Comment on "Inflation expectations, uncertainty, and monetary policy" by Christopher Sims

Comment by Athanasios Orphanides, Governor of the Central Bank of Cyprus, at the 7th BIS Annual Conference: Whither Monetary Policy? Monetary Policy Challenges In The Decade Ahead

Lucerne, 26 June 2008

This is a rich paper to discuss with numerous exciting ideas. It touches upon areas ranging from the questioning of the traditional application of rational expectations in macro modelling, the role of disagreement in expectations for asset price determination and implications of the new thinking on modelling expectations for monetary policy. The paper also presents a view of the history of the Phillips curve and discusses the determination of inflation without a Phillips curve. The Phillips curve was at the centre of a related paper (Sims, 2008), which Chris presented a couple weeks ago at the Federal Reserve Bank of Boston's 53rd economic conference. Given that I had the chance to talk about that paper at that conference, today I will focus on some of the other ideas.

Chris starts his paper by recalling that in earlier work he had criticised central bank macro modellers for not following the rational expectations literature to the letter, for example, in jointly specifying and estimating the system of equations in their models (Sims, 2002). I would like to express my relief that the modellers did not take on board Chris' suggestions to put all of the cross equation restrictions in place at that time. In light of the central argument in this paper, this would have been a terrible idea. Indeed, the main theme of the present paper is to explain why traditional application of rational expectations can be problematic and why it is high time modellers incorporated more realistic models of expectations in their analysis.

The rational expectations assumption has proven extremely powerful over the past few decades and there are some good reasons for its broad theoretical appeal. But there are also numerous limitations that should be kept in mind. Chris highlights as a problem the fact that only one probability distribution is in play in any model---what Tom Sargent recently described as the "communism" of rational expectations(1). The rational expectations assumption imposes strict discipline on a model, requiring all agents in the model, the economists and policymakers outside the model, as well as nature, to share a unique set of beliefs. But it is incomplete as it does not explain how people come to hold these common beliefs. It is also inconsistent with the presence of numerous models,

each one imposing a different and yet unique set of beliefs. Importantly, as Chris emphasizes, it fails to address the presence and the role of heterogeneity of beliefs in the real world we observe around us.

These limitations matter for policy analysis. Oversimplifying the expectations formation process in analytical models can lead to a misreading of the workings of the economy and to false policy recommendations. At present, by and large, traditional modelling imposes rational expectations in a world with fixed and perfectly known structures, including known and stable policy preferences. Under such assumptions, the monetary policy problem seems trivial---and misleadingly so. In particular, anchoring inflation expectations, which is widely acknowledged as a central element of a successful policy strategy---especially in the environment of adverse supply shocks we presently face---becomes a simple matter of policy adopting and adhering to a stable policy rule. Furthermore, in a world of perfect knowledge there is essentially no role for central bank communication, a gap that Alan Blinder also pointed out in his contribution earlier today (Blinder, 2008).

In recognition of the serious limitations of the traditional application of rational expectations, recent work has explored various avenues for improvement. Alternative models of expectations have been developed using the concepts of inattention or learning as refinements of the rational expectations assumption. Chris focuses on the virtues of the rational inattention alternative but I find both approaches promising in different ways. A common element in these approaches is the acknowledgement of the presence of imperfections in the formation of expectations relative to simplistic rational expectations models. Both approaches stress the limited cognitive capacity of humans and by doing so better capture the inherent limitations in gathering and processing information.

The rational inattention models are successful in staying closer to the microfoundations of decision making but are rather hard to work with. Learning models provide a simpler approach for deviating from traditional rational expectations by asking private agents to act as econometricians, respecifying and reestimating forecasting models with limited data to form expectations. The implications, in many respects, are broadly similar, but much more work is needed to figure out whether the simple mechanisms embedded in learning models approximate well the microfoundations of the inattention models.

The paper develops an interesting example of an economy where differences of opinion about the course of inflation could have real economic effects. The example is helpful for

explaining one potential reason why policymakers are interested in monitoring expectations and, also, for establishing that disagreement in expectations may be a determinant of financial asset returns. Motivated by the theoretical example, I thought it would be informative to provide two illustrations of the potential empirical relevance of the story by looking at survey measures of expectations.

The first illustration draws on the Bluechip survey of financial forecasts in the United States. Twice a year since the mid-1980s, this survey presents information on the dispersion of long-horizon expectations of inflation and Treasury-bill rates (in addition to the consensus forecasts). As can be seen in Figure 1 (reproduced from Kim and Orphanides, 2007), the dispersion in these forecasts is positively correlated with the term-premia embedded in long-term government bonds. This suggests that heterogeneity of beliefs may be influencing the expected returns on long-term assets. The theoretical example Chris develops in his paper can be seen as an attempt towards explaining such empirical regularities.

The second illustration regards the usefulness of monitoring inflation expectations---including measures of disagreement---for monetary policy analysis. Survey information on inflation expectations can be a valuable input into monitoring how well-anchored inflation expectations are. There are multiple dimensions that can be examined in this regard. One is to examine whether average expectations over various horizons remain close to the policymaker's price stability objective. Another is to assess the degree of consensus among forecasters about the outlook of inflation---that is lack of disagreement. The ECB Survey of Professional Forecasters provides data that can be useful in illustrating these two dimensions for the euro area. The survey, which, among other things, collects information on the outlook of HICP inflation in the euro area, has been conducted quarterly since 1999.

Figure 2 shows the evolution of HICP inflation in the euro area together with the average one-year and five-year-ahead inflation expectations from the survey. Looking at actual HICP, one can easily understand why inflation has been the cause of great concern in recent months. Looking at the evolution of the one-year-ahead expectations, however, one sees much less volatility. Although one-year ahead expectations have edged up somewhat in the past two quarters, the figure suggests that professional forecasters do not expect the current deterioration in inflation to persist. More importantly, five-year-ahead expectations have remained remarkably stable and very close to the ECB objective of keeping HICP inflation close but below 2 percent, despite the increase in actual inflation observed in the last few quarters.

The evidence from the evolution of average expectations presented in Figure 2 confirms the credibility of the ECB. But another piece of evidence of the ECB's successful policy strategy during its ten years of existence, and the one I find most striking, can be found in examining the evolution of disagreement regarding inflation expectations over the past decade. Figure 3 presents the standard deviation of the individual forecasters' one- and five-year ahead inflation expectations since 1999. Notice that when the ECB started operating in 1999, there was greater disagreement about where inflation would be five years out than one year out. Over the past ten years, disagreement about inflation expectations one year out has moved about but without exhibiting any particular trend. By contrast, disagreement about five year-ahead expectations has trended down and is now considerably below that of one-year-ahead inflation. Relative to a decade ago, there is now very little disagreement about either the ECB's inflation objective or about the determination of the Governing Council to succeed in keeping inflation close to its objective over time.

By presenting these two empirical illustrations, I want to support Chris' emphasis on the need to refine the treatment of expectations in the models we use for understanding the macroeconomy and for formulating policy advice. I now wish to turn my attention to fleshing out some key implications of refining models of expectations formation[(2)]. Many of these implications coincide with those in the paper but some differ somewhat, perhaps because I draw more on models of learning rather than rational inattention. I will discuss, in turn, implications about inflation dynamics, about monetary policy, and about central bank communication.

Regarding inflation dynamics, one implication stands out. Refining expectations formation by introducing learning behaviour introduces a rich layer of non-linear dynamics in otherwise linear economies. This induces time-variation on the formation of expectations and thereby in the structure of the economy, even when fundamental regime changes are absent. In turn, this complicates empirical modelling of fixed-coefficient linear models and the use of such models for forecasting. Indeed, this may explain the difficulties econometricians tend to face when trying to fit simplistic linear expectations models. A model may be otherwise well specified (except for the imposition of rational expectations), and the rational expectations assumption alone may be the source of the problem.

In large part, because of its effect on inflation dynamics, the refinement of expectations behaviour away from rational expectations also has first order effects on monetary policy. For a given monetary policy rule, learning behaviour in the formation of

expectations imparts additional persistence to the inflation process relative to rational expectations, thereby diminishing policymakers' ability to stabilise business cycle fluctuations in addition to maintaining price stability. In my view, this is why policy should focus primarily on price stability as a means of achieving, not only price stability, but also overall economic stability over time. This provides a rationale for the wisdom of stating central bank mandates as if they arise from lexicographic preferences, with price stability recognized as the primary policy objective and other objectives being pursued without prejudice to the primary objective of price stability. Another implication is that learning behaviour may generate endogenous inflation scares that can be particularly damaging to the economy in the absence of forceful policy response. This, in turn, provides an explanation as to why policymakers monitor inflation expectations so closely and place a premium on striving to maintain well-anchored inflation expectations at all times.

As already mentioned, acknowledgment of imperfections in expectations formation introduces a role for central bank communication that is absent in traditional rational expectations models. The precise implications are not as clear-cut as the implications for policy, however. To the extent that central bank communication can facilitate the formation of more accurate inflation expectations, it can prove valuable for improving overall policy outcomes. Under this light, clarity regarding the central bank's price stability objective can be particularly helpful. Nevertheless, care is needed to respect the cognitive limits of economic agents. Thus, as Alan Blinder reported earlier today, "saying more is not always better for a central bank"[(3)]. In this case, the quality of information may be a determining factor as to whether its communication is warranted. At the same time, the presentation of information may also matter, since easier to use information is likely to be more useful than information that is harder to digest.

To conclude, I reiterate that this is a rich paper with numerous exciting ideas. Not all of these ideas are completely fleshed out, but this would have been an impossible task in one paper. By walking us through some of the implications of moving from traditional application of rational expectations to newer refinements for modelling inflation expectations, this paper reminds both policy modellers and policymakers that our understanding of the determinants of inflation expectations is far from perfect. By showing us glimpses of the prospect of further progress in our understanding, it also wets our appetite for what comes next.

ENDNOTES:

(1) See Evans and Honkapohja (2005).

(2) Here I draw on some joint work with John C. Williams (see e.g. Orphanides and Williams, 2005 and 2006).

(3) In recent work with Par Osterholm and Spencer Dale (Dale et al, 2008), we reach a similar conclusion: providing more information is not always necessarily better.

REFERENCES:

Blinder, Alan (2008) "Talking about monetary policy: the virtues (and vices?) of central bank communication", paper presented at the Seventh BIS Annual Conference, Lucerne, Switzerland, 26-27 June.

Dale, Spencer, Athanasios Orphanides and Par Osterholm (2008) "Imperfect central bank communication: information versus distraction", Central Bank of Cyprus Working Paper 2008-1.

Evans, George and Seppo Honkapohja (2005) "An interview with Thomas J. Sargent", *Macroeconomic Dynamics*, 9(4): 561-583.

Goodfriend, Marvin (2003) "Inflation targeting in the United States?", NBER Working Paper W9981.

Kim, Don and Athanasios Orphanides (2007) "The bond market premium: what is it and how can we measure it?", *BIS Quarterly Review,* June: 28-40.

Orphanides, Athanasios and John C. Williams (2005) "Imperfect knowledge, inflation expectations, and monetary policy", in B. S. Bernanke and M. Woodford (eds) *The Inflation-Targeting Debate*, Chicago: University of Chicago Press.

Orphanides, Athanasios and John C. Williams (2006) "Monetary policy with imperfect knowledge", *Journal of the European Economic Association*, 4(2-3): 366-375.

Sims, Christopher (2002) "The role of models and probabilities in the monetary policy process", *Brookings Papers on Economic Activity*, 33(2): 1–62.

Sims, Christopher (2008) "Inflation expectations, uncertainty, the Phillips curve, and monetary policy", paper presented at the Federal Reserve Bank of Boston's 53rd Annual Economic Conference, Chatham, Massachusetts, 10 June.

FIGURE 1: Disagreement in Expectations and 10-Year Bond Term Premia

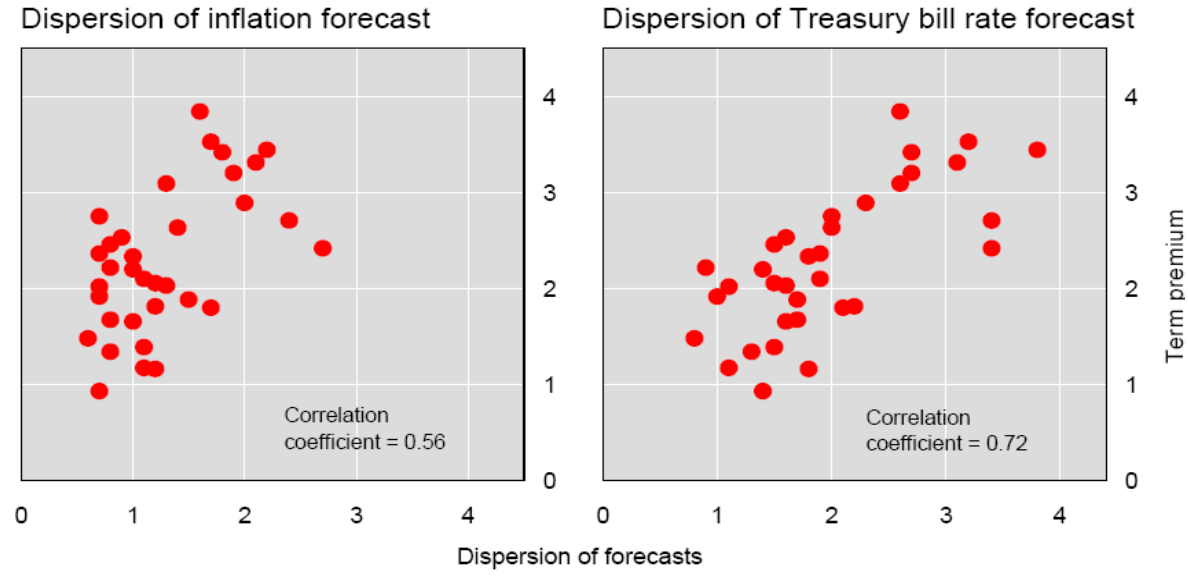

Sources: *Blue Chip Financial Forecasts*; authors' calculations. Graph 3

Note: Reproduced from Kim and Orphanides (2007).

FIGURE 2: Inflation Expectations in the Euro Area

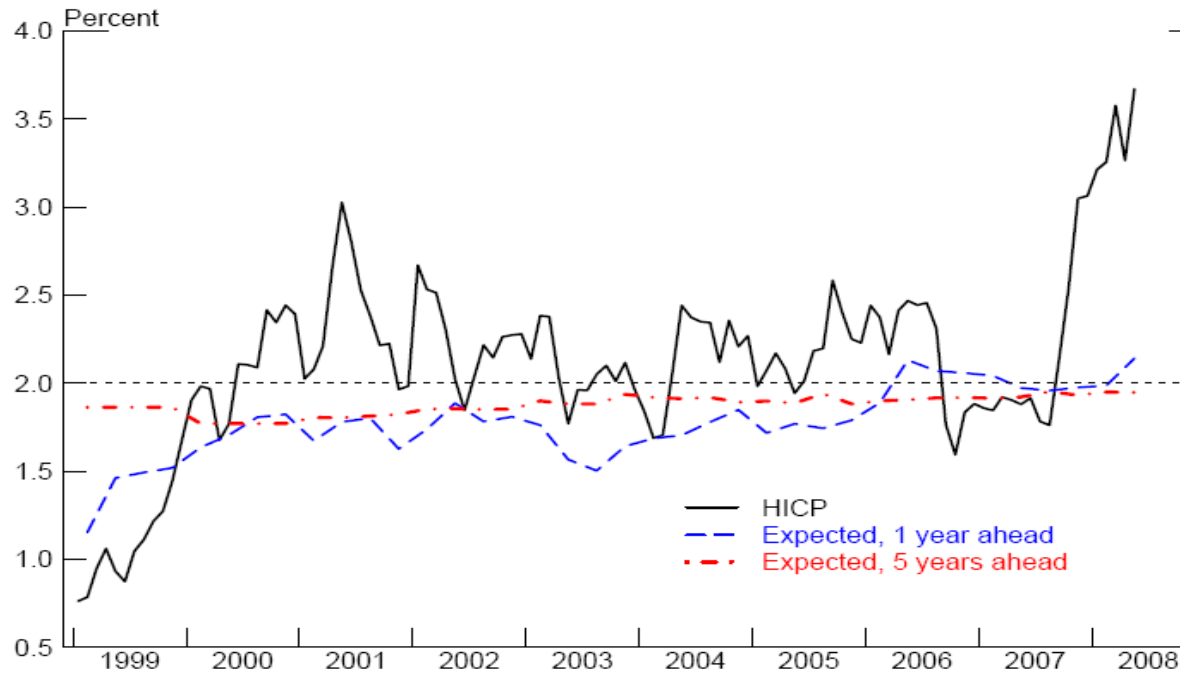

Source: European Central Bank.

FIGURE 3: Disagreement about Inflation Expectations

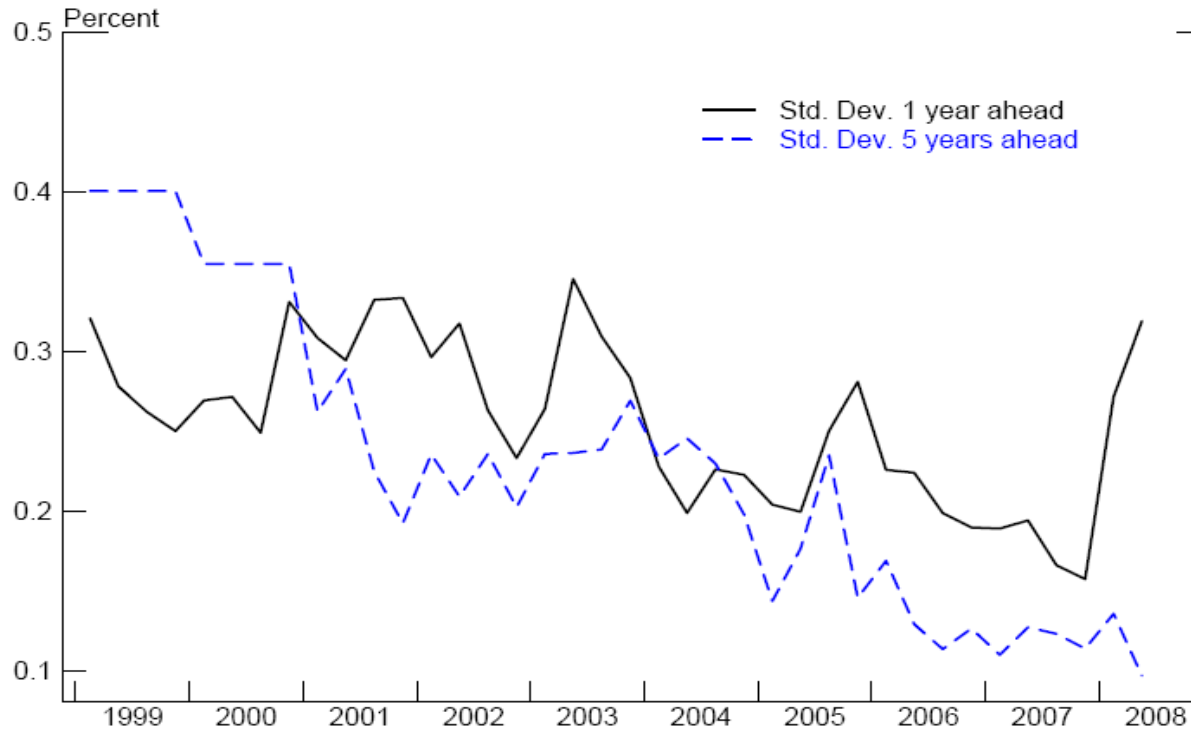

Source: European Central Bank.