risk-free rate puzzles of section 9 are failures of models that impose the restriction.

Finance and market microstructure. Fama’s (1970) efficient markets hypothesis, that prices fully reflect available information, is a statement about equilibrium beliefs. Testing it requires a model of equilibrium expected returns, because predictable returns can reflect risk premia that vary over time. In Kyle (1985), a market maker sets prices knowing the equilibrium relation between order flow and an insider’s private information (section 5.3). Milgrom and Stokey (1982) showed that when agents share a common prior, start from a Pareto efficient allocation, and it is common knowledge that they are rational, private information alone cannot generate trade. Explaining the large volume of trade in actual markets therefore requires relaxing one of those assumptions; an efficient initial allocation, for example, leaves no hedging or liquidity motive to trade. Harrison and Kreps (1978) showed how short-sale constraints and heterogeneity among traders’ statistical models can support speculative trade and asset prices that exceed what any trader would pay to hold the asset forever.

Game theory and mechanism design. A Nash equilibrium (Nash, 1950) can be read as a communism of beliefs about strategies: each player best responds to the strategies that the other players actually use. Conversely, a competitive rational expectations equilibrium with a continuum of agents can be read as a Nash equilibrium of a nonatomic game (Schmeidler, 1973), in which each player’s payoff depends on his or her own actions and on aggregate outcomes, such as prices, that no single player can affect. Lasry and Lions (2007) call such games mean field games and derive their equilibrium conditions as limits of Nash equilibria of games with N players as $N \rightarrow \infty$. Harsanyi (1967) turned games of incomplete information into games of imperfect information by giving players a common prior over types, and Aumann (1987) showed that correlated equilibrium expresses Bayesian rationality under a common prior. In an extensive-form game, a Nash equilibrium already requires each player to predict correctly how the others would play at information sets that are never reached in equilibrium. Subgame perfection adds that those off-path plans be optimal. Self-confirming equilibrium (Fudenberg and Levine, 1993) relaxes the requirement of correct beliefs off the equilibrium path (section 14.1). In mechanism design, Myerson’s (1981) optimal auction assumes that the seller and every bidder know the distribution from which bidders’ values are drawn.

Industrial organization. In a Markov perfect equilibrium (Maskin and Tirole, 2001), each firm’s beliefs about its rivals’ state-contingent policies equal those policies. Ericson and Pakes (1995) built a framework for empirical work on industry dynamics on that equilibrium concept. In Rust (1987), a decision maker who knows the transition law of bus mileage chooses when to replace bus engines. The econometrician estimates that transition law and attributes it to the decision maker. The two-step estimators of Hotz and Miller (1993) and Bajari et al. (2007) first estimate transition laws and policy functions from data and then assume that agents use those estimated objects to form their expectations.

Labor economics. A worker in McCall’s (1970) search model knows the distribution of wage offers, an instance of rational expectations about an exogenous process. In Mortensen and Pissarides (1994), the job-finding rate is an equilibrium object that firms’ free entry determines, and workers and firms know it; that requires a fixed point. With aggregate shocks and heterogeneous workers and firms, the distribution of matches determines whom a worker or firm will meet, so the fixed point includes a law of motion for that distribution. Payne, Rebei, and Yang (2026) compute such equilibria globally with deep neural networks. Keane and Wolpin (1997) estimated a model of career choice in which young men know the distributions of the shocks to their future rewards, the same distributions that the econometricians estimate.

Public finance and policy. Ramsey problems (Lucas and Stokey, 1983), Mirrlees problems (Mirrlees, 1971), time-consistency problems (Kydlund and Prescott, 1977; Calvo, 1978), and models of reputation (Chari and Kehoe, 1990; Stokey, 1991) all assume that private agents know the government’s policy rule.

International economics. In Dornbusch (1976), asset markets clear instantly, goods prices adjust slowly, and the model imposes rational expectations, so the exchange rate overshoots its long-run response to a monetary expansion. In sovereign debt models that descend from Eaton and Gersovitz (1981), such as Arellano (2008), lenders price a government’s bonds using default probabilities implied by the government’s own optimal default policy.

Information and higher-order beliefs. In Townsend’s (1983) model of “forecasting the forecasts of others,” in the beauty contests of Morris and Shin (2002), and in Huo and Pedroni (2020) and Huo and Takayama (2025), agents have different information but share a joint probability distribution (sections 2 and 5). Global games use such dispersed information to select a unique equilibrium in coordination games that have several when fundamentals are common knowledge (section 7.3). In models of rational inattention (Sims, 2003, 2006; Turmuhambetova, 2005), variables that interest a decision maker are polluted with measurement errors whose distribution the decision maker can choose at a cost. Sims (2003, 2006) constrains or penalizes the Shannon mutual information between the variables that a decision maker cares about and the signals that he or she observes. Turmuhambetova (2005) lets a decision maker choose an additional signal process and penalizes that choice according to the signal’s information content. In a rational expectations equilibrium of such a model, decision makers have different information sets because each chooses his or her information according to his or her own objectives. In Turmuhambetova’s (2005) portfolio choice model, for example, endogenous information induces heterogeneity in consumption and portfolio allocations across agents with different amounts of wealth.

Experimental economics. Rational expectations equilibria supply predictions that laboratory markets can test. Plott and Sunder (1982) found that in experimental security markets with insiders, prices came, with experience, to reveal the insiders’ information, as a rational expectations equilibrium predicts. Plott and Sunder (1988) found that markets aggregated traders’ diverse information and attained rational expectations equilibria when traders could trade a complete set of Arrow–Debreu securities or had identical preferences, but not in a third design. Whether a market attains a rational expectations equilibrium can depend on its institutions and trading instruments. In laboratory coordination games,

subjects use threshold strategies, and with private information their thresholds are closer to the global-game prediction than with common information, although the differences between the two treatments are small (Heinemann et al., 2004).

Empirical work in these fields estimates an objective model and attributes it to agents as their subjective model: Hansen and Sargent (1980) in macroeconomics, Rust (1987) and Hotz and Miller (1993) in industrial organization, and Keane and Wolpin (1997) in labor economics.

7.3 Rational expectations in modern macroeconomic models

Expectations management. In New Keynesian models, households and firms respond today to the entire expected future path of interest rates. Rational expectations makes credible announcements about that path affect current decisions. In the standard New Keynesian model with rational expectations, a promise to hold the nominal interest rate low for a long time has large effects on current output and inflation, effects that do not shrink as the promised policy recedes into the future. Del Negro et al. (2023) call those implausibly large effects the “forward guidance puzzle.” Proposed resolutions alter other ingredients, such as the completeness of markets (McKay et al., 2016), or alter beliefs, for example by making agents myopic about distant events (Gabaix, 2020) or by limiting their depth of reasoning (Farhi and Werning, 2019).

Determinacy. Determinacy of rational expectations equilibrium (section 7.1) organizes debates about monetary and fiscal policy. Leeper (1991) showed that in his model a unique equilibrium requires one policy to be “active” and the other “passive.” Monetary policy is active when it obeys the Taylor principle; fiscal policy is passive when it adjusts primary surpluses to stabilize government debt. The fiscal theory of the price level (Sims, 1994; Cochrane, 2023) studies equilibria in which fiscal policy is active and the price level adjusts so that the real value of government debt equals the present value of expected primary surpluses. Both versions of the determinacy argument rely on rational expectations of future policies. Replacing rational expectations with other belief specifications can change determinacy conditions: in Gabaix’s (2020) model, equilibrium can be determinate even when monetary policy violates the Taylor principle.

Runs and financial fragility. In Diamond and Dybvig (1983), depositors who expect others to withdraw early withdraw too, so a run and the absence of a run are both rational expectations equilibria. Selecting between them with a sunspot, as Cole and Kehoe (2000)

do for sovereign debt crises, leaves the probability of a crisis as a free parameter that indexes beliefs, the kind of parameter that the restriction otherwise eliminates (section 5). Global games eliminate it (Carlsson and van Damme, 1993; Morris and Shin, 1998, 2003). Each depositor observes a noisy private signal of the bank’s fundamentals. With small noise, a unique equilibrium survives: depositors run when their signals fall below a threshold, so a run occurs when fundamentals fall below a corresponding threshold. Goldstein and Pauzner (2005) used such an equilibrium to compute how the probability of a run depends on the terms of the deposit contract, and Rochet and Vives (2004) used one to analyze a lender of last resort. Global games have become a standard tool in models of bank runs, currency attacks, and market freezes (Morris and Shin, 2012), and outflows from mutual funds display the strategic complementarities that those models emphasize (Chen et al., 2010).

The common prior that the rational expectations restriction imposes selects the equilibrium. Whether that prior is correct does not matter. Depositors share a joint probability distribution of fundamentals and signals, but not information (section 5). Iterated deletion of dominated strategies uses that distribution at every round to compute the probability that other depositors’ signals lie below a candidate threshold. With common knowledge of depositors’ rationality and of their common prior, iterated dominance then implies correct beliefs about other depositors’ strategies. As the noise vanishes, the marginal depositor’s belief about the fraction of others who run becomes uniform (Morris and Shin, 2003), and Morris et al. (2016) trace uniqueness to approximate common certainty of such “rank beliefs,” beliefs about where one’s own signal ranks among others’ signals. Uniqueness does not need a common prior, but without one the selection changes: in Izmalkov and Yildiz (2010), investors’ sentiments, free parameters of the kind that a sunspot leaves, move the threshold even as noise vanishes.

Transitions after unanticipated shocks. The “MIT shock,” a complete surprise to the agents in a model (section 8.4), has become a computational device. Heterogeneous agent models are often solved by computing perfect foresight paths of prices after an unanticipated aggregate shock. Agents assign probability zero to the shock beforehand and know the entire subsequent path of prices afterward. To first order in aggregate shocks, the resulting impulse responses coincide with those of a rational expectations equilibrium with aggregate risk, so they deliver its linearized law of motion (Boppart et al., 2018; Auclert et al., 2021). Reiter (2009) instead linearized around the stationary cross-section distribution. These methods make rational expectations equilibria of heterogeneous agent models computable for small aggregate shocks; section 12.4 describes global methods for large ones.

Estimation. Bayesian estimation of rational expectations models such as Smets and Wouters (2007) is routine at central banks (Herbst and Schorfheide, 2016; Sargent, 2025). It imposes cross-equation restrictions like those of section 5.1. Auclert et al. (2021) use sequence-space methods to estimate heterogeneous agent models. Payne et al. (2026) solve the direct problem of section 6 for a whole region of structural parameters at once, by treating the parameters as additional state variables of the master equation (section 12.4), and then estimate their search and matching models by the simulated method of moments. Solving and estimating one of their models takes about three hours.

8 Policy under rational expectations

Counterfactual policies produce outcomes that historical data have never displayed. A rational expectations model of policy must take a stand on what people inside the model know about situations they have not experienced.

8.1 Counterfactuals

Marschak described the challenge:

The economist’s objectives are similar to those of an engineer but his data are like those of a meteorologist. The economist is often required to estimate the effects of a given (intended or expected) change in the “economic structure”, i.e., in the very mechanism that produced his data. None of these changes can he produce beforehand, as in a laboratory experiment; and since some of the changes envisaged have never happened before, the economist often has to estimate the results of changes he has never observed. (Marschak, 1950, p. 2)

A counterfactual question asks what a *stochastic process* would have been had government policy been different. A rational expectations model tells how agents’ predictive densities depend on government policies. Policy parameters, such as (A, G) in the Cagan model, are components of θ_s , so $h(\theta_s)$ shows how expectations respond to policy. More generally, let v be a policy parameter and write fixed point (4) as $\theta = \bar{T}(\theta; v)$. On a differentiable branch of equilibria,

$$\frac{\partial \theta}{\partial v} = [I - D_\theta \bar{T}]^{-1} D_v \bar{T}, \quad (19)$$

provided that the inverse exists. The term $D_v \bar{T}$ measures how a policy change affects outcomes at fixed beliefs; the inverse adds the feedback from outcomes to beliefs and back. A model

with “behaviorally biased” agents must explain how those biases depend on government policies. Models that contain robust decision makers do that (section 10.1).

A model’s responses to new policies come from its controlled laws. Two controlled laws, $y_{t+1} = u_t + \varepsilon_{t+1}$ and $y_{t+1} = 2u_t + \varepsilon_{t+1}$, illustrate the point. If policy has always set $u_t = 0$, the two laws generate the same data and the same correct forecasts. They disagree about setting $u_t = 1$. A rational expectations model resolves the disagreement by assuming a controlled law. Evidence from other environments or from experiments must support that assumption.

Counterfactual policies are of two kinds. A policy can be *possible*: historically imagined, but never tried. Or it can be *unimagined*: historically not even contemplated, an “innovation.” Ramsey, Mirrlees, and other optimal policy design problems take a heroic stand: they treat every policy in an admissible set as possible, and they assume that agents know the consequences of each. A planner in such a problem knows agents’ responses to each possible plan. That is expectations management on steroids.

A counterfactual also requires the model builder to decide which equilibria to admit under the new policy. In the Cagan model (6), the model builder rules out bubbles $c \left(\frac{1+\alpha}{\alpha} \right)^t$ that also satisfy the equilibrium difference equation (Sargent, 2024). A Ramsey planner can instead design policies that rule out such bubbles.

8.2 Time inconsistency and credibility

Agents can solve Bellman equations under any subjective model. Rational expectations pins down the probabilities in the nested Bellman equations, “dynamic programming squared,” that analyze time inconsistency and credibility. Examples are Ramsey plans for monetary-fiscal policies (Lucas and Stokey, 1983), repeated moral hazard and enforcement problems, and Mirrlees problems (Mirrlees, 1971) and their dynamic extensions. An agent’s value function, a backward-forward looking object, appears as an argument of a principal’s Bellman equation.

Kydland and Prescott (1977) interpreted time inconsistency as rendering dynamic programming inapplicable. Kydland and Prescott (1980) rescinded that claim by redefining the state that confronts a planner (Sargent, 2024). Chari et al. (2020) used a redefined state to reinterpret time inconsistency. Their Ramsey planner in a growth model must deliver households an exogenous value W of initial wealth measured in utility. The Ramsey problem then has a recursive representation with state (k, W) , where k is capital. Suppose that each period’s government chooses current policies and the value of wealth that its successor must respect, and nothing else. The Ramsey outcome is a Markov equilibrium of that partial commitment environment. Partial commitment to the value of wealth implements the full

commitment outcome. Partial commitment to next period’s tax rates typically does not. The state variable, a promised value, carries the commitment.

8.3 Actions, forecasts, and promises

A communism of models sometimes requires one variable to play multiple roles. In the Calvo–Chang model studied by Sargent (2024), the inflation rate π_t plays three roles: it is the actual rate of inflation; it is the public’s expected rate of inflation; and it is a promised rate of inflation chosen by a Ramsey planner at time 0. In a model of a credible plan, a government’s sequence of money creation rates $\vec{\mu}$ is both a sequence of actions chosen by the government and a sequence of private agents’ forecasts of those actions. Both models are deterministic. In a stochastic model in which the government’s policy rule is part of the model, the public forecasts a policy action by its conditional expectation given the rule, so the policy rule is a private sector forecasting rule. Managing expectations and choosing policies are the same activity.

8.4 Regimes

One response to the challenge posed by counterfactuals is to put prospective changes in policy inside the model, so that people inside the model have probabilities over them.

A *regime* is a fixed government policy, i.e., a strategy, a sequence of possibly history-dependent functions. A regime can be represented recursively in terms of some state, and different theories of government policy bring different states: an exogenous Markov state; a promised inflation rate; or a promised inflation rate together with continuation values (Sargent, 2024). Finding the state is an art. A *regime-switching model* is a collection of regimes, each indexed by a Markov-switching state j_t . Parameters describe regimes. Hyper-parameters describe switches.

Lucas’s parallel universes. Lucas (1976) compared two rational expectations equilibria, each describing outcomes under a *permanently operating* policy rule. The comparison says nothing about transitions between regimes. Policy analysis becomes a comparison of stationary distributions. Policy advice emerges as a recommendation to engineer a complete surprise, i.e., an “MIT shock.” That treats the new policy as an “innovation” (section 8.1) to which agents assigned probability zero. Sargent (1982) studied the ends of four big inflations after World War I as laboratories for the study of regime changes. In Austria, Hungary, Poland, and Germany, inflation stopped abruptly rather than gradually once governments created independent central banks and simultaneously altered their fiscal policy regimes,

experiences that are difficult to reconcile with views that attribute momentum to inflation. In a footnote, Sargent (1982, n. 39) recorded a deep objection to that interpretation: in a rational expectations model of a single economy, what look like changes in regime are realizations of events under a single, more complex regime.

Complete rational expectations models. Sargent and Wallace (1976), Sims (1982), and Sargent (1984) refined the Lucas critique when they posited that prospective switches of government policy belong *inside* a historically realistic rational expectations equilibrium, because proposals for new policies do not emerge in a vacuum. Sims reinterpreted Lucas's examples as neglected nonlinearities: the population limit of a vector autoregression fitted to data from a hidden Markov regime-switching model is an information projection of that model onto a manifold of linear models (section 14.1). Svensson and Williams (2008, 2009) provide a way to put prospective switches inside an equilibrium. Their Markov-jump linear-quadratic dynamic programming problems have N regimes, one for each Markov-jump state, and N coupled Riccati equations. Because the Riccati equations are coupled through the probabilities of switching among regimes, agents' beliefs about prospective switches are part of the equilibrium. A change of policy regime then alters equilibrium allocations through those beliefs, which update as regimes switch, rather than through an unanticipated break. Rational expectations econometrics can estimate the switching probabilities that agents use, e.g., by fitting a hidden Markov model.

9 Puzzles and the hard core

Rational expectations models are coherent, but rarely comprehensive. Users of these models accumulate “puzzles,” i.e., patterns that their models miss. Examples are the equity premium puzzle, the risk-free rate puzzle, and other asset pricing anomalies, and, in growth theory, the failure of incomes in poor countries to converge to those in rich countries.

A standard response is to insist on rational expectations and to modify specifications of

- preferences: Kreps and Porteus (1978) recursive preferences; hyperbolic discounting (Laibson, 1997); Gul and Pesendorfer (2001) temptation preferences; habits; reference dependence (Barberis et al., 2001);
- technologies: Romer (1986) externalities; adjustment costs;
- shock processes: stochastic volatility; long-run risks (Bansal and Yaron, 2004; Hansen et al., 2008).

In their sophisticated versions, each of these responses retains the rational expectations restriction. Naive hyperbolic discounters, by contrast, misforecast their own future behavior.

In the language of Lakatos (1970), the rational expectations restriction functions as the hard core of a research program, and specifications of preferences, technologies, and shock processes form its protective belt. The restriction has earned that status. Preferences, technologies, and shock processes can be disciplined independently: preferences by introspection and experiments, technologies by engineering evidence, and shock processes by financial and other time series. Beliefs about endogenous variables are different: under the restriction they are equilibrium outcomes that change when policies change. Beliefs fitted to survey data (section 12.2) can be disciplined by data, but a model that used them would need another theory of how they change under a new policy. Another response modifies beliefs while keeping a shared approximating model (section 10).

10 Rational expectations with uncertain agents

10.1 Equilibria with a shared approximating model

In a rational expectations equilibrium with uncertain agents, all agents share an *approximating*, or *baseline*, model, which is the objective model that the equilibrium generates, but no agent fully trusts it. Each agent surrounds the approximating model with a set of nearby models whose divergences from it are measured by relative entropy, and chooses the decision rule whose worst-case performance over that set is best. Because agents' objectives differ, the worst-case models under which they evaluate decisions differ too.

Three examples illustrate such equilibria. Hansen and Sargent (2008, ch. 12) construct competitive equilibria with robustness in a pure endowment economy. They formulate a robust planning problem, decentralize its allocation with robust consumers who trade Arrow securities, and give the robust planning problem an *ex post* Bayesian interpretation. The lecture “Robust Markov Perfect Equilibrium” in Sargent and Stachurski (2026a), at https://python-advanced.quantecon.org/rob_markov_perf.html, provides another example. Two duopolists share a baseline model of the state transition law but doubt it. Each firm chooses a decision rule while a fictitious adversary distorts the transition law, subject to a relative entropy penalty. The firms' worst-case transition laws differ, so the firms hold different beliefs about future output and prices. Hansen et al. (2026) provide an example in a heterogeneous agent economy. Consumers in their Bewley economy share a baseline model of their endowments but differ in their discount factors and in how much they distrust that model. Along a locus of those two parameters, any distribution of consumers supports the same interest rate and the same joint distribution of consumption, debt, and income as an economy of consumers who trust the baseline model, while consumers' worst-case forecasts

and welfare evaluations differ.

With a multiplier penalty ξ_i on KL divergence, agent i 's worst-case density twists the approximating model:

$$\tilde{f}_i(x_{t+1} \mid x_t) \propto \exp\left(-\frac{V_i(x_{t+1})}{\xi_i}\right) f_\theta(x_{t+1} \mid x_t),$$

where $V_i(x_{t+1})$ is agent i 's continuation value at next period's state. The penalty that produces this twist is $\xi_i KL(\tilde{f}_i(\cdot \mid x_t), f_\theta(\cdot \mid x_t))$, a divergence between one-step-ahead conditional densities at the current state. It measures the divergence of the worst-case model relative to the approximating model, the reverse of the order in (3). The two orders answer different questions. Maximum likelihood (section 6) and the information projections of section 14.1 average log likelihood ratios under the model that generates the data. A robust agent averages them under the worst-case model, the model under which the agent evaluates decisions. That order delivers the exponential twist above and links robustness to the risk-sensitive control of Jacobson (1973). With agent i 's twisted model written as $S_i(\theta)$, an equilibrium is a fixed point

$$\theta = T(S_1(\theta), S_2(\theta), \dots, S_I(\theta)). \quad (20)$$

In the formulations of Hansen (2014) and Hansen and Sargent (2022), robustness adds two parameters per agent: one measures doubts about a prior over a family of likelihoods (ambiguity), and the other, related to ξ_i , measures doubts about the likelihoods themselves (misspecification). The twist above is the case with misspecification concerns only. Calibrating the two parameters with detection error probabilities keeps worst-case models hard to distinguish statistically from the approximating model. Because worst-case probabilities have an *ex post* Bayesian interpretation, robustness provides a structural model of “behavioral biases” and of belief heterogeneity. It puts agents in the econometrician's position: both doubt the approximating model.¹⁵ Estimation proceeds essentially as in rational expectations econometrics (Szöke, 2022). For applications to asset pricing, including the equity premium puzzle of section 9, see Hansen et al. (1999), Tallarini (2000), Barillas et al. (2009), and Cogley and Sargent (2008). For fiscal policy, see Karantounias (2013) and Ferriere and Karantounias (2019). For a textbook treatment, see Hansen and Sargent (2008); for code, see the lecture “Robust Permanent Income and Pricing” in Sargent and Stachurski (2026b).

¹⁵After I presented a joint paper with Lars Hansen about robustness at the Federal Reserve Bank of Minneapolis, Bob Lucas asked me “why should the people in our models be like us?” (Sargent, 2025).

10.2 Ambiguity, misspecification, and uncertainty prices

Hansen and Sargent (2024) separate two kinds of doubt by pulling apart the two ingredients of the Bayesian econometrician’s predictive density in section 6.2: a family of likelihoods $\{f(Y \mid \theta(\theta_s))\}_{\theta_s \in \Theta_s}$ and a prior over θ_s . An agent inside a rational expectations model knows θ_s . An agent inside a Bayesian model knows the family and a prior (section 6.2). An agent who knows the family but does not know or trust a prior over its members faces *ambiguity*. An agent who distrusts the likelihoods themselves, suspecting that every member of the family is misspecified, faces *misspecification*. Hansen and Sargent (2021) call members of the family *structured* models,¹⁶ and they use *uncertainty* for ambiguity and misspecification together, ignorance about the probability distributions of the risks that the models describe.

Hansen and Sargent (2021) formalize these ideas in continuous time. Each alternative to a baseline diffusion driven by a Brownian motion B is represented by a positive martingale M^v , a likelihood ratio with $dM_t^v = M_t^v v_t \cdot dB_t$. Under the alternative model, the increment dB_t acquires the drift $v_t dt$. Structured models have drift distortions s_t that lie in sets Ξ_t generated by alternative parameter values. A representative investor ranks consumption processes with continuation values

$$V_t = \min_{\{v_{t+\tau}\}} E_t \left[\int_0^\infty \beta^\tau \frac{M_{t+\tau}^v}{M_t^v} \left(\log c_{t+\tau} + \frac{\xi}{2} \min_{s \in \Xi_{t+\tau}} |v_{t+\tau} - s|^2 \right) d\tau \right]. \quad (21)$$

Minimization over structured models expresses ambiguity. The relative entropy penalty on the distance between v and the set of structured models expresses misspecification.¹⁷ The size of that set and the penalty ξ are the two parameters of section 10.1. As $\xi \rightarrow \infty$, the investor has max–min preferences over structured models. When the set of structured models shrinks to the baseline model, only misspecification concerns remain, as in Hansen and Sargent (2001).

Equilibrium prices equal shadow prices of a robust planning problem. The vector of local prices of exposures to dB_t is a risk price, which comes from the curvature of utility, plus an uncertainty price $-v_t^*$, where v_t^* is the worst-case drift distortion. With s_t^* the drift distortion of the worst-case structured model, Hansen and Sargent (2021, pp. 239–40) decompose the uncertainty price as

$$\underbrace{-v_t^*}_{\text{uncertainty price}} = \underbrace{-s_t^*}_{\text{ambiguity price}} \underbrace{-(v_t^* - s_t^*)}_{\text{misspecification price}} \quad (22)$$

¹⁶Structured models are “more or less tightly parameterized statistical models” that “are not what econometricians working in the tradition either of the Cowles commission or of rational expectations econometrics would call ‘structural’ models” (Hansen and Sargent, 2021, p. 222, n. 1).

¹⁷I write v , s , B , and c for their U , S , W , and C , β^τ for their $\exp(-\delta\tau)$, and ξ for their $\theta\delta$.

The first component comes from not knowing which structured model is best, the second from fearing that all of them are misspecified.¹⁸ Uncertainty prices enter the stochastic discount factor through the worst-case likelihood ratio. In the Euler equation (17), a robust investor’s E_t^Q is an expectation under the worst-case model. Measured with objective expectations, the stochastic discount factor is therefore a marginal rate of substitution times a likelihood ratio. An econometrician who imposes rational expectations sets that likelihood ratio to one and must attribute uncertainty prices to risk aversion. Earlier quantitative work had found that “investors’ responses to modest amounts of model ambiguity can replace the implausibly large risk aversions during economic downturns that are required to explain observed market prices of risk” (Hansen and Sargent, 2021, p. 223).

Four features of this work matter here. First, agents, the model builder, and the econometrician share the baseline model and the family of structured models, and uncertainty prices are measured relative to the baseline model. Hansen and Sargent estimated the baseline model from time series of macroeconomic quantities alone. They deliberately ignored asset prices and so did not impose the cross-equation restrictions that their asset pricing theory implies (Hansen and Sargent, 2021, p. 243). Second, detection error probabilities discipline the two new parameters. Statistical tests would take a long time to tell the worst-case models apart from the baseline model or from each other: half-lives of mistake probabilities range from 60 to 668 quarters (Hansen and Sargent, 2021, Table 1). Misspecification contributes roughly a third to a half of long-horizon uncertainty prices (Hansen and Sargent, 2021, p. 242 and Table 2). Third, uncertainty prices vary over time without the stochastic volatility that section 9 lists among modifications of shock processes, because the investor fears high persistence of low growth rates and low persistence of high growth rates. Fourth, a self-confirming equilibrium (section 14.1) admits only models that agree with the model that generates the data along the equilibrium path, so that no history, however long, tells them apart; Hansen and Sargent instead admit models with small positive relative entropy, which only long histories can tell apart (Hansen and Sargent, 2021, p. 234, n. 16).¹⁹

¹⁸Earlier work surrounded a single baseline model with a relative entropy neighborhood. The “market price of uncertainty” that Anderson et al. (2003, p. 68) linked to “a bound on the error in statistically discriminating between an approximating and a worst case model” is a misspecification price when the set of structured models contains only the baseline model (Hansen and Sargent, 2021, p. 226, n. 9).

¹⁹Hansen and Sargent (2010, p. 130) study a discrete-time precursor in which a representative consumer learns by Bayes’ law about two models of consumption growth and their parameters, and “doubts the Bayesian model-mixing probabilities as well as the specification of each model.” Those doubts put countercyclical uncertainty premia into prices of risk.

Part III. Defending rational expectations and taking stock

11 What rational expectations asks of people

11.1 The “as if” stance

A rational expectations model is a coherent model of *what* future outcomes people inside the model forecast. It is not a model of *how* people think about forecasting. Lucas and Prescott (1971, p. 664) took that stance:

It remains for us to link the anticipated price sequence to the actual price sequence—also a sequence of functions on E^t . Typically, this is done by postulating the process by which firms actually develop forecasts from actual past values and other information. In this paper, we shall not take this route but rather postulate a property of the outcome of this (unspecified) process. Specifically, we assume that expectations of firms are rational, or that the anticipated price at time t is the same function of $(u_1, \dots, u_t)$ as is the actual price. That is, we assume that firms know the true distribution of prices for all future periods.

That stance resembles Newton’s refusal to explain the cause of gravity (“hypotheses non fingo”), and it is a version of the “as if” reasoning that physicists use in the principle of least action and that Friedman and Savage (1948, p. 298) and Friedman (1953, p. 21) used to predict the shots taken by an expert billiards player. No particle solves the variational problem in the principle of least action. No billiards player solves the equations of mechanics. Nobody inside a rational expectations model computes its equilibrium. That is the model builder’s and the econometrician’s task.

11.2 What agents need to know

What must people inside a rational expectations model know? Rational expectations econometrics attributes knowledge of θ_s to them (section 6). Their decisions require less. An agent’s decisions depend on the objective model only through the conditional distributions of variables that enter his or her own decision problem, typically exogenous shocks and prices. Restriction (5) requires those distributions to be correct. An agent need not know other agents’ preferences, how goods are produced, or the equations that a model builder

uses to compute an equilibrium.²⁰ As Prescott and Mehra noted (section 4), knowing H , p , and $F(\zeta' | \zeta)$ suffices for forming predictive distributions of future prices. Those predictive distributions are what a household needs. The functions H and p are one way to represent them.

Arrow's (1964) construction (section 5) proves for his economy that agents need forecast only prices. Where trading claims to goods in states at time 0 needs SC markets, his securities-then-spot-markets arrangement needs $S + C$: a consumer observes the S security prices and must forecast, state by state, the C spot prices that tell him or her what a dollar delivered in that state will buy. Arrow (1964, p. 94) observed that this "permits economizing on markets." Forecasts pay for the saving in markets: consumers must forecast spot prices correctly.²¹

Hayek (1945, pp. 526–27) made the same point about a price system: "The most significant fact about this system is the economy of knowledge with which it operates, or how little the individual participants need to know in order to be able to take the right action." The mean field games of section 7.2 share that economy. Lasry and Lions (2007, p. 229) describe "a very large number of 'rational players' with a limited information (or visibility) on the 'game,'" each of whom "chooses his optimal strategy in view of the global (or macroscopic) informations that are available to him and that result from the actions of all players."

Hellwig and Venkateswaran (2025) formalize Hayek's point in a rational expectations model. Monopolistically competitive firms observe only their own sales and wages. Those signals reveal each firm's current marginal cost and revenue. They can be poor indicators of aggregate shocks. With flexible prices, equilibrium prices and quantities equal those of an economy in which firms observe all aggregate and firm-specific shocks. With sticky prices, a firm must forecast its future marginal costs. The equivalence then requires the firm's current marginal cost to be a sufficient statistic for forecasting its future marginal costs, that is, marginal costs must follow a first-order Markov process. That happens only when firm-specific and aggregate shocks are equally persistent and pricing decisions are not strategic complements. Firms then "do not need to know whether their marginal costs are driven by firm-specific shocks, aggregate shocks, or the other firms' pricing decisions" (Hellwig and

²⁰Fernando Alvarez tells first-year macroeconomics students that a competitive equilibrium is like living through COVID. Isolated from other people, you buy and sell things on Amazon. You need to know prices. You need not know how things are produced or who produces them.

²¹Cole and Prescott (1997, proposition 5) prove a counterpart for economies in which people form households, firms, and other clubs. An equilibrium in which consumers trade lotteries over consumption and club memberships at a single date is equivalent to one in which they first make fair wealth gambles and then trade deterministically in markets for memberships. The second arrangement trades a smaller and more natural set of commodities, and because the gambles are chosen at the second-stage prices, it asks consumers to forecast those prices correctly.

Venkateswaran, 2025, p. 2). Otherwise, dispersed information slows the response of prices to nominal shocks.

11.3 Histograms and stationarity

Christopher Sims responded to widespread early-1970s statements like “I don’t believe that real people know enough to have rational expectations.”²² Sims gave what I still find a telling answer: people who live inside a rational expectations equilibrium need only form joint histograms. Under stationarity and ergodicity, a law of large numbers makes those histograms converge to the joint probability distribution that the equilibrium generates. That joint distribution is an optimal forecasting model. Sims’s argument is nonparametric apart from a choice of bin width and of the variables and lags over which to tabulate, in effect a finite-dimensional state.²³

It rests on foundations of time series statistics required to justify applying a law of large numbers (an ergodic theorem), such as stationarity; ergodicity; and transformations that attain stationarity and ergodicity in the midst of growth, e.g., cointegration and scalings by additive and multiplicative functionals (Hansen and Scheinkman, 2009). Rational expectations econometrics also requires those foundations.

Lucas and Prescott (1971, p. 664, n. 9) had rested their defense of rational expectations on the same foundations:

... we can think of no objection to this assumption which is not better phrased as an objection to our hypothesis that the stochastic component of demand has a regular, stationary structure. If the demand shift assumption is reasonable, then expectations rational in our sense are surely more plausible than any simple, adaptive scheme. If the demand assumption is unreasonable, then adopting an alternative expectations hypothesis will certainly not improve matters.

Rational expectations itself does not require stationarity: perfect foresight models and models with known time-varying laws of motion impose it without a stationary distribution.

Sims’s histograms are an account of how people come to know the objective law of motion $f(Y | \theta)$, which reveals θ only at states that the equilibrium visits. They need not reveal the structural parameters θ_s . Section 6.2 said that agents inside a rational expectations model already know θ_s , and hence θ , or, in Bayesian versions, a prior over θ_s . The two accounts are

²²Chris Sims told me this in the early 1970s in the hallway at the Department of Economics of the University of Minnesota. I had just returned from presenting Sargent (1973) at the Brookings Institution and had told Chris about some of the skepticism and criticisms the paper had provoked.

²³Section 12.4 describes models in which no finite-dimensional state reproduces the forecasts of an equilibrium in which agents observe the whole cross-section distribution.

consistent only in the limit. Sims’s histograms tabulate data generated by an equilibrium in which people already know θ . Whether people who are still learning drive the data toward that equilibrium is a separate question. Bray and Kreps (1987) and Marcet and Sargent (1989) studied it.

Knowing θ suffices for forecasting within the stationary process that generated the histograms. Would agents also need θ_s to forecast after a change in government policy? Sims had in mind people with a very long sample, during which the economy has moved among all possible policy regimes, which have switched often. Switches of regime are then themselves part of the stationary joint distribution that the histograms tabulate, as in the complete rational expectations models of section 8.4. A joint histogram over a long enough sample approximates that joint distribution well, including the probabilities of regime switches and the behavior of outcomes after them. Ljungqvist and Sargent (2027, ch. 5) describe the technicalities that underlie Sims’s argument. The argument does not cover policies that have never been tried, whether possible or unimagined (section 8.1), and it requires regime switches to be frequent enough that histograms fill in samples of realistic length (section 12.1).

11.4 Adaptation and market selection

Two other defenses appeal to long experience. Lucas (1986) described economics as studying decision rules that are steady states of adaptive processes, and explored what happens when decision makers recursively update least squares estimates of the parameters of a statistical model. Under suitable conditions, stochastic approximation results make their estimates and decision rules converge to the beliefs and decision rules of what section 14.1 calls a restricted perceptions equilibrium, which is a rational expectations equilibrium when their model is correctly specified and equilibrium-path data single out the correct member of their family (Sargent, 2025). Market selection, an argument that Alchian (1950) and Friedman (1953) made for profit maximization, offers another defense: traders whose beliefs are less accurate lose wealth to traders whose beliefs are more accurate. Blume and Easley (2006) show that when allocations are Pareto optimal, survival depends only on discount factors and beliefs, so that among traders with equal discount factors those whose beliefs have the smallest KL divergence (3) survive, but that with incomplete markets the market selection hypothesis can fail. Selection therefore supports the restriction when markets are complete and some trader whose beliefs are correct, or become correct as a Bayesian’s do when the truth is in the support of his or her prior, is at least as patient as the others.

12 Where the restriction strains

The roles described so far rest on assumptions that do not always hold.

12.1 Rare and off-path events

When decisions depend on probabilities of rare but consequential events, the data say little about the probabilities that matter most. The equilibrium joint probability distribution depends on decision makers' beliefs about conditional probabilities of events that rarely occur in equilibrium. Examples are peso problems, rare disasters (Rietz, 1988; Barro, 2006), and run probabilities in Diamond and Dybvig (1983). A rational expectations model assumes that agents know the consequences of such events. Those consequences are learned slowly, if at all, from observations generated by the model. A rational expectations model therefore expresses faith in an extraneous source of deduction. How much can people deduce, possibly aided by instruction or by experience in related environments? *Rare* events occur along equilibrium paths, but infrequently. Histograms eventually reveal their probabilities, but slowly. *Off-path* events never occur in equilibrium. No amount of equilibrium-path data reveals their consequences. Adopting a policy that has been imagined but never tried is a rare event in a model that assigns it positive probability, as the complete models of section 8.4 do, and an off-path event in a model that assigns it probability zero, as an MIT shock does. A policy that nobody has imagined enters a model only when a model builder adds it, usually as an MIT shock (sections 8.1 and 8.4). Runs in Diamond and Dybvig (1983) can be off-path events: in an equilibrium without runs, agents' beliefs about what a run would bring are never tested. In global games (section 7.3), runs are on-path events that occur when fundamentals are weak, so beliefs about them are tested, but rarely: in data for nearly two hundred countries since 1800, runs occur in about 2 percent of country-years (Jamilov et al., 2026). Self-confirming equilibria (section 14.1) allow beliefs about off-path events to be wrong.

12.2 Survey evidence

Surveys of expectations record disagreement among forecasters and forecast errors that can be predicted. Such evidence is not new. Tests of rational expectations with survey data began in the 1970s (Pesando, 1975), and Lovell (1986) reviewed many of them. Some of that evidence speaks against models in which all agents share the same information, not against the rational expectations restriction. Coibion and Gorodnichenko (2015) interpret predictable errors in average forecasts as evidence of information rigidities, which rational expectations models

with dispersed or costly information, such as those in section 7.2, generate. Bordalo et al. (2020) find that individual forecasters overreact to their own news while average forecasts underreact. Individual overreaction is harder to reconcile with forecasts that are objective conditional expectations. Let $\hat{y}_{i,t}$ be forecaster i 's forecast at t of a variable y_T . If the forecast is the objective conditional expectation given information that grows over time, the law of iterated expectations implies

$$E[(y_T - \hat{y}_{i,t})(\hat{y}_{i,t} - \hat{y}_{i,t-1})] = 0,$$

so the error of such a forecaster cannot be predicted from his or her own forecast revision. A test with survey data tests this implication jointly with assumptions about what respondents report. Hansen et al. (2026) show that optimizing consumers who doubt their model can generate both overreaction and underreaction. Their worst-case forecasts overreact to news in samples in which consumers' distances from a bliss point of consumption are concentrated and underreact in samples in which those distances are dispersed. Those are forecasts of the consumers' own consumption, and reading a worst-case forecast as a reported belief is one of the assumptions about what respondents report.

Survey questions often ask about aggregate conditions. In the flexible-price economy of section 11.2, firms' beliefs about aggregate conditions are widely dispersed, but "the extent of this dispersion in beliefs has absolutely no effects on allocations or aggregate dynamics" (Hellwig and Venkateswaran, 2025, p. 7). Survey evidence therefore tests the restriction only jointly with assumptions about agents' information, about what respondents report, and about which beliefs matter for allocations.

12.3 Existence and selection

Hart (1975) showed that an equilibrium can fail to exist and can fail to be unique in a two-period general equilibrium model with two consumers, two states, two goods, strictly concave utility functions, strictly positive endowments, no transaction costs, and no frictions of any other kind. His examples arise because his securities pay goods (section 5). A forecast of spot prices then determines the span of the securities' payoffs, that span determines consumers' budget sets, and those budget sets determine the prices being forecast. Budget correspondences therefore fail to be upper semicontinuous at prices at which the span drops, so the standard fixed-point proof of existence fails.

In one example the fixed point has two solutions. Under one forecast the two securities have proportional dollar payoffs, so no wealth can be moved across states, consumers trade only on each state's spot markets, and the resulting spot prices are the forecast ones. Under

another forecast the payoffs are linearly independent, so the securities span, the allocation is the complete markets allocation, and the spot prices that support it are again the forecast ones. Both are rational expectations equilibria, they differ in which markets are effectively open, and every consumer strictly prefers the second. Changing only the securities' dividends produces an example with no rational expectations equilibrium at all: every candidate forecast refutes itself. Forecasts under which the securities fail to span imply spot prices under which they do span, and forecasts under which they span imply spot prices under which they do not.

Three things follow. First, an equilibrium is a fixed point, not the maximum of a welfare criterion, so equilibria can be Pareto ranked, and selecting among them is not welfare neutral. Second, the multiplicity here is of a different kind from the one that the determinacy literature of section 7.1 studies. There, a linear model has many stable paths and policy is judged by whether it leaves exactly one; here, what beliefs select is the set of markets that function. Third, the restriction can be unsatisfiable: imposing $f_i = f$ on a model that is otherwise well behaved can define an empty set.

Hart's examples are knife edge, in the sense that small perturbations of preferences, endowments, or dividends restore uniqueness in the first and existence in the second. Failures of existence are not generic (Duffie and Shafer, 1985), but they occur in well-behaved economies. Failures of Hart's kind cannot occur in a two-period economy with one good, or with one money-denominated security for each state, as in Arrow's construction. With fewer money-denominated securities than states, equilibria exist, but forecasts of price levels change real payoffs; when there are more households than securities and nothing else in the model determines price levels, those forecasts generically index a continuum of equilibrium allocations (Balasko and Cass, 1989; Geanakoplos and Mas-Colell, 1989), free parameters that index beliefs (section 5). With one good and three dates, a forecast interest rate can play the role of a forecast relative price, and a third example in Hart (1975) has Pareto-ranked equilibria of the same kind.

Selection by information has limits. The global games of section 7.3 select a unique equilibrium only when private information is precise relative to public information. Precise public signals, including those that equilibrium prices supply, can restore multiplicity (Hellwig, 2002; Angeletos and Werning, 2006), so agents' inference from equilibrium outcomes can undo the selection. Hellwig et al. (2006) make the point in a market setting. When an interest rate is determined in a noisy rational expectations equilibrium, every trader's optimal strategy is unique, yet multiple market-clearing interest rates can survive the introduction of noisy private signals, because demand for domestic bonds bends backward. The model of Morris and Shin (1998) is the special case in which bond supply is perfectly elastic at a fixed interest

rate, so that the interest rate conveys no information. As in Hart’s examples, uniqueness depends on a model’s information structure and on how its markets work.

12.4 Large state spaces

Heterogeneous agent models with aggregate risk add a challenge of their own: their equilibrium prices depend on the cross-section distribution of wealth and income. Two systems of equations can characterize an equilibrium. One couples a backward Bellman equation for each agent’s value, in which the rest of the economy enters only through the stochastic process of prices, with a forward equation for the cross-section distribution. The other makes the distribution a state variable, as K is in section 4, and folds both equations into a single Bellman equation over an infinite-dimensional state. That equation, the “master equation” of mean field games, has been nicknamed the “Monster equation.” Under regularity conditions, the two systems characterize the same equilibrium (Cardaliaguet et al., 2019).

The master equation is a device that a model builder can use to compute an equilibrium. Agents need correct predictive distributions of the prices and shocks that enter their decision problems (section 11.2). With aggregate risk, each date’s price is a function of the history of aggregate shocks, given the initial distribution. Lucas and Prescott (1971, p. 664) assumed that firms know such functions (section 11.1), and sequence-space methods compute them to first order (Auclert et al., 2021). Model builders also compute such equilibria by higher-order perturbation (Bilal, 2023) and globally, with deep neural networks that approximate solutions of master equations (Gu et al., 2026; Payne et al., 2026). Whether agents can learn the predictive distributions of prices that the equilibrium implies is a separate question, the one that Sims’s histograms answer in stationary environments (section 11.3).

Section 3.5 of Moll (2026) implies a limit to Sims’s argument. Because equilibrium prices in these models are not Markov (their conditional distribution depends on the entire cross-section distribution), the histogram that would reveal θ must be taken over the infinite-dimensional state, or over long histories of prices. A law of large numbers still applies, but it fills such histograms very slowly. In the sticky-price economy of section 11.2, the same property determines how much firms need to know: their own market signals suffice when their marginal costs are first-order Markov.

Lucas’s student and coauthor Moll (2026) is not convinced by the “as if” metaphor of section 11.1. He contends that heterogeneous agent models that impose rational expectations attribute implausible forecasting behavior to the people inside them:

... 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? ... Instead of solving Monster equations we should replace the rational expectations assumption and solve the simpler equations corresponding to their actual price-forecasting behaviour. (Moll, 2026, p. 1174)

Section 14.3 describes what alternatives to rational expectations about such prices would keep and lose.

13 The roles collected

The roles that sections 5–10 described are:

1. *Discipline on beliefs.* The restriction eliminates free parameters that would otherwise describe agents' subjective models, all of them when the equilibrium is unique (sections 5.2 and 5.4).
2. *Identification and estimation.* Cross-equation restrictions link structural parameters to the parameters of descriptive models and can make that link invertible. Estimators attribute an estimated objective model to agents as their subjective model (sections 5.1, 6, 7.2, and 7.3).
3. *Policy evaluation.* Agents' beliefs respond to policy rules, so a model builder can compute how outcomes respond to policies (sections 7.1, 7.3, and 8.1).
4. *Policy design and credibility.* Ramsey plans, time consistency, and reputations become problems that dynamic programming can solve (sections 8.2 and 8.3).
5. *Inference from equilibrium outcomes.* Agents know how prices and other agents' actions depend on private information, so they can extract information from them. That underlies theories of information aggregation, no-trade theorems, and higher-order beliefs (sections 5.3, 7.2, 7.3, and 12.3).
6. *Welfare comparisons.* Agents' subjective expected utilities equal objective ones, which settles whose probabilities welfare comparisons should use (section 7.1).
7. *Determinacy.* Whether a policy rule delivers a unique equilibrium becomes a criterion for choosing it, and dispersed private information can select a unique equilibrium where common knowledge allows several (sections 7.1, 7.3, 8.1, and 12.3).

8. *Benchmark.* Departures from rational expectations, including “behavioral biases,” are measured from it (sections 6.2, 7.1, 10.2, and 14).
9. *Computation.* Equilibria are fixed points, and relaxation and stochastic approximation algorithms can find the stable ones. Perturbation and deep learning methods compute them in models with large state spaces (sections 5.5, 7.3, and 12.4).

Each role also uses ingredients that the restriction does not supply. Identification requires rank conditions. Counterfactuals require primitives that stay fixed across policies and an account of what agents know about a new policy. Welfare comparisons require a criterion. Every role requires an equilibrium that exists and a rule for selecting one when there are several (section 12.3).

The nine roles are also criteria.²⁴ A proposed alternative equilibrium concept can be judged by how many of them it keeps.

Part IV. Alternatives and an accounting

14 Alternatives anchored to the restriction

There is a wilderness of models that do not impose rational expectations.²⁵ Most proposed replacements for rational expectations have long histories. Adaptive expectations, a forecasting heuristic, came first (section 5.4). Modern heuristics, such as the myopia and limited depth of reasoning proposed to resolve the forward guidance puzzle (section 7.3), came later. Temporary equilibrium theory (Grandmont, 1977) takes agents’ forecasting rules as given. In the internal rationality of Adam and Marcet (2011), agents optimize under subjective beliefs about prices. Least squares learning (Bray, 1982; Marcet and Sargent, 1989), reinforcement learning by artificially intelligent agents (Marimon et al., 1990; Sargent, 1993), restricted perceptions and self-confirming equilibria (Fudenberg and Levine, 1993; Sargent, 1999; Evans and Honkapohja, 2001), and doubts about models (Hansen and Sargent, 2008) have been studied for decades. Those literatures have already worked out much of what each departure keeps and loses.

Departures come in two classes. Those in the first share a template. An *outcome map* T sends a list of agents’ subjective models into the objective model that their decisions generate.

²⁴Fernando Alvarez regrets that there are only nine. Ten could have been commandments.

²⁵Sims (1980, p. 4) used the word for a different set of models, those in which prices adjust sluggishly: “If instead we take the view that prices themselves may adjust sluggishly, we enter the wilderness of ‘disequilibrium economics.’ This phrase must, it seems to me, denote a situation in which we cannot suppose that business behavior is invariant under changes in the public’s tastes.”

Table 1: Belief maps and equilibrium concepts

Concept	Agent i 's family	Belief map $B_i(\theta)$	Equilibrium
Rational expectations	$\{f(\cdot \mid \theta)\}_{\theta \in \Theta}$	θ	$\theta = \mathsf{T}(\theta, \dots, \theta)$
Robustness	approximating $f(\cdot \mid \theta)$ or structured models; relative entropy neighborhood	worst-case twist $S_i(\theta)$	$\theta = \mathsf{T}(S_1(\theta), \dots, S_I(\theta))$
Restricted perceptions	$\{g_\delta\}_{\delta \in \Delta_i}$, possibly misspecified	projection $\Pi_i(\theta)$, (25)	$\theta = \mathsf{T}(\Pi_1(\theta), \dots, \Pi_I(\theta))$
Self-confirming	$\{f_i(\cdot \mid \theta_i)\}_{\theta_i \in \Theta_i}$	set $\mathcal{C}_i(\theta)$: correct on path	$\theta \in \mathsf{T}(\mathcal{C}_1(\theta), \dots, \mathcal{C}_I(\theta))$
Adaptive learning	$\{g_\delta\}_{\delta \in \Delta_i}$	recursion (15)	$\theta_t = \mathsf{T}_t(\hat{\delta}_{1t}, \dots, \hat{\delta}_{It})$; no fixed point

A *belief map* B_i sends the objective model into agent i 's subjective model. An equilibrium concept is a choice of subjective families $\mathcal{F}_i$ together with belief maps B_i , and an equilibrium is a fixed point

$$\theta \in \mathsf{T}_{\mathcal{F}}(B_1(\theta), \dots, B_I(\theta)). \quad (23)$$

The outcome map encodes the same preferences, technologies, information, and market clearing under every concept; only the families $\mathcal{F}_i$ from which its arguments are drawn differ. The belief maps distinguish the concepts, as Table 1 shows. Departures in the second class, such as adaptive learning along a transition path (the last row of the table), have no fixed point. Agents' beliefs are state variables, and along a transition they are inconsistent with the outcomes that they generate. Rational expectations is a special case of each of rows 2–4 of Table 1. Robustness (section 10) reduces to it as $\xi_i \rightarrow \infty$ when there is no ambiguity, that is, when the only structured model is the approximating model (section 10.2). A restricted perceptions equilibrium is self-confirming when each family $\{g_\delta\}_{\delta \in \Delta_i}$ contains f_θ , and it is a rational expectations equilibrium when, in addition, every member of the family that agrees with f_θ on the equilibrium path agrees with it at every state, because then $g_{\Pi_i(\theta)} = f_\theta$ everywhere. Every rational expectations equilibrium is self-confirming when each agent's family contains the objective family, because then $\theta \in \mathcal{C}_i(\theta)$ for every i .

Measured state by state by KL divergences between one-step-ahead conditional densities, rational expectations sets them to zero in every state, including states reached only after deviations; a self-confirming equilibrium sets them to zero in states visited on the equilibrium path; and a restricted perceptions equilibrium minimizes (3) over a possibly misspecified family. Robustness instead penalizes the divergence of a worst-case model relative to the

approximating model, the reverse order. Hansen et al. (2026) show how much freedom the choice of a reference model leaves and what bounds it.²⁶

14.1 Restricted perceptions and self-confirming equilibria

In a restricted perceptions equilibrium, agent i 's family of subjective models $\{g_\delta\}_{\delta \in \Delta_i}$ need not contain the objective model f_θ .

Definition 14.1. *A restricted perceptions equilibrium is a $(\theta, \delta_1, \dots, \delta_I)$ that satisfies*

$$\theta = \mathsf{T}(\delta_1, \dots, \delta_I), \tag{24}$$

$$\delta_i = \arg \min_{\delta \in \Delta_i} KL(f_\theta, g_\delta), \quad i = 1, \dots, I. \tag{25}$$

The minimizer in (25) is a population maximum likelihood estimator (section 6), the information projection of f_θ onto agent i 's family (Sargent, 2025). With that projection written as $\Pi_i(\theta)$, a restricted perceptions equilibrium is a fixed point $\theta = \mathsf{T}(\Pi_1(\theta), \dots, \Pi_I(\theta))$ of T composed with information projections. When agent i 's family consists of linear regressions with Gaussian disturbances, $\Pi_i(\theta)$ is a population least squares regression. The Krusell–Smith algorithm of section 7.2 computes such a regression, and its households treat the fitted rule as exact. The Berk–Nash equilibrium of Esponda and Pouzo (2016) generalizes this concept to games in which agents entertain sets of models. In a self-confirming equilibrium (Fudenberg and Levine, 1993, 2009), agent i 's parameter need only lie in the set $\mathcal{C}_i(\theta)$ of parameters θ_i for which $f_i(\cdot \mid \theta_i)$ and f_θ assign the same probabilities to events on the equilibrium path. Self-confirming equilibria typically come in sets. An agent's model can misrepresent outcomes after historically unprecedented policies, so a government that uses such a model can adopt inferior policies.

14.2 Adaptive learning

Least squares learning is the subject of a literature of the 1980s and 1990s that includes Marcet and Sargent (1989), Sargent (1993), and Sargent (1999). That literature asks whether learning converges to a rational expectations equilibrium. Game theorists ask how players might come to play a Nash equilibrium (Fudenberg and Levine, 1998). Lectures in Sargent

²⁶In the Bewley economy of section 10.1, quantity data identify neither a consumer's type nor the reference model from which a bias is measured. Measured from a trusting benchmark, robust consumers look pessimistic and impatient; measured from a point near the boundary of admissible robustness, they look optimistic and patient. Four things limit that freedom: bounds on admissible discount factors, detection error probabilities, variation in the environment, and an Euler-equation restriction that elicited forecasts and discount factors can test.

and Stachurski (2026b), beginning with the section “Bounded Rationality in Macroeconomics” at https://python.quantecon.org/bounded_rationality.html, describe that literature and implement models with escape dynamics.

Models that replace rational expectations with forecasting rules describe a temporary equilibrium

$$\theta_t = \mathsf{T}_t(\hat{\delta}_{1t}, \dots, \hat{\delta}_{It}), \quad (26)$$

in which agents’ beliefs $\hat{\delta}_{it}$ are taken as given and T_t maps them into the objective law of motion of outcomes at t . In a learning model, agent i updates $\hat{\delta}_{it}$ with a recursion like (15), a stochastic counterpart of the relaxation iterations of section 5.5. Marcet and Sargent (1989) showed that the limiting behavior of beliefs is governed by the mean dynamics

$$\frac{d\delta}{d\tau} = \Pi(\mathsf{T}(\delta)) - \delta, \quad (27)$$

where Π projects the actual law of motion onto agents’ family of perceived laws of motion, as in (25). Rest points of (27) are restricted perceptions equilibria.²⁷ With a decreasing gain, least squares learning can converge to a rational expectations equilibrium when agents’ families contain the actual law of motion, equilibrium-path data single it out, and the equilibrium is stable under (27), a property that Evans and Honkapohja (2001) call E-stability. Results of that kind rely on what I call *prompting* by the model builder: a correctly specified family, suitable initial beliefs, and devices that keep beliefs in a domain of attraction, such as the projection facility of Marcet and Sargent (1989). With a constant gain, beliefs keep fluctuating, and in some models escapes from restricted perceptions equilibria recur (Cho et al., 2002).²⁸

Learning models add free parameters: families, gains, and initial beliefs (Sargent, 1993, ch. 7, p. 168). The family determines the possible limits, which are restricted perceptions equilibria. With a decreasing gain and a unique limit, gains and initial beliefs affect only transient dynamics, so the limiting distribution of the data contains no information about them. When there are several possible limits, initial beliefs affect which one is reached. With a constant gain and an ergodic belief process, initial beliefs do not affect the stationary distribution. The gain governs both the size of belief fluctuations around a restricted perceptions equilibrium and the frequency of escapes from it. The fluctuations contain

²⁷Sargent (1999) and Cho et al. (2002) call them self-confirming equilibria, a usage that differs from that of section 14.1.

²⁸Applications include escapes from Nash inflation (Cho et al., 2002), recurrent currency crises (Kasa, 2004; Cho and Kasa, 2005), illusory algorithmic collusion (Cho and Williams, 2024), recurrent Latin American stabilizations (Marcet and Nicolini, 2003; Sargent et al., 2009), and U.S. inflation-unemployment dynamics (Sims, 1988; Chung, 1990; Sargent, 1999).

information about the gain, but escapes are rare, so for them the rare-events problem of section 12.1 reappears. By contrast, a model of Bayesian learning in which nature draws parameters from the agents' prior is a rational expectations model with priors as additional parameters (section 6.2; Bray and Kreps, 1987).

In 1993 I described the adaptive learning program this way:

Bounded rationality is a movement to make model agents behave more like econometricians. Despite the compliment thereby paid to their kind, macroeconometricians have shown very little interest in applying models of bounded rationality to data. Within the economics profession, the impulse to build models populated by econometricians has come primarily from theorists with different things on their minds from most econometricians. (Sargent, 1993, ch. 7, pp. 167–68)

Exceptions include Chung (1990), who estimated such a model. Moll (2026) renews the theorists' impulse, with a difference: he asks that models of beliefs be disciplined by survey evidence on expectations (section 12.2), data that are far richer now than in 1993.

14.3 What alternatives keep and lose

Abandoning rational expectations everywhere would put at risk the equilibrium concepts of game theory and mechanism design and the structural estimators of industrial organization and labor economics (section 7.2). Moll (2026) proposes something narrower: not imposing the restriction on equilibrium prices in heterogeneous agent models with aggregate risk, while keeping it for exogenous processes and keeping optimization and market clearing (Moll, 2026, p. 1189, n. 22). He proposes three criteria for a replacement: computational tractability, consistency with empirical evidence on expectations, and endogeneity of beliefs to “model reality,” including policy changes. The third criterion keeps two of the roles collected in section 13: discipline on beliefs, because in stationary environments price beliefs should be approximately consistent with equilibrium prices; and policy evaluation, because price beliefs should respond to policy changes (Moll, 2026, pp. 1192–93). The alternative he prefers is close to the Krusell and Smith (1998) algorithm of section 7.2, whose fixed point is a restricted perceptions equilibrium. Moll would treat that fixed point as a description of how people forecast and would discipline it with survey evidence (Moll, 2026, pp. 1177 and 1188).

Where such proposals apply, Table 2 compares what three alternatives keep and lose of the roles collected in section 13. The phrase “discipline on beliefs” means different things in each column of Table 2, as its first row shows.

Discipline on beliefs, identification, and welfare comparisons are at risk under both restricted perceptions and adaptive learning. Adaptive learning also puts policy evaluation

Table 2: What three alternatives keep and lose

Role	Price beliefs as restricted perceptions	Robustness	Adaptive learning
Discipline on beliefs	parameters pinned down by projections; families free	partly kept: worst-case twists of a shared baseline model, with two parameters per agent bounded by detection-error probabilities	lost: families, gains, initial beliefs
Identification and estimation	families must be specified	baseline model as with rational expectations; quantity data alone may not separate doubts from discount factors (section 10.1)	identification of gains and initial beliefs needs its own analysis
Policy evaluation	beliefs respond through the projection	beliefs respond through the worst-case twist	beliefs respond only gradually
Policy design and credibility	requires modeling how agents read policy	kept	possible, but plans hinge on gains and initial beliefs
Inference from equilibrium outcomes	inference through a best-fitting, possibly misspecified price function	kept, with worst-case inference	agents must first learn how prices reflect information (Bray, 1982)
Welfare comparisons	whose beliefs count?	agents' own robust criteria; welfare evaluations differ across types that quantity data cannot distinguish	whose beliefs count?
Determinacy	a fixed point remains, so the criterion applies	a fixed point remains, so the criterion applies	replaced by stability under learning
Benchmark	kept: a special case when families contain f_θ	kept: the limit as both kinds of doubt vanish	kept: a possible limit of learning
Computation	a fixed point remains	kept: a fixed point remains, at least as hard to compute as under rational expectations	no fixed point in beliefs; each period's temporary equilibrium must still be computed

and policy design at risk: optimal policy under learning can be computed, but the plans depend on gains and initial beliefs (Gaspar et al., 2006; Molnár and Santoro, 2014). Robustness, which retains a communism of approximating models, keeps most roles, although quantity data alone may not separate doubts from impatience. Departures also change a model's predictions. In the search model of Payne et al. (2026), unemployment falls about 25 percent faster during the recovery from a shock calibrated to COVID-19 when workers and firms forecast the distribution of matches than when they act as if it stays at its long-run average (section 7.2).

Two other responses keep the restriction, fully or partially. One starts from forecasting rules that survey evidence suggests people use and, as Muth (1960) did for adaptive expectations, reverse engineers environments in which those rules are objective conditional expectations given the forecasters' information. That is harder in heterogeneous agent models with aggregate risk than in Muth's univariate example, and impossible for rules whose errors are predictable from the forecasters' own revisions (section 12.2). The other imposes the restriction only partially, as generalized method of moments estimation of Euler equations does (section 6.3).

15 Concluding remarks

Rational expectations is a property of a model's equilibrium, not of the people inside it. A rational expectations equilibrium is a fixed point $\theta = T(\theta, \dots, \theta)$ of a mapping from subjective to objective statistical models. That communism of models plays the nine roles collected in section 13. It appears, often unnamed, throughout economics: as equilibrium price expectations, Euler-equation orthogonality conditions, correct beliefs about strategies, common priors from which nature draws types, and known transition laws in structural econometrics. In current macroeconomics, it makes credible announcements of future policies affect current outcomes, the basis of forward guidance; its common prior lets global games select among equilibria in models of runs; and it underlies the computational methods that make heterogeneous agent models tractable. When models that impose it fail, economists usually modify preferences, technologies, and shock processes, not the restriction.

The restriction is hardest to defend where stationarity and a law of large numbers give it the least support: for rare events and off-path events. Agents need to forecast only the prices and shocks that enter their decision problems. Computing the equilibrium is the model builder's task. The fixed point can also have several solutions or none, even in well-behaved economies (section 12.3). Disciplined alternatives remain anchored to the restriction. Equilibria in which agents share but distrust an approximating model keep a communism of

approximating models, and they add ambiguity and misspecification components to prices measured relative to that model. The roles are also criteria for judging proposed alternative equilibrium concepts. The accounting of section 14.3 shows that abandoning the restriction would cost most in discipline on beliefs, identification, and welfare comparisons, and that a leading proposed alternative would treat the fixed point of the Krusell–Smith algorithm as a description of how people forecast. I hope that collecting the roles and the accounting in one place helps students and researchers decide when to impose the restriction and when to depart from it. Rational expectations remains the benchmark from which departures are measured.

References

- Achdou, Yves, Jiequn Han, Jean-Michel Lasry, Pierre-Louis Lions, and Benjamin Moll. 2022. Income and Wealth Distribution in Macroeconomics: A Continuous-Time Approach. *Review of Economic Studies* 89 (1):45–86.
- Adam, Klaus and Albert Marcet. 2011. Internal rationality, imperfect market knowledge and asset prices. *Journal of Economic Theory* 146 (3):1224–1252.
- Aiyagari, S. Rao. 1994. Uninsured Idiosyncratic Risk and Aggregate Saving. *Quarterly Journal of Economics* 109 (3):659–684.
- Alchian, Armen A. 1950. Uncertainty, Evolution, and Economic Theory. *Journal of Political Economy* 58 (3):211–221.
- Anderson, Evan W., Lars Peter Hansen, and Thomas J. Sargent. 2003. A Quartet of Semigroups for Model Specification, Robustness, Prices of Risk, and Model Detection. *Journal of the European Economic Association* 1 (1):68–123.
- Angeletos, George-Marios and Iván Werning. 2006. Crises and Prices: Information Aggregation, Multiplicity, and Volatility. *American Economic Review* 96 (5):1720–1736.
- Arellano, Cristina. 2008. Default Risk and Income Fluctuations in Emerging Economies. *American Economic Review* 98 (3):690–712.
- Arrow, Kenneth J. 1964. The Role of Securities in the Optimal Allocation of Risk-bearing. *Review of Economic Studies* 31 (2):91–96. Read at the Colloque sur les Fondements et Applications de la Théorie du Risque en Économétrie, Centre National de la Recherche Scientifique, Paris, May 13, 1952, and published in French in the proceedings of that colloquium, *Économétrie*, 1953.

- Auclert, Adrien, Bence Bardóczy, Matthew Rognlie, and Ludwig Straub. 2021. Using the Sequence-Space Jacobian to Solve and Estimate Heterogeneous-Agent Models. *Econometrica* 89 (5):2375–2408.
- Aumann, Robert J. 1987. Correlated Equilibrium as an Expression of Bayesian Rationality. *Econometrica* 55 (1):1–18.
- Bajari, Patrick, C. Lanier Benkard, and Jonathan Levin. 2007. Estimating Dynamic Models of Imperfect Competition. *Econometrica* 75 (5):1331–1370.
- Balasko, Yves and David Cass. 1989. The Structure of Financial Equilibrium with Exogenous Yields: The Case of Incomplete Markets. *Econometrica* 57 (1):135–162.
- Bansal, Ravi and Amir Yaron. 2004. Risks for the Long Run: A Potential Resolution of Asset Pricing Puzzles. *Journal of Finance* 59 (4):1481–1509.
- Barberis, Nicholas, Ming Huang, and Tano Santos. 2001. Prospect Theory and Asset Prices. *The Quarterly Journal of Economics* 116 (1):1–53.
- Barillas, Francisco, Lars Peter Hansen, and Thomas J. Sargent. 2009. Doubts or Variability? *Journal of Economic Theory* 144 (6):2388–2418.
- Barro, Robert J. 2006. Rare Disasters and Asset Markets in the Twentieth Century. *Quarterly Journal of Economics* 121 (3):823–866.
- Berk, Robert H. 1966. Limiting Behavior of Posterior Distributions When the Model Is Incorrect. *Annals of Mathematical Statistics* 37 (1):51–58.
- Bhandari, Anmol. 2015. Reverse Engineering Heterogeneous Beliefs. Unpublished manuscript, University of Minnesota.
- Bilal, Adrien. 2023. Solving Heterogeneous Agent Models with the Master Equation. National Bureau of Economic Research working paper.
- Blanchard, Olivier Jean and Charles M. Kahn. 1980. The Solution of Linear Difference Models under Rational Expectations. *Econometrica* 48 (5):1305–1311.
- Blume, Lawrence and David Easley. 2006. If You’re So Smart, Why Aren’t You Rich? Belief Selection in Complete and Incomplete Markets. *Econometrica* 74 (4):929–966.
- 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* 89:68–92.

- Bordalo, Pedro, Nicola Gennaioli, Yueran Ma, and Andrei Shleifer. 2020. Overreaction in Macroeconomic Expectations. *American Economic Review* 110 (9):2748–2782.
- Bray, Margaret. 1982. Learning, Estimation, and the Stability of Rational Expectations. *Journal of Economic Theory* 26 (2):318–339.
- Bray, Margaret and David M Kreps. 1987. Rational Learning and Rational Expectations. In *Arrow and the Ascent of Modern Economic Theory*, 597–625. Springer.
- Brayton, Flint, Andrew Levin, Ralph Tryon, and John C. Williams. 1997. The Evolution of Macro Models at the Federal Reserve Board. *Carnegie-Rochester Conference Series on Public Policy* 47:43–81.
- Brunnermeier, Markus K., Alp Simsek, and Wei Xiong. 2014. A Welfare Criterion for Models with Distorted Beliefs. *Quarterly Journal of Economics* 129 (4):1753–1797.
- Cagan, Philip. 1956. The Monetary Dynamics of Hyperinflation. In *Studies in the Quantity Theory of Money*, edited by Milton Friedman, 25–117. Chicago: University of Chicago Press.
- Calvo, Guillermo A. 1978. On the Time Consistency of Optimal Policy in a Monetary Economy. *Econometrica* 46 (6):1411–1428.
- Campbell, John Y. and Robert J. Shiller. 1991. Yield Spreads and Interest Rate Movements: A Bird’s Eye View. *Review of Economic Studies* 58 (3):495–514.
- Cardaliaguet, Pierre, François Delarue, Jean-Michel Lasry, and Pierre-Louis Lions. 2019. *The Master Equation and the Convergence Problem in Mean Field Games*. No. 201 in Annals of Mathematics Studies. Princeton: Princeton University Press.
- Carlsson, Hans and Eric van Damme. 1993. Global Games and Equilibrium Selection. *Econometrica* 61 (5):989–1018.
- Chari, V. V., Juan Pablo Nicolini, and Pedro Teles. 2020. Optimal Capital Taxation Revisited. *Journal of Monetary Economics* 116:147–165.
- Chari, Varadarajan V and Patrick J Kehoe. 1990. Sustainable Plans. *Journal of Political Economy* 98 (4):783–802.
- Chen, Qi, Itay Goldstein, and Wei Jiang. 2010. Payoff Complementarities and Financial Fragility: Evidence from Mutual Fund Outflows. *Journal of Financial Economics* 97 (2):239–262.

- Cho, In-Koo and Kenneth Kasa. 2005. Model Validation Dynamics and Endogenous Currency Crises. Unpublished.
- Cho, In-Koo and Noah Williams. 2024. Collusive Outcomes Without Collusion. ArXiv:2403.07177.
- Cho, In-Koo, Noah Williams, and Thomas J. Sargent. 2002. Escaping Nash Inflation. *Review of Economic Studies* 69:1–40.
- Chung, Heetaik. 1990. Did Policy Makers Really Believe in the Phillips Curve. PhD thesis, University of Minnesota.
- Cochrane, John H. 2023. *The Fiscal Theory of the Price Level*. Princeton, NJ: Princeton University Press.
- Cogley, Timothy and Thomas J Sargent. 2008. The Market Price of Risk and the Equity Premium: A Legacy of the Great Depression? *Journal of Monetary Economics* 55 (3):454–476.
- Coibion, Olivier and Yuriy Gorodnichenko. 2015. Information Rigidity and the Expectations Formation Process: A Simple Framework and New Facts. *American Economic Review* 105 (8):2644–2678.
- Cole, Harold L. and Timothy J. Kehoe. 2000. Self-Fulfilling Debt Crises. *Review of Economic Studies* 67 (1):91–116.
- Cole, Harold L. and Edward C. Prescott. 1997. Valuation Equilibrium with Clubs. *Journal of Economic Theory* 74 (1):19–39.
- Constantinides, George M. and Darrell Duffie. 1996. Asset Pricing with Heterogeneous Consumers. *Journal of Political Economy* 104 (2):219–240.
- Csiszár, Imre and František Matúš. 2003. Information Projections Revisited. *IEEE Transactions on Information Theory* 49 (6):1474–1490.
- Del Negro, Marco, Marc P. Giannoni, and Christina Patterson. 2023. The Forward Guidance Puzzle. *Journal of Political Economy Macroeconomics* 1 (1):43–79.
- Diamond, Douglas W. and Philip H. Dybvig. 1983. Bank Runs, Deposit Insurance, and Liquidity. *Journal of Political Economy* 91 (3):401–419.
- Dornbusch, Rudiger. 1976. Expectations and Exchange Rate Dynamics. *Journal of Political Economy* 84 (6):1161–1176.

- Duffie, Darrell and Wayne Shafer. 1985. Equilibrium in Incomplete Markets: I. A Basic Model of Generic Existence. *Journal of Mathematical Economics* 14 (3):285–300.
- Eaton, Jonathan and Mark Gersovitz. 1981. Debt with Potential Repudiation: Theoretical and Empirical Analysis. *Review of Economic Studies* 48 (2):289–309.
- Ericson, Richard and Ariel Pakes. 1995. Markov-Perfect Industry Dynamics: A Framework for Empirical Work. *Review of Economic Studies* 62 (1):53–82.
- Esponda, Ignacio and Demian Pouzo. 2016. Berk–Nash Equilibrium: A Framework for Modeling Agents with Misspecified Models. *Econometrica* 84 (3):1093–1130.
- Evans, George W and Seppo Honkapohja. 2001. *Learning and Expectations in Macroeconomics*. Princeton University Press.
- Fama, Eugene F. 1970. Efficient Capital Markets: A Review of Theory and Empirical Work. *Journal of Finance* 25 (2):383–417.
- Farhi, Emmanuel and Iván Werning. 2019. Monetary Policy, Bounded Rationality, and Incomplete Markets. *American Economic Review* 109 (11):3887–3928.
- Ferriere, Axelle and Anastasios G Karantounias. 2019. Fiscal Austerity in Ambiguous Times. *American Economic Journal: Macroeconomics* 11 (1):89–131.
- Friedman, Milton. 1953. The Methodology of Positive Economics. In *Essays in Positive Economics*, 3–43. Chicago: University of Chicago Press.
- . 1957. *A Theory of the Consumption Function*. Princeton: Princeton University Press.
- Friedman, Milton and L. J. Savage. 1948. The Utility Analysis of Choices Involving Risk. *Journal of Political Economy* 56 (4):279–304.
- Fudenberg, Drew and David K. Levine. 1993. Self-Confirming Equilibrium. *Econometrica* 61 (3):523–545.
- . 1998. *The Theory of Learning in Games*. Cambridge, MA: MIT Press.
- . 2009. Self-Confirming Equilibrium and the Lucas Critique. *Journal of Economic Theory* 144 (6):2354–2371.
- Gabaix, Xavier. 2020. A Behavioral New Keynesian Model. *American Economic Review* 110 (8):2271–2327.

- Gallant, A Ronald and George Tauchen. 1996. Which Moments to Match? *Econometric Theory* 12 (4):657–681.
- Gaspar, Vítor, Frank Smets, and David Vestin. 2006. Adaptive Learning, Persistence, and Optimal Monetary Policy. *Journal of the European Economic Association* 4 (2–3):376–385.
- Geanakoplos, John and Andreu Mas-Colell. 1989. Real Indeterminacy with Financial Assets. *Journal of Economic Theory* 47 (1):22–38.
- Goldstein, Itay and Ady Pauzner. 2005. Demand–Deposit Contracts and the Probability of Bank Runs. *Journal of Finance* 60 (3):1293–1327.
- Gouriéroux, Christian, Alain Monfort, and Eric Renault. 1993. Indirect Inference. *Journal of Applied Econometrics* 8 (S1):S85–S118.
- Grandmont, Jean-Michel. 1977. Temporary General Equilibrium Theory. *Econometrica* 45 (3):535–572.
- Grossman, Sanford J. and Joseph E. Stiglitz. 1980. On the Impossibility of Informationally Efficient Markets. *American Economic Review* 70 (3):393–408.
- Grunberg, Emile and Franco Modigliani. 1954. The Predictability of Social Events. *Journal of Political Economy* 62 (6):465–478.
- Gu, Zhouzhou, Mathieu Laurière, Sebastian Merkel, and Jonathan Payne. 2026. Global Solutions to Master Equations for Continuous Time Heterogeneous Agent Macroeconomic Models. Forthcoming in *Journal of Political Economy Macroeconomics*.
- Gul, Faruk and Wolfgang Pesendorfer. 2001. Temptation and Self-Control. *Econometrica* 69 (6):1403–1435.
- Hall, Robert E. 1978. Stochastic Implications of the Life Cycle-Permanent Income Hypothesis: Theory and Evidence. *Journal of Political Economy* 86 (6):971–987.
- Hansen, Lars Peter. 1982. Large Sample Properties of Generalized Method of Moments Estimators. *Econometrica* 50 (4):1029–1054.
- . 2014. Nobel Lecture: Uncertainty Outside and Inside Economic Models. *Journal of Political Economy* 122 (5):945–987.
- Hansen, Lars Peter and Thomas J. Sargent. 1980. Formulating and Estimating Dynamic Linear Rational Expectations Models. *Journal of Economic Dynamics and Control* 2:7–46.

- . 1993. Seasonality and Approximation Errors in Rational Expectations Models. *Journal of Econometrics* 55 (1-2):21–55.
- . 2001. Robust Control and Model Uncertainty. *American Economic Review* 91 (2):60–66.
- Hansen, Lars Peter and Thomas J Sargent. 2008. *Robustness*. Princeton: Princeton University Press.
- Hansen, Lars Peter and Thomas J. Sargent. 2010. Fragile Beliefs and the Price of Uncertainty. *Quantitative Economics* 1 (1):129–162.
- . 2021. Macroeconomic Uncertainty Prices When Beliefs Are Tenuous. *Journal of Econometrics* 223 (1):222–250.
- . 2022. Structured Ambiguity and Model Misspecification. *Journal of Economic Theory* 199:105165.
- . 2024. Risk, Ambiguity, and Misspecification: Decision Theory, Robust Control, and Statistics. *Journal of Applied Econometrics* 39 (6):969–999.
- Hansen, Lars Peter and José A. Scheinkman. 2009. Long-Term Risk: An Operator Approach. *Econometrica* 77 (1):177–234.
- Hansen, Lars Peter and Kenneth J Singleton. 1982. Generalized Instrumental Variables Estimation of Nonlinear Rational Expectations Models. *Econometrica* 50 (5):1269–1286.
- . 1991. Computing Semi-Parametric Efficiency Bounds for Linear Time Series Models. In *Nonparametric and Semiparametric Methods in Econometrics and Statistics*, edited by William A. Barnett, James Powell, and George Tauchen, 387–412. Cambridge: Cambridge University Press.
- Hansen, Lars Peter, Thomas J. Sargent, and Thomas D. Tallarini. 1999. Robust Permanent Income and Pricing. *Review of Economic Studies* 66 (4):873–907.
- Hansen, Lars Peter, John C Heaton, and Nan Li. 2008. Consumption Strikes Back? Measuring Long-Run Risk. *Journal of Political Economy* 116 (2):260–302.
- Hansen, Lars Peter, Francesco Lippi, and Thomas J. Sargent. 2026. Behavioral Bias, Relative to What? Robustness and Discounting in a Bewley Economy. Working paper.
- Harrison, J Michael and David M Kreps. 1978. Speculative Investor Behavior in a Stock Market with Heterogeneous Expectations. *The Quarterly Journal of Economics* 92 (2):323–336.

- Harsanyi, John C. 1967. Games with Incomplete Information Played by “Bayesian” Players, I–III: Part I. The Basic Model. *Management Science* 14 (3):159–182.
- Hart, Oliver D. 1975. On the Optimality of Equilibrium when the Market Structure is Incomplete. *Journal of Economic Theory* 11 (3):418–443.
- Hayek, F. A. 1945. The Use of Knowledge in Society. *American Economic Review* 35 (4):519–530.
- Heinemann, Frank, Rosemarie Nagel, and Peter Ockenfels. 2004. The Theory of Global Games on Test: Experimental Analysis of Coordination Games with Public and Private Information. *Econometrica* 72 (5):1583–1599.
- Hellwig, Christian. 2002. Public Information, Private Information, and the Multiplicity of Equilibria in Coordination Games. *Journal of Economic Theory* 107 (2):191–222.
- Hellwig, Christian and Venky Venkateswaran. 2025. Dispersed Information, Nominal Rigidities and Monetary Business Cycles: A Hayekian Perspective. *Journal of Monetary Economics* 155:103843.
- Hellwig, Christian, Arijit Mukherji, and Aleh Tsyvinski. 2006. Self-Fulfilling Currency Crises: The Role of Interest Rates. *American Economic Review* 96 (5):1769–1787.
- Herbst, Edward P. and Frank Schorfheide. 2016. *Bayesian Estimation of DSGE Models*. Princeton, NJ: Princeton University Press.
- Hood, William C and Tjalling C Koopmans, eds. 1953. *Studies in Econometric Method*. New York: Wiley.
- Hotz, V. Joseph and Robert A. Miller. 1993. Conditional Choice Probabilities and the Estimation of Dynamic Models. *Review of Economic Studies* 60 (3):497–529.
- Huo, Zhen and Marcelo Pedroni. 2020. A Single-Judge Solution to Beauty Contests. *American Economic Review* 110 (2):526–568.
- Huo, Zhen and Naoki Takayama. 2025. Rational Expectations Models with Higher-Order Beliefs. *Review of Economic Studies* 92 (5):3138–3173.
- İmrohoroğlu, Selahattin. 1993. Testing for Sunspot Equilibria in the German Hyperinflation. *Journal of Economic Dynamics and Control* 17 (1–2):289–317.
- Izmalkov, Sergei and Muhamet Yildiz. 2010. Investor Sentiments. *American Economic Journal: Microeconomics* 2 (1):21–38.

- Jacobson, David H. 1973. Optimal Stochastic Linear Systems with Exponential Performance Criteria and Their Relation to Deterministic Differential Games. *IEEE Transactions on Automatic Control* 18 (2):124–131.
- Jamilov, Rustam, Tobias König, Karsten Müller, and Farzad Saidi. 2026. Two Centuries of Systemic Bank Runs. CEPR Discussion Paper 19382, revised July 2026.
- Kaplan, Greg, Benjamin Moll, and Giovanni L. Violante. 2018. Monetary Policy According to HANK. *American Economic Review* 108 (3):697–743.
- Karantounias, Anastasios G. 2013. Managing Pessimistic Expectations and Fiscal Policy. *Theoretical Economics* 8 (1):193–231.
- Kasa, Kenneth. 2000. Forecasting the Forecasts of Others in the Frequency Domain. *Review of Economic Dynamics* 3 (4):726–756.
- . 2004. Learning, Large Deviations, and Recurrent Currency Crises. *International Economic Review* 45:141–173.
- Keane, Michael P. and Kenneth I. Wolpin. 1997. The Career Decisions of Young Men. *Journal of Political Economy* 105 (3):473–522.
- Kihlstrom, Richard E. 2019. The Role of Securities in the Allocation of Risk: Arrow’s 1964 REStud Paper. Class notes, Wharton School, University of Pennsylvania.
- Koopmans, Tjalling C. 1947. Measurement Without Theory. *The Review of Economics and Statistics* 29 (3):161–172.
- Koopmans, Tjalling C, ed. 1950. *Statistical Inference in Dynamic Economic Models*. New York: Wiley.
- Kreps, David M. and Evan L. Porteus. 1978. Temporal Resolution of Uncertainty and Dynamic Choice Theory. *Econometrica* 46 (1):185–200.
- Krusell, Per and Anthony A. Smith, Jr. 1998. Income and Wealth Heterogeneity in the Macroeconomy. *Journal of Political Economy* 106 (5):867–896.
- Kydland, Finn E. and Edward C. Prescott. 1977. Rules Rather than Discretion: The Inconsistency of Optimal Plans. *Journal of Political Economy* 85 (3):473–491.
- . 1980. Dynamic Optimal Taxation, Rational Expectations and Optimal Control. *Journal of Economic Dynamics and Control* 2:79–91.

- . 1982. Time to Build and Aggregate Fluctuations. *Econometrica* 50 (6):1345–1370.
- Kyle, Albert S. 1985. Continuous Auctions and Insider Trading. *Econometrica* 53 (6):1315–1335.
- Laibson, David. 1997. Golden Eggs and Hyperbolic Discounting. *Quarterly Journal of Economics* 112 (2):443–477.
- Lakatos, Imre. 1970. Falsification and the Methodology of Scientific Research Programmes. In *Criticism and the Growth of Knowledge*, edited by Imre Lakatos and Alan Musgrave, 91–196. Cambridge: Cambridge University Press.
- Lasry, Jean-Michel and Pierre-Louis Lions. 2007. Mean Field Games. *Japanese Journal of Mathematics* 2 (1):229–260.
- Leeper, Eric M. 1991. Equilibria under ‘Active’ and ‘Passive’ Monetary and Fiscal Policies. *Journal of Monetary Economics* 27 (1):129–147.
- Ljungqvist, Lars and Thomas J. Sargent. 2027. *Recursive Macroeconomic Theory*. Cambridge, MA: MIT Press, fifth ed. Forthcoming.
- Lovell, Michael C. 1986. Tests of the Rational Expectations Hypothesis. *American Economic Review* 76 (1):110–124.
- Lucas, Robert E., Jr. 1972a. Econometric Testing of the Natural Rate Hypothesis. In *The Econometrics of Price Determination*, edited by Otto Eckstein, 50–59. Washington, DC: Board of Governors of the Federal Reserve System.
- . 1972b. Expectations and the Neutrality of Money. *Journal of Economic Theory* 4 (2):103–124.
- . 1976. Econometric Policy Evaluation: A Critique. *Carnegie-Rochester Conference Series on Public Policy* 1 (1):19–46.
- . 1978. Asset Prices in an Exchange Economy. *Econometrica* 46 (6):1429–1445.
- . 1980. Methods and Problems in Business Cycle Theory. *Journal of Money, Credit and Banking* 12 (4):696–715.
- . 1986. Adaptive Behavior and Economic Theory. *Journal of Business* 59 (4):S401–S426.
- . 1987. *Models of Business Cycles*. Oxford: Basil Blackwell.

- Lucas, Robert E., Jr. and Edward C. Prescott. 1971. Investment Under Uncertainty. *Econometrica* 39 (5):659–681.
- Lucas, Robert E., Jr. and Thomas J Sargent. 1981. *Rational Expectations and Econometric Practice*. Minneapolis: University of Minnesota Press.
- Lucas, Robert E., Jr. and Nancy L Stokey. 1983. Optimal Fiscal and Monetary Policy in an Economy without Capital. *Journal of Monetary Economics* 12 (1):55–93.
- Marcet, Albert and Juan P Nicolini. 2003. Recurrent Hyperinflations and Learning. *American Economic Review* 93 (5):1476–1498.
- Marcet, Albert and Thomas J Sargent. 1989. Convergence of Least Squares Learning Mechanisms in Self-Referential Linear Stochastic Models. *Journal of Economic Theory* 48 (2):337–368.
- Marimon, Ramon, Ellen McGrattan, and Thomas J. Sargent. 1990. Money as a Medium of Exchange in an Economy with Artificially Intelligent Agents. *Journal of Economic Dynamics and Control* 14:329–373.
- Marschak, Jacob. 1950. Statistical Inference in Economics: An Introduction. In *Statistical Inference in Dynamic Economic Models*, edited by Tjalling C Koopmans, 1–50. New York: Wiley.
- Maskin, Eric and Jean Tirole. 2001. Markov Perfect Equilibrium: I. Observable Actions. *Journal of Economic Theory* 100 (2):191–219.
- McCall, John J. 1970. Economics of Information and Job Search. *Quarterly Journal of Economics* 84 (1):113–126.
- McKay, Alisdair, Emi Nakamura, and Jón Steinsson. 2016. The Power of Forward Guidance Revisited. *American Economic Review* 106 (10):3133–3158.
- Meyn, Sean. 2022. *Control Systems and Reinforcement Learning*. Cambridge University Press.
- Milgrom, Paul and Nancy Stokey. 1982. Information, Trade and Common Knowledge. *Journal of Economic Theory* 26 (1):17–27.
- Mirrlees, James A. 1971. An Exploration in the Theory of Optimum Income Taxation. *Review of Economic Studies* 38 (2):175–208.

- Moll, Benjamin. 2026. The Trouble with Rational Expectations in Heterogeneous Agent Models: A Challenge for Macroeconomics. *The Economic Journal* 136 (676):1173–1216.
- Molnár, Krisztina and Sergio Santoro. 2014. Optimal Monetary Policy When Agents Are Learning. *European Economic Review* 66:39–62.
- Morgenstern, Oskar. 1928. *Wirtschaftsprognose: Eine Untersuchung ihrer Voraussetzungen und Möglichkeiten*. Vienna: Julius Springer.
- Morris, Stephen and Hyun Song Shin. 1998. Unique Equilibrium in a Model of Self-Fulfilling Currency Attacks. *American Economic Review* 88 (3):587–597.
- . 2002. Social Value of Public Information. *American Economic Review* 92 (5):1521–1534.
- . 2003. Global Games: Theory and Applications. In *Advances in Economics and Econometrics: Theory and Applications, Eighth World Congress*, vol. 1, edited by Mathias Dewatripont, Lars Peter Hansen, and Stephen J. Turnovsky, 56–114. Cambridge: Cambridge University Press.
- . 2012. Contagious Adverse Selection. *American Economic Journal: Macroeconomics* 4 (1):1–21.
- Morris, Stephen, Hyun Song Shin, and Muhamet Yildiz. 2016. Common Belief Foundations of Global Games. *Journal of Economic Theory* 163:826–848.
- Mortensen, Dale T. and Christopher A. Pissarides. 1994. Job Creation and Job Destruction in the Theory of Unemployment. *Review of Economic Studies* 61 (3):397–415.
- Muth, John F. 1960. Optimal Properties of Exponentially Weighted Forecasts. *Journal of the American Statistical Association* 55 (290):299–306.
- . 1961. Rational Expectations and the Theory of Price Movements. *Econometrica* 29 (3):315–335.
- Myerson, Roger B. 1981. Optimal Auction Design. *Mathematics of Operations Research* 6 (1):58–73.
- Nash, John F. 1950. Equilibrium Points in n -Person Games. *Proceedings of the National Academy of Sciences* 36 (1):48–49.

- Nimark, Kristoffer P. 2017. Dynamic Higher Order Expectations. Society for Economic Dynamics 2017 Meeting Paper 1132; earlier version Universitat Pompeu Fabra Working Paper, 2007.
- Payne, Jonathan, Adam Rebei, and Yucheng Yang. 2026. Deep Learning for Search and Matching Models. Accepted at *Econometrica*, conditional on revisions.
- Pesando, James E. 1975. A Note on the Rationality of the Livingston Price Expectations. *Journal of Political Economy* 83:849–858.
- Plott, Charles R. and Shyam Sunder. 1982. Efficiency of Experimental Security Markets with Insider Information: An Application of Rational-Expectations Models. *Journal of Political Economy* 90 (4):663–698.
- . 1988. Rational Expectations and the Aggregation of Diverse Information in Laboratory Security Markets. *Econometrica* 56 (5):1085–1118.
- Prescott, Edward C. and Rajnish Mehra. 1980. Recursive Competitive Equilibrium: The Case of Homogeneous Households. *Econometrica* 48 (6):1365–1379.
- Radner, Roy. 1972. Existence of Equilibrium of Plans, Prices, and Price Expectations in a Sequence of Markets. *Econometrica* 40 (2):289–303.
- . 1979. Rational Expectations Equilibrium: Generic Existence and the Information Revealed by Prices. *Econometrica* 47 (3):655–678.
- Reiter, Michael. 2009. Solving Heterogeneous-Agent Models by Projection and Perturbation. *Journal of Economic Dynamics and Control* 33 (3):649–665.
- Rietz, Thomas A. 1988. The Equity Risk Premium: A Solution. *Journal of Monetary Economics* 22 (1):117–131.
- Rochet, Jean-Charles and Xavier Vives. 2004. Coordination Failures and the Lender of Last Resort: Was Bagehot Right After All? *Journal of the European Economic Association* 2 (6):1116–1147.
- Romer, Paul M. 1986. Increasing Returns and Long-Run Growth. *Journal of Political Economy* 94 (5):1002–1037.
- Rust, John. 1987. Optimal Replacement of GMC Bus Engines: An Empirical Model of Harold Zurcher. *Econometrica* 55 (5):999–1033.

- Sargent, Thomas, Noah Williams, and Tao Zha. 2009. The Conquest of South American Inflation. *Journal of Political Economy* 117 (2):211–256.
- Sargent, Thomas J. 1971. A Note on the “Accelerationist” Controversy. *Journal of Money, Credit and Banking* 3 (3):721–725.
- Sargent, Thomas J. 1973. Rational Expectations, the Real Rate of Interest, and the Natural Rate of Unemployment. *Brookings Papers on Economic Activity* 1973 (2):429–480.
- Sargent, Thomas J. 1977. The Demand for Money During Hyperinflations under Rational Expectations: I. *International Economic Review* 18 (1):59–82.
- . 1979. A Note on Maximum Likelihood Estimation of the Rational Expectations Model of the Term Structure. *Journal of Monetary Economics* 5 (1):133–143.
- Sargent, Thomas J. 1982. The Ends of Four Big Inflations. In *Inflation: Causes and Effects*, edited by Robert E. Hall, 41–98. Chicago: University of Chicago Press.
- Sargent, Thomas J. 1984. Autoregressions, Expectations, and Advice. *The American Economic Review* 74 (2):408–415.
- . 1993. *Bounded Rationality in Macroeconomics: The Arne Ryde Memorial Lectures*. Oxford: Oxford University Press.
- . 1999. *The Conquest of American Inflation*. Princeton: Princeton University Press.
- . 2024. Critique and Consequence. *Journal of Monetary Economics* 141:2–13.
- Sargent, Thomas J. 2025. Macroeconomics After Lucas. *Journal of Political Economy* 133 (11):3390–3417.
- Sargent, Thomas J. and John Stachurski. 2026a. Advanced Quantitative Economics with Python. QuantEcon lecture series, <https://python-advanced.quantecon.org/intro.html>. Accessed October 1, 2026.
- . 2026b. Intermediate Quantitative Economics with Python. QuantEcon lecture series, <https://python.quantecon.org/intro.html>. Accessed October 1, 2026.
- Sargent, Thomas J and Neil Wallace. 1973. The Stability of Models of Money and Growth with Perfect Foresight. *Econometrica* 41 (6):1043–1048.
- . 1976. Rational Expectations and the Theory of Economic Policy. *Journal of Monetary Economics* 2 (2):169–183.

- Sargent, Thomas J. and Neil Wallace. 1987. Inflation and the Government Budget Constraint. In *Economic Policy in Theory and Practice*, edited by Assaf Razin and Efraim Sadka, 170–207. London: Macmillan.
- Schmeidler, David. 1973. Equilibrium Points of Nonatomic Games. *Journal of Statistical Physics* 7 (4):295–300.
- Simon, Herbert A. 1956. Dynamic Programming Under Uncertainty with a Quadratic Criterion Function. *Econometrica* 24 (1):74–81.
- Sims, Christopher A. 1980. Macroeconomics and Reality. *Econometrica* 48 (1):1–48.
- . 1982. Policy Analysis with Econometric Models. *Brookings Papers on Economic Activity* 1982 (1):107–164.
- . 1988. Projecting Policy Effects with Statistical Models. *Revista de Análisis Económico* 3 (3).
- Sims, Christopher A. 1993. Rational Expectations Modeling with Seasonally Adjusted Data. *Journal of Econometrics* 55 (1-2):9–19.
- . 1994. A Simple Model for Study of the Determination of the Price Level and the Interaction of Monetary and Fiscal Policy. *Economic Theory* 4 (3):381–399.
- . 2003. Implications of Rational Inattention. *Journal of Monetary Economics* 50 (3):665–690.
- . 2006. Rational Inattention: Beyond the Linear-Quadratic Case. *American Economic Review* 96 (2):158–163.
- Smets, Frank and Rafael Wouters. 2007. Shocks and Frictions in US Business Cycles: A Bayesian DSGE Approach. *American Economic Review* 97 (3):586–606.
- Stokey, Nancy L. 1991. Credible Public Policy. *Journal of Economic Dynamics and Control* 15 (4):627–656.
- Svensson, Lars E. O. and Noah Williams. 2008. Optimal Monetary Policy under Uncertainty: A Markov Jump-Linear-Quadratic Approach. *Federal Reserve Bank of St. Louis Review* 90 (4):275–293.
- Svensson, Lars E.O. and Noah Williams. 2009. Optimal Monetary Policy under Uncertainty in DSGE Models: A Markov Jump-Linear-Quadratic Approach. In *Monetary Policy under*

- Uncertainty and Learning*, vol. 13 of *Central Banking, Analysis, and Economic Policies Book Series*, edited by Klaus Schmidt-Hebbel and Carl E. Walsh, chap. 3, 077–114. Central Bank of Chile.
- Szőke, Bálint. 2022. Estimating Robustness. *Journal of Economic Theory* 199:105225.
- Tallarini, Thomas D. 2000. Risk-Sensitive Real Business Cycles. *Journal of Monetary Economics* 45 (3):507–532.
- Theil, Henri. 1957. A Note on Certainty Equivalence in Dynamic Planning. *Econometrica* 25 (2):346–349.
- Townsend, Robert M. 1983. Forecasting the Forecasts of Others. *Journal of Political Economy* 91:546–588.
- Turmuhambetova, Gauhar. 2005. *Decision Making in an Economy with Endogenous Information*. PhD dissertation, Department of Economics, University of Chicago.
- Walker, Todd B. 2007. How Equilibrium Prices Reveal Information in a Time Series Model with Disparately Informed, Competitive Traders. *Journal of Economic Theory* 137 (1):512–537.
- White, Halbert. 1982. Maximum Likelihood Estimation of Misspecified Models. *Econometrica* 50 (1):1–25.
- Whiteman, Charles H. 1983. *Linear Rational Expectations Models: A User's Guide*. Minneapolis: University of Minnesota Press.
- Woodford, Michael. 2003. *Interest and Prices: Foundations of a Theory of Monetary Policy*. Princeton, NJ: Princeton University Press.