



Uncertainty Outside and Inside Economic Models*

Prize Lecture, December 8, 2013

by Lars Peter Hansen

University of Chicago, Chicago, IL, USA.

We must infer what the future situation would have been without our interference, and what change will be wrought in it by our action.

Fortunately or unfortunately, none of these processes is infallible, or indeed ever accurate and complete. Knight (1921: 201–202)

1 INTRODUCTION

Asset pricing theory has long recognized that financial markets compensate investors who are exposed to some components of uncertainty. This is where macroeconomics comes into play. The economy-wide shocks, the primary concern of macroeconomists, by their nature are not diversifiable. Exposures to these shocks cannot be averaged out with exposures to other shocks. Thus, returns on assets that depend on these macroeconomic shocks reflect “risk” premia and are a linchpin connecting macroeconomic uncertainty to financial markets. A risk premium reflects both the price of risk and the degree of exposure to risk. I will be particularly interested in how the exposures to macroeconomic impulses are priced by decentralized security markets.

* I thank Manuel Arellano, Amy Boonstra, Philip Barrett, Xiaohong Chen, John Cochrane, Maryam Farboodi, Eric Ghysels, Itzhak Gilboa, Massimo Marinacci, Nan Li, Monika Piazzesi, Eric Renault, Scott Richard, Larry Samuelson, Enrique Sentana, José Scheinkman, Martin Schneider, Stephen Stigler, Harald Uhlig, Amir Yaron an anonymous referee and especially Jaroslav Borovička, James Heckman, Thomas Sargent and Grace Tsiang for helpful comments.

How do we model the dynamic evolution of the macroeconomy? Following the tradition initiated by Slutsky (1927, 1937) and Frisch (1933), I believe it is best captured by stochastic processes with restrictions; exogenous shocks repeatedly perturb a dynamic equilibrium through the model's endogenous transmission mechanisms. Bachelier (1900), one of the developers of Brownian motion, recognized the value of modeling financial prices as responses to shocks.¹ It took economists fifty years to discover and appreciate his insights. (It was Savage who alerted Samuelson to this important line of research in the early 1950s.) Prior to that, scholars such as Yule (1927), Slutsky (1927, 1937) and Frisch (1933) had explored how linear models with shocks and propagation mechanisms provide attractive ways of explaining approximate cyclical behavior in macro time series. Similarities in the mathematical underpinnings of these two perspectives opened the door to connecting macroeconomics and finance.

Using random processes in our models allows economists to capture the variability of time series data, but it also poses challenges to model builders. As model builders, we must understand the uncertainty from two different perspectives. Consider first that of the econometrician, standing outside an economic model, who must assess its congruence with reality, inclusive of its random perturbations. An econometrician's role is to choose among different parameters that together describe a family of possible models to best mimic measured real world time series and to test the implications of these models. I refer to this as *outside uncertainty*. Second, agents inside our model, be it consumers, entrepreneurs, or policy makers, must also confront uncertainty as they make decisions. I refer to this as *inside uncertainty*, as it pertains to the decision-makers within the model. What do these agents know? From what information can they learn? With how much confidence do they forecast the future? The modeler's choice regarding insiders' perspectives on an uncertain future can have significant consequences for each model's equilibrium outcomes.

Stochastic equilibrium models predict risk prices, the market compensations that investors receive for being exposed to macroeconomic shocks. A challenge for econometric analyses is to ascertain if their predictions are consistent with data. These models reveal asset pricing implications via stochastic discount factors. The discount factors are stochastic to allow for exposures to alternative

¹ See Davis and Etheridge (2006) for a translation and commentary and Dimson and Mussavian (2000) for an historical discussion of the link between Bachelier's contribution and subsequent research on efficient markets.

macroeconomic random outcomes to be discounted differently. Moreover, the compounding of stochastic discount factors shows how market compensations change with the investment horizon. Stochastic discount factors thus provide a convenient vehicle for depicting the empirical implications of the alternative models. I will initially describe the methods and outcomes from an econometrician *outside* the model.

Stochastic discount factors are defined with respect to a probability distribution relevant to investors inside the model. Lucas (1972) and others imposed rational expectations as an equilibrium concept, making the probability distribution relevant to investors inside the model coincide with the probability distribution implied by the solution to the model. It is an elegant response for how to model agents *inside* the model, but its application to the study of asset pricing models has resulted in empirical puzzles as revealed by formal econometric methods that I will describe. These and other asset pricing anomalies have motivated scholars to speculate about investor beliefs and how they respond to or cope with uncertainty. In particular, the anomalies lead me and others to explore specific alternatives to the rational expectations hypothesis.

In this essay I will consider alternatives motivated in part by a decision theory that allows for distinctions between three alternative sources of uncertainty: i) risk conditioned on a model, ii) ambiguity about which is the correct model among a family of alternatives, and iii) potential misspecification of a model or a family of possible models. These issues are pertinent to outside econometricians, but they also may be relevant to inside investors. I will elaborate on how the distinctions between uncertainty components open the door to the investigation of market compensations with components other than more narrowly defined risk prices. Motivated by empirical evidence, I am particularly interested in uncertainty pricing components that fluctuate over time.

Why is it fruitful to consider model misspecification? In economics and as in other disciplines, models are intended to be revealing simplifications, and thus deliberately are not exact characterizations of reality; it is therefore specious to criticize economic models merely for being wrong. The important criticisms are whether our models are wrong in having missed something essential to the questions under consideration. Part of a meaningful quantitative analysis is to look at models and try to figure out their deficiencies and the ways in which they can be improved. A more subtle challenge for statistical methods is to explore systematically potential modeling errors in order to assess the quality of the model predictions. This kind of uncertainty about the adequacy of a model or model family is not only relevant for econometricians outside the model but potentially also for agents inside the models.

This essay proceeds as follows. In Section 2, I review the development of time series econometric modeling, including the initiation of rational expectations econometrics. In Section 3, I review my contributions to the econometric study of partially specified models, adapting to the study asset pricing and macroeconomic uncertainty. I describe methods and approaches to the study of fully specified models based on asset pricing considerations in Section 4. In Section 5, I explore the consequences for asset pricing models when investor beliefs are not in full accord with an underlying model, which can result in investor behavior that resembles extreme risk aversion. In Section 6, I review perspectives on model ambiguity which draw on work by decision theorists and statisticians to revisit the framework that I sketch in Section 5. I draw some conclusions in Section 7.

2 RATIONAL EXPECTATIONS ECONOMETRICS

Rational expectations econometrics explores structural stochastic models of macroeconomic time series with the ambition to be a usable tool for policy analysis. It emerged in response to a rich history of modeling and statistical advances. Yule (1927) and Slutsky (1927, 1937) provided early characterizations of how time series models can generate interesting cyclical behavior by propagating shocks. Yule (1927) showed that a second-order autoregression could reproduce intriguing patterns in the time series. He fit this model to sunspot data, known to be approximately but not exactly periodic. The model was built using independent and identically distributed (iid) shocks as building blocks. The model produced a damped periodic response to random impulses. Similarly, Slutsky (1927, 1937) constructed models that were moving-averages of iid shocks and showed how such processes could be arbitrarily close to exact periodic sequences.² He also demonstrated how moving-average type models could account for British business cycle data.

Frisch (1933), who shared the first Sveriges Riksbank Prize in Economics with Tinbergen, pushed this agenda further by exploring how to capture dynamic economic phenomenon through probability models with explicit economic underpinnings. Frisch discussed propagation from initial conditions and described an important role for random impulses building in part on the work of Yule (1927) and Slutsky (1927, 1937). In effect, Frisch (1933) introduced

²I cite two versions of Slutsky's paper. The first one was published in Russian. The second one was published in English a decade later with a more comprehensive set of results. English translations of the first paper were circulated well in advance of 1937.

impulse response functions to economics as a device to understand the intertemporal impact of shocks on economic variables. Haavelmo (1944) took an additional step by providing foundations for the use of statistical methods to assess formally the stochastic models. This literature set the foundation for a modern time series econometrics that uses economics to interpret evidence in a mathematically formal way. It featured important interactions among economics, mathematics and statistics and placed a premium on formal model building.³ Haavelmo (1944) confronts uncertainty as an econometrician *outside* the model that is to be estimated and tested.

Investment and other decisions are in part based on people's views of the future. Once economic decision makers are included into formal dynamic economic models, their expectations come into play and become an important ingredient to the model. This challenge was well appreciated by economists such as Pigou, Keynes and Hicks, and their suggestions have had durable impact on model building. Thus, the time series econometrics research agenda had to take a stand on how people inside the model made forecasts. Alternative approaches were suggested including static expectations, adaptive expectations or appeals to data on beliefs; but these approaches left 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*. This hypothesis appeals to long histories of data to motivate the modeling expectations. The Law of Large Numbers gives an approximation whereby parameters that are invariant over time are revealed by data, and this revelation gives a model builder a way to formalize the expectations of economic investors inside our models.⁴ This approach to completing the specification of a stochastic equilibrium model was initiated within macroeco-

³ Frisch, in particular, nurtured this ambitious research agenda by his central role in the foundational years of the Econometric Society. His ambition is reflected in the 1933 mission statement he wrote for the journal *Econometrica*: “. . . Experience has shown that each of these three viewpoints, that of statistics, economic theory, and mathematics, is a necessary, but not by itself a sufficient, condition for a real understanding of the quantitative relations in modern economic life. It is the unification of all three that is powerful. And it is this unification that constitutes econometrics.” Frisch (1933b).

⁴ More than three hundred years ago, Jacob Bernoulli proved a result that implied a Law of Large Numbers. He was motivated in part by social problems for which probabilities had to be estimated empirically, in contrast to typical gambling problems. Bernoulli's result initiated an enduring discussion of both the relevance of his simple model specification and of the approximation he established. See Stigler (2014) for an interesting retrospective on Bernoulli's contribution.

nomics by Muth (1961) and Lucas (1972). Following Lucas (1972) in particular, rational expectations became an integral part of an equilibrium for a stochastic economic model.

The aim of structural econometrics is to provide a framework for policy analysis and the study of counterfactuals. This vision is described in Marschak (1953) and articulated formally in the work of Hurwicz (1962). While there are a multitude of interesting implications of the rational expectations hypothesis, perhaps the most important one is its role in policy analysis. It gives a way to explore policy experiments or hypothetical changes that are not predicated on systematically fooling people. See Sargent and Wallace (1975) and Lucas (1976) for a discussion.⁵

From an econometric standpoint, rational expectations introduced important cross-equation restrictions. These recognize that parameters governing the dynamic evolution of exogenous impulses to the model must also be present in decision rules and equilibrium relations. These restrictions reflect how decision-makers within the model are forward-looking. For instance, an investment choice today depends on the beliefs about how profitable such investments will be in the future. Investors forecast the future, and the rational expectations hypothesis predicts how they do this. The resulting cross-equation restrictions add a new dimension to econometric analysis; but these restrictions are built on the premise that investors have figured much out about how the future will evolve. See Sargent (1973), Wallis (1980) and my first published paper, Hansen and Sargent (1980), for characterizations of these restrictions.⁶ To implement this approach to rational expectations econometrics, a researcher is compelled to specify correctly the information sets of economic actors.⁷ When building actual stochastic models, however, it is often not clear what information should be presumed on the part of economic agents, how they should use it, and how much confidence they have in that use.

The introduction of random shocks as impulses to a dynamic economic model in conjunction with the assumption of rational expectations is an example of uncertainty *inside* a model. Under a rational expectations equilibrium, an

⁵ To be clear, rational expectations offers an approach for comparing distinct stochastic equilibria but not the transitions from one to another. For an interesting extension that allows for clustering of observations near alternative *self-confirming equilibria* in conjunction with escapes from such clusters see Sargent (1999).

⁶ While this was my first publication of a full length paper, this was not my first publication. My first was a note published in *Economic Letters*.

⁷ See Sims (2012) for a discussion of the successes and limitations of implementing the Haavelmo (1944) agenda to the study of monetary policy under rational expectations.

investor inside the model knows the model-implied stochastic evolution for the state variables relevant for decision making and hence the likely consequences of the impulses. An econometrician also confronts uncertainty *outside* a model because of his or her lack of knowledge of parameters or maybe even a lack of confidence with the full model specification. There is an asymmetry between the inside and the outside perspectives found in rational expectations econometrics that I will turn to later. But first, I will discuss an alternative approach to imposing rational expectations in econometric analyses.

3 ROBUST ECONOMETRICS UNDER RATIONAL EXPECTATIONS

My econometrics paper, Hansen (1982b), builds on a long tradition in econometrics of “doing something without having to do everything.” This entails the study of partially specified models—that is, models in which only a subset of economic relations are formally delineated. I added to this literature by analyzing such estimation problems in greater generality, giving researchers more flexibility in modeling the underlying time series while incorporating some explicit economic structure. I studied formally a family of Generalized Method of Moments (GMM) estimators, and I adapted these methods to applications that study linkages between financial markets and the macroeconomy.⁸ By allowing for partial specification, these methods gain a form of robustness. They are immune to mistakes in how one might fill out the complete specification of the underlying economic model.

The approach is best thought of as providing initial steps in building a time series econometric model without specifying the full econometric model. Consider a research program that studies the linkages between the macroeconomy and financial markets. One possibility is to construct a fully specified model of

⁸ My exposure to using GMM estimators as a vehicle to represent a broad family of estimators originally came from Christopher Sims’ lectures. As a graduate student I became interested in central limit approximations that allow for econometric error terms to possess general types of temporal dependence by using central limit approximations of the type demonstrated by Gordin (1969). I subsequently established formally large sample properties for GMM estimators in such circumstances. Interestingly, *Econometrica* chose not to publish many of the formal proofs for results in my paper. Instead they were published thirty years later by the *Journal of Econometrics*, see Hansen (2012). Included in my original submission and in the published proofs is a Uniform Law of Large Numbers for stationary ergodic processes. See Hansen (2001) and Ghysels and Hall (2002) for further elaborations and discussion about the connection between GMM and related statistics literatures. See Arellano (2003) for a discussion of applications to panel data.

the macroeconomy including the linkages with financial markets that are presumed to exist. This is a lot to ask in early stages of model development. Of course, an eventual aim is to produce a full model of stochastic equilibrium.

The econometric tools that I developed are well suited to study a rich family of asset pricing models, among other things. Previously, Ross (1978) and Harrison and Kreps (1979) produced mathematical characterizations of asset pricing in frictionless asset pricing markets implied by the absence of arbitrage. Their work provides a general way to capture how financial markets value risky payoffs. My own research, and that with collaborators, built on this conceptual approach, but with an important reframing. Our explicit consideration of stochastic discounting, left implicit in the Ross (1978) and Harrison and Kreps (1979) framework, opened the door to new ways to conduct empirical studies of asset pricing models using GMM and related econometric methods. I now describe these methods.

3.1 A GMM Approach to Empirical Asset Pricing

A productive starting point in empirical asset pricing is

$$E \left[\left(\frac{S_{t+\ell}}{S_t} \right) Y_{t+\ell} \mid \mathcal{F}_t \right] = Q_t \quad (1)$$

where $S > 0$ is a stochastic discount factor (SDF) process. In formula (1), $Y_{t+\ell}$ is a vector of payoffs on assets at time $t + \ell$, and Q_t is a vector of corresponding asset prices. The event collection (sigma algebra), \mathcal{F}_t , captures information available to an investor at date t . The discount factor process is stochastic in order to adjust market values for risk. Each realized state is discounted differently and this differential discounting reflects investor compensation for risk exposure. Rational expectations is imposed by presuming that the conditional expectation operator is consistent with the probability law that governs the actual data generation. With this approach a researcher does not specify formally that probability law and instead “lets the data speak.”

Relations of type (1) are premised on investment decisions made in optimal ways and are fundamental ingredients in stochastic economic models. The specification of a SDF process encapsulates some economics. It is constructed from the intertemporal marginal rates of substitution of marginal investors. Investors consider the choice of consuming today or investing to support opportunities to consume in the future. There are a variety of investment opportunities with differential exposure to risk. Investors’ risk aversion enters the SDF and influences the nature of the investment that is undertaken. While I have used the language of financial markets, this same formulation applies to investments in physical

and human capital. In a model of a stochastic equilibrium, this type of relation holds when evaluated at equilibrium outcomes. Relation (1) by itself is typically not sufficient to determine fully a stochastic equilibrium, so focusing on this relation alone leads us to a partially specified model. Additional modeling ingredients are required to complete the specification. The presumption is that whatever those details might be, the observed time series come from a stochastic equilibrium that is consistent with an equation of the form (1).

Implications of relation (1), including the role of SDFs and the impact of conditioning information used by investors, were explored systematically in Hansen and Richard (1987). But the origins of this empirically tractable formulation traces back to Rubinstein (1976), Lucas (1978) and Grossman and Shiller (1981), and the conceptual underpinnings to Ross (1978) and Harrison and Kreps (1979).⁹ To implement formula (1) as it stands, we need to specify the information set of economic agents correctly. The Law of Iterated Expectations allows us to understate the information available to economic agents.¹⁰ For instance let $\widehat{\mathcal{F}}_t \subset \mathcal{F}_t$ denote a smaller information set used by an external analyst. By averaging over the finer information set \mathcal{F}_t conditioned on the coarser information set $\widehat{\mathcal{F}}_t$, I obtain

$$E \left[\left(\frac{S_{t+\ell}}{S_t} \right) (Y_{t+\ell})' - (Q_t)' \mid \widehat{\mathcal{F}}_t \right] = 0. \tag{2}$$

I now slip in conditioning information through the “back door” by constructing a conformable matrix Z_t with entries in the reduced information set (that are $\widehat{\mathcal{F}}_t$ measurable). Then

$$E \left[\left(\frac{S_{t+\ell}}{S_t} \right) (Y_{t+\ell})' Z_t - (Q_t)' Z_t \mid \widehat{\mathcal{F}}_t \right] = 0.$$

⁹ The concept of a SDF was first introduced in Hansen and Richard (1987). Stochastic discount factors are closely connected to the “risk-neutral” probabilities used in valuing derivative claims. This connection is evident by dividing the one-period SDF by its conditional mean and using the resulting random variable to define a new one-period conditional probability distribution, the risk neutral distribution.

¹⁰ In his study of interest rates, Shiller (1972) in his PhD dissertation suggested omitted information as a source of an “error term” for an econometrician. In Hansen and Sargent (1980), we built on this insight by contrasting implications for a “Shiller error-term” as a disturbance term to processes that are unobserved to an econometrician and enter structural relations. In Hansen and Sargent (1991) we show how to allow for omitted information in linear or log-linear time series models using quasi-likelihood methods.

Under an asset pricing interpretation, $(Y_{t+\ell})'Z_t$ is a synthetic payoff vector with a corresponding price vector $(Q_t)'Z_t$. Finally, we may form the unconditional expectation by averaging over the coarser conditioning information set $\widehat{\mathcal{F}}_t$:

$$E\left[\left(\frac{S_{t+\ell}}{S_t}\right)(Y_{t+\ell})' - (Q_t)' \mid \widehat{\mathcal{F}}_t\right] = 0. \quad (3)$$

This becomes an estimation problem once we parameterize the SDF in terms of observables and unknown parameters to be estimated.

Hansen and Singleton (1982) is an initial example of this approach.¹¹ In that work we consider the case in which the SDF process can be constructed from observables along with some unknown parameters. Economics comes into play in justifying the construction of the SDF process and sometimes in the construction of returns to investment. From an econometric perspective, time series versions of Laws of Large Numbers and Central Limit Theorems give us approximate ways to estimate parameters and test restrictions as in Hansen (1982b).

In Hansen (1982b), I also studied statistical efficiency for a class of GMM estimators given a particular choice of Z in a manner that extends an approach due to Sargan (1958, 1959).¹² When (3) has more equations than unknown parameters, multiple GMM estimators are the outcome of using (at least implicitly) alternative linear combinations of these equations equal to the number of parameters. Since there are many possible ways to embark on this construction, there is a family of GMM estimators. This family of estimators has an attainable efficiency bound derived and reported in Hansen (1982b).¹³ When the number

¹¹ An earlier application of GMM inference is found in my work Hansen and Hodrick (1980). In that paper we studied the empirical relationship between the logarithm of a future spot exchange and the logarithm of the current forward rate and other possible predictors. We applied ordinary least squares in our work, but with corrected standard errors. Others were tempted to (and in fact did) apply generalized least squares (GLS) to “correct for” serial correlation, but applied in this setting GLS is statistically inconsistent. The counterpart to the moment conditions studied here are the least squares orthogonality conditions. The contract interval played the role of ℓ in this least squares analysis and was typically larger than one. In subsequent work, Hansen and Hodrick (1983), we used a SDF formulation to motivate further empirical characterizations, which led us to confront over-identification. See also Bilson (1981) and Fama (1984) who featured a cross-currency analysis.

¹² See Arellano (2002) for a nice discussion relating GMM estimation to the earlier work of Sargan.

¹³ See Hansen (2007b) for a pedagogical discussion of GMM estimation including discussions of large sample statistical efficiency and tests.

of equations exceeds the number of free parameters, there is also a direct way to test equations not used formally in estimation. While nesting estimators into a general GMM framework has great pedagogical value, I was particularly interested in applying a GMM approach to problems requiring new estimators as in many of the applications to financial economics and elsewhere.¹⁴

Notice that the model, as written down in equation (3), is only partially specified. Typically we cannot invert this relation, or even its conditional counterpart, to deduce a full time series evolution for economic aggregates and financial variables.¹⁵ Other relations would have to be included in order to obtain a full solution to the problem.

3.2 Further Econometric Challenges

I now digress temporarily and discuss some econometric extensions that I and others contributed to.

3.2.1 SEMIPARAMETRIC EFFICIENCY

Since the model is only partially specified, the estimation challenge leads directly to what is formally called a semiparametric problem. Implicitly the remainder of the model can be posed in a nonparametric manner. This gives rise to a problem with a finite-dimensional parameter vector of interest and an infinite-dimensional “nuisance” parameter vector representing the remainder of the model. This opens the door to the study of semiparametric efficiency of a large class of estimators as will be evident from the discussion that follows. In typical GMM problems, the actual introduction of the nuisance parameters can be sidestepped.

Relation (2) conditions on the information set of economic agents. We have great flexibility in choosing the matrix process Z . The entries of Z_t should be in the $\widehat{\mathcal{F}}_t$ information set, but this still leaves many options when building a Z process. This flexibility gives rise to an infinite class of estimators. In Hansen (1982b), I studied statistical efficiency given a particular choice of Z . This

¹⁴ Other econometricians have subsequently found value in unifying the treatment of GMM estimators into a broader type of extremum estimators. This, however, misses some of the special features of statistical efficiency within a GMM framework and does not address the issue of how to construct meaningful estimators from economic models.

¹⁵ For those reluctant to work with partially specified models, Lucas (1978) showed how to close a special case of this model by considering an endowment economy. But from an empirical standpoint, it is often not necessary to take the endowment nature of the economy literally. The consumption from the endowment economy may be conceived of as the equilibrium outcome of a model with production and preserves the same pricing relations.

approach, however, understates the class of possible GMM estimators in a potentially important way. Hansen (1985) shows how to construct an efficiency bound for the much larger (infinite dimensional) class of GMM estimators. This efficiency bound is a greatest lower bound on the asymptotic efficiency of the implied GMM estimators. Not surprisingly, it is more challenging to attain this bound in practice. For some related but special (linear) time series problems, Hansen and Singleton (1996) and West et al. (2009) discuss implementation strategies.

There is a more extensive literature exploring these and closely related questions in an iid (independent and identically distributed) data setting, including Chamberlain (1987), who looks at an even larger set of estimators. By connecting to an extensive statistics literature on semiparametric efficiency, he shows that this larger set does not improve the statistical efficiency relative to the GMM efficiency bound. Robinson (1987), Newey (1990), and Newey (1993) suggest ways to construct estimators that attain this efficiency bound for some important special cases.¹⁶ Finally, given the rich array of moment restrictions, there are opportunities for more flexible parameterizations of, say, a SDF process. Suppose the conditional moment restrictions contain a finite-dimensional parameter vector of interest along with an infinite-dimensional (nonparametric) component. Chamberlain (1992) constructs a corresponding efficiency bound and Ai and Chen (2003) extend this analysis and estimation for such problems. While these richer efficiency results have not been shown in the time series environment I consider, I suspect that they can indeed be extended.

3.2.2 MODEL MISSPECIFICATION

The approaches to GMM estimation that I have described so far presume a given parameterization of a SDF process. For instance, the analysis of GMM efficiency in Hansen (1982b) and Hansen (1985) and related literature presumes that the model is correctly specified for one value of the unknown (to the econometrician) parameter. Alternatively, we may seek to find the best choice of a parameter value even if the pricing restrictions are only *approximately* correct. In our paper, Hansen and Jagannathan (1997), we suggest a modification of GMM estimation in which appropriately scaled pricing errors are minimized. We propose this as a way to make model comparisons in economically meaningful ways. Recently, Gosh et al. (2012) adopt an alternative formulation of model

¹⁶ Relatedly, Zhang and Gijbels (2003), Kitamura et al. (2004) and Antoine et al. (2007) studied methods based on restricting nonparametric estimates of conditional density functions to attain Chamberlain (1987)'s efficiency bound in an estimation environment with independent and identically distributed data generation.

misspecification extending the approach of Stutzer (1995) described later. This remains an interesting and important line of investigation that parallels the discussion of model misspecification in other areas of statistics and econometrics. I will return to this topic later in this essay.

3.2.3 NONPARAMETRIC CHARACTERIZATION

A complementary approach to building and testing new parametric models is to treat the SDF process as unobserved by the econometrician. It is still possible to deduce empirical characterizations of such processes implied by asset market data. This analysis provides insights into modeling challenges by showing what properties a valid SDF process must possess.

It turns out that there are potentially many valid stochastic discount factors over a payoff horizon ℓ :

$$s \equiv \frac{S_{t+\ell}}{S_t}$$

that will satisfy either (2) or the unconditional counterpart (3). For simplicity, focus on (3).¹⁷ With this in mind, let

$$\begin{aligned} y' &= (Y_{t+\ell})' Z_t \\ q' &= (Q_t)' Z_t \end{aligned}$$

where for notational simplicity, I omit the time subscripts on the left-hand side of this equation. In what follows I will assume some form of a Law of Large Numbers so that we can estimate such entities. See Hansen and Richard (1987) for a discussion of such issues. Rewriting (3) with this simpler notation:

$$E[sy' - q'] = 0. \tag{4}$$

This equation typically implies many solutions for a positive $s > 0$. In our previous discussion of parametric models, we excluded many solutions by adopting a parametric representation in terms of observables and an unknown parameter vector. In practice this often led to a finding that there were *no* solutions, that is no values of s solving (4), within the parametric family assumed for s . Using Hansen (1982b), this finding was formalized as a test of the pricing restrictions. The finding alone left open the question: rejecting the parametric restrictions

¹⁷ For conditional counterparts to some of the results I summarize see Gallant et al. (1990) and Cochrane and Hansen (1992).

for what alternative? Thus a complementary approach is to characterize properties of the family of s 's that do satisfy (4). These solutions might well violate the parametric restriction.

The interesting challenge is how to characterize the family of SDFs that solve (4) in useful ways. Here I follow a general approach that is essentially the same as that in Almeida and Garcia (2013). I choose this approach both because of its flexibility and because it includes many interesting special cases used in empirical analysis. Consider a family of convex functions ϕ defined on the positive real numbers.¹⁸

$$\phi(r) = \frac{1}{\theta(1+\theta)} [(r)^{1+\theta} - 1] \tag{5}$$

for alternative choices of the parameter θ . The specification $\theta = 1$ is commonly used in empirical practice, in which case ϕ is quadratic. We shall look for lower bounds on the

$$E \left[\phi \left(\frac{s}{Es} \right) \right]$$

by solving the convex optimization problem:¹⁹

$$\lambda = \inf_{s>0} E \left[\phi \left(\frac{s}{Es} \right) \right] \text{ subject to } E[sy' - q'] = 0. \tag{6}$$

By design we know that

$$E \left[\phi \left(\frac{s}{Es} \right) \right] \geq \lambda.$$

Notice that $E \left[\phi \left(\frac{s}{Es} \right) \right]$ hence λ are nonnegative by Jensen's Inequality because ϕ is convex and $\phi(1) = 0$. When $\theta = 1$,

¹⁸ This functional form is familiar from economists' use of power utility (in which case we use $-\phi$ to obtain a concave function), from statisticians' use of F-divergence measures between two probability densities, the Box-Cox transformation, and the applications in the work of Cressie and Read (1984).

¹⁹ Notice that the expectation is also an affine transformation of the moment generating function for $\log s$.

$$\sqrt{2E\left[\phi\left(\frac{s}{Es}\right)\right]}$$

is the ratio of the standard deviation of s to its mean and $\sqrt{2\lambda}$ is the greatest lower bound on this ratio.

From the work of Ross (1978) and Harrison and Kreps (1979), arbitrage considerations imply the economically interesting restriction $s > 0$ with probability one. To guarantee a solution to optimization problem (6), however, it is sometimes convenient to include s 's that are zero with positive probability. Since the aim is to produce bounds, this augmentation can be justified for mathematical and computational convenience. Although this problem optimizes over an infinite-dimensional family of random variables s , the dual problem that optimizes over the Lagrange multipliers associated with the pricing constraint (4) is often quite tractable. See Hansen et al. (1995) for further discussion.

Inputs into this calculation are contained in the pair (y, q) and a hypothetical mean Es . If we have time series data on the price of a unit payoff at date $t + \ell$, Es can be inferred by averaging the date t prices over time. If not, by changing Es we can trace out a frontier of solutions. An initial example of this is found in Hansen and Jagannathan (1991) where we constructed mean-standard deviation tradeoffs for SDFs by setting $\theta = 1$.^{20,21}

While a quadratic specification of ϕ ($\theta = 1$) has been the most common one used in empirical practice, other approaches have been suggested. For instance, Snow (1991) considers larger moments by setting θ to integer values greater than one. Alternatively, setting $\theta = 0$ yields

$$E\left[\phi\left(\frac{s}{Es}\right)\right] = \frac{E[s(\log s - \log Es)]}{Es}$$

²⁰ This literature was initiated by a discussion in Shiller (1982) and my comment on that discussion in Hansen (1982a). Shiller argued why a volatility bound on the SDF is of interest, and he constructed an initial bound. In my comment, I showed how to sharpen the volatility bound, but without exploiting that $s > 0$. Neither Shiller nor I explored mean-standard deviation tradeoffs that are central in Hansen and Jagannathan (1991). In effect, I constructed one point on the frontier characterized in Hansen and Jagannathan (1991).

²¹ When θ is one, the function ϕ continues to be well defined and convex for negative real numbers. As noted in Hansen and Jagannathan (1991), if the negative choices of s are allowed in the optimization problem (which weakens the bound), there are quasi-analytical formulas for the minimization problems with simple links to Sharpe ratios commonly used in empirical finance.

which Stutzer (1995) featured this in his analysis. When $\theta = -1$,

$$E\left[\phi\left(\frac{s}{Es}\right)\right] = -E \log s + \log Es$$

and use of this specification of ϕ gives rise to a bound that has been studied in several papers including Bansal and Lehmann (1997), Alvarez and Jermann (2005), Backus et al. (2011), and Backus et al. (2014). These varying convex functions give alternative ways to characterize properties of SDFs that work through bounding their stochastic behaviour.²² He and Modest (1995) and Luttmer (1996) further extended this work by allowing for the pricing equalities to be replaced by pricing inequalities. These inequalities emerge when transaction costs render purchasing and selling prices distinct.²³

3.3 The Changing Price of Uncertainty

Empirical puzzles are only well defined within the context of a model. Hansen and Singleton (1982, 1983) and others documented empirical shortcomings of macroeconomic models with power utility versions of investor preferences. The one-period SDF of such a representative consumer is:

$$\frac{S_{t+1}}{S_t} = \exp(-\delta) \left(\frac{C_{t+1}}{C_t}\right)^{-\rho} \tag{7}$$

where C_t is consumption, δ is the subjective rate of discount and $\frac{1}{\rho}$ is the intertemporal elasticity of substitution. Hansen and Singleton and others were the bearers of bad news: the model didn't match the data even after taking account of statistical inferential challenges.²⁴

²²The continuous-time limit for the conditional counterpart results in one-half times the local variance for all choices of ϕ for Brownian information structures.

²³There has been some work on formal inferential methods associated with these methods. For instance, see Burnside (1994), Hansen et al. (1995), Peñaranda and Sentana (2011) and Chernozhukov et al. (2013).

²⁴Many scholars make reference to the “equity premium puzzle.” Singleton and I showed how to provide statistically rigorous characterizations of this and other empirical anomalies. The puzzling implications coming from this literature are broader than the expected return differential between an aggregate stock portfolio and bonds and extend to differential returns across a wide variety of securities. See, for instance, Fama and French (1992) for empirical evidence on expected return differences, and see Cochrane (2008) and the discussion by Hansen (2008) for an exchange about the equity premium and related puzzles.

This empirical work nurtured a rich literature exploring alternative preferences and markets with frictions. Microeconomic evidence was brought to bear that targeted financial market participants when constructing the SDFs. These considerations and the resulting modeling extensions led naturally to alternative specifications on SDFs and suggestions for how they might be measured.

The nonparametric methods leading to bounds also added clarity to the empirical evidence. SDFs encode compensations for exposure to uncertainty because they discount alternative stochastic cash flows according to their sensitivity to underlying macroeconomic shocks. Thus, empirical evidence about SDFs sheds light on the “risk prices” that investors need as compensations for being exposed to aggregate risk. Using these nonparametric methods, the empirical literature has found that the risk price channel is a fertile source for explaining observed variations in securities prices and asset returns. SDFs are highly variable (Hansen and Jagannathan (1991)). The unconditional variability in SDFs could come from two sources: on-average conditional variability or variation in conditional means. As argued by Cochrane and Hansen (1992), it is really the former. Conditional variability in SDFs implies that market-based compensations for exposure to uncertainty are varying over time in important ways. Sometimes this observation about time variation gets bundled into the observation about time-varying risk premia. Risk premia, however, depend both on the compensation for being exposed to risk (the price of risk) and on how big that exposure is to risk (the quantity of risk). Price variability, exposure variability or a combination of the two could be the source of fluctuations in risk premia. Deducing the probabilistic structure of SDFs from market data thus enables us to isolate the price effect. In summary, this empirical and theoretical literature gave compelling reasons to explore sources of risk price variation not previously captured, and provided empirical direction to efforts to improve investor preferences and market structures within these models.

Campbell and Cochrane (1999) provided an influential specification of investor preferences motivated in part by this empirical evidence. Consistent with the view that time variation in uncertainty prices is vital for understanding financial market returns, they constructed a model in which SDFs are larger in magnitude in bad economic times than good. This paper is prominent in the asset pricing literature precisely because it links the time series behavior of risk prices to the behavior of the macroeconomy (specifically aggregate consumption), and it suggests one preference-based mechanism for achieving this variation. Under the structural interpretation provided by the model, the implied risk aversion is very large in bad economic times and modest in good times as measured by the history of consumption growth. This work successfully avoided

the need for large risk aversion in all states of the world, but it did not avoid the need for large risk aversion in some states. The statistician in me is intrigued by the possibility that observed incidents of large risk aversion might be proxying for investor doubts regarding the correctness of models. I will have more to say about that later.

4 ECONOMIC SHOCKS AND PRICING IMPLICATIONS

While the empirical methods in asset pricing that I described do not require that an econometrician identify the fundamental macroeconomic shocks pertinent to investors, this shortcut limits the range of questions that can be addressed. Without accounting for shocks, we can make only an incomplete assessment of the consequences for valuation of macroeconomic uncertainty. To understand fully the pricing channel, we need to know how the SDF process itself depends on fundamental shocks. This dependence determines the equilibrium compensations to investors that are exposed to shocks. We may think of this as valuation accounting at the juncture between the Frisch (1933) vision of using shock and impulses in stochastic equilibrium models and the Bachelier (1900) vision of asset values that respond to the normal increments of a Brownian motion process. Why? Because the asset holders exposed to the random impulses affecting the macroeconomy require compensation, and the equilibrating forces affecting borrowers and lenders interacting in financial markets determine those compensatory premia.

In what follows, I illustrate two advantages to a more complete specification of the information available to investors that are reflected in my work.

4.1 Pricing Shock Exposure over Alternative Horizons

First, I explore more fully how a SDF encodes risk compensation over alternative investment horizons. I suggest a way to answer this question by describing valuation counterparts to the impulse characterizations advocated by Frisch (1933) and used extensively in quantitative macroeconomics since Sims (1980) proposed a multivariate and empirical counterpart for these characterizations. Recall that an impulse response function shows how alternative shocks tomorrow influence future values of macroeconomic variables. These shocks also represent alternative exposures to macroeconomic risk. The market-based compensations for these exposures may differ depending on the horizon over which a cash flow is realized. Many fully specified macroeconomic models proliferate shocks, including random changes in volatility, as a device for matching time series. While the additional shocks play a central role in fitting time series, eventually we must seek better answers to what lies within the black box of candidate

impulses. Understanding their role within the models is central to opening this black box in search of the answers. Empirical macroeconomists' challenges for identifying shocks for the macroeconomy also have important consequences for financial markets and the role they play in the transmission of these shocks. Not all types of candidate shocks are important for valuation.

I now discuss how we may distinguish which shock exposures command the largest market compensation and the impact of these exposures over alternative payoff horizons. I decompose the risk premia into *risk prices* and *risk exposures* using sensitivity analyses on underlying asset returns. To be specific, let X be an underlying Markov process and W a vector of shocks that are random impulses to the economic model. The state vector X_t depends on current and past shocks. I take as given a solved stochastic equilibrium model and reveal its implications for valuation. Suppose that there is an implied stochastic factor process S that evolves as:

$$\log S_{t+1} - \log S_t = \psi_s (X_t, W_{t+1}). \tag{8}$$

Typically economic models imply that this process will tend to decay over time because of the role that S plays as a discount factor. For instance, for the yield on a long-term discount bond to be positive,

$$\lim_{t \rightarrow \infty} \frac{1}{t} \log E \left[\frac{S_t}{S_0} \mid X_0 = x \right] < 0.$$

Specific models provide more structure to the function ψ_s relating the stochastic decay rate of S to the current state and next period shock. In this sense, (8) is a reduced form-relation. Similarly, consider a one-period, positive cash-flow G that satisfies:

$$\log G_{t+1} - \log G_t = \psi_g (X_t, W_{t+1}). \tag{9}$$

The process G could be aggregate consumption, or it could be a measure of aggregate corporate earnings or some other process. The logarithm of the expected one-period return of a security with this payoff is:

$$v_t = \log E \left[\frac{G_{t+1}}{G_t} \mid \mathcal{F}_t \right] - \log E \left[\frac{S_{t+1} G_{t+1}}{S_t G_t} \mid \mathcal{F}_t \right]. \tag{10}$$

So-called risk return tradeoffs emerge as we change the exposure of the cash flow to different components of the shock vector W_{t+1} .

Since cash flow growth $\frac{G_{t+1}}{G_t}$ depends on the components of W_{t+1} as a source of risk, *exposure* is altered by changing how the cash flow depends on

the underlying shocks. When I refer to *risk prices*, formally I mean the sensitivity of the logarithm of the expected return given on the left-hand side of (10) to change in cash-flow risk. I compute *risk prices* from measuring how v_t changes as we alter the cash flow, and compute *risk exposures* from examining the corresponding changes in the logarithm of the expected cash-flow growth: $\log E \left[\frac{G_{t+1}}{G_t} \middle| \mathcal{F}_t \right]$ (the first-term on the right-hand side of (10)).

These calculations are made operational by formally introducing changes in the cash-flows and computing their consequences for expected returns. When the changes are scaled appropriately, the outcomes of both the price and exposure calculations are elasticities familiar from price theory. To operationalize the term *changes*, I must impose some additional structure that allows a researcher to compute a derivative of some type. Thus I must be formal about changes in $\frac{G_{t+1}}{G_t}$ as a function of W_{t+1} . One way to achieve this formality is to take a continuous-time limit when the underlying information structure is that implied by an underlying Brownian motion as in the models of financial markets as originally envisioned by Bachelier (1900). This reproduces a common notion of a risk price used in financial economics. Another possibility is to introduce a perturbation parameter that alters locally the shock exposure, but maintains the discrete-time formulation.

These one-period or local measures have multi-period counterparts obtained by modeling the impact of small changes in the components of W_{t+1} on cash flows in future time periods, say $\frac{G_{t+\tau}}{G_t}$, for $\tau \geq 1$. Proceeding in this way, we obtain a valuation counterpart to impulse response functions featured by Frisch (1933) and by much of the quantitative macroeconomics literature. They inform us which exposures require the largest compensations and how these compensations change with the investment horizon. I have elaborated on this topic in my Fisher-Schultz Lecture paper (Hansen (2011)), and I will defer to that and related papers for more specificity and justification.²⁵ My economic interpretation of these calculations presumes a full specification of investor information as is commonly the case when analyzing impulse response functions.

4.2 A Recursive Utility Model of Investor Preferences

Next I consider investor preferences that are particularly sensitive to the assumed available information. These preferences are constructed recursively

²⁵ See Hansen et al. (2008), Hansen and Scheinkman (2009), Borovička et al. (2011), Hansen and Scheinkman (2012) and Borovička and Hansen (2014).

using continuation values for prospective consumption processes, and they are featured prominently in the macro-asset pricing literature. With these preferences the investor cares about intertemporal composition of risk as in Kreps and Porteus (1978) with these preferences. As a consequence, general revisions of the recursive utility model make investor preferences potentially sensitive to the details of the information available in the future. As I will explain, this feature of investor preferences makes it harder to implement a “do something without doing everything” approach to econometric estimation and testing.

The more general recursive utility specification nests the power utility model commonly used in macroeconomics as a special case. Interest in a more general specification was motivated in part by some of the statistical evidence that I described previously. Stochastic equilibrium models appealing to recursive utility featured in the asset pricing literature were initially advocated by Epstein and Zin (1989) and Weil (1990). They provide researchers with a parameter to alter risk preferences in addition to the usual power utility parameter known to determine the intertemporal elasticity of substitution. The one-period SDF measured using the intertemporal marginal rate of substitution is:

$$\frac{S_{t+1}}{S_t} = \exp(-\delta) \left(\frac{C_{t+1}}{C_t} \right)^{-\rho} \left[\frac{V_{t+1}}{\mathcal{R}_t(V_{t+1})} \right]^{\rho-\gamma} \tag{11}$$

where C_t is equilibrium consumption, δ is the subjective rate of discount, $\frac{1}{\rho}$ is the elasticity of intertemporal substitution familiar from power utility models, V_t is the forward-looking continuation value of the prospective consumption process, and $\mathcal{R}_t(V_{t+1})$ is the risk adjusted continuation value:

$$\mathcal{R}_t(V_{t+1}) = \left(E \left[(V_{t+1})^{1-\gamma} \mid \mathcal{F}_t \right] \right)^{\frac{1}{1-\gamma}}.$$

The parameter γ governs the magnitude of the risk adjustment. The presence of the forward-looking continuation values in the stochastic discount factor process adds to the empirical challenge in using these preferences in an economic model. When $\rho = \gamma$, the forward-looking component drops out from the SDFs and the preferences become the commonly used power utility model as is evident by comparing (7) and (11). Multi-period SDFs are the corresponding products of single period discount factors.

The empirical literature has focused on what seems to be large values for the parameter γ that adjusts for the continuation value risk. Since continuation values reflect all current prospective future consumption, increasing γ enhances

the aversion of the decision maker to consumption risk. Applied researchers have only been too happy to explore this channel. A fully solved out stochastic equilibrium model represents C and V as part of the model solution. For instance $\log C$ might have an evolution with the same form as $\log G$ as specified in (9) along a balanced stochastic growth trajectory. Representing S as in (8) presumes a solution for V_t or more conveniently $\frac{V_t}{C_t}$ as a function of X_t along with a risk adjusted counterpart to V_t and these require a full specification of investor information.

For early macro-finance applications highlighting the computation of continuation values in equilibrium models, see Hansen et al. (1999) and Tallarini (2000). The subsequent work of Bansal and Yaron (2004) showed how these preferences in conjunction with forward looking beliefs about stochastic growth and volatility have a potentially important impact on even one-period (in discrete time) or instantaneous (in continuous time) risk prices through the forward-looking channel. Hansen (2011) and Borovička et al. (2011) show that the prices of growth rate shocks are large for all payoff horizons with recursive utility and when γ is much larger than ρ . By contrast, for power utility models with large values of $\rho = \gamma$, the growth rate shock prices start off small and only eventually become large as the payoff horizon increases. The analyses in Hansen et al. (2008) and Restoy and Weil (2011) also presume that one solves for the continuation values of consumption plans or their equivalent. This general approach to the use of recursive utility for investor preferences makes explicit use of the information available to investors and hence does not allow for the robustness that I discussed in section 3.²⁶

Sometimes there is a way around this sensitivity to the information structure when conducting an econometric analysis. The empirical approach of Epstein and Zin (1991) assumes that an aggregate equity return measures the return on an aggregate wealth portfolio. In this case the continuation value relative to a risk-adjusted counterpart that appears in formula (11) is revealed by the return on the wealth portfolio for alternative choices of the preference parameters. Thus there is no need for an econometrician to compute continuation values provided that data are available on the wealth portfolio return. Epstein and Zin (1991) applied GMM methods to estimate preference parameters and test model restrictions by altering appropriately the approach in Hansen and

²⁶ Similarly, many models with heterogenous consumers/investors and incomplete markets imply pricing relation (1) for marginal agents defined as those who participate in the market over the relevant investment period. Such models require either microeconomic data and/or equilibria solutions computed using numerical methods.

Singleton (1982). Given that the one-period SDF can be constructed from consumption and return data, the full investor information set does not have to be used in the econometric implementation.²⁷ Campbell (1993) and Campbell and Vuolteenaho (2004) explored a related approach using a log-linear approximation, but this research allowed for market segmentation. Full participation in financial markets is not required because the econometric specification that is used to study the risk-return relation avoids having to use aggregate consumption. Like Epstein and Zin (1991), this approach features the return on the wealth portfolio as measured by an aggregate equity return, but now prospective beliefs about that return also contribute to the (approximate) SDF.

4.3 A Continuing Role for GMM-based Testing

Even when fully specified stochastic equilibria are formulated and used as the basis for estimation, there remains the important task of assessing the performance of the pricing implications remains. SDFs constructed from fully specified and estimated stochastic equilibrium models can be constructed *ex post* and used in testing the pricing implications for a variety of security returns. These tests can be implemented formally using direct extensions of the methods that I described in section 3. Thus the SDF specification remains an interesting way to explore empirical implications, and GMM-style statistical tests of pricing restrictions remain an attractive and viable way to analyze models.

In the remainder of this essay I will speculate on the merits of one productive approach to addressing empirical challenges based in part on promising recent research.

5 MISSPECIFIED BELIEFS

So far I have focused primarily on uncertainty *outside the model* by exploring econometric challenges, while letting risk averse agents inside the model have rational expectations. Recall that rational expectations uses the model to construct beliefs about the future.²⁸ I now consider the consequences of altering

²⁷ In contrast to recursive utility models with $\rho \neq \gamma$, often GMM-type methods can be applied to habit persistence models of the type analyzed by Sundaresan (1989), Constantinides (1990) and Heaton (1995) without having to specify the full set of information available to investors.

²⁸ A subtle distinction exists between two efforts to implement rational expectations in econometric models. When the rational expectations hypothesis is imposed in a fully specified stochastic equilibrium model, this imposition is part of an internally consistent

beliefs inside the model for two reasons. First, investor beliefs may differ from those implied by the model even if other components of the model are correctly specified. For instance, when historical evidence is weak, there is scope for beliefs that are different from those revealed by infinite histories of data. Second, if some of the model ingredients are not correct but only approximations, then the use of model-based beliefs based on an appeal to rational expectations is less compelling. Instead there is a rationale for the actors inside the model to adjust their beliefs in face of potential misspecification.

For reasons of tractability and pedagogical simplicity, throughout this and the next section I use a baseline probability model to represent conditional expectations, but not necessarily the beliefs of the people inside the model. Presuming that economic actors use the baseline model with full confidence would give rise to a rational expectations formulation, but I will explore departures from this approach. I present a tractable way to analyze how varying beliefs will alter this baseline probability model. Also, I will continue my focus on the channel by which SDFs affect asset values. A SDF and the associated risk prices, however, are only well-defined relative to a baseline model. Alterations in beliefs affect SDFs in ways that can imitate risk aversion. They also can provide an additional source of fluctuations in asset values.

My aim in this section is to study whether statistically *small* changes in beliefs can imitate what appears to be a *large* amount of risk aversion. While I feature the role of statistical discipline, explicit considerations of both learning and market discipline also come into play when there are heterogeneous consumers. For many environments there may well be an intriguing interplay between these model ingredients, but I find it revealing to narrow my focus. As is evident from recent work by Blume and Easley (2006), Kogan et al. (2011) and Borovička (2013), distorted beliefs can sometimes survive in the long run. Presumably when statistical evidence for discriminating among models is weak, the impact of market selection, whereby there is a competitive advantage of confidently

model specification a model. A model builder may impose these restrictions prior to looking at the data. The expectations become “rational” once the model is fit to data, assuming that the model is correctly specified. I used GMM and related methods to examine only a portion of the implications of a fully specified, fully solved model. In such applications, an empirical economist is not able to use a model solution to deduce the beliefs of economic actors. Instead these methods presume that the beliefs of the economic actors are consistent with historical data as revealed by the Law of Large Numbers. This approach presumes that part of the model is correctly specified, and the data are used as part of the implementation of the rational expectations restrictions.

knowing the correct model, will at the very least be sluggish. In both this and the next section, I am revisiting a theme considered by Hansen (2007a).

5.1 Martingale Models of Belief Perturbations

Consider again the asset pricing formula but now under an altered or perturbed belief relative to a baseline probability model:

$$\tilde{E}\left[\left(\frac{\tilde{S}_{t+\ell}}{\tilde{S}_t}\right)Y_{t+\ell} \mid \mathcal{F}_t\right] = Q_t \tag{12}$$

where the \tilde{E} is used to denote the perturbed expectation operator and \tilde{S} is the SDF derived under the altered expectations. Mathematically, it is most convenient to represent beliefs in an intertemporal environment using a strictly positive (with probability one) stochastic process M with a unit expectation for all $t \geq 0$. Specifically, construct the altered conditional expectations via the formula:

$$\tilde{E}[B_\tau \mid \mathcal{F}_t] = E\left[\left(\frac{M_\tau}{M_t}\right)B_\tau \mid \mathcal{F}_t\right]$$

for any bounded random variable B_τ in the date $\tau \geq t$ information set \mathcal{F}_τ . The martingale restriction imposed on M is necessary for the conditional expectations for different calendar dates to be consistent.²⁹

Using a positive martingale M to represent perturbed expectations we rewrite (12) as:

$$E\left[\left(\frac{M_{t+\ell}\tilde{S}_{t+\ell}}{M_t\tilde{S}_t}\right)Y_{t+\ell} \mid \mathcal{F}_t\right] = Q_t$$

which matches our original pricing formula (1) provided that

$$S = M\tilde{S}. \tag{13}$$

²⁹ The date zero expectation of random variable B_t that is in the \mathcal{F}_t information set may be computed in multiple ways

$$\tilde{E}[B_t \mid \mathcal{F}_0] = E\left[\left(\frac{M_t}{M_0}\right)B_t \mid \mathcal{F}_0\right] = E\left[\left(\frac{M_t}{M_0}\right)B_t \mid \mathcal{F}_0\right]$$

for any $\tau \geq t$. For this equality to hold for all bounded random variables B_t in the date t information set, $E(M_\tau \mid \mathcal{F}_t) = M_t$. This verifies that M is a martingale relative to $\{\mathcal{F}_t : t \geq 0\}$.

This factorization emerges because of the two different probability distributions that are in play. One comes from the baseline model and another is that used by investors. The martingale M makes the adjustment in the probabilities. Risk prices relative to the $\tilde{\nu}$ distribution are distinct from those relative to the baseline model. This distinction is captured by (13).

Investor models of risk aversion are reflected in the specification of \tilde{S} . For instance, example (7) implies an \tilde{S} based on consumption growth.³⁰ The martingale M would then capture the belief distortions including perhaps some of the preferred labels in the writings of others such as “animal spirits,” “over-confidence,” “pessimism,” *etc.* Without allowing for belief distortions, many empirical investigations resort to what I think of as “large values of risk aversion.” We can see, however, from factorization (13) that once we entertain belief distortions it becomes challenging to disentangle risk considerations from belief distortions.

My preference as a model builder and assessor is to add specific structure to these belief distortions. I do not find it appealing to let M be freely specified. My discussion that follows suggests a way to use some tools from statistics to guide such an investigation. They help us to understand if statistically small belief distortions in conjunction with seemingly more reasonable (at least to me) specifications of risk aversion can explain empirical evidence from asset markets.

5.2 Statistical Discrepancy

I find it insightful to quantify the statistical magnitude of a candidate belief distortion by following in part the analysis in Anderson et al. (2003). Initially, I consider a specific alternative probability distribution modeled using a positive martingale M with unit expectation and I ask if this belief distortion could be detected easily from data. Heuristically when the martingale M is close to one, the probability distortion is small. From a statistical perspective we may think of M as a relative likelihood process of a perturbed model *vis a vis* a baseline probability model. Notice that M_t depends on information in \mathcal{F}_t , and can be viewed as a “data-based” date t relatively likelihood. The ratio $\frac{M_t + 1}{M_t}$ has conditional

³⁰ When $\rho \neq \gamma$ in (11), continuation values come into play; and they would have to be computed using the distorted probability distribution. Thus M would also play a role in the construction of \tilde{S} . This would also be true in models with investor preferences that displayed “habit persistence” that is internalized when selecting investment plans. Chabi-Yo et al. (2008) nest some belief distortions inside a larger class of models with state-dependent preferences and obtain representations in which belief distortions also have an indirect impact on SDFs.

expectation equal to unity, and this term reflects how new data that arrive between dates t and $t + 1$ are incorporated into the relative likelihood.

A variety of statistical criteria measure how close M is to unity. Let me motivate one such model by bounding probabilities of mistakes. Notice that for a given threshold η ,

$$\log M_t - \eta \geq 0$$

implies that

$$[M_t \exp(-\eta)]^\alpha \geq 1 \tag{14}$$

for positive values of α . Only α 's that satisfy $0 < \alpha < 1$ interest me because only these α 's provide meaningful bounds. From (14) and Markov's Inequality,

$$Pr\{\log M_t \geq \eta | \mathcal{F}_0\} \leq \exp(-\eta\alpha) E[(M_t)^\alpha | \mathcal{F}_0]. \tag{15}$$

The left-hand side gives the probability that a log-likelihood formed with a history of length t exceeds a specified threshold η . Given inequality (15),

$$\frac{1}{t} \log Pr\{\log M_t \geq \eta | \mathcal{F}_0\} \leq -\frac{\eta\alpha}{t} + \frac{1}{t} \log E[(M_t)^\alpha | \mathcal{F}_0]. \tag{16}$$

The right-hand side of (16) gives a bound for the log-likelihood ratio to exceed a given threshold η for any $0 < \alpha < 1$. The first term on the right-hand side converges to zero as t gets large but often the second term does not and indeed may have a finite limit that is negative. Thus the negative of the limit bounds the decay rate in the probabilities as they converge to zero. When this happens we have an example of what is called a large deviation approximation. More data generated under the benchmark model makes it easier to rule out an alternative model. The decay rate bound underlies a measure of what is called Chernoff (1952) entropy. Dynamic extensions of Chernoff entropy are given by first taking limits as t gets arbitrarily large and then optimizing by the choice of α :

$$\kappa(M) = - \inf_{0 < \alpha < 1} \limsup_{t \rightarrow \infty} \frac{1}{t} \log E[(M_t)^\alpha | \mathcal{F}_0].$$

Newman and Stuck (1979) characterize Markov solutions to the limit used in the optimization problem. Minimizing over α improves the sharpness of the

bound. If the minimized value is zero, the probability distortion vanishes and investors eventually settle on the benchmark model as being correct.

A straightforward derivation shows that even when we change the roles of the benchmark model and the alternative model, the counterpart to $\kappa(M)$ remains the same.³¹ Why is Chernoff entropy interesting? When this common decay rate is small, even long histories of data are not very informative about model differences.³² Elsewhere I have explored the connection between this Chernoff measure and Sharpe ratios commonly used in empirical finance, see Anderson et al. (2003) and Hansen (2007a).³³ The Chernoff calculations are often straightforward when both models (the benchmark and perturbed models) are Markovian. In general, however, it can be a challenge to use this measure in practice without imposing considerable *a priori* structure on the alternative models.

In what follows, I will explore discrepancy measures that are similar to this Chernoff measure but are arguably more tractable to implement. What I describe builds directly on my discussion of GMM methods and extensions. Armed with factorization (13), approaches that I suggested for the study of SDFs can be adapted to the study of belief distortions. I elaborate in the discussion that follows.

5.3 Ignored Belief Distortions

Let me return to GMM estimation and model misspecification. Recall that the justification for GMM estimation is typically deduced under the premise that the underlying model is correctly specified. The possibility of permanent belief distortions, say distortions for which $\kappa(M) > 0$, add structure to the model misspecification. But this is not enough structure to identify fully the belief distortion unless an econometrician uses sufficient asset payoffs and prices to reveal the SDF. Producing bounds with this extra structure can still proceed along the lines of those discussed in Section 3.2.3 with some modifications. I sketch below one such approach.

³¹ With this symmetry and other convenient properties of $\kappa(M)$, we can interpret the measure as a metric over (equivalence classes of) martingales.

³² Bayesian and max-min decision theory for model selection both equate decay rates in type I and type II error rates.

³³ The link is most evident when a one-period (in discrete time) or local (in continuous time) measure of statistical discrimination is used in conjunction with a conditional normal distribution, instead of the large t measure described here.

Suppose the investors in the model are allowed to have distorted beliefs, and part of the estimation is to deduce the magnitude of the distortions. How big would these distortions need to be in a statistical sense in order to satisfy the pricing restrictions? What follows makes some progress in addressing this question. To elaborate, consider again the basic pricing relation with distorted beliefs written as unconditional expectation:

$$E \left[\left(\frac{M_{t+\ell} \tilde{S}_{t+\ell}}{M_t \tilde{S}_t} \right) (Y_{t+\ell})' Z_t - (Q_t)' Z_t \right] = 0.$$

As with our discussion of the study of SDFs without parametric restrictions, we allow for a multiplicity of possible martingales and impose bounds on expectations of convex functions of the ratio $\frac{M_{t+\ell}}{M_t}$.

To deduce restrictions on M , for notational simplicity I drop the t subscripts and write the pricing relation as:

$$\begin{aligned} E(m\tilde{y}' - q') &= 0 \\ E(m - 1) &= 0. \end{aligned} \tag{17}$$

To bound properties of m solve

$$\inf_{m > 0} E[\phi(m)] \tag{18}$$

subject to (17) where ϕ is given by equation (5). This formulation nests many of the so-called F -divergence measures for probability distributions including the well known Kullback-Leibler divergence ($\theta = -1, 0$). A Chernoff-type measure can be imputed by computing the bound for $-1 < \theta < 0$ and optimizing after an appropriate rescaling of the objective by $\theta(1 + \theta)$. As in the previous analysis of Section 3.2.3, there may be many solutions to the equations given in (17). While the minimization problem selects one of these, I am interested in this optimization problem to see how small the objective can be in a statistical sense. If the infimum of the objective is small, then statistically small changes in distributions suffice to satisfy the pricing restrictions. Such departures allow for “behavior biases” that are close statistically to the benchmark probabilities used in generating the data.

I have just sketched an *unconditional* approach to this calculation by allowing conditioning information to be used through the “back door” with the specification of Z but representing the objective and constraints in terms

of unconditional expectations. It is mathematically straightforward to study a conditional counterpart, but the statistical implementation is more challenging. Application of the Law of Iterated Expectations still permits an econometrician to condition on less information than investors, so there continues to be scope for robustness in the implementation. By omitting information, however, the bounds are weakened.

By design, this approach allows for the SDF to be misspecified, but in a way captured by distorted beliefs. If the SDF \tilde{S} depends on unknown parameters, say subjective discount rates, intertemporal elasticities of substitution or risk aversion parameters, then the parameter estimation can be included as part of the minimization problem. Parameter estimation takes on a rather different role in this framework than in GMM estimation. The large sample limits of the resulting parameter estimators will depend on the choice of θ unless (as assumed in much of existing econometrics literature) there are no distortions in beliefs.³⁴ Instead of featuring these methods as a way to get parameter estimators, they have potential value in helping applied econometricians infer how large probability distortions in investor beliefs would have to be from the vantage point of statistical measures of discrepancy. Such calculations would be interesting precursors or complements to a more structured analysis of asset pricing with distorted beliefs.³⁵ They could be an initial part of an empirical investigation and not the ending point as in other work using bounds in econometrics.

Martingales are present in SDF processes, even without resort to belief distortions. Alvarez and Jermann (2005), Hansen and Scheinkman (2009), Hansen

³⁴ Extensions of a GMM approach have been suggested based on an empirical likelihood approach following Qin and Lawless (1994) and Owen (2001) ($\theta = -1$), a relative-entropy approach of Kitamura and Stutzer (1997) ($\theta = 0$), a quadratic discrepancy approach of Antoine et al. (2007) ($\theta = 1$) and other related methods. Interestingly, the quadratic ($\theta = 1$) version of these methods coincides with a “continuously updating” GMM estimator of Hansen et al. (1996). Empirical likelihood methods and their generalizations estimate a discrete data distribution given the moment conditions such as pricing restrictions. From the perspective of parametric efficiency, Newey and Smith (2004) show these methods provide second-order asymptotic refinements to what is often a “second-best” efficiency problem. Recall that the statistical efficiency problem studied in Hansen (1982b) took the unconditional moment conditions as given and did not seek to exploit the flexibility in their construction giving rise to a second-best problem. Perhaps more importantly, these methods sometimes have improvements in finite sample performance but also can be more costly to implement. The rationales for such methods typically abstract from belief distortions of the type featured here and typically focus on the case of iid data generation.

³⁵ Although Gosh et al. (2012) do not feature belief distortions, with minor modification and reinterpretation their approach fits into this framework with $\theta = 0$.

(2011) and Bakshi and Chabi-Yo (2012) all characterize the role of martingale components to SDF's and their impact on asset pricing over long investment horizons. Alvarez and Jermann (2005), Bakshi and Chabi-Yo (2012) and Borovička et al. (2014a) suggest empirical methods that bound this martingale component using a very similar approach to that described here. Since there are multiple sources for martingale components to SDF's, adding more structure to what determines other sources of long-term pricing can play an essential role in quantifying the martingale component attributable to belief distortions.

In summary, factorization (13) gives an abstract characterization of the challenge faced by an econometrician outside the model trying to disentangle the effects of altered beliefs from the effects of risk aversion on the part of investors inside the model. There are a variety of ways in which beliefs could be perturbed. Many papers invoke “animal spirits” to explain lots of empirical phenomenon in isolation. However, these appeals alone do not yield the formal modeling inputs needed to build usable and testable stochastic models. Adding more structure is critical to scientific advancement if we are to develop models that are rich enough to engage in the type of policy analysis envisioned by Marschak (1953), Hurwicz (1962) and Lucas (1976). What follows uses decision theory to motivate some particular constructions of the martingale M .³⁶

Next I explore one strategy for adding structure to the martingale alterations to beliefs that I introduced in this section.

6 UNCERTAINTY AND DECISION THEORY

Uncertainty often takes a “back seat” in economic analyses using rational expectations models with risk averse agents. While researchers have used large and sometimes state dependent risk aversion to make the consequences of exposure to risk more pronounced, I find it appealing to explore uncertainty in a conceptually broader context. I will draw on insights from decision theory to suggest ways to enhance the scope of uncertainty in dynamic economic modeling. Decision theorists, economists and statisticians have wrestled with uncertainty for a very long time. For instance, prominent economists such as Keynes (1921) and Knight (1921) questioned our ability to formulate uncertainty in terms of

³⁶ An alternative way to relax rational expectations is to presume that agents solve their optimization problems using the expectations measured from survey data. See Piazzesi and Schneider (2013) for a recent example of this approach in which they fit expectations to time series data to produce the needed model inputs.

precise probabilities. Indeed Knight (1921) posed a direct challenge to time series econometrics:

We live in a world full of contradiction and paradox, a fact of which perhaps the most fundamental illustration is this: that the existence of a problem of knowledge depends on the future being different than the past, while the possibility of the solution of the problem depends on the future being like the past.

While Knight's comment goes to the heart of the problem, I believe the most productive response is not to abandon models but to exercise caution in how we use them. How might we make this more formal? I think we should use model misspecification as a source of uncertainty. One approach that has been used in econometric model-building is to let approximation errors be a source for random disturbances to econometric relations. It is typically not apparent, however, where the explicit structure comes from when specifying such errors; nor is it evident that substantively interesting misspecifications are captured by this approach. Moreover, this approach is typically adopted for an outside modeler but not for economic actors inside the model. I suspect that investors or entrepreneurs inside the models we build also struggle to forecast the future.

My co-authors and I, along with many others, are reconsidering the concept of uncertainty and exploring operational ways to broaden its meaning. Let me begin by laying out some constructs that I find to be helpful in such a discussion. When confronted with multiple models, I find it revealing to pose the resulting uncertainty as a two-stage lottery. For the purposes of my discussion, there is no reason to distinguish unknown models from unknown parameters of a given model. I will view each parameter configuration as a distinct model. Thus a model, inclusive of its parameter values, assigns probabilities to all events or outcomes within the model's domain. The probabilities are often expressed by shocks with known distributions and outcomes are functions of these shocks. This assignment of probabilities is what I will call *risk*. By contrast there may be many such potential models. Consider a two-stage lottery where in stage one we select a model and in stage two we draw an outcome using the model probabilities. Call stage one model ambiguity and stage two risk that is internal to a model.

To confront model ambiguity, we may assign *subjective* probabilities across models (including the unknown parameters). This gives us a way of averaging model implications. This approach takes a two-stage lottery and *reduces it* to a single lottery through subjective averaging. The probabilities assigned by each of a family of models are averaged using the subjective probabilities. In a dynamic setting in which information arrives over time, we update these probabilities

using Bayes' Rule. de Finetti (1937) and Savage (1954) advocate this use of subjective probability. It leads to an elegant and often tractable way to proceed. While both de Finetti (1937) and Savage (1954) gave elegant defenses for the use of subjective probability, in fact they both expressed some skepticism or caution in applications. For example, de Finetti (as quoted by Dempster (1975) based on personal correspondence) wrote:³⁷

Subjectivists should feel obligated to recognize that any opinion (so much more the initial one) is only vaguely acceptable . . . So it is important not only to know the exact answer for an exactly specified initial problem, but what happens changing in a reasonable neighborhood the assumed initial opinion.

Segal (1990) suggested an alternative approach to decision theory that avoids reducing a two-stage lottery into a single lottery. Preserving the two-stage structure opens the door to decision making in which the behavioral responses for risk (stage two) are distinct from those for what I will call *ambiguity* (stage one). The interplay between uncertainty and dynamics adds an additional degree of complexity into this discussion, but let me abstract from that complexity temporarily. Typically there is a recursive counterpart to this construction that incorporates dynamics and respects the abstraction that I have just described. It is the first stage of this lottery that will be the focus of much of the following discussion.

6.1 Robust Prior Analysis and Ambiguity Aversion

One possible source of *ambiguity*, in contrast to risk, is in how to assign subjective probabilities across the array of models. Modern decision theory gives alternative ways to confront this ambiguity from the first stage in ways that are tractable. Given my desire to use formal mathematical models, it is important to have conceptually appealing and tractable ways to represent preferences in environments with uncertainty. Such tools are provided by decision theory. Some of the literature features axiomatic development that explores the question of what is a “rational” response to uncertainty.

The de Finetti quote suggests the need for a prior sensitivity analysis. When there is a reference to a decision problem, an analysis with multiple priors can

³⁷ Similarly, Savage (1954) wrote: “No matter how neat modern operational definitions of personal probability may look, it is usually possible to determine the personal probabilities of events only very crudely.” See Berger (1984) for further discussion.

deduce bounds on the expected utility consequences of alternative decisions, and more generally a mapping from alternative priors into alternative expected outcomes. Building on discussions in Walley (1991) and Berger (1994), there are multiple reasons to consider a family of priors. This family could represent the views of alternative members of an audience, but they could also capture the *ambiguity* to a single decision maker struggling with which prior should be used. Ambiguity aversion as conceived by Gilboa and Schmeidler (1989) and others confronts this latter situation by minimizing the expected utility for each alternative decision rule. *Max-min* utility gives a higher rank to a decision rule with the larger expected utility outcome of this minimization.³⁸

Max-min utility has an extension whereby the minimization over a set of priors is replaced by a minimization over priors subject to penalization. The penalization limits the scope of the prior sensitivity analysis. The penalty is measured relative to a benchmark prior used as a point of reference. A discrepancy measure for probability distributions, for instance some of the ones I discussed previously, enforce the penalization. See Maccheroni et al. (2006) for a general analysis and Hansen and Sargent (2007) for implications using the relative entropy measure that I already mentioned. Their approach leads to what is called *variational preferences*.

For either form of ambiguity aversion, with some additional regularity conditions, a version of the Min-Max Theorem rationalizes a worst-case prior. The chosen decision rule under ambiguity aversion is also the optimal decision rule if this worst-case prior were instead the single prior of the decision maker. Dynamic counterparts to this approach do indeed imply a martingale distortion when compared to a benchmark prior that is among the set of priors that are entertained by a decision maker. Given a benchmark prior and a dynamic formulation, this worst-case outcome implies a positive martingale distortion of the type that I featured in Section 5. In equilibrium valuation, this positive martingale represents the consequences of ambiguity aversion on the part of investors inside the model. This martingale distortion emerges endogenously as a way to confront multiple priors that is ambiguity averse or robust. In sufficiently simple environments, the decision maker may in effect learn the model that generates the data in which case the martingale may converge to unity.

There is an alternative promising approach to ambiguity aversion. A decision theoretic model that captures this aversion can be embedded in the analysis of Segal (1990) and Davis and Pate-Cornell (1994), but the application to

³⁸ See Epstein and Schneider (2003) for a dynamic extension that preserves a recursive structure to decision making.

ambiguity aversion has been developed more fully in Klibanoff et al. (2005) and elsewhere. It is known as a *smooth ambiguity* model of decision making. Roughly speaking, distinct preference parameters dictate behavior responses to two different sources of uncertainty. In addition to aversion to risk given a model captured by one concave function, there is a distinct utility adjustment for ambiguity aversion that emerges when weighting alternative models using a Bayesian prior. While this approach does not in general imply a martingale distortion for valuation, as we note in Hansen and Sargent (2007), such a distortion will emerge with an exponential ambiguity adjustment. This exponential adjustment can be motivated in two ways, either as a penalization over a family of priors as in variational preferences or as a smooth ambiguity behavioral response to a single prior.

6.2 Unknown Models and Ambiguity Aversion

I now consider an approach with an even more direct link to the analysis in Section 5. An important initiator of statistical decision theory, Wald (1939), explored methods that did not presume *a priori* weights could be assigned across models. Wald (1939)'s initial work generated rather substantial literatures in statistics, control theory and economics. I am interested in such an approach as a structured way to perform an analysis of robustness. The alternative models represented as martingales may be viewed as ways in which the benchmark probability model can be misspecified. To explore robustness, I start with a family of probability models represented as martingales against a benchmark model. Discrepancy measures are most conveniently expressed in terms of convex functions of the martingales as in Section 5. Formally the ambiguity is over models, or potential misspecifications of a benchmark model.

What about learning? Suppose that the family of positive martingales with unit expectations is a convex set. For any such martingale M in this set and some $0 < \omega < 1$, construct the mixture $\omega M + (1 - \omega)1$ is a positive martingale with unit expectations. Notice that

$$\frac{\omega M_{t+\tau} + (1 - \omega)1}{\omega M_t + (1 - \omega)1} = \frac{\omega M_t \left(\frac{M_{t+\tau}}{M_t} \right) + (1 - \omega)1}{\omega M_t + (1 - \omega)1}.$$

The left-hand side is used to represent the conditional expectations operator between dates $t + \tau$ and t . If we interpret ω as the prior assigned to model M and $(1 - \omega)$ as the prior assigned to a benchmark model, then the right-hand side

reveals the outcome of Bayes' rule conditioning on date t information where M_t is a date t likelihood ratio between the two original models. Since all convex combinations are considered, we thus allow all priors including point priors. Here I have considered mixtures of the two models, but the basic logic extends to a setting with more general *a priori* averages across models.

Expected utility minimization over a family of martingales provides a tractable way to account for this form of ambiguity aversion, as in max-min utility. Alternatively the minimization can be subject to penalization as in variational preferences. Provided that we can apply the Min-Max Theorem, we may again produce a (constrained or penalized) worst-case martingale distortion. The ambiguity averse decision maker behaves as if he or she is optimizing using the worst-case martingale as the actual probability specification. This same martingale shows up in first-order conditions for optimization and hence in equilibrium pricing relationships. With this *as if* approach I can construct a distorted probability starting from a concern about model misspecification. The focus on a worst-case distortion is the outcome of a concern for robustness to model misspecification.

Of course there is no "free lunch" for such an analysis. We must limit the family of martingales to obtain interesting outcomes. The idea of conducting a sensitivity analysis would seem to have broad appeal, but of course the "devil is in the details." Research from control theory as reflected in Basar and Bernhard (1995) and Petersen et al. (2000), Hansen and Sargent (2001) and Hansen et al. (2006) and others has used discrepancies based on discounted versions of relative entropy measured by $E[M_t \log M_t | \mathcal{F}_0]$. For a given date t this measure is the expected log-likelihood ratio under the M probability model and lends itself to tractable formulas for implementation.³⁹ Another insightful formulation is given by Chen and Epstein (2002), which targets misspecification of transition densities in continuous time. Either of these approaches requires additional parameters that restrict the search over alternative models. The statistical discrepancy measures described in Section 5 provide one way to guide this choice.⁴⁰

As Hansen and Sargent (2007) emphasize, it is possible to combine this multiple models approach with a multiple priors approach. This allows simultaneously for multiple benchmark models and potential misspecification. In addition there is ambiguity in how to weight the alternative models.

³⁹ See Strzalecki (2011) for an axiomatic analysis of associated preferences.

⁴⁰ See Anderson et al. (2003) for an example of this approach.

6.3 What Might We Achieve?

For the purposes of this essay, the important outcome of this discussion is the ability to use ambiguity aversion or a concern about model misspecification as a way to generate what looks like distorted beliefs. In an application, Chamberlain (2000) studied individual portfolio problems from the vantage point of an econometrician (who could be placed inside a model) using max-min utility and featuring calculations of the endogenously determined worst-case models under plausible classes of priors. These worst-case models give candidates for the distorted beliefs mentioned in the previous section. A worst-case martingale belief distortion is part of the equilibrium calculation in the macroeconomic model of Ilut and Schneider (2014). These authors study simultaneously production and pricing using a recursive max-min formulation of the type advocated by Epstein and Schneider (2003) and introduce ambiguity shocks as an exogenous source of fluctuations.

Ambiguity aversion with unknown models provides an alternative to assuming large values of risk aversion parameters. This is evident from the control theoretic link between what is called risk sensitivity and robustness, noted in a variety of contexts including Jacobson (1973), Whittle (1981) and James (1992). Hansen and Sargent (1995) and Hansen et al. (2006) suggest a recursive formulation of risk sensitivity and link it to recursive utility as developed in the economics literature. While the control theory literature features the equivalent interpretations for decision rules, Hansen et al. (1999), Anderson et al. (2003), Maenhout (2004) and Hansen (2011) consider its impact on security market prices. This link formally relies on the use of relative entropy as a measure of discrepancy for martingales, but more generally I expect that ambiguity aversion often will have similar empirical implications to (possibly extreme) risk aversion for models of asset pricing. Formal axiomatic analyses can isolate behaviorally distinct implications. For this reason I will not overextend my claims of the observational similarity between risk and ambiguity. Axiomatic distinctions, however, are not necessarily present in actual empirical evidence.

The discussion so far produces an ambiguity component to prices in asset markets in addition to the familiar risk prices. There is no endogenous rationale for market compensations fluctuating over time. While exogenously specified stochastic volatility commonly used in asset pricing models also delivers fluctuations, this is a rather superficial success that leaves open the question of what the underlying source is for the implied fluctuations. The calculations in Hansen (2007a) and Hansen and Sargent (2010) suggest an alternative mechanism. Investors concerned with the misspecification of multiple models view these

models differently in good versus bad times. For instance, persistence in economic growth is welcome in good times but not in bad times. Given ambiguity about how to weight models and aversion to that ambiguity, investors' worst-case models shift over time leading to changes in ambiguity price components.

Introducing uncertainty about models even with a unique prior will amplify risk prices, although for local risk prices this impact is sometimes small (see Hansen and Sargent (2010) for a discussion). Introducing ambiguity aversion or a concern about model misspecification will lead to a different perspective on both the source and magnitude of the market compensations for exposure to uncertainty. Moreover, by entertaining multiple models and priors over those models there is additional scope for variation in the market compensations as investors may fear different models depending on the state of the economy.⁴¹

A framework for potential model misspecification also gives a structured way to capture "over-confidence." Consider an environment with multiple agents. Some express full commitment to a benchmark model. Others realize the model is flawed and explore the consequences of model misspecification. If indeed the benchmark model is misspecified, then agents of the first type are over-confident in the model specification. Such an approach offers a novel way to capture this form of heterogeneity in preferences.

What is missing in my discussion of model misspecification is a prescription for constructing benchmark models and/or benchmark priors. Benchmarks are important for two reasons in this analysis. They are used as a reference point for robustness and as a reference point for computing ambiguity prices. I like the transparency of simpler models especially when they have basis in empirical work, and I view the ambition to construct the perfect model to be unattainable.

7 CONCLUSION

I take this opportunity to make four concluding observations.

1. The first part of my essay explored formal econometric methods that are applicable to a researcher outside the model when actors inside the model possess rational expectations. I showed how to connect GMM estimation methods with SDF formulations of stochastic discount factors

⁴¹ See Collin-Dufresne et al. (2013) for a Bayesian formulation with parameter learning that generates interesting variation in risk prices. Given that recursive utility and a preference for robustness to model misspecification have similar and sometimes identical implications for asset pricing in other settings, it would be of interest to see if this similarity carries over to the parameter learning environments considered by these authors.

to estimate and assess asset pricing models with connections to the macroeconomy. I also described how to use SDF formulations to assess the empirical implications of asset pricing models more generally. I then shifted to a discussion of investor behavior inside the model, perhaps even motivated by my own experiences as an applied econometrician. More generally these investors may behave as if they have distorted beliefs. I suggested statistical challenges and concerns about model misspecification as a rationale for these distorted beliefs.

2. I have identified ways that a researcher might alter beliefs for the actors within a model, but I make no claim that this is the only interesting way to structure such distortions. Providing structure, however, is a prerequisite to formal assessment of the resulting models. I have also suggested statistical measures that extend the rational expectations appeal to the Law of Large Numbers for guiding the types of belief distortions that are reasonable to consider. This same statistical assessment should be a valuable input into other dynamic models within which economic agents have heterogeneous beliefs.
3. How best to design econometric analysis in which econometricians and agents formally acknowledge this misspecification is surely a fertile avenue for future research. Moreover, there remains the challenge of how best to incorporate ambiguity aversion or concerns about model misspecification into a Marschak (1953), Hurwicz (1962) and Lucas (1972) style study of counterfactuals and policy interventions.
4. Uncertainty, generally conceived, is not often embraced in public discussions of economic policy. When uncertainty includes incomplete knowledge of dynamic responses, we might well be led away from arguments that “complicated problems require complicated solutions.” When complexity, even formulated probabilistically, is not fully understood by policy makers, perhaps it is the simpler policies that are more prudent. This could well apply to the design of monetary policy, environmental policy and financial market oversight. Enriching our toolkit to address formally such challenges will improve the guidance that economists give when applying models to policy analysis.

REFERENCES

- Ai, Chunrong and Xiaohong Chen. 2003. Efficient Estimation of Models with Conditional Moment Restrictions Containing Unknown Functions. *Econometrica* 71 (6):1795–1843.

- Almeida, Caio and Rene Garcia. 2013. Robust Economic Implications of Nonlinear Pricing Kernels. Tech. rep., Social Science Research Network. <http://ssrn.com/abstract=1107997>.
- Alvarez, Fernando and Urban J. Jermann. 2005. Using Asset Prices to Measure the Persistence of the Marginal Utility of Wealth. *Econometrica* **73** (6):1977–2016.
- Anderson, Evan W., Lars Peter Hansen, and Thomas J. Sargent. 2003. A Quartet of Semigroups for Model Specification, Robustness, Prices of Risk, and Model Detection. *Journal of the European Economic Association* **1** (1):68–123.
- Antoine, Bertille, Helene Bonnal, and Eric Renault. 2007. On the Efficient Use of Informational Content of Estimating Equations: Implied Probabilities and Euclidean Empirical Likelihood. *Journal of Econometrics* **138**:461–487.
- Arellano, Manuel. 2002. Sargan's Instrumental Variables Estimation and the Generalized Method of Moments. *Journal of Business and Economic Statistics* **20** (4):450–459.
- . 2003. *Panel Data Econometrics*. Advanced Texts in Econometrics. Oxford University Press.
- Bachelier, Louis. 1900. Theorie de la Speculation. *Annales Scientifiques de l' Ecole Normale Supérieure* **17**:21–86.
- Backus, David, Mikhail Chernov, and Ian Martin. 2011. Disasters Implied by Equity Index Options. *The Journal of Finance* **66** (6):1969–2012.
- Backus, David K., Mikhail Chernov, and Stanley E. Zin. 2014. Sources of Entropy in Representative Agent Models. *Journal of Finance* **69** (1):51–99.
- Bakshi, Gurdip and Fousseni Chabi-Yo. 2012. Variance Bounds on the Permanent and Transitory Components of Stochastic Discount Factors. *Journal of Financial Economics* **105** (1): 191–208.
- Bansal, Ravi and Bruce N. Lehmann. 1997. Growth-Optimal Portfolio Restrictions on Asset Pricing Models. *Macroeconomic Dynamics* **1** (02):333–354.
- Bansal, Ravi and Amir Yaron. 2004. Risks for the Long Run: A Potential Resolution of Asset Pricing Puzzles. *Journal of Finance* **59** (4):1481–1509.
- Basar, Tamer and Pierre Bernhard. 1995. *H[∞]-Optimal Control and Related Minimax Design Problems: A Dynamic Game Approach*. Birkhauser.
- Berger, James O. 1984. The Robust Bayesian Viewpoint. In *Robustness of Bayesian Analysis*, edited by Joseph B. Kadane, 63–144. North-Holland.
- . 1994. An Overview of Robust Bayesian Analysis (with discussion). *Test* **3** (1):5–124.
- Bilson, John F. O. 1981. The “Speculative Efficiency” Hypothesis. *The Journal of Business* **54** (3):435–451.
- Blume, Lawrence and David Easley. 2006. If You're so Smart, why Aren't You Rich? Belief Selection in Complete and Incomplete Markets. *Econometrica* **74** (4):929–966.
- Borovička, Jaroslav. 2013. Survival and Long-Run Dynamics with Heterogeneous Beliefs Under Recursive Preferences. Tech. rep., New York University.
- Borovička, Jaroslav and Lars Peter Hansen. 2014. Examining Macroeconomic Models Through the Lens of Asset Pricing. *Journal of Econometrics* forthcoming.

- Borovička, Jaroslav, Lars Peter Hansen, Mark Hendricks, and José A. Scheinkman. 2011. Risk Price Dynamics. *Journal of Financial Econometrics* **9**:3–65.
- Borovička, Jaroslav, Lars Peter Hansen, and José A. Scheinkman. 2014a. Misspecified Recovery. SSRN Working Paper.
- , 2014b. Shock Elasticities and Impulse Response Functions. *Mathematics and Financial Economics* Forthcoming.
- Burnside, Craig. 1994. Hansen-Jagannathan Bounds as Classical Tests of Asset-Pricing Models. *Journal of Business and Economic Statistics* **12** (1):57–79.
- Campbell, John Y. 1993. Intertemporal Asset Pricing Without Consumption Data. *The American Economic Review* **83** (3):487–512.
- Campbell, John Y. and John Cochrane. 1999. Force of Habit: A Consumption-Based Explanation of Aggregate Stock Market Behavior. *Journal of Political Economy* **107** (2):205–251.
- Campbell, John Y and Thomo Vuolteenaho. 2004. Bad Beta, Good Beta. *American Economic Review* **94** (5):1249–1275.
- Chabi-Yo, Fousseni, Rene Garcia, and Eric Renault. 2008. State Dependence Can Explain the Risk Aversion Puzzle. *Review of Financial Studies* **21** (2):973–1011.
- Chamberlain, Gary. 1987. Asymptotic Efficiency in Estimation with Conditional Moment Restrictions. *Journal of Econometrics* **34** (3):305–334.
- . 1992. Efficiency Bounds for Semiparametric Regression. *Econometrica* **60** (3):567–96.
- . 2000. Econometric Applications of Maxmin Expected Utility. *Journal of Applied Econometrics* **15** (6):625–644.
- Chen, Zengjing and Larry Epstein. 2002. Ambiguity, Risk, and Asset Returns in Continuous Time. *Econometrica* **70** (4):1403–1443.
- Chernoff, Herman. 1952. A Measure of Asymptotic Efficiency for Tests of a Hypothesis Based on the Sum of Observations. *The Annals of Mathematical Statistics* **23** (4):493–507.
- Chernozhukov, Victor, Emre Kocaturum, and Konrad Menzel. 2013. Inference on Sets in Finance. Unpublished, MIT, DRW Trading Group and NYU.
- Cochrane, John. 2008. *Financial Markets and the Real Economy*, 237–322. Elsevier.
- Cochrane, John H. and Lars Peter Hansen. 1992. Asset Pricing Explorations for Macroeconomics. *NBER Macroeconomics Annual* **7**:115–165.
- Collin-Dufresne, Pierre, Michael Johannes, and Lars A. Lochstoer. 2013. Parameter Learning in General Equilibrium: The Asset Pricing Implications. Tech. rep., Columbia University, 3022 Broadway, New York, NY 10027.
- Constantinides, George M. 1990. Habit Formation: A Resolution of the Equity Premium Puzzle. *Journal of Political Economy* **98** (3):519–543.
- Cressie, Noel and Timothy R. C. Read. 1984. Multinomial Goodness-of-Fit Tests. *Journal of the Royal Statistical Society. Series B (Methodological)* **46** (3):440–464.
- Davis, Donald B. and M.-Elisabeth Pate-Cornell. 1994. A Challenge to the Compound Lottery Axiom: A Two-Stage Normative Structure and Comparison to Other Theories. *Theory and Decision* **37** (3):267–309.

- Davis, Mark and Alison Etheridge. 2006. *Louis Bachelier's Theory of Speculation: The Origins of Modern Finance*. Princeton University Press.
- Dempster, A.P. 1975. A Subjectivist Look at Robustness. *Bulletin of the International Statistical Institute* **46**:349–374.
- Dimson, Elroy and Massoud Mussavian. 2000. Market Efficiency. In *The Current State of Business Disciplines*, vol. 3, 959–970. Spellbound Publications.
- Epstein, Larry and Stanley E. Zin. 1989. Substitution, Risk Aversion and the Temporal Behavior of Consumption and Asset Returns: A Theoretical Framework. *Econometrica* **57** (4):937–969.
- Epstein, Larry G. and Martin Schneider. 2003. Recursive Multiple-Priors. *Journal of Economic Theory* **113** (1):1–31.
- Epstein, Larry G and Stanley E Zin. 1991. Substitution, Risk Aversion, and the Temporal Behavior of Consumption and Asset Returns: An Empirical Analysis. *Journal of Political Economy* **99** (2):263–86.
- Fama, Eugene F. 1984. Forward and Spot Exchange Rates. *Journal of Monetary Economics* **14** (3):319–338.
- Fama, Eugene F and Kenneth R. French. 1992. The Cross-Section of Expected Returns. *Journal of Finance* **47** (2):427–465.
- de Finetti, Bruno. 1937. La Prevision: Ses Lois Logiques, ses Sources Subjectives. *Annales de l'Institut Henri Poincaré* **7**:1–68.
- Frisch, Ragnar. 1933. Propagation Problems and Impulse Problems in Dynamic Economics. In *Economic Essays in Honour of Gustav Cassel*, 171–205. Allen and Unwin.
- Frisch, Ragnar. 1933b. Editor's note. *Econometrica* **1** (1): 1–4.
- Gallant, A. Ronald, Lars Peter Hansen, and George Tauchen. 1990. Using Conditional Moments of Asset Payoffs to Infer the Volatility of Intertemporal Marginal Rates of Substitution. *Journal of Econometrics* **45**:141–179.
- Ghysels, Eric and Alastair Hall. 2002. Interview with Lars Peter Hansen. *Journal of Business and Economic Statistics Twentieth Anniversary Issue on the Generalized Method of Moments* **20** (4):442–447.
- Gilboa, Itzhak and David Schmeidler. 1989. Maxmin Expected Utility with Non-Unique Prior. *Journal of Mathematical Economics* **18** (2):141–153.
- Gordin, Mikhail I. 1969. The Central Limit Theorem for Stationary Processes. *Soviet Mathematics Doklady* **10**:1174–1176.
- Gosh, Anisha, Christian Julliard, and Alex Taylor. 2012. What is the Consumption-CAPM Missing? An Information-Theoretic Framework for the Analysis of Asset Pricing Models. Working paper, Tepper School of Business, Carnegie Mellon University.
- Grossman, Sanford J. and Robert J. Shiller. 1981. The Determinants of the Variability of Stock Market Prices. *The American Economic Review* **71** (2):222–227.
- Haavelmo, Trygve. 1944. The Probability Approach in Econometrics. *Econometrica* **12** supplement:1–115.
- Hansen, Lars Peter. 1982a. Consumption, Asset Markets, and Macroeconomic Fluctuations: A Comment. *Carnegie-Rochester Conference Series on Public Policy* **17**:239–250.

- . 1982b. Large Sample Properties of Generalized Method of Moments Estimators. *Econometrica* **50** (4):1029–1054.
- . 1985. A Method for Calculating Bounds on the Asymptotic Covariance Matrices of Generalized Method of Moments Estimators. *Journal of Econometrics* **30** (1):203–238.
- . 2001. Generalized Method of Moments Estimation: A Time Series Perspective (published title Method of Moments). In *International Encyclopedia of the Social and Behavior Sciences*, edited by S.E. Fienberg and J.B. Kadane. Pergamon: Oxford. <http://www.larspeterhansen.org/generalized-method-of-moments-estimation-32.html>.
- . 2007a. Beliefs, Doubts and Learning: Valuing Macroeconomic Risk. *American Economic Review* **97** (2):1–30.
- . 2007b. Generalized Method of Moments Estimation. In *The New Palgrave Dictionary of Economics*, edited by Steven N. Durlauf and Lawrence E. Blume. Palgrave Macmillan.
- . 2008. *Discussion of: Financial Markets and the Real Economy*, by J. Cochrane, 326–329. Elsevier.
- . 2011. Dynamic Valuation Decomposition Within Stochastic Economies. *Econometrica* **80** (3):911–967. Fisher-Schultz Lecture at the European Meetings of the Econometric Society.
- . 2012. Proofs for Large Sample Properties of Generalized Method of Moments Estimators. *Journal of Econometrics* **170** (2):325–330.
- Hansen, Lars Peter and Robert J. Hodrick. 1980. Forward Exchange Rates as Optimal Predictors of Future Spot Rates: An Econometric Analysis. *Journal of Political Economy* **88** (5):829–853.
- . 1983. Risk Averse Speculation in the Forward Foreign Exchange Market: An Econometric Analysis of Linear Models. In *Exchange Rates and International Macroeconomics*, NBER Chapters, 113–152. National Bureau of Economic Research, Inc.
- Hansen, Lars Peter, C.A. Helt, and D. Peled. 1978. A note on first degree stochastic dominance *Economics Letters* **1** (4): 315–319.
- Hansen, Lars Peter and Ravi Jagannathan. 1991. Implications of Security Market Data for Models of Dynamic Economies. *Journal of Political Economy* **99** (2):225–262.
- . 1997. Assessing Specification Errors in Stochastic Discount Factor Models. *The Journal of Finance* **52** (2):557–590.
- Hansen, Lars Peter and Scott F. Richard. 1987. The Role of Conditioning Information in Deducing Testable Restrictions Implied by Dynamic Asset Pricing Models. *Econometrica* **50** (3):587–613.
- Hansen, Lars Peter and Thomas Sargent. 2010. Fragile Beliefs and the Price of Uncertainty. *Quantitative Economics* **1** (1):129–162.
- Hansen, Lars Peter and Thomas J. Sargent. 1980. Formulating and Estimating Dynamic Linear Rational Expectations Models. *Journal of Economic Dynamics and Control* **2**:7–46.

- . 1991. Exact Linear Rational Expectations Models: Specification and Estimation. In *Rational Expectations Econometrics: Specification and Estimation*, edited by Lars Peter Hansen and Thomas J. Sargent, 45–76. Westview Press.
- . 1995. Discounted Linear Exponential Quadratic Gaussian Control. *IEEE Transactions on Automatic Control* **40** (5):968–971.
- . 2001. Robust Control and Model Uncertainty. *The American Economic Review* **91** (2):60–66.
- . 2007. Recursive Robust Estimation and Control Without Commitment. *Journal of Economic Theory* **136** (1):1–27.
- Hansen, Lars Peter and Jose Scheinkman. 2009. Long-Term Risk: An Operator Approach. *Econometrica* **77** (1):117–234.
- . 2012. Pricing Growth-Rate Risk. *Finance and Stochastics* **16** (1):1–15.
- Hansen, Lars Peter and Kenneth J. Singleton. 1982. Generalized Instrumental Variables Estimation of Nonlinear Rational Expectations Models. *Econometrica* **50** (5):1269–1286.
- . 1983. Stochastic Consumption, Risk Aversion, and the Temporal Behavior of Asset Returns. *Journal of Political Economy* **91** (2):249–265.
- . 1996. Efficient Estimation of Linear Asset-Pricing Models with Moving Average Errors. *Journal of Business and Economic Statistics* **14** (1):53–68.
- Hansen, Lars Peter, John Heaton, and Erzo G. J. Luttmer. 1995. Econometric Evaluation of Asset Pricing Models. *The Review of Financial Studies* **8** (2):237–274.
- Hansen, Lars Peter, John Heaton, and Amir Yaron. 1996. Finite-Sample Properties of Some Alternative GMM Estimators. *Journal of Business and Economic Statistics* **14** (3):262–280.
- Hansen, Lars Peter, Thomas J. Sargent, and Jr. Tallarini, Thomas D. 1999. Robust Permanent Income and Pricing. *The Review of Economic Studies* **66** (4):873–907.
- Hansen, Lars Peter, Thomas J. Sargent, Gauhar A. Turmuhambetova, and Noah Williams. 2006. Robust Control and Model Misspecification. *Journal of Economic Theory* **128** (1):45–90.
- Hansen, Lars Peter, John C. Heaton, and Nan Li. 2008. Consumption Strikes Back?: Measuring Long Run Risk. *Journal of Political Economy* **116** (2):260–302.
- Harrison, J. Michael and David M. Kreps. 1979. Martingales and Arbitrage in Multiperiod Securities Markets. *Journal of Economic Theory* **20** (3):381–408.
- He, Hua and David M. Modest. 1995. Market Frictions and Consumption-Based Asset Pricing. *Journal of Political Economy* **103** (1):94–117.
- Heaton, John C. 1995. An Empirical Investigation of Asset Pricing with Temporally Dependent Preference Specifications. *Econometrica* **63** (3):681–717.
- Hurwicz, Leonid. 1962. On the Structural Form of Interdependent Systems. *Logic, Methodology and Philosophy of Science: Proceedings of the 1960 International Congress* 232–239.
- Ilut, Cosmin and Martin Schneider. 2014. Ambiguous Business Cycles. *American Economic Review* forthcoming.

- Jacobson, David H. 1973. Optimal Stochastic Linear Systems with Exponential Performance Criteria and Their Relation to Deterministic Differential Games. *IEEE Transactions for Automatic Control* **AC-18** (2):1124–131.
- James, Matthew R. 1992. Asymptotic Analysis of Nonlinear Stochastic Risk-Sensitive Control and Differential Games. *Mathematics of Control, Signals and Systems* **5** (4):401–417.
- Keynes, John Maynard. 1921. *A Treatise on Probability*. London: Macmillan.
- Kitamura, Yuichi and Michael Stutzer. 1997. An Information-Theoretic Alternative to Generalized Method of Moments Estimation. *Econometrica* **65** (4):861–874.
- Kitamura, Yuichi, Gautam Tripathi, and Hyungtaik Ahn. 2004. Empirical Likelihood-Based Inference in Conditional Moment Restriction Models. *Econometrica* **72** (6):1667–1714.
- Klibanoff, Peter, Massimo Marinacci, and Sujoy Mukerji. 2005. A Smooth Model of Decision Making Under Uncertainty. *Econometrica* **73** (6):1849–1892.
- Knight, Frank H. 1921. *Risk, Uncertainty, and Profit*. Houghton Mifflin.
- Kogan, Leonid, Stephen A. Ross, Jiang Wang, and Mark M. Westerfield. 2011. Market Selection. Tech. rep., Massachusetts Institute of Technology.
- Kreps, David M. and Evan L. Porteus. 1978. Temporal Resolution of Uncertainty and Dynamic Choice. *Econometrica* **46** (1):185–200.
- Lucas, Robert E. 1972. Econometric Testing of the Natural Rate Hypothesis. In *The Econometrics of Price Determination*, edited by O. Eckstein, 50–59. Board of Governors of the Federal Reserve Board.
- . 1976. Econometric Policy Evaluation: A Critique. *Carnegie-Rochester Conference Series on Public Policy* **1** (0):19–46.
- . 1978. Asset Prices in an Exchange Economy. *Econometrica* **46** (6):1429–1445.
- Luttmer, Erzo G. J. 1996. Asset Pricing in Economies with Frictions. *Econometrica* **64** (6):1439–1467.
- Maccheroni, Fabio, Massimo Marinacci, and Aldo Rustichini. 2006. Ambiguity Aversion, Robustness, and the Variational Representation of Preferences. *Econometrica* **74** (6):1147–1498.
- Maenhout, Pascal J. 2004. Robust Portfolio Rules and Asset Pricing. *Review of Financial Studies* **17**:951–983.
- Marschak, Jacob. 1953. Economic Measurements for Policy and Prediction. In *Studies in Econometric Method*, edited by Tjalling Charles Koopmans and William C. Hood, 1–26. John Wiley and Sons.
- Muth, John H. 1961. Rational Expectations and the Theory of Price Movements. *Econometrica* **29** (3):315–335.
- Newey, Whitney K. and Richard J. Smith. 2004. Higher Order Properties of GMM and Generalized Empirical Likelihood Estimators. *Econometrica* **72** (1):219–255.
- Newey, W.K. 1990. Efficient Instrumental Variables Estimation of Nonlinear Models. *Econometrica* **58** (4):809–837.

- . 1993. Efficient Estimation of Models with Conditional Moment Restrictions. In *Handbook of Statistics*, vol. 11, edited by G.S. Maddala, C.R. Rao, and H.D. Vinod, 809–837. Amsterdam: North-Holland.
- Newman, C. M. and B. W. Stuck. 1979. Chernoff Bounds for Discriminating Between Two Markov Processes. *Stochastics* 2 (1–4):139–153.
- Owen, Art B. 2001. *Empirical Likelihood*, vol. 92 of *CRC Monographs on Statistics and Applied Probability*. Chapman and Hall.
- Peñaranda, Francisco and Enrique Sentana. 2011. Inferences about Portfolio and Stochastic Discount Factor Mean Variance Frontiers. Working Paper, UPF and CEMFI.
- Petersen, I.R., M.R. James, and P. Dupuis. 2000. Minimax Optimal Control of Stochastic Uncertain Systems with Relative Entropy Constraints. *Automatic Control, IEEE Transactions on* 45 (3):398–412.
- Piazzesi, Monika and Martin Schneider. 2013. Inflation and the Price of Real Assets. Unpublished, Stanford University.
- Qin, Jin and Jerry Lawless. 1994. Empirical Likelihood and General Estimating Equations. *The Annals of Statistics* 22 (1):300–325.
- Restoy, Fernando and Philippe Weil. 2011. Approximate Equilibrium Asset Prices. *Review of Finance* 15 (1):1–28.
- Robinson, P. M. 1987. Asymptotically Efficient Estimation in the Presence of Heteroskedasticity of Unknown Form. *Econometrica* 55 (4):875–891.
- Ross, Stephen A. 1978. A Simple Approach to the Valuation of Risky Streams. *The Journal of Business* 51 (3):453–75.
- Rubinstein, Mark. 1976. The Valuation of Uncertain Income Streams and the Pricing of Options. *The Bell Journal of Economics* 7:407–425.
- Sargan, J. D. 1958. The Estimation of Economic Relationships Using Instrumental Variables. *Econometrica* 26 (3):393–415.
- . 1959. The Estimation of Relationships with Autocorrelated Residuals by the Use of Instrumental Variables. *Journal of the Royal Statistical Society. Series B (Methodological)* 21 (1):91–105.
- Sargent, Thomas J. 1973. Rational Expectations, the Real Rate of Interest, and the Natural Rate of Unemployment. *Brookings Papers in Economic Activity* 4 (2):429–480.
- . 1999. *The Conquest of American Inflation*. Princeton, New Jersey: Princeton University Press.
- Sargent, Thomas J. and Neil Wallace. 1975. "Rational" Expectations, the Optimal Monetary Instrument, and the Optimal Money Supply Rule. *Journal of Political Economy* 83 (2):241–254.
- Savage, Leonard J. 1954. *The Foundations of Statistics*. Wiley Publications in Statistics.
- . 1961. The Foundations of Statistics Reconsidered. In *Proceedings of the Fourth Berkeley Symposium on Mathematical Statistics and Probability*, edited by Jerzy Neyman, 575–586.
- Segal, Uzi. 1990. Two-Stage Lotteries Without the Reduction Axiom. *Econometrica* 58 (2):349–377.
- Shiller, Robert. 1972. *Rational Expectations and the Structure of Interest Rates*. Ph.D. thesis, M.I.T.

- . 1982. Consumption, Asset Markets and Macroeconomic Fluctuations. *Carnegie Rochester Conference Series on Public Policy* **17**:203–238.
- Sims, Christopher A. 1980. Macroeconomics and Reality. *Econometrica* **48** (1):1–48.
- . 2012. Statistical Modeling of Monetary Policy and Its Effects. *The American Economic Review* **102** (4):1187–1205.
- Slutsky, Eugen. 1927. The Summation of Random Causes as the Source of Cyclic Processes. In *Problems of Economic Conditions*, vol. 3. Moscow: The Conjunction Institute.
- . 1937. The Summation of Random Causes as the Source of Cyclic Processes. *Econometrica* **5** (2):105–146.
- Snow, Karl N. 1991. Diagnosing Asset Pricing Models Using the Distribution of Asset Returns. *The Journal of Finance* **46** (3):955–983.
- Stigler, Stephen. 2014. Soft Questions, Hard Answers: Jacob Bernoulli's Probability in Historical Context. *International Statistical Review* forthcoming.
- Strzalecki, Tomasz. 2011. Axiomatic Foundations of Multiplier Preferences. *Econometrica* **79**:47–73.
- Stutzer, Michael. 1995. A Bayesian Approach to Diagnosis of Asset Pricing Models. *Journal of Econometrics* **68** (2):367–397.
- Sundaresan, Suresh M. 1989. Intertemporally Dependent Preferences and the Volatility of Consumption and Wealth. *The Review of Financial Studies* **2** (1):73–89.
- Tallarini, Thomas D. 2000. Risk-Sensitive Real Business Cycles. *Journal of Monetary Economics* **45** (3):507–532.
- Wald, Abraham. 1939. Contributions to the Theory of Statistical Estimation and Testing Hypotheses. *The Annals of Mathematical Statistics* **10** (4):299–326.
- Walley, Peter. 1991. *Statistical Reasoning with Imprecise Probabilities*. CRC Monographs on Statistics and Applied Probability. Chapman and Hall.
- Wallis, Kenneth F. 1980. Econometric Implications of the Rational Expectations Hypothesis. *Econometrica* **48** (1):49–73.
- Weil, Philippe. 1990. Unexpected Utility in Macroeconomics. *The Quarterly Journal of Economics* **105** (1):29–42.
- West, Kenneth, Ka fu Wong, and Stanislav Anatolyev. 2009. Instrumental Variables Estimation of Heteroskedastic Linear Models Using All Lags of Instruments. *Econometric Reviews* **28** (5):441–467.
- Whittle, Peter. 1981. Risk Sensitive Linear Quadratic Gaussian Control. *Advances in Applied Probability* **13** (4):764–777.
- Yule, G. Udny. 1927. On a Method of Investigating Periodicities in Disturbed Series, with Special Reference to Wolfer's Sunspot Numbers. *Philosophical Transactions of the Royal Society of London. Series A, Containing Papers of a Mathematical or Physical Character* **226**:267–298.
- Zhang, Jian and Irene Gijbels. 2003. Sieve Empirical Likelihood and Extensions of the Generalized Least Squares. *Scandinavian Journal of Statistics* **30** (1):1–24.

Portrait photo of Lars Peter Hansen by photographer Alexander Mahmoud.