



Robust identification of investor beliefs

Xiaohong Chen^{a,1}, Lars Peter Hansen^{b,c,d,1,2}, and Peter G. Hansen^{e,1}

^aDepartment of Economics, Yale University, New Haven, CT 06520; ^bDepartment of Economics, University of Chicago, Chicago, IL 60637; ^cDepartment of Statistics, University of Chicago, Chicago, IL 60637; ^dBooth School of Business, University of Chicago, Chicago, IL 60637; and ^eMIT Sloan School of Management, Cambridge, MA 02142

Contributed by Lars Peter Hansen, October 26, 2020 (sent for review September 24, 2020; reviewed by Anmol Bhandari and Monika Piazzesi)

This paper develops a method informed by data and models to recover information about investor beliefs. Our approach uses information embedded in forward-looking asset prices in conjunction with asset pricing models. We step back from presuming rational expectations and entertain potential belief distortions bounded by a statistical measure of discrepancy. Additionally, our method allows for the direct use of sparse survey evidence to make these bounds more informative. Within our framework, market-implied beliefs may differ from those implied by rational expectations due to behavioral/psychological biases of investors, ambiguity aversion, or omitted permanent components to valuation. Formally, we represent evidence about investor beliefs using a nonlinear expectation function deduced using model-implied moment conditions and bounds on statistical divergence. We illustrate our method with a prototypical example from macro-finance using asset market data to infer belief restrictions for macroeconomic growth rates.

subjective beliefs | asset pricing | intertemporal divergence | bounded rationality | large deviation theory

Prices in asset markets reflect a combination of investor beliefs and their risk preferences. Researchers, as well as policymakers, look to asset market data as a barometer of public beliefs. Derivative claims prices potentially enrich what we can infer about conditional probability distributions of future events, but events of interest often entail components of macroeconomic uncertainty for which there will be a paucity of information along some dimensions. Moreover, since a central tenet of asset pricing is that investors must be compensated for exposure to macroeconomic shocks that are not diversifiable, beliefs about future macroeconomic performance are of paramount importance to understanding asset prices.

To disentangle the contributions of risk aversion from beliefs, many empirical approaches in the last few decades have focused on models of investor preferences by assuming rational expectations. Using the implied moment conditions of the investor's portfolio choice problem in conjunction with this restriction gives a directly applicable and tractable approach for estimating and testing alternative model specifications. This approach, however, often leads to risk prices in some time periods that are attributed to an arguably extreme level of investor risk aversion or a rejection of the model. Risk aversion and belief formulation are intertwined. Rational expectation as a model of belief formation is meant to be a simplifying approximation. It can serve as an elegant and powerful modeling choice when appropriate. In a complex environment, however, it can be challenging to make statistical inferences pertinent to forward-looking decision making. In such settings, we find the presumption that model-dwelling investors and entrepreneurs know the true data-generating process to be tenuous and worthy of relaxation. We are not alone in this view.

Some researchers have explored mechanisms that could account for this evidence via a different channel, namely beliefs which differ from rational expectations. It is sometimes argued, but typically not justified formally, that these alternatives are small departures from rational expectations. These "belief distortions" relative to rational expectations alternatively could

reflect the lack of investor confidence about the assignment of probabilities to future events. This has been modeled and captured formally as ambiguity aversion or concerns about model misspecification.

This paper proposes a formal methodology for analyzing models that imply conditional moment restrictions where the restrictions are presumed to hold under a distorted probability measure. It extends a previous econometrics literature that represents the statistical implications of asset pricing implications as conditional moment restrictions under rational expectations. Rational expectations on the part of individuals or enterprises can be motivated by a law of large numbers approximation used to pin down the beliefs of these economic agents "inside the model." Once we relax the rational expectations, there are typically many choices of investor beliefs that satisfy the conditional moment restrictions. Rather than imposing a specific alternative to rational expectations, we restrict the family of investor probabilities to satisfy discrepancy bounds, which gives us a way of relaxing the rational expectations hypothesis based on pushing back from a law of large numbers approximation. We then use the conditional moment restrictions along with statistical discrepancy bounds, to characterize families of probabilities that satisfy the conditional moment restrictions.

Our approach provides a version of "bounded rationality" when assessing empirical evidence. Not only are the bounds we deduce of direct interest; they also can be used as diagnostics for specific models of belief distortions. Additionally, we show how to include survey data on subjective beliefs. Such data are typically sparse and not sufficient to pin down full probabilistic characterizations of beliefs. Given these data limitations, we

Significance

Prices in asset markets reflect a combination of investor beliefs and their risk preferences. Researchers, as well as policymakers, look to asset market data as a barometer of public beliefs. Such data are informative because investors must be compensated for exposure to macroeconomic shocks, and thus beliefs about future macroeconomic performance are encoded in asset prices. This paper develops a method informed by data and models to recover information about investor beliefs. Our approach uses information embedded in forward-looking asset prices in conjunction with asset pricing models. We entertain families of potential belief distortions bounded by a statistical measure of divergence. Additionally, our method allows for the direct use of sparse survey evidence to make these bounds more informative.

Author contributions: X.C., L.P.H., and P.G.H. designed research and wrote the paper.

Reviewers: A.B., University of Minnesota; and M.P., Stanford University.

This open access article is distributed under [Creative Commons Attribution-NonCommercial-NoDerivatives License 4.0 \(CC BY-NC-ND\)](https://creativecommons.org/licenses/by-nc-nd/4.0/).

The authors declare no competing interest.

¹X.C., L.P.H., and P.G.H. contributed equally to this work.

²To whom correspondence may be addressed. Email: lhansen@uchicago.edu.

This article contains supporting information online at <https://www.pnas.org/lookup/suppl/doi:10.1073/pnas.2019910117/-DCSupplemental>.

First published December 14, 2020.

bound probabilities even when certain features of the data may be known through direct evidence.

A common way to represent a probability distribution of a random vector is through how it assigns expectations to functions of that random vector. Since we have multiple probability distributions in play, we represent our bounds by building what is called a “nonlinear expectation” that minimizes expectations over members of the family of probability distortions that we identify. This gives us a formal way to characterize properties of probability distributions that are consistent with model-implied conditional moment restrictions. The choice of minimization is essentially a normalization as we bound expectations of functions of observable variables as well as the negative of such functions. Given the flexibility in our choice of what functions of observables we use when forming expectation bounds, our analysis provides a rich characterization of implications for the belief distortions.

While we use asset pricing applications for motivation, our analysis is more generally applicable to economic models with forward-looking agents. These agents may be groups of individuals making investment or portfolio choices when facing production or financial opportunities that are exposed to uncertainty in different ways. Alternatively, they may be forward-looking enterprises, making decisions today that have important consequences for the future.

In summary, our methodology gives a way to extract information on investor beliefs from asset market and survey data pertinent for both external analysts and policymakers who are looking for evidence to gauge private sector sentiments. In addition, our computations provide revealing diagnostics for model builders that embrace specific formulations of belief distortions as is common in the behavioral economics and finance literatures.

Literature Review. There is a long intellectual history exploring the impact of expectations on investment decisions. As was well appreciated by economists such as in refs. (1–3), investment decisions are in part based on people’s views of the future. Alternative approaches for modeling expectations of economic actors were suggested including static expectations, ref. 4’s extrapolative expectations, ref. 5’s adaptive expectations, or appeals to data on beliefs; but these approaches leave open how to proceed when using dynamic economic models to assess hypothetical policy interventions. A productive approach to this modeling challenge has been to add the hypothesis of rational expectations. Motivated by long histories of data, this hypothesis pins down beliefs by equating the expectations of agents inside the model to those implied by the data-generating distribution. This approach to completing the specification of a stochastic equilibrium model was initiated by ref. 6 and developed fully in ref. 7.

Recently there has been a renewed interest in alternative belief distortions within the asset pricing literature. See, for example, refs. 8–12. Relatedly, since survey evidence on investor beliefs is typically rather sparse and not able to produce entire predictive distributions, refs. 13 and 14 fitted time series models to the observed beliefs that can be distinct from the actual data evolution. Our approach is different from these literatures, but complementary to them. We focus on the construction of the implied bounds for expectations for functions of the stochastic process of interest that could provide empirical targets of tests of parametric models of subjective beliefs fitted to time series.

There is similarity in motivation and overlaps in the methods we use to the study of robust optimization (see, for instance, refs. 15 and 16) and robust Markov chain modeling (see, for instance, ref. 17). But our aim is different. While the robust optimization research features decision makers that confront multiple probabilities, our perspective is that of analyst seeking information about the beliefs of economic decision makers from

observed financial market data or survey data through the lens of a dynamic economic model.

Refs. 18 and 19 describe and implement econometric methods for confronting conditioning information under correct model specification under rational expectations, within a generalized method of moments framework. Refs. 20 and 21 give extensions of the measure of model misspecification proposed by ref. 22 to accommodate conditioning information. Similarly, the models we consider are misspecified under rational expectations. This misspecification is induced as it is in precursors (23, 24) by investor belief distortions (these two papers abstract from the role of conditioning information). Our innovation is to propose and justify a dynamic formulation with belief uncertainty that 1) accommodates conditioning and 2) uses the recursive structure of multiperiod likelihoods to characterize families of beliefs that are consistent with alternative divergence thresholds.

Outline of the Paper. 1. *Asset Pricing with Distorted Beliefs* introduces the framework we use for moment restrictions implied by an asset pricing model. 2. *Data Generation and Probability Divergence* specifies the probabilistic environment that underlies our computations and gives a dynamic version of the divergence with a built-in recursive structure. 3. *Moment Bounds* presents and justifies our recursive formulation of the functional equation used to compute the bounds along with some special cases that are of particular interest. This section provides a more complete characterization of the solution for the familiar relative entropy divergence and discusses the relation to results from large deviation theory. Finally, 3. *Moment Bounds* characterizes a nonlinear expectation as a way to represent the bounds on the subjective probabilities. 4. *Illustration* presents an empirical illustration of our methodology. 5. *Bounding Other Probabilities* shows how to apply our approach to extract information about the one-period, risk neutral measure and the long-term counterpart without assuming the existence of data on a complete set of Arrow–Debreu securities. Both of these probability measures are of interest in their own right. 6. *Conclusions* concludes.

1. Asset Pricing with Distorted Beliefs

In standard economic applications, moment conditions are justified via an assumption of rational expectations. This assumption equates population expectations with those used by economic agents inside the model. These expectations are therefore presumed to be revealed by the law of large numbers applied to time series data.

Let (Ω, \mathcal{G}, P) denote the underlying probability space and $\mathcal{I} \subset \mathcal{G}$ represent information available to investors. The original moment equations under rational expectations are of the form

$$\mathbb{E}[f(X, \theta) | \mathcal{I}] = 0, \quad [1]$$

where the function, f , captures the parameter dependence, θ , of either the payoff or the stochastic discount factor along with a random vector, X , of variables observed by the econometrician and used to construct the payoffs, prices, and the stochastic discount factor.

A typical asset pricing example is as follows: Let R denote an n -dimensional vector of gross returns corresponding to payoffs on financial or physical assets over some investment horizon, let S denote the corresponding stochastic discount factor for this horizon, and let \mathcal{I} denote the investor information set. The stochastic nature of the stochastic discount factor captures the market compensations for exposure to uncertainty.

The underlying asset pricing equation is

$$\mathbb{E}[SR - \mathbf{1}_n | \mathcal{I}] = 0,$$

where $\mathbf{1}_n$ is an n -dimensional vector of ones. Both the stochastic discount factor and the return vector R may depend on unknown

parameters, giving rise to Eq. 1. The vector of returns can be parameter dependent when the investment is in a physical asset with an unobserved return. While we posed this as an asset-pricing relation, restrictions of the same form can be derived from investment or asset demand equations with potentially differential exposures to uncertainty. Even in such models, there is a counterpart to stochastic discount factors that are perhaps different across agent types. More generally, these equations feature the forward-looking behavior of individuals or enterprises as they depend on the perceptions of the future.

A. Market Beliefs. We allow for the beliefs that are revealed by the market to differ from the rational expectations beliefs implied by (infinite) histories of data. We represent what we call “market beliefs” by introducing a positive random variable N with a unit conditional expectation. Thus, we consider moment restrictions of the form

$$\mathbb{E}[Nf(X, \theta) | \mathcal{J}] = 0. \quad [2]$$

The random variable N provides a flexible change in the probability measure and is sometimes referred to as a Radon–Nikodym derivative or a likelihood ratio. The dependence of N on random variables not in the information captured by \mathcal{J} defines a relative density that informs how rational expectations are altered by market beliefs.

By using N to represent a potential market belief, we require that any event that depends on the realization of X and has conditional probability measure zero under the rational expectation distribution will continue to have conditional probability zero under this change in distribution. We will, however, sometimes allow for N to be zero on events that have positive conditional probability under the original measure, at least as a possible limiting case. Such events would be a complete surprise to market participants.

The introduction of N into the analysis is seemingly an innocuous change in formulating the observable implications. But it has rather dramatic consequences for econometric analyses. Specifically, we consider estimation environments in which a researcher uses data only on a limited set of asset returns. With observations on a complete set of asset returns and a prespecified stochastic discount factor, we could identify uniquely the belief distortion, N . Given our interest in macroeconomic risk compensation, we presume a more modest set of data is available to use as empirical inputs. As a consequence, even with a known stochastic discount factor, there may be an extensive family of beliefs that is consistent with the underlying pricing restrictions expressed as conditional moments. To elaborate, we suppose that Eq. 1 may not have solutions for any θ under rational expectations. Once we relax rational expectations by introducing N , Eq. 2 will in general be satisfied for an infinite-dimensional set of possible N s for each value of θ . Thus, the parameter vector θ and the corresponding N fail to be point identified in a rather spectacular way. The set of N s associated with a given value of θ will be of particular interest to us.

Given our interest in the set of N s, we are led to deduce implied bounds on moments. Consider, for instance, the relation

$$\mathbb{E}(NSR | \mathcal{J}) - \mathbf{1}_n = 0. \quad [3]$$

Then the proportional risk premia from the perspective of the altered probability are

$$\log \mathbb{E}(NR | \mathcal{J}) + \mathbf{1}_n \log \mathbb{E}(NS | \mathcal{J}).$$

The first term is the logarithm altered expectation of R and the second term is the negative of the logarithm of the risk-free return. Our methods allow us to compare the rational expectations version of the risk compensations to bounds on these proportional compensations as implied by market data.

Two classes of asset pricing models that have received considerable attention provide motivation for our analysis. One class allows for subjective beliefs to differ from those implied by rational expectations because of “market psychology.” Alternative models of expectations from behavioral finance imply alternative specifications of N . Another class includes models with investors that are ambiguity averse. Associated with many such models are belief specifications that emerge as altered probabilities encoded in asset prices. These distortions reflect some form of caution, depending on modeling details. While both literatures derive counterparts to N , our methods put very modest structure on the beliefs beyond potentially small statistical departures from rational expectations and can provide revealing diagnostics for assessing models that impose specific distortions in expectations.

Since the form of our pricing equation applies to investment or asset demand equations, there is a direct extension of our analysis to the case in which there are distinct classes of economic agents with potentially different subjective beliefs or concerns about ambiguity aversion.

B. Incorporating Survey Evidence. When constructing our moment conditions, we could also include direct data on investor expectations to help inform the direction and magnitude of the subjective belief distortion from historical evidence. This would entail augmenting the moment conditions used to constrain beliefs to include the variable being forecasted minus the observed forecast all scaled by N .

Suppose we have data D on beliefs that reflect subjective expectations of \tilde{X} . These data could include survey responses or analyst forecasts. We may include this in our analysis by imposing the conditional moment condition:

$$\mathbb{E}(N\tilde{X} | \mathcal{J}) = D. \quad [4]$$

In words, this restriction says that D is the best forecast of \tilde{X} under the subjective belief measure. Note that we can incorporate probabilistic forecasts into our framework by letting \tilde{X} be an indicator function.*

Remark 1.1: Time series of survey data are often shorter relative to data on returns or macroeconomic variables. This can be accommodated in our framework provided that there is sufficient time series variation for these data to add nontrivial incremental information to the analysis.

2. Data Generation and Probability Divergence

In this section we construct and use the dynamic counterpart to the statistical divergence measure. We focus initially on relative entropy as a measure of statistical divergence, a measure which frequently arises in the analysis of large deviations of stochastic processes with temporal dependence (see, for instance, refs. 27 or 28). While we use relative entropy as a starting point, we go much farther by extending the so-called ϕ divergence measures in ways that have a recursive structure that is very similar to that of relative entropy.

While the applications that interest us use Markov formulations, we relax this assumption to entertain non-Markov distortions. For this reason, we initially consider a stationary and

*See ref. 25 and the published comments for an overview and discussion of the use of survey data in macroeconomics and ref. 26 for a probe into the impact of heterogeneity in the study of overreaction.

ergodic formulation that nests stationary, ergodic Markov processes. This is the same environment used by ref. 29 to study bounds on statistical efficiency and ref. 30 to study testable implications of asset pricing models, both in the presence of conditional moment restrictions. The ref. 30 analysis imposes rational expectations in contrast to the analysis in this paper.

We start with a baseline probability triple $(\Omega, \mathfrak{G}, \mathbb{P})$ and a measurable one-to-one transformation \mathbb{U} which is measure preserving and ergodic under \mathbb{P} . We use \mathbb{U} to construct stochastic processes and filtrations.[†]

Let $\mathfrak{J}_0 \subset \mathfrak{G}$ depict information available at date zero. We use the transformation \mathbb{U} to capture the information available at future dates via the recursion

$$\mathfrak{J}_t = \{\Lambda \in \mathfrak{G} : \mathbb{U}^{-1}\Lambda \in \mathfrak{J}_{t-1}\} = \{\Lambda \in \mathfrak{G} : \mathbb{U}^{-t}\Lambda \in \mathfrak{J}_0\}.$$

We presume that information accumulates,

$$\mathfrak{J}_t \subset \mathfrak{J}_{t+1},$$

which in turn implies that $\{\mathfrak{J}_t : -\infty < t < +\infty\}$ is a filtration. Similarly, for any random variable B_0 that is \mathfrak{J}_0 measurable, we form B_t recursively:

$$B_t(\omega) = B_{t-1}[\mathbb{U}(\omega)] = B_0[\mathbb{U}^t(\omega)].$$

Thus for each initial random vector B_0 , there is a corresponding stochastic process $\{B_t : t \geq 0\}$ that is adapted to the filtration $\{\mathfrak{J}_t : t \geq 0\}$. Since \mathbb{U} is measure preserving, the process $\{B_t : t \geq 0\}$ is stationary.

A. Alternative Probabilities. In what follows, we hold fixed the transformation \mathbb{U} while considering alternative probability measures. Let \mathbb{Q} denote an alternative probability distribution on (Ω, \mathfrak{G}) that is measure preserving and ergodic, and let \mathbb{Q}_t be the restriction of \mathbb{Q} to \mathfrak{J}_t . We consider only \mathbb{Q} s for which there exists an $N_1 \geq 0$ that is \mathfrak{J}_1 measurable and satisfies

$$\int B_1 d\mathbb{Q}_1 = \int \mathbb{E}(N_1 B_1 | \mathfrak{J}_0) d\mathbb{Q}_0 \quad [5]$$

for all bounded \mathfrak{J}_1 measurable random variables B_1 . This N_1 necessarily satisfies $\mathbb{E}(N_1 | \mathfrak{J}_0) = 1$. When \mathbb{Q} is distinct from \mathbb{P} , there is a bounded random variable for which the \mathbb{Q} and \mathbb{P} distributions differ. Because both are ergodic, the law of large numbers is applicable to both with distinct almost sure limit points. For the purposes of this analysis, the probability measure \mathbb{Q} encodes the limits implied by the law of large numbers.[‡]

Form the product

$$M_T = \prod_{t=1}^T N_t. \quad [6]$$

Then under \mathbb{Q} , the date- T conditional expectation of a bounded, \mathfrak{J}_T random variable B_T is

$$\mathbb{E}(M_T B_T | \mathfrak{J}_0).$$

We think of M_T as a relative likelihood between two models over horizon T constructed recursively through the familiar likelihood factorization. We further restrict \mathbb{Q} to imply stochastic stability:[§]

[†]A common specification of \mathbb{U} is the shift transformation applied to the space of infinite sequences of vectors of real numbers.

[‡]From a measure-theoretic perspective, \mathbb{Q} and \mathbb{P} cannot be equivalent. There exist events based on limits for which \mathbb{Q} assigns probability one and \mathbb{P} assigns probability zero and conversely.

[§]Stochastic stability as defined here is satisfied when the process is beta mixing (or absolutely regular); see, e.g., theorem 3.29 in ref. 31.

Definition 2.1: We say that \mathbb{Q} induces stochastic stability if for any B_0 that is \mathfrak{J}_0 measurable and satisfies $\int |B_0| d\mathbb{Q}_0 < \infty$,

$$\lim_{T \rightarrow \infty} \mathbb{E}(M_T B_T | \mathfrak{J}_0) = \int B_0 d\mathbb{Q}_0.$$

Definition 2.2: The set \mathcal{N} contains all N_1 s for which there is a corresponding probability \mathbb{Q} satisfying Eq. 5 and is stochastically stable.

We presume that $N_1 = 1$ is in this set and hence \mathbb{P} is stochastically stable.

B. Intertemporal Divergences. First consider the Kullback–Leibler divergence. Represent the expected log-likelihood ratio as a sum of contributions for each date by using the recursive structure of a relative likelihood:

$$\mathbb{E}(M_T \log M_T | \mathfrak{J}_0) = \mathbb{E}\left(M_T \sum_{t=1}^T \log N_t \mid \mathfrak{J}_0\right) \geq 0.$$

Dividing by T and taking limits gives

$$\begin{aligned} \mathcal{R}(N_1) &= \lim_{T \rightarrow \infty} \frac{1}{T} \mathbb{E}(M_T \log M_T | \mathfrak{J}_0) \\ &= \lim_{T \rightarrow \infty} \frac{1}{T} \mathbb{E}\left(M_T \sum_{t=1}^T \log N_t \mid \mathfrak{J}_0\right) \\ &= \int \mathbb{E}(N_1 \log N_1 | \mathfrak{J}_0) d\mathbb{Q}_0, \end{aligned}$$

which is the measure of relative entropy that we will use in our analysis. Note that there is an explicit connection between N_1 and \mathbb{Q}_0 , which gives rise to a restriction that we impose when computing bounds.

Remark 2.3: The relative entropy measure $\mathcal{R}(N_1)$ is the discrete-time analog to the relative entropy measure that is used in the Donsker–Varadhan large deviation theory applied to Markov processes [see refs. 28, 32, and 33 (chap. 3)].

Finite relative entropy restricts substantially the tail behavior of the probability distributions. For this reason we consider other divergences, but modified to exploit the recursive structure implied by the likelihood factorization. We use the conditional version of what is commonly called a ϕ divergence as an important building block:

$$\mathbb{E}[\phi(N_t) | \mathfrak{J}_{t-1}]$$

for a strictly convex function ϕ defined on $(0, \infty)$ with $\phi(1) = 0$.[¶] There is an equivalent way to represent this divergence that we use. Construct a function ψ such that

$$n\psi\left(\frac{1}{n}\right) = \phi(n).$$

It may be shown that ψ is also strictly convex with $\psi(1) = 0$. By design,

$$\mathbb{E}[\phi(N_t) | \mathfrak{J}_{t-1}] = \mathbb{E}\left[N_t \psi\left(\frac{1}{N_t}\right) \mid \mathfrak{J}_{t-1}\right]. \quad [7]$$

On the right-hand side, we use the conditional probability associated with N_t to compute expectations. The random variable $\frac{1}{N_t}$ is the Radon–Nikodym derivative of the baseline conditional

[¶]For some ϕ divergences the domain can be extended to include zero.

probability with respect to the N_t conditional probability distribution. The function ψ in Eq. 7 plays a role analogous to ϕ once we change the probability we use in the expectation.

We use ψ to represent an intertemporal divergence measure:

$$\begin{aligned} \mathcal{R}(N_1) &= \lim_{T \rightarrow \infty} \frac{1}{T} \sum_{t=1}^T \mathbb{E} (M_{t-1} \mathbb{E} [\phi(N_t) | \mathcal{J}_{t-1}] | \mathcal{J}_0) \\ &= \lim_{T \rightarrow \infty} \frac{1}{T} \sum_{t=1}^T \mathbb{E} \left[M_t \psi \left(\frac{1}{N_t} \right) \middle| \mathcal{J}_0 \right] \\ &= \lim_{T \rightarrow \infty} \frac{1}{T} \mathbb{E} \left[M_T \sum_{t=1}^T \psi \left(\frac{1}{N_t} \right) \middle| \mathcal{J}_0 \right]. \end{aligned} \quad [8]$$

The limiting version of this measure as implied by the law of large numbers for stationary, ergodic processes is

$$\int \mathbb{E} \left[N_1 \psi \left(\frac{1}{N_1} \right) \middle| \mathcal{J}_0 \right] d\mathbf{Q}_0 = \int \mathbb{E} [\phi(N_1) | \mathcal{J}_0] d\mathbf{Q}_0.$$

Note that we use the conditional version of what is commonly called a ϕ divergence measure averaged using the altered stationary probability \mathbf{Q} . This coincides with the relative entropy divergence when $\phi(n) = n \log n$. *SI Appendix, section A* shows that this intertemporal divergence extension of ϕ divergence is convex.

Alternatively, there has been considerable interest in members of the Wasserstein class of divergences, including in the study of robust Markov chains and machine learning. Wasserstein divergences could also be used as conditional divergences that are averaged using the altered stationary probability \mathbf{Q} .[#]

Remark 2.4: Empirical likelihood methods (35, 36) use a divergence measure $\phi(n) = -\log n$ for independent and identically distributed data. The corresponding intertemporal divergence measure is

$$\mathbb{E} (-\log N_1). \quad [9]$$

Note that this measure uses the original baseline probability measure implied by \mathbf{P}_0 and not \mathbf{Q}_0 to average over the conditioning information. Were we to average using \mathbf{Q}_0 in the divergence, then our analysis would apply. Elsewhere, we and others have argued that this strictly decreasing ϕ is not well suited for detecting departures from baseline probabilities restricted by moment conditions.

Remark 2.5: Using positive random variables, M_T , to depict alternative probabilities for date- T events imposes absolute continuity (conditioned on date-zero information). As we already noted, this same absolute continuity will not be true over the infinite future, however. Our division by T when constructing intertemporal divergence Eq. 8 allows for the altered probability to have different limits under the law of large numbers.

Remark 2.6: Our analysis assumes that the underlying processes are stationary under the alternative \mathbf{Q} s. However, the results we obtain will still apply under certain classes of nonstationarity processes. While we restrict \mathbf{Q} to be measure preserving and ergodic, any probability measure that is absolutely continuous with respect to \mathbf{Q} will obey the same laws of large numbers even though it may not be measure preserving. Such measures may also imply the same intertemporal divergence measure Eq. 8. Moreover, specific examples that could be of interest are probabilities associated with Markov processes that eventually “escape” from their dependence on the initialization. While we

do not explore such cases formally, these observations suggest that the bounds that we compute could be justified under an even broader set of probability measures.

3. Moment Bounds

We next present the recursive approach we use for computing probability bounds. To represent these bounds, we entertain a rich family of functions g and compute sharp lower bounds on the expectation $g(X_1)$. As a special case, $g(X_1)$ could be an indicator of an event with its expectation being the probability of that event. We will be interested in other choices of g in our empirical illustration. We compute an upper bound on $g(X_1)$ by finding the negative of the lower bound on the expectation $-g(X_1)$. Formally, we are interested in solving

$$\inf_{N_1 \in \mathcal{N}} \int \mathbb{E} \left(N_1 g(X_1) \middle| \mathcal{J}_0 \right) d\mathbf{Q}_0$$

subject to the constraints

$$\begin{aligned} \mathcal{R}(N_1) &\leq \kappa \\ \mathbb{E}[N_1 f(X_1) | \mathcal{J}_0] &= 0. \end{aligned}$$

To compute a bound for a given choice of g , we will borrow an idea from the robust control literature. See, for instance, refs. 15 and 37. We will initially solve a problem with a discrepancy penalty indexed by a parameter $\xi > 0$ and show how to solve a problem given ξ . We will then treat ξ as a Lagrange multiplier and trace out the implied discrepancies for each such ξ . In this way, ξ may be chosen to enforce the constraint $\mathcal{R}(N_1) \leq \kappa$. For notational simplicity we suppress the parameter dependence in the moments of $f(X_t)$ and $g(X_t)$.

In much of what follows, we aim to solve the following:

Problem 3.1.

$$\begin{aligned} \mu^* &= \inf_{N_1 \in \mathcal{N}} \lim_{T \rightarrow \infty} \frac{1}{T} \mathbb{E} \left(M_T \sum_{t=1}^T \left[g(X_t) + \xi \psi \left(\frac{1}{N_t} \right) \right] \middle| \mathcal{J}_0 \right) \\ &= \inf_{N_1 \in \mathcal{N}} \int \mathbb{E} \left(N_1 \left[g(X_1) + \xi \psi \left(\frac{1}{N_1} \right) \right] \middle| \mathcal{J}_0 \right) d\mathbf{Q}_0 \end{aligned}$$

subject to $\mathbb{E}[N_1 f(X_1) | \mathcal{J}_0] = 0$.

It suffices to optimize over N_1 because this choice determines the N_t for all $t \geq 1$ through the transformation \mathbb{U}^{t-1} and M_T as constructed in Eq. 6.

A. Martingale Construction.. While we solve *Problem 3.1* via recursive methods, as a precursor to this, we use a convenient martingale construction used in connection with the objective function. Suppose that for $N_1 \in \mathcal{N}$, we find a random variable v_0 such that

$$\mathbb{E} \left(N_1 \left[g(X_1) + \xi \psi \left(\frac{1}{N_1} \right) - \mu + v_1 \right] \middle| \mathcal{J}_0 \right) - v_0 = 0, \quad [10]$$

where

$$\mu = \int \mathbb{E} \left(N_1 \left[g(X_1) + \xi \psi \left(\frac{1}{N_1} \right) \right] \middle| \mathcal{J}_0 \right) d\mathbf{Q}_0,$$

and where $v_1(\omega) = v_0[\mathbb{U}(\omega)]$ and v_0 is \mathcal{J}_0 measurable. Then iterating on Eq. 10, it follows that

$$\begin{aligned} \mathbb{E} \left(M_T \sum_{t=1}^T \left[g(X_t) + \xi \psi \left(\frac{1}{N_t} \right) - \mu \right] \middle| \mathcal{J}_0 \right) \\ + \mathbb{E} (M_T v_T | \mathcal{J}_0) - v_0 = 0. \end{aligned}$$

[#]For example, see ref. 34. Ref. 34 deduces a Laplace principle for large deviation theory and other approximations with an overlapping class of intertemporal divergences. While ref. 34 poses a problem that is mathematically related, the formulation does not nest ours.

In fact,

$$\sum_{t=1}^T \left[g(X_t) + \xi \psi \left(\frac{1}{N_t} \right) \right] - T\mu + v_T - v_0$$

is a martingale under the probability measure \mathbb{Q} with expectation zero. Moreover,

$$v_0 - \int v_0 d\mathbb{Q}_0 \tag{11}$$

$$= \lim_{T \rightarrow \infty} \mathbb{E} \left(M_T \sum_{t=1}^T \left[g(X_t) + \xi \psi \left(\frac{1}{N_t} \right) - \mu \right] \middle| \mathcal{J}_0 \right)$$

provided that the almost sure limit on the right-hand side is finite. The random variable v_0 is in fact well defined only up to a constant translation. Our recursive approach will simultaneously solve for v_0 used in the martingale construction and μ computed at the minimizing solution for N_1 .

B. Recursive Formulation. We use a rearranged version of Eq. 10 when posing the recursive formulation of the optimization problem.

Problem 3.2. Find a pair (μ, v) that satisfies

$$\mu = \inf_{N_1 \in \mathcal{N}} \mathbb{E} \left(N_1 \left[g(X_1) + \xi \psi \left(\frac{1}{N_1} \right) + v_1 \right] \middle| \mathcal{J}_0 \right) - v_0$$

subject to the constraint

$$\mathbb{E} [N_1 f(X_1) | \mathcal{J}_0] = 0,$$

where $v_1(\omega) = v_0[\mathbb{U}(\omega)]$ and v_0 is \mathcal{J}_0 measurable and μ is a finite number. This optimization problem determines the constant μ and the random variable v_0 up to a translation by a constant.

This problem is recognizable as a fixed point problem in (μ, v_0) captured by the relation between v_1 and v_0 . While we posed this problem in terms of date zero and date one, given our presumed stationary data generation the problem could equivalently be stated in terms of date t and date $t + 1$ for $t > 0$. The minimization over N_1 can be solved using convex duality methods familiar from the analysis of ϕ divergence measures. We let (μ^*, v_0^*) denote the solution to Problem 3.2 and N_1^* be the corresponding minimizer. The following objects are of interest from this problem:

- the moment bound, $\int \mathbb{E} [N_1^* g(X_1) | \mathcal{J}_0] d\mathbb{Q}_0^*$;
- the corresponding conditional moment, $\mathbb{E} [N_1^* g(X_1) | \mathcal{J}_0]$;
- the implied divergence,

$$\int \mathbb{E} [\phi(N_1^*) | \mathcal{J}_0] d\mathbb{Q}_0^* = \int \mathbb{E} \left[N_1^* \psi \left(\frac{1}{N_1^*} \right) \middle| \mathcal{J}_0 \right] d\mathbb{Q}_0^*,$$

where N_1^* solves Problem 3.2 and \mathbb{Q}_0^* is an implied stationary distribution.^{||}

There are three features of Problem 3.2 that require further comment. First, the minimization problem includes “continuation-value” adjustments depicted by v_0 and its next period counterpart v_1 along with the numerical value μ . We include these adjustments to account for the fact that the choice of N_1 has implications for future time periods. Second, limit Eq. 11 is not finite for all measure-preserving and ergodic probabilities measures \mathbb{Q} . For the minimum to be attained, we

^{||}In contrast to our analysis, ref. 38 incorporates conditioning information to identify a subjective belief by minimizing the divergence given in Eq. 11 under \mathbb{P} .

presume that this limit is finite at the minimizing solution. Third, $\xi \mathbb{E} \left(N_1 \psi \left(\frac{1}{N_1} \right) \middle| \mathcal{J}_0 \right)$ acts as a per-period divergence penalty, but we may equivalently think of ξ as a Lagrange multiplier and subsequently maximize over ξ . Thus, we use ξ to index alternative problems that penalize the increment to the intertemporal divergence. The value μ^* of this objective depends on ξ , leading us to write $\mu^*(\xi)$. To impose a specific divergence constraint κ , we solve

$$\sup_{\xi > 0} \mu^*(\xi) - \xi \kappa \tag{12}$$

$$= \sup_{\xi > 0} \int \mathbb{E} \left(N_1^* \left[g(X_1) + \xi \psi \left(\frac{1}{N_1^*} \right) - \xi \kappa \right] \middle| \mathcal{J}_0 \right) d\mathbb{Q}_0^*.$$

Alternatively, we can back out the implied κ for each ξ by computing the derivative $\frac{d\mu^*}{d\xi}$ or by computing directly the divergence associated with N_1^* . To determine the ξ sensitivity, many versions of the optimization problem could be solved in parallel using convex duality and dynamic programming methods. The computed bound $\mu^*(\xi^*)$ is for the unconditional expectation of $g(X_1)$, although we find the conditional expectation, $\mathbb{E} [N_1^* g(X_1) | \mathcal{J}_0]$ that is a central part of the calculation, to be of interest in its own right. Finally, $\lim_{\xi \rightarrow \infty} \mu^*(\xi)/\xi$ reveals the minimum possible divergence subject to the conditional moment restrictions. This limit gives a lower bound on magnitude of κ used in our analysis. It can be computed directly by solving a counterpart to Problem 3.2.

C. Nonlinear Expectation. Consider now the set \mathcal{B} of bounded Borel measurable functions g to be evaluated at alternative realizations of the random vector X_1 . Given a divergence bound κ , we construct a mapping \mathbb{K} from functions g in \mathcal{B} into the real line that assigns the computed bound. Formally, define

$$\mathbb{K}(g) = \int \mathbb{E} [N_1^* g(X_1) | \mathcal{J}_0] d\mathbb{Q}_0^*,$$

where the right-hand side is computed using the value of ξ that solves Eq. 12. This mapping can be thought of as a nonlinear expectation, as formalized in the following proposition:

Proposition 3.3. The mapping $\mathbb{K} : \mathcal{B} \rightarrow \mathbb{R}$ given by $\int \mathbb{E} [N_1^* g(X_1) | \mathcal{J}_0] d\mathbb{Q}_0^*$ implied by Problem 3.2 for each $g \in \mathcal{B}$ has the following properties**:

- 1) If $g_2 \geq g_1$, then $\mathbb{K}(g_2) \geq \mathbb{K}(g_1)$;
- 2) if g is constant, then $\mathbb{K}(g) = g$;
- 3) $\mathbb{K}(rg) = r\mathbb{K}(g)$, for a scalar $r \geq 0$;
- 4) $\mathbb{K}(g_1) + \mathbb{K}(g_2) \leq \mathbb{K}(g_1 + g_2)$.

All four properties follow from the definition of \mathbb{K} . Property 4 includes an inequality instead of an equality because we compute by solving a minimization problem, and the N_1 s that solve this problem can differ depending on g .

Remark 3.4: While $\mathbb{K}(g)$ gives a lower bound on the expectation of $g(X)$, by replacing g with $-g$, we construct an upper bound on the expectation of $g(X)$. The upper bound will be given by $-\mathbb{K}(-g)$. The interval

$$[\mathbb{K}(g), -\mathbb{K}(-g)]$$

captures the set of possible values for the distorted expectation of $g(X)$ consistent with divergence less than or equal to κ .

**The first two of these properties are taken to be the definition of a nonlinear expectation by ref. 39. Properties (3) and (4) are referred to as “positive homogeneity” and “superadditivity.”

Remark 3.5: In our statement of *Problem 3.2*, we suppressed the dependence of the function f on an unknown parameter vector θ in a parameter space Θ . But many applications will necessarily include parameter uncertainty. Provided that we can compute the solution to this problem quickly using duality, we can assess parameter sensitivity by perhaps solving this problem many times in parallel. Bounds that are robust to parameter sensitivity could be obtained by minimizing over the set Θ or over a family of probability distributions over Θ .

Remark 3.6: For some applications, it is of interest to bound ratios of expectations of functions g_1 and g_2 of X_1 . Expectations conditioned on discrete events are such ratios and proportional risk compensations are logarithms of such ratios. As we elaborate in *SI Appendix, section E*, bounds of ratios can be computed by first bounding expectations of $g_1(X_1) - \zeta g_2(X_1)$ for alternative choices of the real number ζ and then searching over ζ for the smallest ratio.

D. Dual Problem. We show that the dual problem for the relative entropy divergence is equivalent to a principal eigenvalue problem. This gives a representation of the distorted measures that underlies the bounds and a revealing link to large deviation theory.

By a direct application of duality for $\phi(n) = n \log n$,

$$\begin{aligned} & \mu + v_0 \\ &= \max_{\lambda_0} -\xi \log \mathbb{E} \left(\exp \left[-\frac{1}{\xi} g(X_1) + \lambda_0 \cdot f(X_1) - \frac{1}{\xi} v_1 \right] \middle| \mathcal{J}_0 \right), \end{aligned}$$

where the random vector λ_0 is restricted to be \mathcal{J}_0 measurable and is the vector of Lagrange multipliers for the conditional moment restriction. See *SI Appendix, section C* for a more complete development of the primal and the dual problems. Let $\epsilon = \exp\left(-\frac{\mu}{\xi}\right)$ and $e_0 = \exp\left(-\frac{v_0}{\xi}\right)$. Then an equivalent statement of the dual problem is

$$\epsilon = \min_{\lambda_0} \mathbb{E} \left(\exp \left[-\frac{1}{\xi} g(X_1) + \lambda_0 \cdot f(X_1) \right] \left(\frac{e_1}{e_0} \right) \middle| \mathcal{J}_0 \right).$$

In this optimization problem λ_0 is again restricted to be a \mathcal{J}_0 measurable random vector and e_0 is restricted to be positive as is the real number ϵ .

When the state space is not discrete, this eigenvalue problem can have multiple solutions. While there could be multiple solutions to this eigenvalue problem, the next result identifies the eigenvalue of interest.

Lemma 3.7. *When there are multiple positive eigenvalue solutions for a given λ_0 , at most one of them induces a probability measure that is stochastically stable.*

See *SI Appendix, section B* for a proof.^{††}

Proposition 3.8. *Problem 3.2 with $\phi(n) = n \log n$ can be solved by finding the answer to*

$$\epsilon = \min_{\lambda_0} \mathbb{E} \left(\exp \left[-\frac{1}{\xi} g(X_1) + \lambda_0 \cdot f(X_1) \right] \left(\frac{e_1}{e_0} \right) \middle| \mathcal{J}_0 \right),$$

where

$$\begin{aligned} \mu &= -\xi \log \epsilon \\ v_0 &= -\xi \log e_0. \end{aligned}$$

In this optimization problem, the random vector λ_0 is restricted to be a \mathcal{J}_0 measurable random vector, the random variable e_0 is restricted to be \mathcal{J}_0 measurable and positive, and $e_1(\omega) = e_0[\mathbb{U}(\omega)]$ with probability one. The real number ϵ is positive. The implied solution for the probability distortion is

$$N_1^* = \frac{\exp \left[-\frac{1}{\xi} g(X_1) + \lambda_0^* \cdot f(X_1) \right] e_1^*}{\epsilon^* e_0^*},$$

where λ_0^* is the optimizing choice for λ_0 and (ϵ^*, e_0^*) are selected so that the resulting Q^* induces stochastically stability. The conditional expectation implied by the bound is

$$\mathbb{E} [N_1^* g(X_1) | \mathcal{J}_0],$$

which in turn implies a bound on the unconditional expectation equal to

$$\int \mathbb{E} [N_1^* g(X_1) | \mathcal{J}_0] dQ_0^*.$$

The implied relative entropy is

$$\int \mathbb{E} (N_1^* \log N_1^* | \mathcal{J}_0) dQ_0^*.$$

Remark 3.9: Our characterization is reminiscent of the results from large deviation theory for Markov processes. As in our analysis, large deviation theory studies an undiscounted limiting problem. See, for instance, refs. 28 and 32 for valuable treatises on large deviation theory.^{††}

A more substantive link to large deviations helps us interpret the relative entropy bounds that we input into our analysis. When using the empirical probability to detect potential departures from the baseline model, there is typically a positive probability that the empirical distribution mistakenly detects a departure. For a fixed criterion, the probability of this mistake becomes increasingly small as the sample size gets large with a well-defined rate characterized by large deviation theory. Under some additional regularity conditions, remarkably, the decay rate can be made to be arbitrarily close to the minimum relative entropy bound that we compute. Moreover, this theory computes excursions, represented probabilistically, that make the decay rate as small as possible. See *SI Appendix, section D* for an elaboration.

While we draw on insights from large deviation theory, our ultimate aim is quite different from that theory. Nevertheless, we find it revealing to compute both relative entropies and related Chernoff entropy as described, for instance, in refs. 42 and 43 as part of a sensitivity analysis.

E. Markov Specification. To proceed in a tractable way, we impose Markovian restrictions on the underlying data-generating processes. Specifically we presume that $\{X_t : t \geq 0\}$ is a time-invariant function of Markov process, appropriately restricted.

Assumption 3.10. $\{(X_t, Z_t) : t = 0, 1, \dots\}$ is a first-order Markov process for which the joint distribution of (X_{t+1}, Z_{t+1}) conditioned on (X_t, Z_t) depends only on Z_t .

Given this assumption, the $\{Z_t\}$ process by itself is a first-order Markov process. We view both X_t and Z_t as observable. The triangular structure for the dynamic evolution allows us to use a more sparse representation of the conditioning information. The alternative probabilities that we explore are not restricted to be Markov, but the solution to the minimization problem will be, for

^{††} Relatedly, ref. 40 proves a counterpart to this result for continuous-time specifications in a Markovian environment. Moreover, ref. 41 extends the ref. 40 analysis by, among other things, relaxing the Markov assumption.

^{**} In particular, the analysis in ref. 28, chaps. 7 and 8 features discrete-time Markov specifications and large sample approximation in the formulation of a Laplace principle.

reasons that are familiar from dynamic programming. With the Markov specification, we solve the recursion.

Problem 3.11. Find a pair (μ, v) that solves

$$\begin{aligned} & \mu + v(z) \\ &= \min_{N_1 \in \mathcal{N}} \mathbb{E} \left(N_1 \left[g(X_1) + \xi \psi \left(\frac{1}{N_1} \right) + v(Z_1) \right] \mid Z_0 = z \right) \end{aligned}$$

subject to the constraint

$$\mathbb{E}[N_1 f(X_1) \mid Z_0 = z] = 0.$$

This optimization problem determines the constant μ and the random variable v up to a translation by a constant.

While the primal problem “imposed” stochastic stability, it suffices to verify the stability of the process that we obtain as our candidate solution. Since it is Markovian, this restriction is satisfied when the process $\{Z_t\}$ is aperiodic and Harris recurrent.

4. Illustration

We illustrate these methods using a familiar asset pricing model with recursive utility investors as in refs. 44 and 45. While much of the asset pricing literature appeals to an arguably large risk aversion in conjunction with rational expectations when confronting data, we constrain risk aversion and instead explore belief distortions as an alternative expectation. We follow ref. 46 by allowing for market segmentation and avoid the direct use of consumption data. We then ask what implications asset market data have for predicted consumption growth rates.

Much of the macroasset pricing literature imposes risk aversion that is arguably large at least for some states of the macroeconomy. Refs. 47 and 48 give alternative rationales for substantially restricting the risk aversion coefficient. While we do not view as a settled issue what precise bounds should be imposed on risk aversion, we assume a unit risk aversion to feature the role of belief distortions when confronting asset pricing evidence. Our choice of unity is admittedly for convenience. As we will see, this choice leads to some particularly simple asset pricing implications.

Let R_t^w denote a presumed observable return on wealth. As noted by ref. 45, the one-period stochastic discount factor under rational expectations is the reciprocal of the gross return on wealth when risk aversion is unity. Thus, under distorted beliefs represented by N_t ,

$$S_t = N_t (R_t^w)^{-1},$$

where S_t is the one-period stochastic discount factor under rational expectations. We use this setup for our illustration.

As ref. 45 notes, the consumption Euler equation for an investor implies

$$\mathbb{E}[N_t \log R_t^w \mid \mathcal{J}_{t-1}] = -\log \beta + (1 - \rho) \mathbb{E}[N_t \log G_t \mid \mathcal{J}_{t-1}]$$

where G_t is the ratio of consumption growth over two adjacent time periods.^{§§} The parameter β is the subjective discount factor, and ρ is the reciprocal of the intertemporal elasticity of substitution. By deducing bounds on the left-hand side, we may infer bounds on ρ times the market expectation of consumption growth of equity market participants expressed in logarithms. Recursive utility preferences are specified in terms of continuation values that determine the rankings of prospective consumption processes. As a rough approximation, when $\rho < 1$,

^{§§}See equations 17 and 18 of ref. 45 except that we allow for more general expectations.

the wealth is positively related to the continuation value, where both are relative to current consumption. Conversely, they are negatively related when $\rho > 1$. Thus, for this model of investor preferences, whether ρ is larger or smaller than one impacts how we interpret the evidence based on conditioning information.

In our illustration, we draw on the literature that suggests returns can be predicted from dividend–price ratios. While there have been debates on how fragile this evidence is, we step aside from that discourse and take the predictability evidence at face value to illustrate our method. Given our direct use of dividend–price measures, we purposefully choose a very coarse conditioning of information and split the dividend–price ratios into three bins using the three empirical terciles. We take the dividend–price terciles to be a three-state Markov process. Dividend–price ratios are known to be persistent, and this will be evident in our calculations.^{¶¶}

We implement our approach using quarterly data from 1954 to 2016. We use the return on the CRSP (Center for the Research on Security Prices) value-weighted index to proxy for the return on wealth. For asset returns, we use the return on a 3-mo treasury bill and the three Fama–French factor excess returns. We impose moment conditions for each return implied by Eq. 3, each scaled by three indicator functions for the terciles of the dividend–price ratio, giving a total of 12 moment conditions. All returns are converted from nominal to real returns using the deflator for nondurables consumption obtained from the Bureau of Labor Statistics. We then apply the methods described in 3. *Moment Bounds* to bound functions of the return on wealth as measured by the value-weighted return.

In Fig. 1, we report the bounds on the beliefs about the expected log return, which under the assumption of the unitary risk aversion coefficient are approximately proportional to the consumption growth rate belief when the subjective discount factor β is very close to one. The conditional expectation of log returns and the unconditional counterpart are all lower than their empirical counterparts. This observation follows by comparing the ●s with the boxes in Fig. 1, where the top and bottom of the boxes are the upper and lower bounds with a relative entropy constraint imposed at a magnitude that is 20% higher than the minimum. The minimum relative entropy rate implies a half-life of about 24 quarters for reducing the probability by 50% of mistakenly rejecting the rational expectations. Increasing this by 20% reduces the half-life by the same percentage to about 20 quarters.^{##} While our choice of a 20% increase is a bit arbitrary and used for illustration purposes, it is straightforward to compute bounds with other choices of divergence thresholds. Across alternative applications, a choice of 20%, independent of magnitude of the minimal possible divergence, would be hard to defend. One nice aspect of relative entropy is that there is an explicit statistical interpretation that we find to be revealing, and thus the magnitude of the divergence has meaning.

Interestingly, it is when we condition on the low value of the dividend–price ratio that we find the box with the largest height (biggest difference between the upper and lower bounds). Also, the bounds on the unconditional distorted expectations are very similar to those we found for the low dividend–price ratio. As a robustness check, we repeated these calculations for a quadratic

^{¶¶}As an alternative starting point, we could use the regime probabilities from Markov switching models of ref. 49 as possible states along with the implied return distributions.

^{##}The Chernoff entropy at the minimum is 0.0096 with a half-life of 72 quarters. The Chernoff measure is motivated by a common decay rate imposed on type I and type II errors of testing one model against another and is expected to be considerably smaller. We computed it using the approach described in ref. 42 for Markov processes. While symmetric, this measure is less tractable to implement and not included in the family of recursive divergences that we describe. We use it merely to provide, ex post, additional information about the magnitude of the bound.




Fig. 1. Expected log market return. The \bullet s are empirical averages and the boxes give the imputed bounds when we inflated the minimum relative entropy by 20%. The minimum relative entropy is 0.0284 with a half-life of 24.4 quarters.

conditional divergence and found only very modest differences (SI Appendix, section F).

Not only are conditional means distorted, but so are the transition probabilities as reported in Table 1. While the implied stationary probabilities are fairly evenly distributed over the three dividend states, essentially by construction, the minimal entropy probabilities down-weight substantially the high dividend-price ratio state and up-weight the low dividend-price state. The high dividend-price state, in particular, has a very small stationary probability under the minimum distorted stationary distribution. Consistent with this, the transition probabilities into this state are lower under the distortion and they are higher for exiting this state. The opposite happens for transitions in and out of the low dividend-price state. Thus, a hypothetical process that behaves in accordance with the minimum entropy distorted Markov transition matrix is likely to spend substantially more time in the low expected log-return state and much less time in the high expected log-return state. When we increase the relative entropy bound by 20%, the implied distorted transition matrices are quite similar to the implied transition matrix recovered by the minimizing relative entropy and depart from the empirical transition matrix in comparable ways.

An alternative approach would be to solve static versions of our analysis for each of the three different specifications of the conditioning information. Under this approach, there would be no reason for distorting the transition probabilities as the dynamical evolution of the conditioning information is ignored. Not surprisingly, this approach does lead to notable differences in bounds. The minimized entropy over the alternative configurations of the conditioning information using the undistorted transition probabilities is 0.047 in comparison to the much smaller 0.028 that we found using our method. We view belief distortions in the transition probabilities are of particular interest in behavioral models and in models with ambiguity aversion and see this as a virtue over methods that analyze the conditional problems separately.

As we mentioned at the outset of this section, there is a substantial asset pricing literature that studies time-varying risk compensation, often appealing to high values of risk aversion. We illustrate how belief distortions can imitate large risk compensations. Thus, consider the bounds on the implied risk compensations when we restrict the risk aversion parameter to be one. We report proportional risk premium using the ex post real return on treasury bills, R^f , as our riskless benchmark in Fig. 2. To construct these results, we compute bounds on $\log \mathbb{E}R^w -$

$\log \mathbb{E}R^f$ by extending the approach of 3. *Moment Bounds* as described in SI Appendix, section E. Restricting investor risk aversion allows for belief distortions to capture the fluctuating empirical compensations for exposure to uncertainty.

Finally, since our empirical method looks at other information from two other Fama-French excess returns, we also report bounds on other risk compensations that we included in our analysis. We convert the excess returns into returns by adding the gross returns on bonds. We report these findings in SI Appendix, section E. Again, the risk compensations are greatly reduced relative to the empirical counterpart, and in most instances they are less sensitive to conditioning information.

Putting aside the empirical debate on return predictability, we see two possible conclusions from these results. One possibility is that the statistical divergence (measured as relative entropy) for the distortions is high enough to challenge a bounded rationality view of the recursive utility model with a unitary risk aversion. The other possibility is that this divergence is defensible, in which case our dynamic implementation reveals the most statistically plausible distortions on the evolution of the dividend-price ratios. It remains a judgment call as to when the resulting statistical bounds we find here are implausible. Researchers that embrace rational expectations do not consider belief distortions, while behavioral finance researchers seldom consider the implied statistical divergence of their modeled beliefs. Neither practice uses tools for assessing statistical approximation as we have done in this paper. We view these computations as providing a valuable empirical complement to the assessment of specific models of belief distortions or ambiguity aversion. Moreover, we have noted how survey evidence can be included within our framework, and we view this as a potentially valuable extension of the illustration presented here.

5. Bounding Other Probabilities

We have motivated our analysis as a method for extracting expectation bounds for subjective beliefs or restricted “worst-case” beliefs that support valuation under ambiguity aversion. These same methods provide expectation bounds for two other probability measures that are of interest in asset valuation. These measures are the one-period risk-neutral probabilities and the long-term forward probabilities. While surveys cease to provide information about these probabilities, even sparsely observed asset values, as we assume here, are revealing. We now comment on how to apply our method for both of these applications.

A. Risk-Neutral Measure. Under risk-neutral pricing, the reciprocal of the gross one-period riskless return acts as a stochastic discount factor. Thus, in this case,

$$N_t S_t = N_t (R_t^f)^{-1}.$$

Refs. 50–52 target the N_t s with the smallest relative entropy divergence to use in pricing derivative claims. We map this type of problem into our analysis by viewing the empirically relevant distribution as the “correct distribution” and the risk-neutral transformation as a way to correct for model misspecification. As in refs. 51 and 52, the measures of particular interest to us are the

Table 1. Empirical and distorted transition probabilities

	Empirical	Minimum entropy
Transition matrix	$\begin{pmatrix} 0.96 & 0.04 & 0 \\ 0.05 & 0.88 & 0.07 \\ 0 & 0.08 & 0.92 \end{pmatrix}$	$\begin{pmatrix} 0.98 & 0.02 & 0 \\ 0.08 & 0.88 & 0.04 \\ 0 & 0.17 & 0.83 \end{pmatrix}$
Stationary probabilities	$[0.42 \quad 0.31 \quad 0.27]$	$[0.76 \quad 0.20 \quad 0.04]$




Fig. 2. Proportional risk compensations computed as $\log \mathbb{E}R^w - \log \mathbb{E}R^f$ scaled to an annualized percentage. The \bullet s are the empirical averages and the boxes give the imputed bounds when we inflated the minimum relative entropy by 20%.

ones with a small divergence, although we explore more probabilities than just the N_t with the minimal divergence. While not our primary motivation, the methods we develop in this paper allow the user to obtain robust bounds on risk-neutral expectations of macroeconomic variables that incorporate information embedded in asset prices.

B. Long-Term Forward Measure. A substantively distinct, but mathematically related, literature studies the martingale decomposition of the stochastic discount factor. This decomposition expresses the stochastic discount factor as the product of a martingale component and a transitory component. The martingale component can be interpreted as a change of probability measures that imposes risk neutrality in valuation over long investment horizons. Refs. 53 and 54 show that the reciprocal of the gross holding-period return on a long-term bond is the stochastic discount factor net of a martingale component.

Since the work of ref. 54, the martingale component is referred to as the permanent component to the cumulative stochastic discount factor process. In providing a more formal mathematical characterization, refs. 40 and 41 find it more revealing to appeal to probabilistic characterization of this component. As emphasized by ref. 55, the probability measure associated with the martingale absorbs long-term risk adjustments for stochastically growing cash flows. This probability measure is the risk-neutral, forward measure that captures long-term risk–return tradeoffs.

Many structural models of asset pricing have stochastic discount factor processes with martingale components that dominate risk prices over long investment horizons. These components can reflect permanent shocks to the macroeconomy or forward-looking components to valuation. These components are present when investors have recursive utility preferences in which the intertemporal composition of risk matters or when they are averse to ambiguity in assigning probabilities to future events.

To relate this to our analysis, suppose that this martingale component is missing from the model specification. In such circumstances, ref. 53 justifies the use of the reciprocal of the gross holding-period return on a long-term bond, R_t^b , as the stochastic discount factor: $S_t = (R_t^b)^{-1}$. When there is a martingale component, ref. 54 advocated bounding its magnitude by, in effect, using this return reciprocal as a misspecified stochastic discount factor. For this application, we use

$$S_t = N_t (R_t^b)^{-1}$$

as the stochastic discount factor where $N_t \geq 0$ has conditional expectation equal to one and thus induces a change of a probability measure that absorbs long-term risk adjustments. Our method applied to this problem complements and extends those of refs. 54 and 56 by providing bounds on expectations implied by this measure.

Remark 5.1: Ref. 57 gives conditions under which investor beliefs can be uniquely identified from asset prices. This identification hinges on a complete panel of state prices being available to the econometrician, as well as an implicit restriction that the stochastic discount factor process does not contain a martingale component. This latter restriction is violated in many standard asset pricing models (see ref. 55 for a discussion). More generally this identification reveals a forward probability measure that absorbs long-term growth rate risk. Our approach avoids the requirement of a complete set of state prices and can be used to identify sets of long-term forward probabilities that are consistent with asset pricing model restrictions and within a small deviation of rational expectations.

6. Conclusions

In this paper, we developed methods designed to extract information on investor beliefs from data on asset prices and investor surveys. Our approach presumes an econometric model of investors or enterprises that could be misspecified under rational expectations. We illustrated how limiting the statistical discrepancy between investor beliefs and rational expectations implies bounds on investors' expectations. Formally, we represented this relationship through a nonlinear expectation function and derived its dual representation.

Going forward, we see two types of applications of our methods. Deducing market expectations about the future from forward-looking asset prices is a common practice in both the public and private sectors. But this is typically done either informally or by targeting so-called risk-neutral probabilities that confound beliefs and risk preferences. Our method provides a formal way to compute and represent information on investor beliefs constrained by a model of risk aversion along with a measure of statistical divergence.

Alternatively, we could use our approach to provide diagnostics for model misspecification under rational expectations. The bounds we deduce will help assess alternative models of subjective beliefs or ambiguity aversion. Implied belief bounds for small or moderate restrictions on the statistical divergence can give suggestive results for model builders as to how to repair potentially misspecified models. By comparing models of subjective beliefs or ambiguity aversion supported by belief distortions to the implied bounds, applied researchers could assess whether such departures from rational expectations could be easily discerned from limited data.

Future applications of our methodology could incorporate information from survey data on investor beliefs. One approach would be to include survey data directly as additional moment conditions when constructing expectation bounds. Another approach would be to compare survey-implied expectations to expectation bounds on the corresponding variables formed without using the survey data as information. The latter approach would provide a check on how plausible the survey data are as a representation of investor beliefs used in decision making.

While we focused on constraining probabilities based on intertemporal measures of divergence, these bounds could be used formally in the design of economic policy. For instance, refs. 58 and 59 pose dynamic policy problems in which a government policymaker fails to have precise knowledge of the beliefs of the private sector when designing a prudent course of action. While these papers explore what impact this imprecision could have on policy, our work looks more systematically at how to extract credible information about the beliefs.

Data Availability. Computer code and computations with standard data sources have been deposited in Github at <https://github.com/lphansen/Beliefs> with computational details on the implementation. All study data are included in this article and *SI Appendix*.

ACKNOWLEDGMENTS. We thank Fernando Alvarez, Orazio Attanasio, Anmol Bhandari, Stéphane Bonhomme, Tim Christensen, Jim Heckman,

Leo De Aparisi Lannoy, Winston Dou, Ralph Koijen, Marco Loseto, Yueran Ma, Andrey Malenko, Stefan Nagel, Diana Petrova, Monika Piazzesi, Eric Renault, Jose Scheinkman, Azeem Shaikh, Ken Singleton, Grace Tsiang, Harald Uhlig, and Xiangyu Zhang for helpful comments and suggestions. We gratefully acknowledge the research assistance of Han Xu and Zhenhuan Xie. We thank the Alfred P. Sloan Foundation (Grant G-2018-11113) for financial support.

1. A. Pigou, *Industrial Fluctuations* (Macmillan and Co. Limited, London, UK, 1927).
2. J. M. Keynes, *The General Theory of Employment, Interest, and Money* (Macmillan, London, UK, 1936).
3. J. Hicks, *Value and Capital* (Oxford, 1939).
4. L. A. Metzler, The nature and stability of inventory cycles. *Rev. Econ. Stat.* **23**, 113–129 (1941).
5. M. Nerlove, Adaptive expectations and cobweb phenomena. *Q. J. Econ.* **72**, 227–240 (1958).
6. J. F. Muth, Rational expectations and the theory of price movements. *Econometrica* **29**, 315–335 (1961).
7. R. E. Lucas, Expectations and the neutrality of money. *J. Econ. Theor.* **4**, 103–124 (1972).
8. A. Fuster, D. Laibson, B. Mendel, Natural expectations and macroeconomic fluctuations. *J. Econ. Perspect.* **24**, 67–84 (2010).
9. D. Hirshleifer, J. Li, J. Yu, Asset pricing in production economies with extrapolative expectations. *J. Monet. Econ.* **76**, 87–106 (2015).
10. N. Barberis, R. Greenwood, L. Jin, A. Shleifer, X-CAPM: An extrapolative capital asset pricing model. *J. Financ. Econ.* **115**, 1–24 (2015).
11. K. Adam, A. Marcet, J. P. Nicolini, Stock market volatility and learning. *J. Finance* **71**, 33–82 (2016).
12. P. Bordalo, N. Gennaioli, R. L. Porta, A. Shleifer, Diagnostic expectations and stock returns. *J. Finance* **74**, 2839–2874 (2019).
13. M. Piazzesi, J. Salomao, M. Schneider, Trend and cycle in bond premia. <https://web.stanford.edu/~piazzesi/trendcycle.pdf> Accessed 8 December 2020.
14. A. Bhandari, J. Borovicka, P. Ho, Survey data and subjective beliefs in business cycle models. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3473120 (22 October 2019).
15. I. R. Petersen, M. R. James, P. Dupuis, Minimax optimal control of stochastic uncertain systems with relative entropy constraints. *IEEE Trans. Automat. Contr.* **45**, 398–412 (2000).
16. J. Duchi, P. Glynn, H. Namkoong, Statistics of robust optimization: A generalized empirical likelihood approach. *Math. Oper. Res.*, in press.
17. D. Škulj, Discrete time Markov chains with interval probabilities. *Int. J. Approx. Reason.* **50**, 1314–1329 (2009).
18. L. P. Hansen, K. J. Singleton, Efficient estimation of linear asset-pricing models with moving average errors. *J. Bus. Econ. Stat.* **14**, 53–68 (1996).
19. S. Nagel, K. J. Singleton, Estimation and evaluation of conditional asset pricing models. *J. Finance* **66**, 873–909 (2011).
20. P. Gagliardini, D. Ronchetti, Comparing asset pricing models by the conditional Hansen-Jagannathan distance. *J. Financ. Econom.* **18**, 333–394 (2020).
21. B. Antoine, K. Proulx, E. Renault, Pseudo-true SDFs in conditional asset pricing models. *J. Financ. Econ.*, 10.1093/jfinec/nby017 (2018).
22. L. P. Hansen, R. Jagannathan, Assessing specification errors in stochastic discount factor models. *J. Finance* **52**, 557–590 (1997).
23. L. P. Hansen, Nobel lecture: Uncertainty outside and inside economic models. *J. Polit. Econ.* **122**, 945–987 (2014).
24. A. Ghosh, C. Julliard, A. P. Taylor, What is the consumption-CAPM missing? An information-theoretic framework for the analysis of asset pricing models. *Rev. Financ. Stud.* **30**, 442–504 (2017).
25. C. F. Manski, Survey measurement of probabilistic macroeconomic expectations: Progress and promise. *NBER Macroecon. Annu.* **32**, 411–471 (2018).
26. P. Bordalo, N. Gennaioli, Y. Ma, A. Shleifer, Over-reaction in macroeconomic expectations. *Am. Econ. Rev.* **110**, 2748–2782 (2020).
27. M. D. Donsker, S. S. Varadhan, Asymptotic evaluation of certain Markov process expectations for large time, I-IV. *Commun. Pure Appl. Math.* **28**, 1–47 (1975).
28. P. Dupuis, R. Ellis, *A Weak Convergence Approach to the Theory of Large Deviations* (John Wiley and Sons, Inc., New York, NY, 1997).
29. L. P. Hansen, A method for calculating bounds on the asymptotic covariance matrices of generalized method of moments estimators. *J. Econom.* **30**, 203–238 (1985).
30. L. P. Hansen, S. F. Richard, The role of conditioning information in deducing testable restrictions implied by dynamic asset pricing models. *Econometrica* **55**, 587–613 (1987).
31. R. Bradley, *Introduction to Strong Mixing Conditions* (Kendrick Press, Heber City, UT, 2007), vol. 1.
32. S. R. Varadhan, Special invited paper: Large deviations. *Ann. Probab.* **36**, 397–419 (2008).
33. A. Dembo, O. Zeitouni, *Large Deviations Techniques and Applications* (Springer Science & Business Media, 2009), vol. 38.
34. S. Eckstein, Extended Laplace principle for empirical measures of a Markov chain. *Adv. Appl. Probab.* **51**, 136–167 (2019).
35. J. Qin, J. Lawless, Empirical likelihood and general estimating equations. *Ann. Stat.* **22**, 300–325 (1994).
36. A. B. Owen, *Empirical Likelihood* (CRC Press, 2001).
37. L. P. Hansen, T. J. Sargent, Robust control and model uncertainty. *Am. Econ. Rev.* **91**, 60–66 (2001).
38. A. Ghosh, G. Roussellet, Identifying beliefs from asset prices. <http://doi.org/10.2139/ssrn.3400005> (17 June 2019).
39. S. Peng, “Nonlinear expectations, nonlinear evaluations and risk measures” in *Stochastic Methods in Finance: Lecture Notes in Mathematics*, M. Frittelli, W. Runggaldier, Eds. (Springer Berlin Heidelberg, Berlin, Heidelberg, Germany, 2004), pp. 165–253.
40. L. P. Hansen, J. A. Scheinkman, Long-term risk: An operator approach. *Econometrica* **77**, 177–234 (2009).
41. L. Qin, V. Linetsky, Long-term risk: A martingale approach. *Econometrica* **85**, 299–312 (2020).
42. C. M. Newman, B. W. Stuck, Chernoff bounds for discriminating between two Markov processes. *Stochastics* **2**, 139–153 (1979).
43. E. W. Anderson, L. P. Hansen, T. J. Sargent, A quartet of semigroups for model specification, robustness, prices of risk, and model detection. *J. Eur. Econ. Assoc.* **1**, 68–123 (2003).
44. D. M. Kreps, E. L. Porteus, Temporal resolution of uncertainty and dynamic choice. *Econometrica* **46**, 185–200 (1978).
45. L. G. Epstein, S. E. Zin, Substitution, risk aversion, and the temporal behavior of consumption and asset returns: An empirical analysis. *J. Polit. Econ.* **99**, 263–286 (1991).
46. J. Campbell, Intertemporal asset pricing without consumption data. *Am. Econ. Rev.* **83**, 487–512 (1993).
47. L. G. Epstein, E. Farhi, T. Strzalecki, How much would you pay to resolve long-run risk? *Am. Econ. Rev.* **104**, 2680–2697 (2014).
48. D. Dillenberger, D. Gottlieb, P. Ortleva, Stochastic impatience and the separation of time and risk preferences. <http://doi.org/10.2139/ssrn.3645071> (8 July 2020).
49. J. M. Maheu, T. H. McCurdy, Identifying bull and bear markets in stock returns. *J. Bus. Econ. Stat.* **18**, 100–112 (2000).
50. M. Stutzer, A Bayesian approach to diagnosis of asset pricing models. *J. Econom.* **68**, 367–397 (1995).
51. M. Stutzer, A simple nonparametric approach to derivative security valuation. *J. Finance* **51**, 1633–1652 (1996).
52. M. Avellaneda, Minimum-relative-entropy calibration of asset-pricing models. *Int. J. Theor. Appl. Finance* **1**, 447–472 (1998).
53. H. B. Kazemi, An intertemporal model of asset prices in a Markov economy with a limiting stationary distribution. *Rev. Financ. Stud.* **5**, 85–104 (1992).
54. F. Alvarez, U. J. Jermann, Using asset prices to measure the persistence of the marginal utility of wealth. *Econometrica* **73**, 1977–2016 (2005).
55. J. Borovicka, L. P. Hansen, J. A. Scheinkman, Misspecified recovery. *J. Finance* **71**, 2493–2544 (2016).
56. G. Bakshi, F. Chabi-Yo, Variance bounds on the permanent and transitory components of stochastic discount factors. *J. Financ. Econ.* **105**, 191–208 (2012).
57. S. Ross, The recovery theorem. *J. Finance* **70**, 615–648 (2015).
58. M. Woodford, Robustly optimal monetary policy with near-rational expectations. *Am. Econ. Rev.* **100**, 274–303 (2010).
59. L. P. Hansen, T. J. Sargent, Three types of ambiguity. *J. Monet. Econ.* **59**, 422–445 (2012).