COMMENT ON WOODFORD'S "ROBUSTLY OPTIMAL MONETARY POLICY WITH NEAR-RATIONAL EXPECTATIONS"

CHRIS SIMS

1. OUTLINE

(1) MinMax robustness as a dual approach to probabilistic decision theory: a method for producing "automatic" priors.
(2) Varieties of irrationality
  (a) Learning
  (b) Rational inattention

2. DUALITY

- MinMax does pick out admissible decision rules — that is it cannot lead us to a rule that can be dominated in every state of the world by some other rule.
- Therefore, under reasonable regularity conditions, it also picks out Bayesian rules.
- Bayesian rules are those that minimize *expected* losses under some probability measure over states of the world.
- So minmax rules generally can be characterized as picking a particular prior and minimizing expected losses. The nature of the prior depends on the feasible set of decision rules.
- Its appeal is that we may well find it easier to assess the consequences of actions in various states of the world than to assess directly a probability measure over states of the world. MinMax maps the former type of assessment into the latter.

3. THE NON-ROBUSTNESS OF MINMAX ROBUST ANALYSIS

- It is usually a lot easier to spot a proposed distribution over states of the world that is ridiculous than it is to specify in detail a distribution that exactly captures one's information and uncertainty.
- The automatic priors emerging from minmax analysis can be ridiculous.
- It is always legitimate to ask that the policy recommendation emerging from such an analysis be accompanied by a description of the supporting prior distribution, and to reject the recommendation if the supporting prior is ridiculous.

4. THIS PAPER'S VARIANT OF MINMAX

- Before applying the minmax criterion, losses are downweighted by distance of the state from some central value, according to a particular (Kullback-Leibler) notion of distance.
- This still selects an admissible rule.

- As you might expect, it will be a rule supported by a prior that puts less probability weight on states of the world farther, in KL distance, from the central state.
- If one is going to do automatic prior generation — and it is a very reasonable thing to do, in fact you probably want to do it several different ways — this paper's approach (and that of its antecedents in Hansen-Sargent and the engineering literature) is appealing.

5.

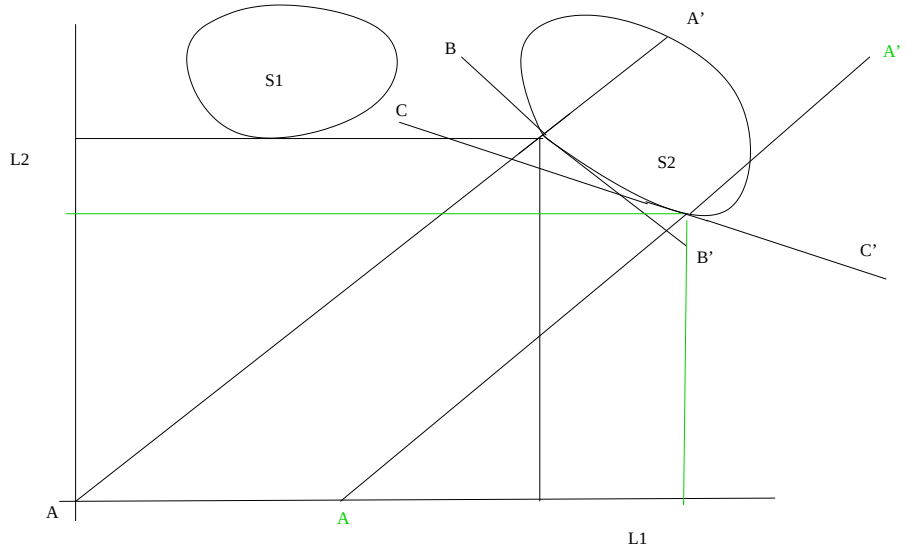

6. Explanation of the figure

This figure represents a decision problem with two states of the world. The vertical axis shows losses in state 2, the horizontal axis losses in state 1. Each point on the graph describes losses in the two states, which is what characterizes a decision rule. A set of feasible decision rules might be represented by the regions labeled S1 or S2 on the graph. Admissible, or undominated, rules are those such that no feasible rule is both below and to the left of them. The forty-five degree line labeled A, A' in black separates a region above, in which the maximum losses occur in state 2, from a region below, in which the maximum losses occur in state 1. A minimax decision rule will always be a point of tangency of a square region like the large black square in the diagram with the feasible set. For S2 this is where the vertex of the square touches the set. For S1 it is the tangency of the top side of the square with the set.

Prior distributions put probability weights on the two states. Level curves of expected losses with a given prior are parallel negatively sloped lines, with the slope defining the relative weight on the two states. A minimax rule like that for S1 implies zero probability on state 1. Lower losses in state 1 could be obtained, in the diagram's S1 case, with negligible increases in state 2 losses. For the S2 case, the prior that supports the minimax rule puts some weight on each state.

Economists should find this reasoning familiar. The admissible rules are the efficient frontier. Points on the efficient frontier are generally supported by tangent planes that can

be interpreted as shadow values in a consumer or producer optimization problem, or as prior probability weights in a decision problem.

This paper uses a modified minmax criterion. Before calculating maximum losses, it reduces losses associated with states distant from the central state. If state 2 is the central state, this amounts in this two-state example to choosing decisions that are a tangency between the feasible set and a rectangle with upper right vertex on the line marked AA' in green, so that for S2 the choice is the tangency with the green rectangle. Note that the supporting prior has a flatter slope, which corresponds to putting more weight on state 2 in the prior.

7. Learning

- The paper's approach seems motivated by a notion of learning by the private sector — plausible deviations from rational expectations are taken to be those it would be hard for the private sector to learn to be incorrect.
- But, as the paper points out, existing approaches to modeling learning involve postulating a particular, usually rather ad hoc, learning mechanism, so that "optimal policy with learning" could seem to involve systematically exploiting the nature of the learning mechanism.
- While this is an advantage of this paper's approach, this paper ignores belief dynamics. If one contemplates a major policy change, an explicit learning model will recognize the costs (or benefits) of the period during which the public revises its beliefs, while this paper ignores them.

8. Rational Inattention

- This is the notion that economic agents, as they translate observation of their environment into decisions and actions, act as finite Shannon-capacity channels.
- It is another form of near-rationality.
- But it has different implications.
  - The most likely deviations are those that involve only crude monitoring of the state of the economy by agents, not those that are hardest to distinguish from RE with perfect observation.
  - If policy succeeds in, say, making fluctuations in the price level so small that they matter little to agents, the agents may monitor them extremely crudely and hence appear in some sense extremely irrational, while if policy makes inflation high and variable, agents may monitor it closely.
  - This paper's approach does not take these factors into account.