Volume Title: NBER International Seminar on Macroeconomics 2007

Volume Author/Editor: Richard Clarida and Francesco Giavazzi, organizers

Volume Publisher: University of Chicago Press

ISSN: 1932-8796

Volume URL: http://www.nber.org/books/clar07-1

Conference Date: June 15-16, 2007

Publication Date: January 2009

Chapter Title: Comment on "Interest Rate Signals and Central Bank Transparency"

Chapter Author: Michael Woodford

Chapter URL: http://www.nber.org/chapters/c2999

Chapter pages in book: (55 - 62)

Comment

Michael Woodford, *Columbia University and NBER*

This chapter makes an important contribution to the literature on the advantages and disadvantages of central-bank transparency. The most influential recent contribution on this topic has been the analysis of Morris and Shin (2002), who use a simple, highly stylized model of central-bank communication to make the point that it is possible for full transparency not to be the optimal policy. In the model of Morris and Shin, the central bank has (imperfectly precise) information about a random fundamental state that is not directly observable by private decision makers, but the value of which is relevant to their decisions; the question posed is how precise a measure of its information it is desirable for the central bank to publicly reveal. Morris and Shin point out that the mere presumption that it would be desirable to increase the precision with which private decisions track the unobserved fundamental does not suffice to answer this question, for one can imagine circumstances under which a more precise announcement by the central bank actually reduces the average conformity of private decision makers' actions with the fundamental, even though it increases the precision of each of their estimates of the fundamental.

If, for strategic reasons, decision makers care about the conformity of their actions with the average actions of other private agents, and not merely about the conformity of their actions with the fundamental, then the fact that the central bank's announcement is common knowledge makes it relevant as an indicator of what *others' estimates* are likely to be, in addition to the information that it provides about the fundamental itself. This leads private decision makers to base their actions more on the central bank announcement than on their own (idiosyncratic) sources of information about the fundamental, which are known not to be observed by others. But excessive reliance upon the central bank announcement about the economy's state has the unfortunate conse-

quence that errors in the central bank's information come to affect everyone's action in the same way, and so cause *aggregate* outcomes to depart from conformity with the fundamental, whereas reliance upon private information would lead to individual errors that would tend to cancel in their consequences for the aggregate outcome. As a result, Morris and Shin show that (on a certain assumption about welfare) it is possible for a less precise announcement by the central bank to achieve a better outcome, because an announcement that is understood to be less informative has less effect on people's actions.

The practical relevance of this analysis has been contested by a number of commentators, but generally by pointing to various special features of the Morris-Shin example; thus, Svensson (2006) argues that the parameter values under which Morris and Shin find transparency to be suboptimal are not the ones most likely to be empirically relevant, while Woodford (2005) argues that the welfare measure under which transparency is suboptimal in their model is unlikely to correspond to the private objectives of the people whose behavior the model is intended to describe. The present chapter instead offers a more trenchant criticism, because it considers an abstract model of information revelation by a central bank, and presents results that depend little on arguments about the empirically relevant magnitudes of particular parameters.

The present chapter extends the analysis of Morris and Shin (2002) in a number of respects. One that is emphasized more than I think it should have been is the idea that the central bank is required to reveal at least *part* of its information about fundamentals, as a consequence of public announcement of its interest-rate target; the authors' central concern then becomes the question whether a central bank should reveal additional information *beyond* what is revealed by the interest rate. There is surely some validity to the observation that central banks must reveal some aspects of their view of current conditions through their policy actions, even if they refrain from communication of other sorts; after all, a central bank is not purely, or even primarily, a news service, though Morris and Shin treat it as if it were. But the interest of this observation derives entirely from the fact that (in the case that full transparency is not optimal) it will create a tension between the considerations that would otherwise determine the optimal interest rate decisions and the central bank's interest in strategic information revelation. Gosselin, Lotz, and Wyplosz assume instead that the interest rate decision has no economic effects; it is purely an announcement of one particular statistic from the central bank's information set.

Since the central bank is also assumed to be free to make this announcement any function whatsoever of its information (there is no meaning at all to its being an interest rate announcement—the meaning of the announcement derives entirely from the correlation that it happens to have with variables observed by the central bank, as a result of the bank's policy with regard to the announcement), there is no meaningful sense in which the necessity of revealing some information through an interest-rate decision represents a constraint on the central bank's communication policy. After all, a central bank that wishes not to reveal any of its information through its interest-rate decision can simply set the interest rate in a way that does not depend on any information that is not already public; if (as assumed in this chapter) the interest-rate decision has no consequences apart from the information that it reveals, there is no cost in acting in this way. This is obscured in the present chapter by the unmotivated assumption that the interest rate must be a *deterministic* function of the central bank's information, with weights summing to one; if either arbitrary randomization of the interest rate, or a constant interest rate (zero weights on each of the central bank's signals) were allowed, as would be possible in reality, it would be obvious that announcement of the interest rate need not reveal any information.

But in fact none of the chapter's interesting results depend on this assumption. The finding that under certain circumstances it is optimal to reveal only the interest rate should more properly be phrased as a result that in these cases it is optimal to reveal only *a single summary statistic*, rather than the central bank's complete information set. (There is nothing important or interesting about the supposition that this statistic should be an interest rate; in the case that the central bank has an interest-rate decision to make that affects the economy, and must reveal it, the result about the optimality of the single statistic will almost certainly not apply.) Under this reinterpretation, the chapter's results still hold, without any need to discuss information revelation by the bank's interest-rate decision.

A more substantive extension of the Morris and Shin analysis is the treatment of the case in which both the central bank and private decision makers have noisy observations of several different fundamental state variables, rather than assuming a one-dimensional fundamental. It might seem that the extension to a multi-dimensional fundamental is only required by the assumption that one linear combination of the central bank's information variables is necessarily revealed by its interest-rate decision; in that case, there only remains a question about the de-

sirable degree of transparency when there is more than a single dimension of information for the central bank to reveal. But in fact, the multidimensional case is of independent interest. For it is only in this case that the authors are able to obtain a result with the generality of their Proposition 1, according to which (in the case that the precision of the various signals is common knowledge) partial transparency is *invariably* preferable to full transparency.

This very clever argument relies on the fact that, in their model, under full transparency there is only a single linear combination of the central bank's information variables that would be used in private decisions. (This in turn follows from functional-form assumptions—quadratic objectives, linear constraints, and normally-distributed disturbances—that make the optimal decision rules linear; from the fact that the decision problems of all private decision makers are identical, so that each uses the information revealed by the central bank in exactly the same way; and from the fact that each private decision maker has only a one-dimensional decision to make.) This means that one can necessarily achieve as good an outcome under partial transparency as under full transparency: it is necessary only to reveal the particular linear combination of central bank information variables that the private decision makers would actually use, rather than also revealing redundant information. But under partial transparency, it is *also* possible to reveal only a single dimension of the bank's information that differs from the one that private decision makers would use under full transparency; this allows the central bank to manipulate the information of private decision makers, should this be desirable.

Thus, partial transparency is a dominant strategy, in the sense that it allows the central bank to achieve a superset of the possible outcomes under full transparency. In the generic case, the equilibrium under full transparency is not fully optimal (owing to the common-knowledge effect identified by Morris and Shin), so that some kind of manipulation of private decision makers' information can improve upon that outcome; one thus concludes that full transparency is strictly inferior in the generic case. As the authors note, this argument does not rely on any particular, special choice of the loss function used to evaluate outcomes; it is only necessary that one be outside any of the special cases in which full transparency leads to precisely the optimal allocation. This is a much stronger result than the original one of Morris and Shin.

Nonetheless, the result is far from conclusive, as far as the practical

question of the desirability of transparency in central banking is concerned. The result still depends (at least for its generality) on quite special assumptions—for example, on the assumptions that private decision makers each have use for only a single dimension of the central bank's information, and that this single dimension is identical in the case of each private decision maker—and one may question the relevance of such assumptions to the actual circumstances of central banks. In fact, Gosselin, Lotz, and Wyplosz themselves provide an effective critique, pointing out that their first result depends on assuming that the *precisions* of the various private observations are common knowledge, even though the *values observed* are not. In particular, it depends on an assumption that the central bank knows enough about the precision of the private information of private decision makers (as well as their understanding of the precision of the information that the central bank chooses to reveal) in order to be able to predict how they will react to the information that it chooses to reveal, and hence to determine the optimal single linear combination of signals to reveal. If the central bank is instead mistaken about the precision of private information, and so incorrectly calculates the optimal partial information to reveal, the result can be worse than would be achieved by full transparency.

This is the other important extension of the Morris-Shin analysis in the present chapter: relaxation of the common-knowledge assumption, first to consider a case in which the central bank misestimates the precision of the information available to private decision makers (though they correctly understand the nature of their incomplete information), and then to consider a further case in which private decision makers also misestimate the precision of their information. In both of the latter cases there are ranges of parameter values for which full transparency will be the optimal policy.

These results are perhaps not surprising; essentially, what is shown is that it may be desirable to constrain the central bank *not* to manipulate information in the way that it would otherwise choose to, when the central bank's choice would be based on *incorrect beliefs*. And here it should be remarked that the analysis would be more interesting if it were to consider what a central bank should choose to reveal that (correctly) *understands that it does not know* the precision of private decision makers' information, rather than assuming that the bank will choose what information to reveal on the basis of incorrect parameter estimates that it treats as known with certainty. (The chapter's frequent references to the

case that it treats as the case in which the central bank is uncertain about private signal precision are somewhat misleading, as the central bank is not treated as being uncertain at all.)

Genuine uncertainty (as opposed to simple erroneous belief) on the part of the central bank could be dealt with in various ways. One might assume a prior over various possible values for the vector of precisions, and choose a communication policy on Bayesian grounds; or one might assume only that the vector of precisions is known to belong to a certain set, and choose a maxmin communication policy relative to this set of possibilities, as in the robust policy analyses of Hansen and Sargent (2007). Even without the assumption that the central bank will naively act on the basis of wrong beliefs if it tries to be clever in its choice of what to reveal, one might well conclude that a bank that recognizes that it does not know exactly what it is that private decision makers need to know about its information should prefer on that ground to reveal more.

But even if the results obtained here do not yet provide a full exploration of the topic, the issues that the authors raise are welcome ones. The standard literature on the optimal degree of central bank transparency, following Morris and Shin, conceives the problem as one in which the central bank must decide whether or not to uncover various pieces of information, the meaning of which will be unambiguous if revealed. In fact, practical debates about central-bank communication policy almost never concern questions such as whether a central bank should reveal statistics that it has collected relating to the state of the economy. The real questions have to do with how much a central bank should reveal about *its own decision processes*—something that the standard literature treats as being so well understood by everyone that they require no discussion.

Gosselin, Lotz, and Wyplosz introduce a concern of this kind when they suppose that the central bank may have its own (mistaken) estimates of the coefficients β_k, and then ask whether it is desirable for the central bank to make these estimates public. This can be beneficial, in their analysis, by allowing private decision makers to better judge the information content of the central bank's announcement, when it chooses to reveal only part of its information set. In reality, private decision makers have not only to guess the meaning (in the sense of the statistical relation to the underlying state of the economy) of central bank announcements of a purely informational character; they also have, above all, to try to forecast future central bank *actions* that affect the economy. But it is certainly true that a central bank cannot take it for

granted that its decision processes are already perfectly transparent to the public, without any need for explanation on its part; thus, the kind of issue that is represented by the discussion here of the desirability of revealing central bank estimates of the β_k is very much one that needs to become more central to analyses of central bank communication policy.

When this is understood to be the real issue, the arguments of Morris and Shin lose much of their force. The question of how the central bank approaches its decision problem—and even the question of how the central bank views the world—is unlikely to be one about which private decision makers have their own, private sources of information, to which they should be paying more attention than they do to what central banks choose to reveal about the matter. Instead, these are preeminent examples of issues about which the central bank is inevitably vastly better informed than anyone on the outside, so that there need be no fear that the common knowledge effect will lead to insufficient use of private information. At the same time, it is clear that better understanding of how central banks make their decisions can help private decision makers make more accurate decisions, owing to the substantial impact of monetary policy on financial markets and on the economy more broadly. For this reason, it is hard to see why transparency on the part of central banks about such matters, to the extent that it is *possible* without interfering with the decision processes themselves, should not be desirable.

The difficult questions about the benefits of transparency arise from doubts about how well attempts at openness will be *understood*, and from questions about how a bank's own decision processes are affected by the need to be able to give a public account of those processes. These are real concerns, and deserve to be carefully weighed, but the analysis of Morris and Shin does not help central banks to think about them. They are not yet present in the analysis here, either. But the present chapter represents at least a step in the direction of greater realism, by introducing the possibility of differences in how the central bank and private decision makers understand the implications of a central bank announcement. One hopes that future discussions of communication policy will go further in this direction.

References

Hansen, L. P., and T. J. Sargent. 2007. *Robustness*. Princeton: Princeton University Press.

Morris, S., and H. S. Shin. 2002. The social value of public information. *American Economic Review* 92: 1521–34.

Svensson, L. E. O. 2006. Social value of public information: Morris and Shin (2002) is actually pro transparency, not con. *American Economic Review Papers and Proceedings* 96: 448–51.

Woodford, M. 2005. Central-Bank communication and policy effectiveness. In *The Greenspan Era: Lessons for the Future*, ed. G. A. Kahn. Kansas City: Federal Reserve Bank of Kansas City.