Discussion of "Robustness and Pricing with Uncertain Growth"

Pascal J. Maenhout
INSEAD

1. Why Robustness?

Hansen and Sargent, and a number of coauthors have recently initiated a research agenda that introduces the notion of robustness to model uncertainty (or a concern with model misspecification) into macroeconomics and finance. Methodologically they modify key techniques from the robust control literature in applied mathematics in order to allow for recursivity, discounting, and stochastic dynamics, which are obvious ingredients of the conventional modeling paradigm in financial economics. The fundamental idea behind robustness is that economic models are at best viewed as stylized approximations of reality rather than perfectly accurate descriptions thereof. It is then natural to entertain the possibility that economic agents making decisions in these models share the concern for model misspecification that economists and econometricians routinely express. When agents use a particular model as guidance in a dynamic decision problem and worry that this model might be misspecified, one would expect them to insist on considering alternative models in order to obtain decision rules that not only work well when the particular baseline model is correct, but that also perform reasonably well when the model is somewhat misspecified. This is the sense in which the decision rule is designed to be robust. The question of course is how agents go about considering alternative models and obtaining a robust decision rule given these alternative models. In this article, as in other work introducing robustness, the agents are assumed to achieve robustness by considering a least-favorable or worst-case model that is similar to and statistically hard to distinguish from the baseline model.

Worrying about a worst-case distortion of the baseline model seems quite drastic. First, it should be emphasized that the least-favorable model is (endogenously) selected from a set of models that are "close" to the baseline model (where distance between models is measured by an entropy or likelihood-type criterion) and in accordance with a preference parameter that measures the strength of the preference for robustness (θ). Effectively the preference parameter θ dictates how much confidence the decision maker has in the baseline model so that $\frac{1}{\theta}$ indicates to what extent she is willing to consider

Address correspondence to Pascal J. Maenhout, Finance Department, INSEAD, Boulevard de Constance, 77305 Fontainebleau Cedex, France, or e-mail: pascal.maenhout@insead.edu.

The Review of Financial Studies Special 2002 Vol. 15, No. 2, pp. 405–411

alternative models that deviate from it. In particular, setting $\theta = \infty$ nests standard expected utility decision making where the decision maker has infinite confidence in the benchmark model and is not considering any alternative model at all. Finite values of θ correspond to situations where the decision maker is skeptical about the benchmark model and considers alternative models that are similar yet pessimistically perturbed.

But why, given a set of alternative models indexed by θ, focus on the worst case in that set? Why not follow standard Bayesian decision making where each model is assigned a probability and model uncertainty is integrated out? A decision-theoretic justification can be based on the work by Gilboa and Schmeidler (1989). They axiomatically characterize maxmin behavior and motivate their analysis by observed behavior that violates the expected utility framework, most notably the Ellsberg paradox. In this spirit, θ becomes a measure of (model) uncertainty aversion. Essentially the idea is that the decision maker is unable to assign probabilities to the alternative models considered. Uncertainty aversion leads her to concentrate on the least-favorable element in the set of models. Interpreted this way, robustness is conceptually similar to the work on ambiguity aversion and multiple prior expected utility of Chen and Epstein (2001), who build directly on the analysis of Gilboa and Schmeidler.

Finally, the pessimistic nature of the maxmin approach, absent from standard Bayesian modeling efforts, is especially promising for understanding important asset pricing phenomena. Mehra and Prescott's (1985) equity premium puzzle is often interpreted as indicative of (overly) cautious and conservative investor behavior. Instead of fully attributing this caution to high-risk aversion, robustness identifies an additional source of conservativism: aversion to (Knightian or model) uncertainty or ambiguity rather than aversion to risk. Along these lines, robustness potentially provides a foundation for the assumed behavior in Cecchetti, Lam, and Mark (2000) and Abel (2001).

2. Main Results

The authors introduce robustness in a standard continuous-time stochastic growth economy with production. The mean growth rate of productivity is not observed and has to be inferred by the representative agent from observations of the level of productivity. Productivity growth is driven by low-frequency regime switches between high- and low-growth states and by high-frequency Brownian increments. Learning about the mean growth rate occurs by means of an appropriate Bayesian updating scheme, that is, the nonlinear Kalman filter. This makes the setup quite similar to recent work of David (1997) and Veronesi (1999). The novelty in this article is to simultaneously allow for robustness in an environment with regime switching and filtering. Methodologically this is a nontrivial contribution.

As described before, agents achieve robustness by considering a worst-case alternative model. Algorithmically this maxmin procedure can be interpreted as a zero-sum game between the decision maker and a fictitious agent ("Nature") who is belligerently deceitful and perturbs the baseline model. This interpretation is useful for understanding the difference between games I and II as developed by the authors. In particular, when combining robustness and learning about the growth rate of the economy, the authors are careful about the information structure for both the decision maker and the fictitious malicious agent. The growth rate of the economy is not observed by the decision maker, but might or might not be known to Nature. This is precisely the difference between game I and game II. In game II, Nature is given an informational advantage relative to the decision maker and is allowed to condition on the actual growth state of the economy when chosing the worst-case model. Model misspecification becomes state dependent, which makes robustness especially potent.

Within this framework, the resource allocation problem of a social planner is solved numerically. A first important result is the extension of the observational equivalence result of Hansen, Sargent, and Tallarini (1999) to the regime-switching environment with filtering of this article. An economy without robustness is shown to be isomorphic in terms of macroeconomic quantities to an economy with robustness, but with more impatience. A concern for model misspecification enhances the precautionary savings motive and makes investors accumulate more capital. This can be offset by increasing the rate of time preference. The importance of this result is twofold. First, it implies that quantity predictions of standard macroeconomic models can be preserved when introducing robustness. Second, it means that robustness will not necessarily affect the risk-free rate in the economy, even though it typically increases the equilibrium equity premium. Explanations of the equity premium puzzle based on robustness are thus able to avoid Weil's (1989) risk-free rate puzzle, unlike explanations based on high-risk aversion in the context of an expected utility isoelastic model [see Maenhout (2001) for an analytical expression of the equilibrium riskfree rate in a Brownian exchange economy with robustness].

Given the solution to the planner's problem, shadow risk premia are computed. As in Anderson, Hansen, and Sargent (2000), Chen and Epstein (2001), and Maenhout (2001), the authors obtain a Breeden-type consumption capital asset pricing model, where the equilibrium Sharpe ratio is the sum of the standard price of risk and a price of model uncertainty or ambiguity aversion. However, in this article the market price of model uncertainty is a nonlinear function of the posterior probability of the low-growth state. Because this posterior probability changes dramatically over time, the equilibrium equity premium becomes time varying.

Finally, the article introduces a second Brownian motion and an earnings process in order to analyze the model's predictions in terms of price-earnings ratios and stock price volatility. While robustness improves the predictions of the model somewhat, the excess volatility puzzle remains to a large extent unexplained.

3. Open Questions and Extensions

A number of questions and challenges can be brought up for discussion concerning the analysis and results in this article. Finally, the article suggests some interesting and promising extensions for further research.

A first important question concerns the updating rule that agents follow in this article. In particular, it is not straightforward that the standard nonlinear Kalman filter is still optimal for agents who worry about model misspecification. This issue is related to the separation result that the authors use. In conventional models of learning, separation between learning or signal extraction and control or optimal dynamic choice readily obtains and substantially simplifies the solution of these models. The original robust control literature analyzed this problem extensively and established some separation results, however, many contributions in this literature are of limited use for applications in financial economics. Tornell (2000) develops robust control tools based on the H_∞ tradition (as opposed to the entropy-based approach used here) and does not obtain separation between filtering and control. The authors tackle the problem in a separate article and make an important contribution by establishing a suitable form of the separation principle.

Second, it can be noted that economies with learning are often observationally equivalent to full information economies with more complicated dynamics for the state variables (i.e., precisely the dynamics induced by the posterior beliefs). Simultaneously, economies with robustness are typically equivalent to an expected utility economy with pessimistically distorted dynamics. What about an economy with both learning and robustness? The authors state that the solution to game I (where neither the agent nor Nature observe the growth rate) is equivalent to a full information economy. Of interest, it seems there is no full information equivalent to game II. This is not surprising given that game II is characterized by an information asymmetry, which would obviously be lacking in any full information economy.

To reduce the number of state variables, the authors introduce stochastic discounting. An alternative analytical simplification could be based on the homothetic version of robustness (obtained by scaling θ appropriately) as proposed by Maenhout. When scaling θ in a natural way to impose homotheticity, robustness becomes isomorphic to stochastic differential utility in the sense of Duffie and Epstein (1992) in the context of a full information, pure Brownian economy. One would conjecture that the same would apply here for game I, but probably not for game II. This issue is complicated by the fact that little is known about stochastic differential utility when combined with filtering and jump diffusions.

A fundamental problem faced by any modeling effort involving robustness or model uncertainty is the calibration of the θ parameter. Introducing a free preference parameter might render a model untestable unless the researcher carefully pins down this parameter. When this parameter measures the strength of the preference for robustness it is not immediately obvious how to impose this type of discipline. The authors refer to the strategy proposed in Anderson, Hansen, and Sargent (2000), where the choice of θ is based on Bayesian detection probabilities. In the setup of the present article, this is more difficult (especially for game II) and it is not straightforward to interpret the relative magnitudes of the entropy measures reported in Figure 4. An alternative and more informal approach analyzes the implied worst-case distortion for the dynamics of the relevant state variable. In the context of a dynamic portfolio choice problem, Maenhout judges the plausibility of a value for θ on the basis of the implied pessimistic assessment of the equity premium. A similar approach could be helpful in this article, especially for game II.

More generally speaking, the quantitative asset pricing predictions of the model deserve further exploration. The time variation in the shadow equity premium is intriguing, but primarily driven by changes in the posterior beliefs of the agent. It is unclear to what extent the time variation in the factor risk price is empirically plausible quantitatively or even qualitatively. The same question arises about second moments of asset returns. The model predicts stochastic volatility. It would be interesting to understand its time-series properties. Similarily the authors point out that the instantaneous risk-free rate is determined by the marginal product of capital. Robustness enhances the precautionary savings motive. In this economy, the precautionary savings motive changes with the state of the economy and with the beliefs of the agent about the state. The capital stock fluctuates accordingly and hence one would expect the risk-free rate to be counterfactually volatile.

Both conceptually and theoretically the article could be extended in a variety of ways. The theoretical model developed by the authors accommodates a finite number of growth states. The illustrative growth economy that is subsequently studied specializes to the case of two growth states. An obvious and promising extension could allow for a third state that is unlikely but severe, in the spirit of Rietz (1988). In particular, the presence of a crash state is very natural when investors worry about robustness and Peso problems. Moreover, a robust investor overstates the likelihood of the crash state and distorts the intensity with which the crash regime occurs (the probability slanting described by the authors is endogenous and depends on the severity of the state). This is a potentially powerful way of making Rietz's solution quantitatively reasonable, even when the observed or estimated likelihood of a crash state is insufficiently large to produce any effect in an expected utility economy [Mehra and Prescott (1988)].

Another promising extension consists of using the framework of the authors to study derivative pricing. The combination of jump diffusions, uncertainty about the state of the economy, and robustness is very relevant for this issue and would introduce a concern for model misspecification in the recent work of David and Veronesi (2000).

A final comment concerns the question of learning about model uncertainty. The authors study model uncertainty in an economy with learning about the growth rate of the economy. While agents attempt to learn the state of the economy, uncertainty about the validity of the model describing the economy does not resolve. Once agents hit upon a benchmark model of the economy they give up on attempts to further refine this model and instead entertain a family of similar models. Presumably the agents lack data to further distinguish between these statistically similar models. The questions of initial model selection and subsequent refinement of this model (or of learning about model uncertainty) are related and ultimately remain unresolved issues that warrant further analysis.

References

Abel, A., 2001, "An Exploration of the Effects of Pessimism and Doubt on Asset Returns," Working Paper 8132, National Bureau of Economic Research, forthcoming in *Journal of Economic Dynamics and Control*.

Anderson, E., L. P. Hansen, and T. Sargent, 2000, "Robustness, Detection and the Price of Risk," working paper, Stanford University.

Cecchetti, S., P. Lam, and N. Mark, 2000, "Asset Pricing with Distorted Beliefs: Are Equity Returns Too Good to be True?," *American Economic Review*, 90, 787–805.

Chen, Z., and L. Epstein, 2001, "Ambiguity, Risk and Asset Returns in Continuous Time," forthcoming in *Econometrica*.

David, A., 1997, "Fluctuating Confidence in Stock Markets: Implications for Returns and Volatility," *Journal of Financial and Quantitative Analysis*, 32, 427–462.

David, A., and P. Veronesi, 2000, "Option Prices with Uncertain Fundamentals: Theory and Evidence on the Dynamics of Implied Volatilities," working paper, University of Chicago.

Duffie, D., and L. Epstein, 1992, "Stochastic Differential Utility," *Econometrica*, 60, 353–394.

Gilboa, I., and D. Schmeidler, 1989, "Maxmin Expected Utility with Non-unique Prior," *Journal of Mathematical Economics*, 18, 141–153.

Hansen, L., T. Sargent, and T. Tallarini, 1999, "Robust Permanent Income and Pricing," *Review of Economic Studies*, 66, 873–907.

Maenhout, P., 2001, "Robust Portfolio Rules, Hedging and Asset Pricing," working paper, INSEAD.

Mehra, R., and E. Prescott, 1985, "The Equity Premium: A Puzzle," *Journal of Monetary Economics*, 15, 145–161.

Mehra, R., and E. Prescott, 1988, "The Equity Premium: A Solution?," *Journal of Monetary Economics*, 22, 133–136.

Rietz, T., 1988, "The Equity Risk Premium: A Solution," *Journal of Monetary Economics*, 22, 117–131.

Tornell, A., 2000, "Robust $\mathscr{H}_\infty$ Forecasting and Asset Pricing Anomalies," working paper, University of California, Los Angeles.

Veronesi, P., 1999, "Stock Market Overreaction to Bad News in Good Times: A Rational Expectations Equilibrium Model," *Review of Financial Studies*, 12, 975–1007.

Weil, P., 1989, "The Equity Premium Puzzle and the Risk-Free Rate Puzzle," *Journal of Monetary Economics*, 24, 401–421.