Pitfalls of a Minimax Approach to Model Uncertainty

Christopher A. Sims*

January 11, 2001

The recent spate of work applying the ideas of "least favorable prior" decision theory (Gilboa and Schmeidler, 1989; Chamberlain, 2000) to macroeconomic models (Hansen and Sargent, 2000; Giannoni, 1999; Onatski and Stock, 1999) has produced thought-provoking results, but there are aspects of the work[1] that I am uncomfortable with:

- it does not always keep the normative analysis of decision making distinct from its descriptive analysis, losing sight of the fact that these methods are appealing shortcut approximations, not improvements on, decision-making based on the Savage axioms;

- in all the existing applications to monetary policy it analyzes uncertainty about relatively unimportant aspects of models, while making strong, but actually uncertain, assumptions about other, more important aspects.

I Normative vs. Descriptive

Least-favorable-prior ideas are appealing in good part because they can explain certain patterns of behavior that are not rational from the perspective of Bayesian decision theory. This appeal of course does not support a claim that these patterns of behavior are in any sense "better" than those predicted by subjective probability and expected utility maximization.

In fact, because it violates the sure-thing principle, most people, and certainly most policy-makers, would be likely to alter behavior fitting the maximin theory if they were shown certain consequences of it. For example, because maximin expected utility behavior, if applied *de novo* to each of a sequence of choice sets, can imply behavior consistent with no single set of probabilistic prior beliefs, it can allow a Dutch Book, a situation where someone agrees to a set of bets that cause him to lose money with probability one.

I can see no reason why we should recommend that policy-makers deliberately violate the sure-thing principle. And some proponents of maximin expected utility theory have been quite clear that it is a shortcut, a way of constructing a prior that may be worth considering in the process of assessing actual prior beliefs or even of using directly if it does not look too outrageous and time is short. But the literature in some cases takes a different viewpoint. Hansen and Sargent (2000), for example, suggest imposing the same maximin prior on the private agents in a model and on the optimizing policy makers, in effect recommending to policy-makers the same sub-rational behavior that they postulate in private agents. But the criteria for acceptable shortcuts in decision-making by a central bank should generally be much stricter than those applying to, say, a consumer buying a new washing machine. And on the other hand, a "representative agent" that summarizes the behavior of many individuals with disparate information sources, coordinated through many markets, may be well modeled as having fewer computational constraints than a monetary policy maker. In either case, the criteria for good descriptive modeling and for good normative policy advice ought to be kept distinct.

Shortcuts are not inherently bad. The fact that a central bank has less need for shortcuts in evaluating the current and future state of the economy than has the washing machine purchaser does not mean that the central bank should disdain all shortcuts. The Savage axioms are normative, and not accurately descriptive of decision making in actual real-world contexts. If we formally go through the steps of Bayesian decision analysis — assessing a prior, looking at evidence, minimizing expected loss — we are likely to use conventional forms for our prior distributions and not to fully assess every possible nonlinearity or cross-dependency among sources of uncertainty. If as is more usual we never formally assess a prior, but instead make decisions in a procedurally reasonable way, we may never realize that our procedures actually imply prior beliefs that we would firmly reject if they were made explicit.

Automatic procedures to generate priors and associated decision rules are therefore potentially useful as part of a decision-making process. They may alert us to forms of prior that, on reflection, do not seem far from what we actually might believe, yet imply decisions very different from what we arrive at by other simple approaches. Maximin expected utility naturally takes a place alongside Jeffreys priors and other methods for generating reference priors in the statistical literature. These are tools that can stimulate our imaginations or help us critique conventional choices of priors or decision rules.

But once we understand the appropriate role for this tool, it should be apparent that its results should always where possible be compared to more direct approaches to assessing prior beliefs. The results may imply prior beliefs that we can see make no sense, in which case they should be discarded. They may also — and this is more likely in the recent implementations in macroeconomics — focus the minimaxing on a narrow, technically convenient, uncontroversial range of deviations from a central model. Then the results will remain close to those of the central model, and the danger is that we will be misled by the rhetoric of robustness into devoting less attention than we should to technically inconvenient, controversial deviations from the central model.

II Forest and Trees

The existing robust control literature focuses on relatively unimportant sources of model uncertainty. The point can be understood from two perspectives, one technical, one broadly conceptual. We take the broadly conceptual point of view first, then the technical in the next section.

What are the most important gaps in our understanding of what makes good monetary policy? Is it really uncertainty about the actual values of coefficients in log-linear local approximations to a particular model? Let me state my own (no doubt controversial, but that is part of the point) views on this.

There are two things we don't understand that are likely to be much more important than our lack of knowledge of coefficients in log-linear approximate models. One is that we do not know whether there are level or growth effects on output from the inflation rate — in other words whether there is a long-run tradeoff between inflation and output. Of course we are by now quite sure that short-run correlations of output and inflation can arise that should not be extrapolated, as a policy menu, to the long run. But long-run effects of inflation on output need not be very large and need not be completely permanent in order to be important. Second-order effects on the steady state of a log-linear approximation are as important for welfare evaluation as first-order effects on its coefficients. There is some weak evidence to suggest that steadily modestly higher inflation reduces output or productivity growth. There is also some weak evidence to suggest the opposite. Since we are talking about long run effects of inflation on a variable subject to a lot of other long-run influences, it will remain unclear what size these effects are, or even whether they exist.

Another thing we don't understand is the danger of deflationary or inflationary spirals that run out of the control of the monetary authority. With inflationary spirals I think the uncertainty is not so great. The process by which deficits feed a spiraling inflation has been observed enough times historically, and the public's distaste for inflation is firmly enough established as a political fact, so that I doubt that there is any trigger, at inflation rates below, say, 15%, above which monetary policy would lose control. Deflationary spirals are another matter. As Benhabib et al. (1998) have shown, a sufficiently benighted fiscal policy can imply that the economy can easily, or even inevitably, fall into a permanent deflationary spiral. It may seem that the fiscal policy required to support this type of equilibrium — rising primary surpluses as the economy spirals downward — is implausible, but we have historical examples (the US in the 30's, Japan recently) in which fiscal policy behaved nearly this perversely. And understanding in policy circles of the connection of fiscal policy to deflation in a liquidity trap is not deep. A monetary authority that targets inflation too low, therefore, runs a risk that a surprising change in the real return on investment or in budget politics could push the nominal rate close to zero. At that point, control of the price level has essentially been handed off to the fiscal authorities. How dangerous a prospect is that in the US? I for one am not sure.

Do we find these uncertainties reflected in macroeconomic robust control exercises? Not at all. All three of the applied papers in this area that I have cited use models that must be interpreted as local log-linear approximations and all three therefore ignore the zero bound on nominal rates and the role of fiscal policy in price determination. Hansen and Sargent and Giannoni use a model that implies a long-run tradeoff, but use an objective function implying that increases in the level of output above the zero-inflation steady state value are as undesirable as decreases. Onatski and Stock also assume increases and decreases of output from its non-inflationary steady state are equally undesirable, and in addition use a model that is constrained to allow no long run tradeoff between inflation and output. Note that it would not help matters to replace local log-linear approximations with nonlinear models if the nonlinear models also have rigidly imposed neutrality properties and assume the optimality of the non-inflationary deterministic steady state.

While it ought to appear unreasonable, this pattern of giving careful attention to minor sources of uncertainty, while ignoring, by making dogmatic assumptions, major sources, is not uncommon in economics. The pattern may reflect the fact that few economists have been taught, or have thought carefully about, the differences and similarities between subjective and objective uncertainty. When a source of uncertainty is a matter of public dispute, we hesitate to apply probabilities to it, and indeed for a variety of reasons may tend to make dogmatic assertions of our own views, or the views of the professional faction with which we are associated, instead. People understand this tendency, so it does not do us great harm in general. But we will get more respect if we can control the tendency, and putting forth a claim that we have mastered "model uncertainty", when in fact the sources of uncertainty that are most important are treated in as *ad hoc* a way as usual, is a step in the wrong direction.

III Local Approximation: First, Second . . . n^{th} order

Note that if the model is constructed in terms of deviations from a non-stochastic steady state, and if the uncertainty in coefficients is taken to apply only to the deviations part of the model, all effects of parameter uncertainty on the joint distribution of shocks and variables are third-order or higher, and all effects on welfare are fourth order or higher. In other words, extending our analysis from a first-order to a second-order approximation to the solution has no effect. This point is perhaps most easily understood in the simplest static problem:

$$\max_x E[-\tfrac{1}{2}(y - ax)^2] \,. \tag{1}$$

As is well known, if $a \,|\, \{y, x\}$ has mean $\bar{a}$ and variance σ_a^2, then the optimal choice of x, given an observation on y, is

$$x = \frac{\bar{a}}{\bar{a}^2 + \sigma_a^2} y \,. \tag{2}$$

If y has mean $\bar{y}$ and variance σ_y^2, the deterministic steady state value of x (i.e. the value when $y = \bar{y}$ and $\sigma_a^2 = 0$) is $\bar{x} = \bar{y}/\bar{a}$. Writing the solution in terms of steady-state values and deviations from steady state ($\tilde{x}$ and $\tilde{y}$), we have

$$\tilde{x} = \frac{\bar{a}}{\bar{a}^2 + \sigma_a^2}\tilde{y} - \frac{\sigma_a^2}{\bar{a}(\bar{a}^2 + \sigma_a^2)} \,. \tag{3}$$

As should be clear from this expression, the effect of uncertainty about $\tilde{a}$ on the coefficient applying to $\tilde{y}$ contributes to the right-hand side of (3) approximately

$$-\frac{\sigma_a^2 \tilde{y}}{\bar{a}} \,. \tag{4}$$

This effect is second-order if the random component of y is thought of as $O(1)$, but in the macroeconomic applications we are discussing all random components have to be thought of as "small" to justify the linearity assumptions, so this term would be third order. The effect of uncertainty about $\tilde{a}$ on the constant term in (3), though, is obviously second order, unless $\bar{y} = 0$. Of course if we had written the model in terms of $\tilde{x}$ and $\tilde{y}$ from the start, this second-order effect on the constant term would have been missed.

The models in the monetary policy robust control literature have been written in terms of deviations from steady states and ignore possible effects of parameter uncertainty on constant terms. Breaking away from these assumptions involves recognizing that we are not sure that the deterministic steady states of stochastic models of the economy ought to be assumed to be optimal and that we are not sure what the tradeoffs between levels and variabilities of inflation and output may be.

IV Where To Go From Here

Some (Onatski and Stock, 1999, for example) have suggested that minimax robust approaches are appealing because they are more tractable than methods that put an explicit probability distribution on the model uncertainty. It is not at all clear that this is true. Rudebusch (1998) takes a straightforward computational approach to evaluating the implications of several kinds of model uncertainty using a probabilistic approach. While the models he uses are small, they are no smaller than those used by existing applications of robust control to monetary policy. Whether there is a difference in scalability of the methods remains to be seen. Furthermore, there is recent work that standardizes methods for taking higher-order local expansions of nonlinear models (Sims, 2000; Kim and Kim, 1999; Collard and Juillard, 2000). These should make it easier to carry out at least local analysis of the effects of model uncertainty using explicit probability distributions and fairly large models.

An approach to robustness that aims at serious evaluation of all dimensions of model uncertainty will give an important place to minimax calculations. Improved methods of carrying them out and applications of them to increasingly realistic models are welcome, therefore, so long as they are seen as a tool for assessing uncertainty, not a replacement for assessing uncertainty.

References

Benhabib, Jess, Stephanie Schmitt-Grohe, and Martin Uribe, "Monetary Policy and Multiple Equilibria," Technical Report, New York University 1998.

Chamberlain, Gary, "Econometric Applications of Maxmin Expected Utility," discussion paper, Harvard University June 2000.

Collard, Fabrice and Michel Juillard, "Perturbation Methods for Rational Expectations Models," Technical Report, CEPREMAP, Paris August 2000. `fabrice.collard@cepremap.cnrs.fr`.

Giannoni, Marc P., "Does Model Uncertainty Justify Caution? Robust Optimal Monetary Policy in a Forward-Looking Model," discussion paper, Princeton University November 1999.

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

Hansen, Lars Peter and Thomas J. Sargent, "Robust Control and Filtering of Forward-Looking Models," discussion paper, Stanford University, Hoover Institution, University of Chicago October 2000.

Kim, Jinill and Sunghyun Kim, "Spurious Welfare Reversals in International Business Cycle Models," Technical Report, Brandeis University, `http://www.people.virginia.edu/~jk9n/` 1999.

Onatski, Alexei and James H. Stock, "Robust Monetary Policy Under Model Uncertainty in a Small Model of the U.S. Economy," manuscript, Harvard University February 1999.

Rudebusch, Glenn D., "Is the Fed Too Timid? Monetary Policy in an Uncertain World," Technical Report, Federal Reserve Bank of San Francisco October 1998.

Sims, Christopher A., "Second Order Accurate Solution of Discrete Time Dynamic Equilibrium Models," discussion paper, Princeton University August 2000. `eco-072399b.princeton.edu/gensys2/`.

Notes

*Department of Economics,Princeton University, sims@princeton.edu. Comments from Marc Giannoni, Lars Hansen, Tom Sargent and Jim Stock have been helpful.

[1]There are applications of minimax robust decision theory to other areas of economics. This paper focuses entirely on the applications to monetary policy.