

Risk Aversion or Mistaken Beliefs? Three Martingales and More*

Thomas J. Sargent[†]

A single mathematical object connects risk-neutral pricing, subjective belief distortions, and concerns about model misspecification, namely, a multiplicative martingale increment that defines a likelihood ratio. I describe likelihood ratios that: (i) encode a representative investor’s risk aversion, following Lucas and Hansen; (ii) represent mistaken beliefs of professional forecasters, following Piazzesi, Salomao, and Schneider; and (iii) capture fears of model misspecification, following Hansen, Szöke, Han, and Sargent. By using a tilted discounted entropy ball that accommodates concerns about particular parametric models featuring long-run risk, the third likelihood ratio can generate state-dependent uncertainty prices that can explain hitherto puzzling movements in the term structure of interest rates.

Key Words: Risk; Uncertainty; Likelihood ratios; Relative entropy; Interest rates.

JEL Classification Numbers: G12, E43, D81, C58.

1. INTRODUCTION

This paper studies some long-standing puzzles about the term structure of interest rates. By “puzzle” I mean a feature of a joint distribution of macroeconomic variables and asset returns that is implied by a “standard” benchmark model but that is contradicted by U.S. time series data.¹ Relative to a “standard” asset pricing model like Lucas (1978) or Cox et al.

* The research presented in this paper was the basis for the 2017 inaugural David Backus Memorial Lecture. Dave was a pioneer in the study of affine term structure models and international risk sharing. His paper with Foresi and Telmer on affine models of currency pricing (Backus et al., 2001) is a direct intellectual ancestor of the framework presented here. Dave was skilled at connecting theory to data in sophisticated ways. He was also extraordinarily generous as a colleague, critic, and teacher. I thank Jonathan Payne, Jiao Shi, and Bálint Szöke for helpful discussions preceding my delivery of the Backus lecture, and Ziyue Yang for criticisms of earlier drafts of this paper.

[†] New York University

¹ Thus, a “puzzle” is always relative to a baseline model.

(1985), three patterns qualify as puzzles: 1) the nominal yield curve usually slopes upward; 2) the long-minus-short yield spread narrows before recessions and widens after them; and 3) long-term bond yields are nearly as volatile as short-term yields even though, under a rational expectations version of the expectations theory of the term structure, they should be much smoother averages of expected future short rates—this is what (Shiller, 1979) called the *volatility puzzle*.²

Hansen and Sargent (1981) tested cross-equation restrictions on vector autoregressions implied by the rational expectations theory of the term structure.³ They found some support for the theory's implications for conditional means, but also evidence of the difficulties with conditional second moments that Shiller emphasized. To resolve the volatility puzzle, risk prices—or something observationally equivalent—must be large, volatile, and state-dependent. What generates those state-dependent risk prices? Risk aversion? Wrong beliefs? Fears of model misspecification? Disentangling these possible sources is the topic of this paper.

To set the stage, I say what I mean by “model,” “rational expectations model,” “rational expectations econometrics,” and “wrong beliefs.” A (statistical) “model” is a joint probability distribution over a sequence of random variables indexed by a vector of known parameters. In a “rational expectations model,” all of the agents inside the model share the model with the economic theorist who created it. “Rational expectations econometrics” means that the econometrician shares the model with the economic theorist. “Wrong beliefs” means that some of the people inside the model don't use the econometrician's model. Thus, “wrong beliefs” are relative to a model that an econometrician has somehow inferred from historical data via a process of induction that I won't describe in this paper. Below we'll sometimes refer to the econometrician's model as the “baseline model” or the econometrician's “approximating model.”

In asset pricing data, consequences of investors' *risk aversion* and *wrong beliefs* are confounded. In a rational expectations model with a risk-averse representative investor, higher mean returns compensate for higher risks. But if the representative investor holds wrong beliefs, belying the rational expectations “communism-of-beliefs” assumption, observed average returns depend on *both* risk aversion *and* misunderstood return distributions. Wrong beliefs can generate what, from the perspective of an econometrician's approximating model, look like “stochastic discount factor shocks.”

²For influential research about the expectations theory of the term structure of interest rates before rational expectations was imposed, see Bierwag and Grove (1967) and Meiselman (1961).

³Also see Sargent (1972, 1979) and Hansen and Sargent (1991), which reprinted Hansen and Sargent (1981).

Can such “shocks” explain the observed countercyclical behavior of risk prices?

This paper organizes ideas around a single concept: the *likelihood ratio*, a non-negative random variable with unit mean that twists one probability distribution into another. Much of the analysis draws on two contributions: Hansen et al. (2020), who construct twisted probability models from tilted discounted entropy balls to price model uncertainty, and Szőke (2022), who estimates a robust control model with cross-equation restrictions implied by a representative agent’s worst-case beliefs and evaluates its implications for the yield curve and survey expectations. Likelihood ratios—equivalently, multiplicative martingale increments—appear in at least four distinct roles in modern asset pricing, summarized in Table 1:

TABLE 1.

Four probability measures and associated likelihood ratios. Each likelihood ratio takes the log-normal form $m_{t+1}^b = \exp(-b'_t \varepsilon_{t+1} - \frac{1}{2} b'_t b_t)$ with $b_t = 0, \lambda_t, w_t$, or a worst-case distortion.

Probability	Likelihood ratio	Describes
Econometric	1	macro risk factors
Risk neutral	m_{t+1}^λ	prices of risks
Mistaken	m_{t+1}^w	experts’ forecasts
Doubtful	$m_{t+1} \in \mathcal{M}$	misspecification fears

The “three martingales” of my subtitle are cumulative likelihood ratio processes $\{M_t = \prod_{s=0}^t m_s\}_{t=0}^\infty$ formed from (i) risk prices λ_t , (ii) belief distortions w_t^* , and (iii) worst-case model distortions $\tilde{w}_t$. The “and more” in my subtitle refers to $\bar{w}_t$, a parametric long-run risk worry that tilts an entropy ball. The risk-price twister λ_t stands apart from the other three: it shifts the price of risk directly, whereas w_t^* , $\bar{w}_t$, and $\tilde{w}_t$ all twist the conditional *mean* of the shocks. Section 5 onward shows how the single λ_t that a rational-expectations econometrician estimates can be decomposed into a smaller risk price plus one of these mean distortions.

The rest of this paper proceeds as follows. Section 2 sets up a statistical framework in which a likelihood ratio twists probabilities while preserving a tractable structure. Section 3 reviews rational-expectations asset pricing and shows how a single likelihood ratio turns risk-neutral pricing formulas into modern asset pricing formulas. Section 4 discusses a statistical parameter identification issue that formalizes how asset pricing data confound risk-aversion and wrong beliefs: the same likelihood ratio can be interpreted as risk aversion or as belief distortion relative to the econometrician’s approximating model. Section 5 describes an approach of Piazzesi et al. (2015), who use survey data on professional forecasters to decompose the likelihood ratio into a smaller risk price and a belief distortion. Section 6

pauses to discuss how rational expectations may not be a good approximation when persistent dynamics make learning converge slowly. Section 7 develops a robust control theory approach, beginning with Hansen's formulation with a discounted entropy ball. Section 8 discusses Hansen and Szóke's tilted entropy ball that accommodates concerns about parametric alternatives with more long-run risk. Section 9 describes how models of robustness impose powerful cross-equation restrictions that resemble those imposed by rational expectations models. Section 10 summarizes empirical results. Section 11 concludes with reflections on why different interpretations and fits of these joint probability distributions matter for macroeconomics. Appendix A.1 describes the multiplier preference formulation of a dubious agent's problem.

2. RISKS AND BELIEFS

This section introduces the paper's central mathematical device: a likelihood ratio that twists one probability distribution into another. In a Gaussian setting, this twist amounts to shifting the mean of the shock distribution by a vector λ while leaving the covariance structure unchanged. The framework sets forth vocabulary and notation that will recur throughout the paper.

Let ε denote a vector of risks to be borne and priced. Suppose that, under the econometrician's probability model, ε has a standard multivariate normal density:

$$\phi(\varepsilon) \propto \exp\left(-\frac{1}{2} \varepsilon' \varepsilon\right), \quad \varepsilon \sim \mathcal{N}(0, I).$$

Define a *likelihood ratio*

$$m(\varepsilon) = \exp\left(-\lambda' \varepsilon - \frac{1}{2} \lambda' \lambda\right) \geq 0,$$

which satisfies $E m(\varepsilon) = 1$ by construction. The *twisted density*

$$\hat{\phi}(\varepsilon) = m(\varepsilon) \phi(\varepsilon) \propto \exp\left(-\frac{1}{2} (\varepsilon + \lambda)' (\varepsilon + \lambda)\right)$$

is a $\mathcal{N}(-\lambda, I)$ density. The *relative entropy* of the twisted density with respect to the baseline density is

$$E[m(\varepsilon) \log m(\varepsilon)] = \frac{1}{2} \lambda' \lambda,$$

a convenient measure of the statistical distance between the two models.

Depending on context, the vector λ represents risk prices, belief distortions, or worst-case mean perturbations under model uncertainty.

3. ASSET PRICING WITH LIKELIHOOD RATIOS

With the likelihood ratio in hand, we now embed it in a dynamic asset pricing framework. The key insight is that the same formulas used to price assets under risk neutrality and rational expectations carry over to settings with risk-averse or belief-distorted investors—one simply replaces the econometrician's probability measure with a twisted one. This section develops that idea in a linear Gaussian state-space setting and shows how a single likelihood ratio generates the risk-neutral dynamics that underlie modern affine term structure models.

3.1. The econometrician's model

The econometrician works with a linear Gaussian state-space system:

$$x_{t+1} = A x_t + C \varepsilon_{t+1}, \quad (1)$$

$$y_{t+1} = D x_t + G \varepsilon_{t+1}, \quad (2)$$

$$\varepsilon_{t+1} \sim \mathcal{N}(0, I),$$

where y_{t+1} collects utility-relevant variables such as the first difference $c_{t+1} - c_t$ of the logarithm of per capita consumption, $r_t = \bar{r} x_t$ is the risk-free one-period interest rate, and $d_t = \bar{d} x_t$ is the payout process from an asset.⁴ The object of interest is the price at date t of a claim to the random payout stream $\{d_{t+j}\}_{j=1}^{\infty}$. Figure 1 illustrates an estimated model of the form (1)–(2) with $y_{t+1} = c_{t+1} - c_t$.

3.2. Risk-neutral rational expectations pricing

Under rational expectations with a risk-neutral representative investor, stock prices satisfy Shiller (1979):


$$p_t = \exp(-r_t) E_t(p_{t+1} + d_{t+1}).$$

The expectations theory of the term structure of interest rates (from Hicks (1939) to Shiller (1979)) prices a zero-coupon risk-free claim to one dollar at time $t + n$ as:

$$\begin{aligned} p_t(1) &= \exp(-r_t), \\ p_t(n+1) &= \exp(-r_t) E_t p_{t+1}(n), \\ p_t(n) &= \exp(B_n x_t). \end{aligned} \quad (3)$$

⁴To accommodate constant terms easily, we assume that the first entry in the state vector is a constant 1 so that $x_t = \begin{bmatrix} 1 \\ \tilde{x}_t \end{bmatrix}$ where $A = \begin{bmatrix} 1 & 0 \\ 0 & \tilde{A} \end{bmatrix}$ and $\tilde{A}$ is a stable matrix, $C = \begin{bmatrix} 0 \\ \tilde{C} \end{bmatrix}$, and the first component of x_0 is 1.

FIG. 1. An econometrician’s model of growth in aggregate consumption: estimated state dynamics.



These formulas work “pretty well” for conditional means but less well for conditional variances—the Shiller “volatility puzzles.”

3.3. Modern asset pricing: adding risk aversion

It would be convenient if versions of the same formulas worked even when investors are risk averse or do not have rational expectations. The likelihood ratio makes this possible.

Let

$$\begin{aligned} m_{t+1}^\lambda &= \exp\left(-\lambda'_t \varepsilon_{t+1} - \frac{1}{2} \lambda'_t \lambda_t\right), & \lambda_t &= \lambda x_t, \\ E_t m_{t+1}^\lambda &= 1, & m_{t+1}^\lambda &\geq 0. \end{aligned} \quad (4)$$

The likelihood ratio m_{t+1}^λ distorts the conditional distribution of ε_{t+1} from $\mathcal{N}(0, I)$ to $\mathcal{N}(-\lambda x_t, I)$. Covariances of returns with m_{t+1}^λ affect mean returns—this is the channel through which risk aversion prices risks.

With this device, modern asset pricing (Lucas (1978) for stocks, Dai and Singleton (2000), Cox et al. (1985), and Backus and Zin (1994) for bonds) takes the form:

$$p_t = \exp(-r_t) E_t(m_{t+1}^\lambda (p_{t+1} + d_{t+1})), \quad (5)$$

and for the term structure:

$$\begin{aligned} p_t(1) &= \exp(-r_t), \\ p_t(n+1) &= \exp(-r_t) E_t(m_{t+1}^\lambda p_{t+1}(n)), \\ p_t(n) &= \exp(\tilde{B}_n x_t). \end{aligned} \tag{6}$$

3.4. Risk-neutral dynamics

The risk-neutral representation implies twisted dynamics:

$$x_{t+1} = (A - C\lambda) x_t + C \tilde{\varepsilon}_{t+1}, \quad \tilde{\varepsilon}_{t+1} \sim \mathcal{N}(0, I).$$

The risk-neutral dynamics assert that the shock distribution ε_{t+1} has conditional mean $-\lambda_t$ instead of 0. The dependence of $\lambda_t = \lambda x_t$ on the state modifies the dynamics relative to the econometrician's model.

3.5. Expectation under a twisted distribution

The mathematical expectation of y_{t+1} under the probability distribution twisted by likelihood ratio m_{t+1} is

$$\tilde{E}_t y_{t+1} = E_t m_{t+1} y_{t+1}.$$

With respect to the risk-neutral dynamics, the term structure theory becomes:

$$\begin{aligned} p_t(1) &= \exp(-r_t), \\ p_t(n+1) &= \exp(-r_t) \tilde{E}_t p_{t+1}(n), \\ p_t(n) &= \exp(\tilde{B}_n x_t). \end{aligned}$$

These are the same formulas as rational-expectations asset pricing, but mathematical expectations are taken with respect to a probability measure *twisted by risk aversion*.

4. THE IDENTIFICATION CHALLENGE

The elegance of the likelihood ratio device comes at a cost: the same mathematical object can represent very different economic forces. This section makes the observational equivalence precise. Because asset pricing formulas depend on λ_t but not on its interpretation, the data alone cannot tell us whether observed risk premia reflect the representative investor's risk aversion or systematically mistaken beliefs. Resolving this issue requires either additional data or additional theory—the subjects of the sections that follow.

The vector λ_t can be interpreted as either:

- a risk price vector expressing the representative agent's risk aversion, or
- the representative agent's belief distortion relative to the econometrician's model.

The asset pricing formulas (5)–(6) are identical under both interpretations, and so are the econometric fits. Relative to the model of a risk-averse representative investor with rational expectations, a model of a risk-neutral investor with appropriately mistaken beliefs produces observationally equivalent predictions. This insight was articulated by Hansen et al. (1999) and Piazzesi et al. (2015) (henceforth PSS).

To distinguish risk aversion from belief distortion, one needs either more information (the PSS approach using survey data) or more theory (the Hansen and Szőke robust control approach), or both (Szőke (2022)).

5. MORE INFORMATION: EXPERTS' FORECASTS

One natural strategy for breaking the identification deadlock is to bring in additional data. Piazzesi et al. (2015) pursue exactly this approach by exploiting surveys of professional forecasters. If forecasters' reported expectations differ systematically from those implied by the econometrician's model, the discrepancy can be attributed to belief distortions, leaving a smaller residual to be explained by risk aversion alone. This section describes their framework and its empirical findings.

5.1. The PSS framework

PSS exploit data on professional forecasters' expectations to decompose the likelihood ratio into risk prices and belief distortions. Their setup posits:

- The representative agent's risk aversion leads him to price risks ε_{t+1} with prices $\lambda_t^* = \lambda^* x_t$.
- The representative agent has twisted beliefs $(A^*, C) = (A - Cw^*, C)$ relative to the econometrician's model (A, C) .
- Professional forecasters use the twisted beliefs (A^*, C) to answer survey questions about their forecasts.

The question PSS assume experts think they are answering is: "What forecasts underlie your investment advice decisions?"

5.2. Estimation strategy

PSS proceed as follows:

1. Use data $\{x_t\}_{t=0}^T$ to estimate the econometrician's model A, C .

2. Project experts' forecasts $\{\hat{x}_{t+1}\}$ on x_t to obtain $\hat{x}_{t+1} = A^* x_t$ and interpret A^* as a belief distortion.
3. Back out the implied shift $w^* x_t = -C^{-1}(A^* - A) x_t$ in the conditional mean of ε_{t+1} .
4. Reinterpret the λ estimated by the rational-expectations econometrician as $\lambda = \lambda^* + w^*$, where $\lambda_t^* = \lambda^* x_t$ is the (smaller) price of risk vector actually charged by the representative agent with distorted beliefs.

An econometrician who mistakenly imposes rational expectations estimates risk prices λ_t that are the sum of two parts:

- smaller risk prices λ_t^* actually charged by the erroneous-beliefs representative agent, and
- conditional mean distortions w_t^* of the risks ε_{t+1} that the twisted-beliefs representative agent's model displays relative to the econometrician's.

5.3. Empirical specification

The key variables in PSS's state-space system are:

- the level and slope of the yield curve (short rate; spread between 5-year and short rate),
- inflation, and
- conditional expectations of these variables.

5.4. PSS successes

PSS find $w^* \neq 0$:

- The experts' model differs systematically from the econometrician's.
- Experts perceive the level and slope of the yield curve to be more persistent than the econometrician's estimates imply.
- Subjective risk prices $\lambda^* x_t$ vary less than the λx_t estimated by the rational-expectations econometrician.

5.5. Why are beliefs distorted?

PSS offer no explanation for why beliefs are distorted. Are they mistakes? Ignorance of good econometrics? And how distorted are they in terms of the divergence of probabilities? These questions motivate a theory of belief distortions rooted in robust control theory.

An intriguing interpretive question lurks here. What questions do professional forecasters think that they are answering when they respond to surveys? Are they reporting their best guess of the future, or the scenario that rationalizes their portfolio decisions? If the latter, their "mistakes"

may not be mistakes at all but instead reflect a pessimistic model that they use precisely because they fear misspecification.⁵ Section 7 develops a theory in which this distinction is formalized: the robust control agent's worst-case model generates forecasts that look like the experts' distorted beliefs, but arise from an optimizing response to model uncertainty rather than from ignorance.

6. RATIONAL EXPECTATIONS AND ITS ALTERNATIVES

Before developing a theory of belief distortions, it is worth pausing to ask why a rational expectations model might not be a good approximation.⁶ The standard justification treats rational expectations as the outcome of learning from a long history—but how long is long enough? This section reviews the conditions under which learning converges to rational expectations and highlights circumstances, particularly those involving persistent dynamics, in which convergence is so slow that model uncertainty becomes economically first-order.

6.1. The 1970s–1980s informal justification

The standard justification for rational expectations treats it as the outcome of learning from an infinite history: least-squares learning converges to rational expectations. The argument requires that:

- agents know correct functional forms, and
- a stochastic approximation argument partitions dynamics into fast parts (justifying a law of large numbers) and slow parts (justifying an ODE).

However, long intertemporal dependencies make *rates of convergence slow*. For example, distinguishing two Gaussian AR(1) processes whose persistence parameters differ by 0.01 (say $\rho = 0.98$ versus $\rho = 0.99$) at the 5% level requires several hundred years of quarterly data. For the long-run risk specifications considered by Bansal and Yaron (2004), the required sample lengths can exceed a millennium. This is not a failure of econometrics—it is the statistical reality that makes model uncertainty economically important.

⁵Ellison and Sargent (2012) applied this idea to interpreting responses of Federal Reserve officials to questions about the forecasts of macroeconomic outcomes that rationalize their decisions.

⁶Sargent (2008) discusses the intellectual history of rational expectations and adaptive learning, framing the tension between them as one between “intelligent design” and “evolution.” The introduction to Lucas and Sargent (1981) lays out the case for rational expectations as a discipline on econometric practice.

6.2. Good econometricians

Good econometricians have limited data and only hunches about functional forms. They fear that their fitted models are incorrect. An agent who is like a good econometrician:

- has a parametric model estimated from limited data,
- acknowledges that many other specifications fit nearly as well—other parameter values, other functional forms, omitted variables, neglected nonlinearities, and history dependencies,
- fears that one of those other models actually prevails, and
- seeks “good enough” decisions under *all* such alternative models—*robustness*.

7. A THEORY OF BELIEF DISTORTIONS: ROBUST CONTROL

Having seen that survey data point to systematic belief distortions and that statistical theory explains why such distortions are hard to eliminate through learning, we now develop a theory that rationalizes them. Robust control theory provides a framework in which a decision maker who acknowledges that her model may be misspecified optimally distorts her probability assessments toward a worst-case scenario. The resulting worst-case beliefs look like the “mistakes” identified by PSS, but arise from a coherent response to model uncertainty rather than from ignorance.

7.1. Hansen’s dubious agent

Inspired by robust control theory, consider a dubious investor who:

- shares the econometrician’s model A, C, D, G ,
- expresses doubts by using a continuum of likelihood ratios to form a *discounted entropy ball* of size η around the econometrician’s model,
- wants a valuation that is good for every model in the entropy ball, and
- constructs a lower bound on values and a worst-case model that attains it.

7.1.1. Why discounted entropy?

Discounted entropy includes models that undiscounted entropy excludes. Undiscounted entropy over infinite sequences admits only models that share laws of large numbers, so it effectively requires alternative models to share the same long-run averages as the baseline and excludes models that differ only in persistent, low-frequency dynamics. But those are precisely the models that are hardest to distinguish with finite data and that matter

most for long-lived asset prices. Discounted entropy, by treating future divergences less severely, admits these statistically elusive but economically important alternatives into the set of models that the dubious agent contemplates.⁷

7.2. Valuation under the econometrician's model

Taking the log consumption process to be linear Gaussian with shocks $\varepsilon_{t+1} \sim \mathcal{N}(0, I)$:

$$c_{t+1} - c_t = D x_t + G \varepsilon_{t+1}, \quad (7)$$

$$x_{t+1} = A x_t + C \varepsilon_{t+1}, \quad (8)$$

the value function is

$$V(x_0, c_0) := E \left[\sum_{t=0}^{\infty} \beta^t c_t \mid x_0, c_0 \right] = c_0 + \beta E[V(x_1, c_1) \mid x_0, c_0].$$

7.3. The sequence problem

The dubious agent solves:

$$\begin{aligned} J(x_0, c_0 \mid \eta) &:= \min_{\{m_{t+1}\}_{t=0}^{\infty}} E \left[\sum_{t=0}^{\infty} \beta^t M_t c_t \mid x_0, c_0 \right] \\ \text{s.t. } &c_{t+1} - c_t = D x_t + G \varepsilon_{t+1}, \\ &x_{t+1} = A x_t + C \varepsilon_{t+1}, \\ &E \left[\sum_{t=0}^{\infty} \beta^t M_t E[m_{t+1} \log m_{t+1} \mid x_t, c_t] \mid x_0, c_0 \right] \leq \eta, \\ &M_{t+1} = M_t m_{t+1}, \quad E[m_{t+1} \mid x_t, c_t] = 1, \quad M_0 = 1. \end{aligned} \quad (9)$$


The likelihood ratio process $\{M_t\}_{t=0}^{\infty}$ is a multiplicative *martingale*. Figure 2 depicts the discounted entropy ball—the set of alternative models, indexed by their likelihood ratios, that the dubious agent contemplates.

7.4. Entropy and the likelihood ratio

With the log-normal likelihood ratio

$$m_{t+1} := \exp \left(-w_t' \varepsilon_{t+1} - \frac{w_t' w_t}{2} \right),$$

⁷This property of discounted entropy is what motivated Hansen and Sargent (1995) to introduce it.

FIG. 2. Discounted entropy ball around the econometrician's model.

conditional entropy takes the simple form

$$E[m_{t+1} \log m_{t+1} \mid x_t, c_t] = \frac{1}{2} w_t' w_t.$$

Substituting this into (9) yields the reformulated problem:

$$\begin{aligned} J(x_0, c_0 \mid \eta) &:= \min_{\{w_t\}_{t \geq 0}} E^w \left[\sum_{t=0}^{\infty} \beta^t c_t \mid x_0, c_0 \right] \\ \text{s.t. } &c_{t+1} - c_t = D x_t + G (\tilde{\varepsilon}_{t+1} - w_t), \quad \tilde{\varepsilon}_{t+1} \sim \mathcal{N}(0, I), \\ &x_{t+1} = A x_t + C (\tilde{\varepsilon}_{t+1} - w_t), \\ &\frac{1}{2} E^w \left[\sum_{t=0}^{\infty} \beta^t w_t' w_t \mid x_0, c_0 \right] \leq \eta. \end{aligned} \tag{10}$$

7.5. Outcome of the Hansen model

The worst-case mean distortion turns out to be a constant vector:

$$w_t = \bar{w}.$$

The consequence is that the contribution of w_t to risk prices is state-independent. This does *not* help explain countercyclical prices of risk (or prices of model uncertainty).

8. TILTING THE ENTROPY BALL

The basic robust control model of the previous section has a limitation: it generates worst-case distortions that are constant across states, and therefore cannot explain the countercyclical variation in risk prices that the data demand. This section shows how Hansen et al. (2020) overcome this limitation by *tilting* the entropy ball—requiring it to include particular parametric alternatives, such as models with more long-run risk. The tilt makes the entropy constraint bind differently across states, producing the state-dependent uncertainty prices needed to match term-structure dynamics.

8.1. Hansen and Szöke’s more refined dubious agent

To generate state-dependent uncertainty prices, Hansen et al. (2020) introduce a more refined dubious agent who:

- shares the econometrician’s model A, C, D, G ,
- expresses doubts by using a continuum of likelihood ratios to form a discounted entropy ball around the econometrician’s model,
- also insists that some martingales representing particular alternative *parametric* models be included in the discounted entropy ball.

The inclusion of those alternative parametric models *tilts* the entropy ball, which affects the worst-case model in a way that can produce countercyclical uncertainty prices. Figure 3 contrasts the resulting tilted balls with the symmetric ball of the untilted problem.

8.2. Concern about other parametric models

The investor wants to include particular alternative models with

$$E_t[\bar{m}_{t+1} \log \bar{m}_{t+1}] = \frac{1}{2} \bar{w}'_t \bar{w}_t = \xi(x_t)$$

and discounted entropy

$$E^{\bar{w}} \left[\sum_{t=0}^{\infty} \beta^t \xi(x_t) \mid x_0, c_0 \right].$$

This is accomplished by replacing the earlier entropy constraint with

$$\frac{1}{2} E^w \left[\sum_{t=0}^{\infty} \beta^t w'_t w_t \mid x_0, c_0 \right] \leq E^w \left[\sum_{t=0}^{\infty} \beta^t \xi(x_t) \mid x_0, c_0 \right]. \quad (11)$$

8.3. Bansal and Yaron's bazooka

Bansal and Yaron (2004) mix long-run risk (LRR) with Epstein–Zin (EZ) preferences.⁸ LRR is statistically difficult to detect and estimate, but an Epstein–Zin agent or a dubious agent really hates LRR. Under certain conditions, the EZ value function is the indirect utility function of a dubious agent. This sets the stage for big risk prices—but not countercyclical ones.

Remark 8.1. [EZ–robust control equivalence] The equivalence between Epstein–Zin preferences and robust control, established by Hansen and Sargent (2001) and Maenhout (2004), implies that EZ preferences with unit elasticity of intertemporal substitution and risk aversion $\gamma > 1$ are observationally equivalent to expected utility with $\gamma = 1$ combined with a penalty parameter $\theta = 1/(\gamma - 1)$ governing robustness. The higher the risk aversion under EZ, the smaller the entropy ball the robust agent is willing to tolerate—she fears a narrower but more intense set of model misspecifications.

8.4. Concern about bigger long-run risk in the Szöke model

Inspired by Bansal and Yaron (2004), an agent fears particular long-run risks expressed by

$$x_{t+1} = \bar{A} x_t + C \tilde{\varepsilon}_{t+1}.$$

This corresponds to $\bar{w}_t = \bar{w} x_t$ with

$$\bar{w} = -C^{-1}(\bar{A} - A),$$

which implies a quadratic ξ function:

$$\xi(x_t) := \frac{1}{2} x_t' \bar{w}' \bar{w} x_t =: \frac{1}{2} x_t' \Xi x_t.$$

8.5. State-dependent contributions to the entropy constraint

The time- t contributions to the right-hand side of (11) relax the discounted entropy constraint in states x_t in which $\xi(x_t)$ is larger. This sets the stage for state-dependent mean distortions in the worst-case model, which in turn ignite countercyclical market prices of uncertainty.

⁸Dave Backus dubbed Bansal and Yaron's idea of combining two hypotheses—the presence of long-run risk in consumption growth together with a representative agent who hates long-run risk—a 'bazooka' because of the powerful way it raises market prices of risk.

8.6. The Szöke agent's sequence problem

The resulting linear-quadratic problem is:

$$\begin{aligned}
 J(x_0, c_0 \mid \Xi) &:= \max_{\tilde{\theta} \geq 0} \min_{\{w_t\}_{t \geq 0}} E^w \left[\sum_{t=0}^{\infty} \beta^t c_t + \tilde{\theta} \frac{1}{2} \sum_{t=0}^{\infty} \beta^t (w_t' w_t - x_t' \Xi x_t) \mid x_0, c_0 \right] \\
 &\quad (12) \\
 \text{s.t.} \quad &c_{t+1} - c_t = D x_t + G (\tilde{\varepsilon}_{t+1} - w_t), \quad \tilde{\varepsilon}_{t+1} \sim \mathcal{N}(0, I), \\
 &x_{t+1} = A x_t + C (\tilde{\varepsilon}_{t+1} - w_t).
 \end{aligned}$$

The worst-case shock mean distortion is⁹

$$\tilde{w}_t = \tilde{W} x_t,$$

and the worst-case model is $(\tilde{A}, C, \tilde{D}, G)$ with

$$\begin{aligned}
 \tilde{A} &= A - C \tilde{W}, \\
 \tilde{D} &= D - G \tilde{W}.
 \end{aligned}$$

9. CROSS-EQUATION RESTRICTIONS AND ESTIMATION

A recurring theme throughout the rational expectations econometrics literature is that equilibrium models impose testable restrictions across the equations of a vector autoregression.¹⁰ This section highlights how the robust control framework preserves that desirable feature: even though the decision maker's beliefs depart from the econometrician's model, the departure is disciplined by the worst-case optimization problem, yielding a new but equally sharp set of cross-equation restrictions that can be taken to the data.


⁹As assumed above, the first entry of the state vector is a constant 1, so that $x_t = \begin{bmatrix} 1 \\ \tilde{x}_t \end{bmatrix}$ with $A = \begin{bmatrix} 1 & 0 \\ 0 & \tilde{A} \end{bmatrix}$ and $\tilde{A}$ stable. Consequently, the worst-case distortion is now affine in the stochastic part $\tilde{x}_t$ of the state vector:

$$\tilde{w}_t = a + \tilde{W} \tilde{x}_t$$

where a is a constant vector determined by the linear terms in the value function. When $\Xi = 0$ (no tilting), the constant a reduces to the Hansen $\bar{w}$ from the untilted problem. However, when $\Xi \neq 0$ the two differ because the quadratic value function contributes an additional $2\beta C' \Xi C$ term to the left-hand side of the first-order condition for the constant part. The state-dependent component $\tilde{W} \tilde{x}_t$ is the new contribution of the tilted entropy ball, and is what generates countercyclical uncertainty prices.

¹⁰See Lucas and Sargent (1981).

FIG. 3. Tilted discounted entropy balls. Including particular parametric alternatives with more long-run risk tilts the entropy ball and generates state-dependent worst-case distortions.



Hansen and Sargent (1991) provided systematic methods for deriving and imposing the cross-equation restrictions on vector autoregressions that are implied by exact linear rational expectations models; those methods set the stage for the estimation strategy used here. As Szőke puts it:

An appealing feature of robust control theory is that it lets us deviate from rational expectations, but still preserves a set of powerful cross-equation restrictions on decision makers' beliefs... Consequently, estimation can proceed essentially as with rational expectations econometrics. The main difference is that now restrictions through which we interpret the data emanate from the decision maker's best response to a worst-case model instead of to the econometrician's model.

10. EMPIRICAL RESULTS

We now confront the theoretical framework with data. This section reviews the key empirical patterns in the U.S. term structure of interest rates—the upward-sloping yield curve, the predictive power of the yield spread for recessions, and Shiller's volatility puzzle—and assesses how well

different models account for them. The centerpiece is Szőke’s (2022) estimation of the tilted entropy ball model, which imposes cross-equation restrictions linking the worst-case model to both bond prices and survey forecasts.

10.1. The U.S. term structure

FIG. 4. U.S. term structure of interest rates.




Figure 4 plots the U.S. term structure of interest rates. Several recognized patterns characterize it:

- The nominal yield curve usually slopes upward.
- The long-minus-short yield spread narrows before U.S. recessions and widens after them.
- Consequently, the slope of the yield curve helps predict aggregate economic activity.
- Long and short yields are almost equally volatile (the “Shiller puzzle”).
- To solve the Shiller puzzle, risk prices (or something observationally equivalent) must depend on volatile state variables.

Specific historical episodes bring these patterns to life. The yield curve inverted before each of the last seven U.S. recessions. In the Szőke model, the state-dependent worry $\xi(x_t)$ ties these episodes together. Late in an expansion, when the agent’s long-run risk worry is muted ($\xi(x_t)$ small), uncertainty prices are compressed and the long-minus-short spread narrows, so that the curve can invert—the “narrows before recessions” half of the

second pattern. As a downturn arrives and uncertainty about long-run consumption growth surges ($\xi(x_t)$ large), as in the 2008–09 financial crisis and the March 2020 COVID shock, uncertainty prices and term premia spike and the spread widens—the “widens after them” half. This is exactly the state dependence that the tilted entropy ball is designed to capture.

10.2. Challenges and responses

Table 2 summarizes how various models perform against these empirical challenges.

TABLE 2.

Challenges and model performances. Affine risk prices with persistent consumption growth can address the second column; the third column requires state-dependent prices of risk. “Built-in” indicates patterns that Piazzesi et al. (2015) match by construction, since they extract belief distortions directly from survey data.

	Average slope	Slopes near recessions	Volatile long yield
Lucas (1978)	no	no	no
Epstein–Zin with LRR (Piazzesi and Schneider (2007), Hansen and Sargent (2001))	maybe	yes	no
Piazzesi et al. (2015)	built-in	built-in	yes
Szöke (2022)	YES	yes	yes

10.3. Szöke’s empirical strategy

Szöke proceeds in two stages:

Stage I: Estimation..

1. Use $\{x_t, c_t\}_{t=0}^T$ to estimate the econometrician’s A, C, D, G .
2. View Ξ as a matrix of additional free parameters and estimate them simultaneously with risk prices $\tilde{\lambda} x_t$ from data $\{p_t(n+1)\}_{n=1}^N, t = 0, \dots, T$, by imposing cross-equation restrictions:

$$\begin{aligned} p_t(n+1) &= \exp(-r_t) E_t[m_{t+1}^{\text{total}} p_{t+1}(n)], \\ m_{t+1}^{\text{total}} &= \exp\left(-(\tilde{w}_t + \tilde{\lambda}_t)' \varepsilon_{t+1} - \frac{1}{2} \|\tilde{w}_t + \tilde{\lambda}_t\|^2\right) \end{aligned} \quad (13)$$

where E_t is taken with respect to the econometrician’s model, $\tilde{w}_t = \tilde{W} x_t$ is the dubious investor’s worst-case mean distortion, and $\tilde{\lambda}_t = \tilde{\lambda} x_t$ is the

ordinary (log-utility) risk price that coexists with that distortion in the total likelihood ratio m_{t+1}^{total} .

Stage II: Assessment..

1. Assess improvements in predicted behavior of the term structure.
2. Use estimated worst-case dynamics to form distorted forecasts $\tilde{x}_{t+1} = (A - C\tilde{W})x_t$ and compare them to those of professional forecasters.
3. Compute discounted relative entropy of the worst-case twisted model $(A - C\tilde{W}, C, D - G\tilde{W}, G)$ relative to the econometrician's model (A, C, D, G) , and use it together with Chernoff (Chernoff, 1952) and Newman–Stuck Newman and Stuck, 1979 entropy measures to assess the difficulty of distinguishing the two models.

10.4. Interpretations of conditional mean distortions

The conditional mean distortions $w x_t$ admit distinct interpretations:

- In PSS, w_t^* is a vector of “mistakes” or “suboptimal forecasts.”
- In Szőke, $\tilde{w}_t$ is a vector of “model uncertainty” prices: compensations the representative agent charges to bear risks ε_{t+1} with unknown probability distribution.

10.5. Insights from Szőke's model

Szőke's framework delivers:

1. A theory of state-dependent belief distortions $\tilde{w}_t = \tilde{W} x_t$.
2. A theory about the question that professional forecasters answer: they respond with their worst-case model because they hear “tell me forecasts that rationalize your (max-min) decisions.”
3. A way to measure the size of belief distortions relative to the econometrician's model.

He uses his estimated Ξ matrix:

- to infer an equivalence class of alternative parametric models parameterized by $\bar{w}$ that concern the representative investor—these models have more long-run risk than the econometrician's model;
- to infer a worst-case mean distortion $\tilde{W} x_t$ whose state dependence causes the term structure to move with x_t in ways that explain hitherto puzzling term-structure movements, including Shiller's “volatility puzzle.”

10.6. Three probability twisters

To summarize, alongside the risk-price twister λ_t of Section 3.3, three distinct *mean* distortions play roles in this analysis:

- w_t^* : Piazzesi, Salomao, and Schneider’s mistaken agent;
- $\bar{w}_t$: Szőke’s particular LRR parametric concern;
- $\tilde{w}_t$: Szőke’s worst-case model.

11. WHO CARES?

I have used likelihood ratios to organize three distinct explanations for observed risk prices—risk aversion, mistaken beliefs, and fears of model misspecification—and have shown how the third, properly formalized, can account for puzzling term-structure dynamics. The question posed in my title matters because the same empirical regularities can call for different responses depending on which interpretation is correct. Joint probability distributions of asset prices and macroeconomic fundamentals therefore matter throughout macroeconomics.

• **Costs of aggregate fluctuations.** Welfare assessments of business cycles depend sensitively on how risks are priced. Lucas (1987) calculated the welfare cost of business cycles at roughly 0.05% of lifetime consumption under standard risk aversion. But if the representative agent fears model misspecification in the sense formalized above, that cost can rise by one or two orders of magnitude (see Hansen et al. (1999)), because the agent evaluates fluctuations under a pessimistic probability model that amplifies downside risk.¹¹

• **Consumption Euler equations.** The “New Keynesian IS curve” is a log-linearized consumption Euler equation whose risk adjustments are controlled by the stochastic discount factor. Whether those adjustments reflect risk aversion, wrong beliefs, or model uncertainty affects how monetary policy transmits to real activity through intertemporal substitution.

• **Optimal taxation and government debt management.** Government bond prices embed risk prices whose state dependence matters for optimal fiscal policy.¹²

• **Central bank expectations management.** Forward guidance works by shifting the term structure, an exercise whose effects depend on the same likelihood ratios studied here.

¹¹See Barillas et al. (2009) for a quantitative exploration of this idea.

¹²See Karantounias (2013) and Ferrier and Karantounias (2019).

• **Long-run risk and secular stagnation.** The Bansal–Yaron long-run risk hypothesis is difficult to detect statistically, yet an agent who fears it in the sense formalized above may behave very differently than one who does not.¹³

Understanding whether observed asset prices reflect risk aversion, mistaken beliefs, or fears of model misspecification—and quantifying each component—thus matters for both positive and normative macroeconomics. Dave Backus spent a lot of time teaching us how to improve connections between theory and evidence. I wish he were still here. In some important ways, he still is.

APPENDIX A

A.1. MULTIPLIER PREFERENCES

The multiplier preference version of the dubious agent’s problem is:

$$W(x_0, c_0 \mid \theta) := \min_{\{m_{t+1}\}_{t \geq 0}} E \left[\sum_{t=0}^{\infty} \beta^t M_t (c_t + \theta m_{t+1} \log m_{t+1}) \mid x_0, c_0 \right], \quad (\text{A.1})$$

$$M_{t+1} = M_t m_{t+1}, \quad E[m_{t+1} \mid x_t, c_t] = 1, \quad M_0 = 1.$$

The recursive formulation is:

$$\begin{aligned} W(x_t, c_t \mid \theta) &= c_t + \min_{m_{t+1}} E[m_{t+1} [\beta W(x_{t+1}, c_{t+1}) + \theta \log m_{t+1}] \mid x_t, c_t] \\ &= c_t - \theta \log E \left[\exp \left(-\frac{\beta W(x_{t+1}, c_{t+1})}{\theta} \right) \mid x_t, c_t \right] \\ &=: c_t + T_t[\beta W(x_{t+1}, c_{t+1})], \end{aligned} \quad (\text{A.2})$$

where the right-hand side is attained by

$$m_{t+1}^* \propto \exp \left(-\frac{\beta W(x_{t+1}, c_{t+1})}{\theta} \right).$$

Relationship between multiplier and constraint problems.. By Lagrange multiplier theory,

$$W(x_t, c_t \mid \tilde{\theta}) = J(x_t, c_t \mid \eta) + \tilde{\theta} \eta.$$

¹³See Hansen and Sargent (2010).

REFERENCES

- Backus, David K. and Stanley E. Zin, 1994. Reverse engineering the yield curve. Working Paper 4676, National Bureau of Economic Research.
- Backus, David K., Silverio Foresi, and Chris I. Telmer, 2001. Affine term structure models and the forward premium anomaly. *Journal of Finance* **56**, 279–304.
- Bansal, Ravi and Amir Yaron, 2004. Risks for the long run: A potential resolution of asset pricing puzzles. *Journal of Finance* **59**, 1481–1509.
- Barillas, Francisco, Lars Peter Hansen, and Thomas J. Sargent, 2009. Doubts or variability? *Journal of Economic Theory* **144**(6), 2388–2418.
- Bierwag, Gerald O. and M. A. Grove, 1967. A model of the term structure of interest rates. *The Review of Economics and Statistics* **49**(1), 50–62.
- Chernoff, Herman, 1952. A measure of asymptotic efficiency for tests of a hypothesis based on the sum of observations. *Annals of Mathematical Statistics* **23**(4), 493–507.
- Cox, John C., Jonathan E. Ingersoll, and Stephen A. Ross, 1985. A theory of the term structure of interest rates. *Econometrica* **53**, 385–407.
- Dai, Qiang and Kenneth J. Singleton, 2000. Specification analysis of affine term structure models. *Journal of Finance* **55**, 1943–1978.
- Ellison, Martin and Thomas J. Sargent, 2012. A defense of the FOMC. *International Economic Review* **53**(4), 1047–1065.
- Ferriere, Axelle and Anastasios G. Karantounias, 2019. Fiscal austerity in ambiguous times. *American Economic Journal: Macroeconomics* **11**(1), 89–131.
- Hansen, Lars Peter and Thomas J. Sargent, 1981. Exact linear rational expectations models: Specification and estimation. Staff Report 71, Federal Reserve Bank of Minneapolis.
- Hansen, Lars Peter and Thomas J. Sargent, 1991. Exact linear rational expectations models: Specification and estimation. In Lars Peter Hansen and Thomas J. Sargent, editors, *Rational Expectations Econometrics*. Westview Press, Boulder, CO.
- Hansen, Lars Peter and Thomas J. Sargent, 1995. Discounted linear exponential quadratic gaussian control. *IEEE Transactions on Automatic Control* **40**(5), 968–971.
- Hansen, Lars Peter and Thomas J. Sargent, 2001. Robust control and model uncertainty. *American Economic Review* **91**, 60–66.
- Hansen, Lars Peter and Thomas J. Sargent, 2010. Fragile beliefs and the price of uncertainty. *Quantitative Economics* **1**(1), 129–162.
- Hansen, Lars Peter, Thomas J. Sargent, and Thomas D. Tallarini, 1999. Robust permanent income and pricing. *Review of Economic Studies* **66**, 873–907.
- Hansen, Lars Peter, Bálint Szőke, Lloyd S. Han, and Thomas J. Sargent, 2020. Twisted probabilities, uncertainty, and prices. *Journal of Econometrics* **216**, 151–174.
- Hicks, John R., 1939. *Value and Capital*. Clarendon Press, Oxford.
- Karantounias, Anastasios G., 2013. Managing pessimistic expectations and fiscal policy. *Theoretical Economics* **8**(1), 193–231.
- Lucas, Robert E., 1978. Asset prices in an exchange economy. *Econometrica* **46**, 1429–1445.
- Lucas, Robert E., 1987. *Models of Business Cycles*. Basil Blackwell, Oxford.
- Lucas, Robert E. and Thomas J. Sargent, editors, 1981. *Rational Expectations and Econometric Practice*. University of Minnesota Press, Minneapolis.

- Maenhout, Pascal J., 2004. Robust portfolio rules and asset pricing. *Review of Financial Studies* **17**, 951–983.
- Meiselman, David I., 1961. The term structure of interest rates. PhD thesis, University of Chicago.
- Newman, Charles M. and Barton W. Stuck, 1979. Chernoff bounds for discriminating between two markov processes. *Stochastics* **2(1–4)**, 139–153.
- Piazzesi, Monika and Martin Schneider, 2007. Equilibrium yield curves. *NBER Macroeconomics Annual* **21**, 389–442.
- Piazzesi, Monika, Juliana Salomao, and Martin Schneider, 2015. Trend and cycle in bond premia. Working Paper, Stanford University, March 2015. URL <https://web.stanford.edu/~piazzesi/trendcycle.pdf>.
- Sargent, Thomas J., 1972. Rational expectations and the term structure of interest rates. *Journal of Money, Credit and Banking* **4**, 74–97.
- Sargent, Thomas J., 1979. A note on maximum likelihood estimation of the rational expectations model of the term structure. *Journal of Monetary Economics* **5**, 133–143.
- Sargent, Thomas J., 2008. Evolution and intelligent design. *American Economic Review* **98**, 5–37.
- Shiller, Robert J., 1979. The volatility of long-term interest rates and expectations models of the term structure. *Journal of Political Economy* **87**, 1190–1219.
- Szőke, Bálint, 2022. Estimating robustness. *Journal of Economic Theory* **199**, 105225.