BANK OF ENGLAND

Speech

Monetary Policy and Asset Prices

Speech given by

John Vickers, Chief Economist and member of the Monetary Policy Committee

At the Money, Macro and Finance Group 31st Annual Conference, Oxford University

22 September 1999

This paper is based on work by Niki Anderson and James Talbot of the Bank's Monetary Instruments and Markets Division, to whom I am most grateful. Thanks also to Bill Allen, Peter Andrews, Andrew Bailey, Hasan Bakhshi, Roger Clews, Roy Cromb, Spencer Dale, Shamik Dhar, Tolga Ediz, Mark Gertler, Charles Goodhart, Simon Hall, Neal Hatch, Nigel Jenkinson, DeAnne Julius, Mervyn King, Jo Paisley, Chris Salmon, Andrew Scott, Clifford Smout and Peter Westaway for helpful comments and conversations. I alone am responsible for the views expressed in the paper, which are not necessarily shared by other MPC members.

Introduction

How should—and how do—asset prices affect monetary policy? This rather basic question can have strikingly different answers. In some regimes, both ancient and modern, monetary policy is entirely about asset prices. Under the gold standard, and in fixed exchange rate regimes, monetary policy is wholly geared to the domestic currency price of a single asset—gold or an international currency such as the dollar. An asset price is then the very anchor of the domestic price level. By contrast, asset prices do not figure at all in simple monetarist rules for the growth rates of monetary quantities. And where the aim—or at least an aim—of monetary policy is directly to control the rate of inflation of goods and services prices, it is not immediately obvious what role, if any, asset prices should play in the setting of policy.

In this paper I shall concentrate on the role of asset prices when monetary policy targets inflation explicitly—as in the UK for the past seven years—or at least implicitly. The discussion will be organised around three main questions:

- Should asset prices be in the measure of inflation targeted by monetary policy?
- What can asset prices tell us directly about monetary policy?
- What do asset prices add to other indicators that inform monetary policy?

Needless to say, these are large questions with large literatures, and my remarks will be selective and in no way comprehensive. In particular, I shall draw some illustrations from asset price analysis in support of monetary policy at the Bank of England.

Before going any further, let me clarify what asset prices I am talking about. These include the prices of both financial assets—e.g. bonds, equities, and the derivatives such as swaps, futures and options based on them—and non-financial assets including residential property. I will not attempt to cover foreign exchange rates or the prices of durable commodities such as gold.

There are lots more asset prices nowadays than there used to be. This is partly because, thanks to financial innovation, there are many new kinds of asset, and also because assets are more traded across markets. For example, securitised debts, unlike bank loans, have readily observable prices. Future and contingent claims markets are far from complete, but they are somewhat less incomplete than they used to be. Nevertheless many financial assets, including most corporate debt as well as the bulk of personal debt in the UK, are not traded. Market prices for those assets are therefore unavailable for monetary policy assessment, though of course there are other measures of credit conditions (both prices and quantities). So while asset-price information has become vastly richer, it must be kept in mind that many of the effects of monetary policy continue to work through channels involving assets such as bank loans that do not have market prices. The prices of traded assets nonetheless have important effects on intermediated finance via loan collateral values and indeed the balance sheets of financial intermediaries.

A theme worth highlighting at the outset is expectations, which are central to any discussion of asset prices. Asset values depend on—and so can be revealing about—expectations of future behaviour, and asset prices are generally determined flexibly in markets with forward-looking, and perhaps far-sighted, participants. In particular, since the value of money depends on monetary policy, expectations about future monetary policy are key to the pricing of at least those assets with returns denominated in nominal terms. Thus the prices of five-year conventional bonds incorporate, among other things including risk and the possibility of the UK joining EMU, market expectations about the outcomes of the next sixty MPC meetings. One of the interesting questions, then, is how monetary policy should take account of asset prices that reflect market expectations about its own future course.

Should asset prices be in the inflation measure?

Asset prices can affect monetary policy in two conceptually distinct ways. They could be part of the *objective* that monetary policy makers pursue, and/or they can be part of the *information* that policy makers look at, and hence an element in the policy ‘reaction function’. Most of what follows concerns the informational role of asset prices, but there is a prior conceptual question about the policy objective—should asset prices be in the targeted measure of inflation? This question is academic, because available price indices, such as those in the RPI family, generally measure the money cost of a basket of goods and services that are for current consumption.[2] Moreover, under UK arrangements the choice of monetary policy target is of course for the Chancellor to make.

If current utility depends on current consumption—which of course includes consumption of the current services yielded by assets—then, subject to well-known measurement problems,[3] movements in those indices can be interpreted as changes in the money cost of achieving a current level of utility for a hypothetical individual—in short, changes in the cost of *living*. However, lifetime utility depends on future consumption as well as current consumption. So movements in price indices based on current consumption will mis-measure changes in the money cost of achieving a given level of lifetime utility—changes in the cost of *life*, as it were—unless the relative prices of current and future consumption happen to stay constant.

On these grounds, Alchian and Klein (1973) argue that a correct measure of inflation should include asset prices because they reflect the current money prices of claims on future, as well as current, consumption.[4] For example, the fall in long-term real interest rates in the UK over the past year or so has significantly increased the current money price of future consumption.[5] Standard inflation measures do not pick this up.

True, but as Alchian and Klein themselves point out, the appropriateness of a price index depends on the question at hand. The issue for monetary policy is price stability. I believe that this should mean stability over time of the money price of current consumption, and *not* stability of the money price of current-and-future consumption.

---

[2] In fact this is not exactly true. For example, for obvious practical reasons the retail price indices include the prices of durable goods such as refrigerators, rather than the prices of the current services that they yield. A measure of house prices is used to measure housing depreciation, which seeks to proxy the user cost of housing. So there is some direct effect of house prices on the RPIX index. And because a rise in house prices tends to increase the stock of mortgage lending, there is also an effect on the mortgage interest payment component of the RPI.

[3] See the Symposium on measuring the U.S. Consumer Price Index in the Winter 1998 issue of the *Journal of Economic Perspectives*.

[4] Goodhart (1999), among others, agrees.

[5] By contrast, falls in short-term *nominal* interest rates *reduce* the RPI by lowering mortgage interest payments (which are excluded from the RPIX measure of inflation that is the target for UK monetary policy).

First, money is the medium of exchange, including inter-temporal exchange—i.e. exchange involving consumption goods (not utility) at different dates. The efficiency with which such exchanges can occur is enhanced the more that the rate of exchange at each date between money and consumption is stable or at least predictable. Achieving that sort of stability—that is, stability of the price level in terms of current consumption—would seem to be the best that monetary policy can hope to do in terms of facilitating efficient inter-temporal exchange.

In particular, suppose realistically that markets do not exist for inter-temporal exchange of consumption goods (including services),[6] but that they do exist for inter-temporal exchange of money—i.e. there are bond markets. So parties wanting to exchange goods-now for goods-at-time-*t* cannot do so directly. They can however exchange money-now for money-at-*t* by trading cash for bonds. This would replicate the (non-existent) opportunity for inter-temporal exchange of goods, which the parties would really like to have, if the rate of exchange between money-at-*t* and goods-at-*t* were known in advance. In that case, the missing market for inter-temporal exchange of goods would effectively come into being.

This suggests that the best that monetary policy can do to facilitate efficient inter-temporal exchange of goods is to make predictable, for each time *t*, the rate of exchange between money-at-*t* and goods-at-*t*. This argument implies that in principle (not just in practice), the price stability objective of monetary policy should concern stability in terms of the money price of current consumption.

Second, introducing inter-temporal considerations into price measurement literally adds a new dimension to the problems for cost-of-utility measurement arising from diversity among economic agents (without which there would be no trade, so no need for a medium of exchange). The concept of *the* cost of lifetime utility is therefore an order of magnitude more problematic than that of *the* current cost of living, which itself is far from straightforward.

Third, even if the concept were well defined, I do not know how we could hope to measure the prices and quantities of the relevant assets, or even decide what they were.[7] Of course

---

[6] Of course there are some index-linked bond markets, but with relatively few (mostly government) issuers.

measurement difficulties—e.g. relating to quality—are serious also for current price indices, but the inter-temporal dimension would surely magnify them enormously. Measurement matters for all sorts of obvious reasons involving the specification and enforcement of index-linked contracts, including the implicit contracts—e.g. embodied in explicit inflation targets—between the monetary authorities and those to whom they are accountable. Moreover, volatility of asset prices—or at least that part of it due to the volatility of real discount rates—would be transmitted to the price measure, further complicating measurement and accountability.

A separate line of argument for including asset prices in the inflation measure used for monetary policy purposes is that it would make the authorities act in a more timely manner to inflationary (or deflationary) pressures. With consumption price indices as an indicator or target of monetary policy, say Alchian and Klein, 'policy changes will often come too late and move too far'. Believers in this view might point to episodes such as the late 1980s inflation of UK house prices or of Japanese asset prices.

Their argument seems to be that policy makers should be given objectives that include asset prices, because they often fail to take timely action when they have objectives that exclude asset prices, with the result that those objectives are missed. Put simply, if policy tends to be behind the curve, and asset prices ahead of it, then policy might be better if its objective included asset prices. Note that this argument is quite different from the claim discussed above that, as a matter of principle, the ultimate policy objective should include future goods prices. It is saying that, as a pragmatic matter, the authorities should be set a target that differs from the ultimate objective (which might just concern current goods price inflation) in order to offset a perceived bias against timely action. But if such a problem exists, a more direct solution involves the way that policy draws on asset price *information*, rather than the second-best solution of putting asset prices in the policy *objective*.

For these reasons I do not subscribe to the view that asset prices should be in the targeted measure of inflation for monetary policy purposes.

---

[7] Alchian and Klein (1973, pages 187-89) discuss measurement problems. They conclude that the marginal cost of improving a price index along their lines is likely to be less than the marginal benefit of improved policy based on 'less misleading indicators of inflation'.

What can asset prices tell us directly about monetary policy?

Asset prices inform monetary policy both directly and more broadly. The direct value of asset price information is given by the answer to the question: what can asset prices on their own tell us about what monetary policy should be, or is expected to be? The more general value of asset price information (see the next section) can be assessed by asking what can asset price information tell us over and above all the other information that is relevant to inflation prospects?

Extraction of information from asset prices typically involves addressing three kinds of question:

- How do observed market data relate to the economic variables of interest?
- Do asset prices reflect fundamental values—i.e. discounted (risk-adjusted) expected returns?
- What is the attitude to risk of market participants?

As usual in economics, there is the dilemma that making more assumptions yields sharper but more questionable results. (Assuming more can-openers opens more cans but their contents get more dubious.) For example, the nominal and real yield curves would provide very sharp information about inflation expectations if one was heroic enough to assume that one had a robust yield curve estimation technique *and* rational expectations *and* risk-neutrality. Given this dilemma, it is important constantly to be aware which results depend on which assumptions, so that, especially when anomalies or puzzles arise, the assumptions (and data and methods) are properly questioned before results are given credence.

For the task of relating economic data to economic variables, some basic tools of inference are yield curves and implied probability distribution (PDF) functions.[8] Yield curves can be estimated for nominal or real interest rates and for spot or implied forward rates.[9] Of particular interest are

---

[8] Thanks to the development of reasonably liquid derivative markets and of techniques of financial market analysis, it is now possible to make inferences about entire probability distributions, not just the means of those distributions. Söderlind and Svensson (1997) provide an excellent recent survey of those techniques.

[9] Campbell (1995) gives an overview of yield curve economics; see also the Annex to this paper. The spot yield at maturity t is the continuously compounded rate of interest on a zero-coupon bond of that maturity. The s-period forward rate at maturity t—or the implied forward if there is no forward market—is the s-period spot yield t periods hence. Yield curves show the term structure—i.e. rates as a function of time t.

yield *spreads*—i.e. differences between types of yield. For example, the spread between nominal and real interest rates gives 'break-even' inflation rates,[10] which under certain assumptions can be interpreted as expected inflation rates. And the spreads between corporate and government yields (or between the yields on common currency debt of different countries) provide measures of relative default risk.

Yield curve (and PDF) analysis involves the technical issue of how to derive curves from market data points, and the economic issue of how to interpret them. In what follows I will simply assume that yield curves (and PDFs) can be observed directly, and so bypass the challenging and fascinating problems of how best to derive them from market data. But the Annex to this paper by Niki Anderson of the Bank's Monetary Instruments and Markets Division contains a brief discussion of some aspects of those problems, and how we are tackling them at the Bank.

Inflation expectations

Measurement of inflation expectations is central to monetary policy analysis. Inflation prospects are influenced by inflation expectations, for example through the process of wage and price setting. Low and stable inflation expectations are key to price stability and the wider economic benefits that flow from it. And inflation expectations provide a way of appraising the MPC—inflation out-turns are unlikely to be exactly on target *ex post*, and inflation expectations are a guide to whether policy is set broadly right *ex ante*, which is the best one can hope for. This point applies to introspection by the MPC as well as to external appraisal.

For example, if the forward break-even inflation rates implied by the nominal and index-linked yield curves are substantially above or below the inflation target of 2.5%, then

(a) the market expects inflation to miss the target, and/or
(b) forward break-even inflation rates do not measure the market's expectation (i.e. mean) of future inflation.

---

[10] The term 'break-even inflation rate' is used specifically in the market to refer to the average rate of inflation which equalises the real return on two comparable gilts: one conventional and one index-linked. The term is used more broadly here to refer also to the difference between the nominal and real (zero-coupon) yield curves derived from the gilt market.

In case (a), it might be that the MPC also expects inflation on average to miss the target—for example if shocks have hit the economy and the horizon is too short sensibly to offset them. Otherwise the interpretation is that the market's expectation about likely economic developments and the MPC's reaction function together imply missing the target on average. In that case one should think hard about why the market's apparent assessment differs from that of the MPC. In that sense 'the market' is like an external forecaster which puts its money where its mouth isn't.[11]

But, as case (b) recognises, the implied break-even inflation rate should not necessarily be equated with the market's expectation of inflation. Short real rates are notoriously hard to measure, institutional or liquidity factors may be affecting prices, and non-zero inflation risk premia are entirely possible.

In part for these reasons, if break-even inflation rates were on target, it would not necessarily mean that policy was being set right. For example, if policy were totally credible, then market-derived measures of inflation expectations would contain no information about inflation expectations beyond being a check that policy remained credible. In the possibly unhealthy situation in which the market was certain that the MPC would do whatever it takes to get expected inflation on target, market-derived inflation expectations would give no indication what to do, but market-based interest rate expectations might still be informative—see below. This is not for a moment to suggest, conversely, that so long as market participants form independent views, policy should follow the market—for example by aiming to get break-even inflation rates on target. The point is simply that market-derived inflation forecasts, like external forecasts generally, potentially have *some* informative value for policy to the extent that they are independent.

Charts 1 and 2 show five and ten year forward nominal and real interest rates derived from the Bank's zero-coupon government yield curve. The charts also show forward break-even inflation rates at the same maturities. Assuming rational expectations and risk neutrality, these break-even inflation rates represent the market's inflation expectation for the six-month period beginning five and ten years from any given point on the chart. Ten-year forward break-even inflation rates have

---

[11] Humanising 'the market' by attributing expectations to it is deeply unsound philosophically, but useful shorthand.

fallen from over 10% in 1982 to 2.6% today. King (1995) used implied forward inflation rates to assess the credibility of monetary policy: a ten year forward rate of 2.6% is consistent with the market expecting future UK inflation to remain broadly in line with the current inflation target. The fall in forward break-even inflation rates may in part reflect lower inflation risk premia. If so, that too is consistent with greater monetary policy credibility.

Chart 1

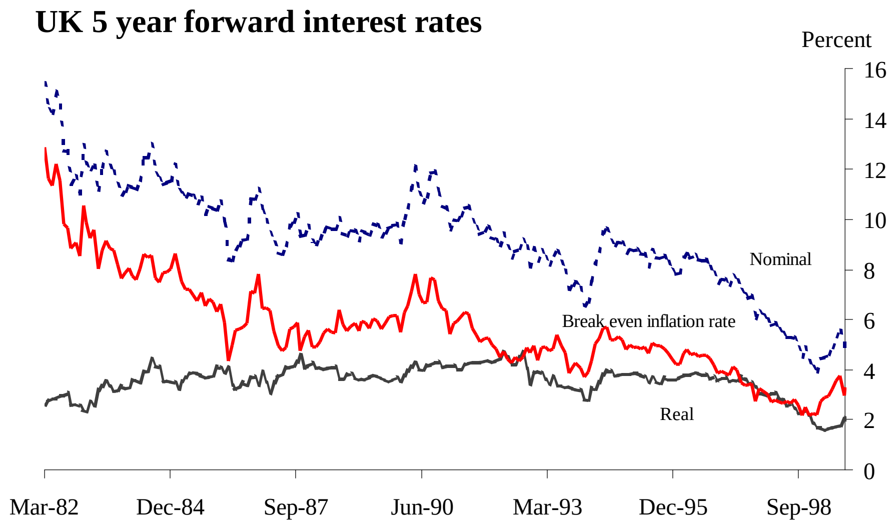

Chart 2

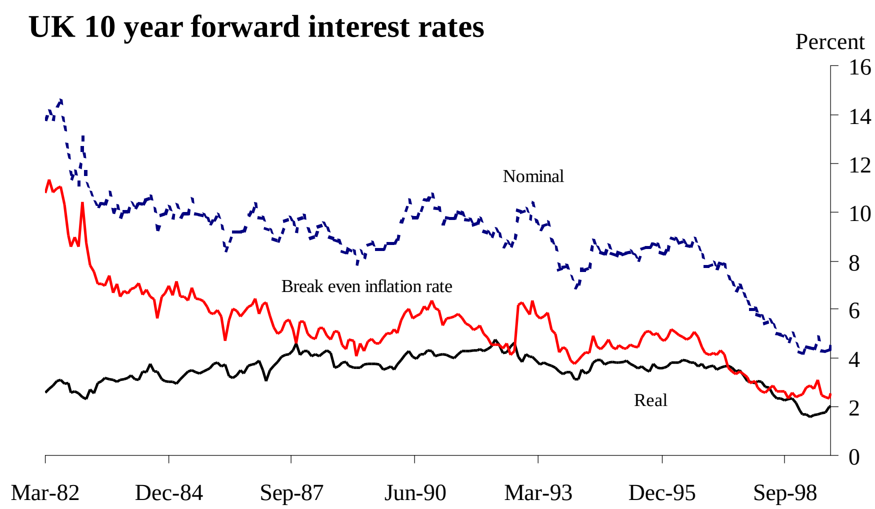

Of course yield curves are not the only source of information about inflation expectations, just as financial market participants are but one of a number of groups in the economy whose inflation expectations matter. In addition there are direct measures such as the Barclays Basix survey—see Table 1, which shows that reported inflation expectations appear generally to have become closer to the 2.5% target over the past year or so. These survey measures also need careful interpretation, but they are a valuable complement to, and cross-check upon, market-based measures.

Table 1: Short-term inflation expectations [(a)]

	1998 Q2	1999 Q2
General public	5.1	4.6
Trade unions	3.8	3.2
Finance directors	3.2	2.4
Business economists	2.9	2.7
Investment analysts	3.3	2.8
Academic economists	3.1	2.6

(a) Expectations of inflation rate one to two years ahead. RPI inflation, except for general public, for which the measure of inflation is not specified.

Source: Barclays Basix Survey.

Interest rate expectations

Yield curves are directly informative about expected future monetary policy, and therefore about market views of inflationary pressures. For example, if the market respectfully took the view that the MPC would do whatever was needed to keep inflation on target, the near-term yield curve would indicate how the market thought monetary policy should be set to deliver the target. A rising/falling curve would indicate a market view of strengthening/easing inflationary pressures.

Short sterling futures and swap rates indicate expected nominal interest rates, subject to the point above about risk premia. However, expected monetary policy—i.e. expected *official* rates—cannot be inferred directly from the short sterling or swap curves because there are credit spreads, which may vary over time.

Thus short sterling futures are contracts on the three-month London inter-bank rate, whereas the MPC sets a two-week repo rate.[12] Normally the short sterling futures rate might be around 20 basis points (hundredths of one per cent) above the corresponding expected official repo rate for reasons of credit risk, liquidity and the difference in maturity of the two rates.[13] But there is no reason to think that this spread is constant—at present there are clear signs of a millennium effect

[12] The repo rate is the official rate of interest chosen by the MPC at which the Bank of England lends to the money market. See further the box on page 5 of Monetary Policy Committee (1999).

[13] See the box on page 331 of the November 1997 *Bank of England Quarterly Bulletin*.

for example—and indeed the analysis of spreads is an important topic of its own (see below). Likewise, since swap rates are based on current and expected future values of the six-month libor rate, these tend to be higher than yields on government bonds of the same maturity.

For any given path for expected nominal interest rates—whether official or market rates—an important issue for monetary policy analysis is the decomposition between expected real rates and expected inflation (and risk premia). A recent practical example of this issue arose in May this year. The short sterling curve, which had been close to flat in the early part of the month, rose quite sharply—see Chart 3.[14] (This phenomenon was not confined to the UK, and indeed it seemed in part to have followed US developments.) Had expected inflation risen sharply? Or expected real rates? Or was something else, not in the textbooks, going on?

Chart 3

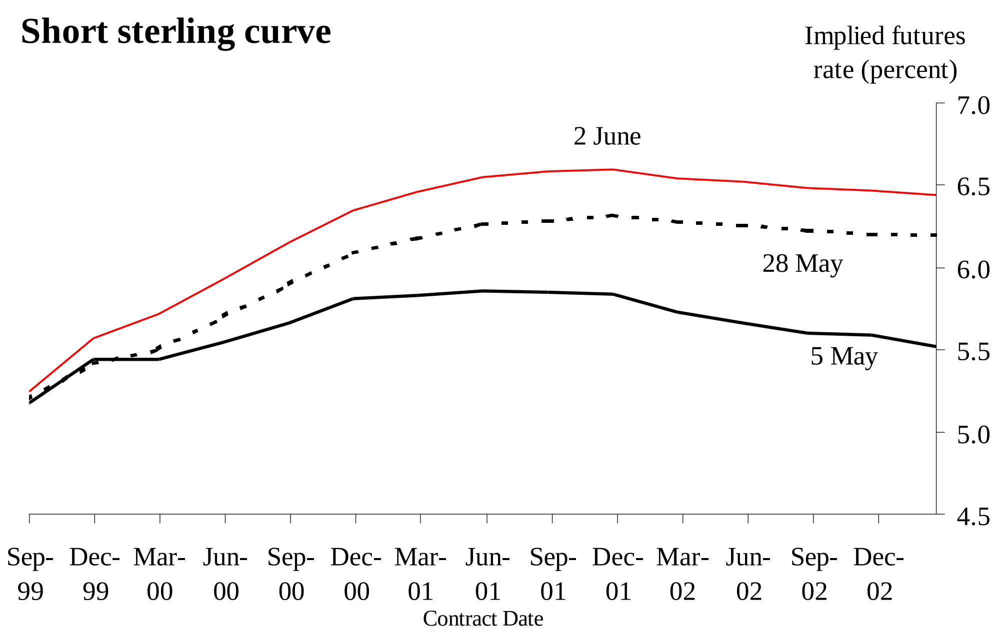

[14] The steeping of the curve went further over the summer. One manifestation of this is the comparison between inflation projections conditional upon a constant interest rate and upon market interest rates. In the May 1999 *Inflation Report*, market rates were consistent with the official rate remaining not far from its then prevailing rate of 5.25% over the two-year forecast horizon. But by the time of the August *Inflation Report*, market rates were implying an official rate two years ahead some 2% above the then prevailing official rate of 5%. In May the inflation projections conditional upon market and constant rates were virtually indistinguishable. Not so in August.

Different explanations have potentially different policy implications. Higher inflation expectations, if justified, might point in the direction of raising the official rate. However, a tightening in market real rates would tend to dampen activity and prospective inflation, which could be a factor pointing towards lowering the official rate, at least in the short term, to keep inflation on track to meet the target. (But note that this would not be appropriate if the higher real rates were anticipating a tightening of monetary policy needed to head off a rise in prospective inflation.)

The evidence considered at the June meeting of the MPC was mixed. At longer maturities, index-linked bond yields had changed little while nominal yields had risen, suggesting that inflation expectations had increased. But it was possible that changes in risk premia (e.g. as the after-effects of LTCM faded) and institutional factors (e.g. pension fund regulations) were distorting inflation expectations as measured by yield curve differentials. Survey evidence and other forecasters' projections showed little rise in short-to-medium term inflation expectations. Overall it was judged implausible that inflation expectations had risen by as much as short maturity forward rates, so short real yields had probably risen.

What do asset prices add to other information?

The fact that asset prices on their own can yield information for monetary policy purposes does not mean that such information is necessarily worth much—beyond its value as a cross-check on other forecasts—on top of all the other information that is available to guide policy. I want to argue, however, that its value added is substantial.

One way to consider this issue is to start from the standard Taylor Rule.[15] This rule sets the interest rate above or below some 'neutral' level to a degree depending on (i) the difference between current inflation and target inflation, and (ii) the difference between actual and potential output—the output gap. Assuming that asset prices are not in the inflation measure, this rule

[15] There are of course many variants of the standard Taylor Rule—for example incorporating expected future inflation rather than current inflation. Bernanke and Gertler (1999) compare policy rules based on expected future inflation with and without an independent role for asset prices. In their simulation experiments, the best policy—in terms of minimising inflation and output volatility—is to focus aggressively on expected inflation and not to pay attention to asset prices except insofar as they matter for expected inflation (which they do).

appears to ignore asset price information except insofar as it is used to determine the neutral level of the interest rate.

Compare this with a regime, such as that in the UK, where monetary policy, which works with lags, is set with a view to getting *prospective* inflation on target. In that context asset prices are helpful both in forecasting—making projections of inflation prospects—and in policy analysis—assessing how policy might affect those prospects.

At a technical level, asset prices feature in a number of the models, including the core macroeconometric model, that are used to help the MPC form its projections for inflation and growth—as described in the recent book *Economic Models at the Bank of England*. And a central theme of the recent MPC paper on *The Transmission Mechanism of Monetary Policy* is the key role of asset prices in the processes whereby monetary policy decisions work through the economy. But I will not attempt to summarise those expositions here.

House prices

Most of the discussion so far has concerned financial assets that are continuously traded in liquid wholesale markets with low transactions costs, where prices reflect market perceptions more or less instantaneously. Markets for residential property are different in nature. Homes, and loans secured on them, are respectively the most important assets and liabilities of the personal sector in the UK. Developments in house prices and activity are therefore a major part of the asset price information relevant for monetary policy. However, the relationships between house prices and inflation are complex and imperfectly understood.

Consider, for example, some aspects of the relationship between house prices and consumer expenditure.[16] Higher house prices increase the (gross) wealth of homeowners, and greater wealth tends to lead to higher consumer spending. Indeed it may be that increases in housing wealth tend to influence consumption more than increases in financial wealth. Property can be used as collateral for household borrowing—through mortgage equity withdrawal—more readily

[16] There are also direct effects of house prices on inflation measures—see footnote 2 above.

than, say, pension fund wealth. House prices may also reflect other influences on consumption, such as the level of, and uncertainty about, expected future income from employment—in short, confidence. And activity in the housing market, which tends to be greater when prices are rising, can have natural direct effects on demand for consumer durables. However, the effects of higher house prices on consumer spending are far from unambiguous. For example, higher house prices may lead non-homeowners—and more generally those who own less housing than their anticipated housing needs—to scale back expenditure on non-housing services.

With these and other various possible influences at work, there can be no precision in judging the implications for aggregate consumption even of *past* developments in house prices. There is the further difficulty of assessing possible *future* house price developments. At times house prices move sharply, not just in particular localities, and there have been episodes such as occurred ten years ago when an apparently self-fuelling house price boom abruptly collapsed. The rise in house prices was followed by a surge in goods and services price inflation, which neared 10% in 1990. A protracted depression of consumer demand followed the fall in house prices, partly on account of the negative equity of a number of indebted homeowners.

The current level of house prices nationally does not appear to be obviously abnormal in relation to income—see Chart 4—but house price indices have accelerated quite sharply in recent months. Annual house price inflation is currently 9% or so, and in the three months to August, house prices rose by about 5% and 4% respectively according to the Halifax and Nationwide measures.

Chart 4

House price to earnings ratio[a]

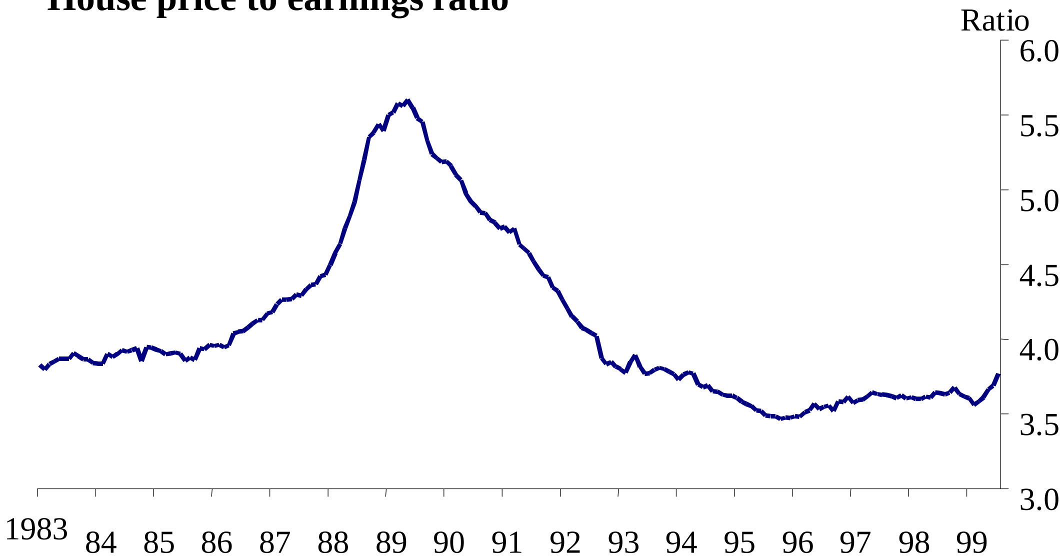

(a) Ratio of Halifax house price index to ONS whole economy earnings

Assessing the possible implications for demand and inflation of house price developments is a highly uncertain matter. But monetary policy requires continuous assessment of evolving balance of risks to prospective inflation. House price developments are a part of this—if only a relatively small part of a larger picture—because, like other asset price developments, they add to the other information available to guide judgments about inflation prospects. So house prices are not an independent concern of monetary policy. Like every other economic indicator, they matter to the extent—and only to the extent—that they say something about overall inflation prospects.

Equity prices and the bubble question

No discussion of asset prices nowadays would be complete without a word on the equity bubble question. Is there a bubble in equity prices? Can monetary policy deflate, or pre-emptively burst, asset price bubbles? If so, should it?

The answer to the first of these questions is that you never know for sure whether market prices do, or do not, reflect economic fundamentals. Most of the techniques for extracting information from asset prices discussed above are based on the bubble-free assumption that asset prices do reflect fundamentals. Sensible use of those tools requires constant awareness that this might not be true. Bubble detection, by contrast, requires independent knowledge of the fundamentals, or at least of their bounds.

Explanations based on fundamentals of the current levels of equity prices, notably in the US, have to stretch in a number of dimensions—for example concerning future profit growth rates, discount rates and risk premia.[17] However, it is by no means certain that there is a large bubble in equity prices, and that a sharp 'correction' will happen. But of course it might.

Can monetary policy burst bubbles? This question is hard to think about, because in theory the ways that monetary policy might naturally affect equity prices involve the fundamentals—short-term discount rates and perhaps profit growth rates—whereas bubbles are departures from the fundamentals. It is quite possible that a sharp monetary policy tightening could be the straw that breaks a bubble's back, but exactly how we do not know.

Supposing, I think unrealistically, that the monetary authorities *could* identify and pre-emptively burst bubbles, should they do so, even if this meant prospective inflation undershooting the target? Assume that bubbles grow before they burst, and that, as they grow so too do the associated uncertainties for future inflation and output. If those uncertainties are undesirable, which is entirely consistent with inflation targeting, then there is a case for tightening policy to burst a growing bubble early on. Even with a symmetric inflation target, expected inflation somewhat under target with moderate inflation uncertainty might be better than expected inflation on target with high inflation uncertainty.[18]

Yes, but there are such doubts about the authorities' ability to detect and burst bubbles, especially in their infancy, that I do not find this a practically persuasive argument for *aiming off* the inflation target on account of possible bubbles. Of course this is quite consistent with taking account of possible bubbles when *aiming at* the target. Moreover, if markets believed that monetary policy was responding, or would respond, directly to particular asset price

---

[17] See Wadhwani (1999) for a discussion of US equity price levels.

[18] Kent and Lowe (1997) develop this sort of argument formally. The logic is like Brainard uncertainty. If the policy-makers' loss function—i.e. loss as a function of the extent to which inflation misses the target—is convex, then inflation uncertainty matters as well as mean inflation. For example, with a quadratic loss function, the mean and variance of inflation both matter.

developments, then monetary policy could itself come to distort those asset prices, with consequent misallocation of resources.

This discussion has echoed the findings in the recent paper by Bernanke and Gertler (1999). The main conclusion from their analysis is best quoted verbatim (but with emphasis added):

'Inflation-targeting provides an effective, unified framework for achieving both general macroeconomic stability and financial stability. Given a strong commitment to stabilizing expected inflation, it is neither necessary nor desirable for monetary policy to respond to changes in asset prices, *except to the extent that they help to forecast inflationary or deflationary pressures*'.

Conclusion

In a number of countries, including the UK, recent times have seen sharply rising asset prices while the inflation of consumption goods and services prices has been subdued. Does asset price inflation mean that monetary policy should be tight even if current inflation is low? And if asset prices were to fall sharply, should monetary policy be loose even if current inflation was higher?

One way of summarising the arguments in this paper is to say that the answers to questions of this kind flow from the straightforward proposition that monetary policy should be set so that *prospective* inflation of consumption prices is on target. If asset prices were a substantial element of the targeted inflation measure, then policy would respond partly automatically to asset price inflation/deflation. But I doubt that they should be. If policy-makers went bubble hunting, then suspicious asset prices could acquire special significance for policy. But that would seem a hazardous pursuit.

Asset prices matter for monetary policy simply because they help inform judgments about inflation prospects. They do this because, in part, asset prices *are* judgments about economic prospects. So asset price makers and interest rate setters have a good deal in common. Both must look ahead, and in doing so they are well advised to keep an eye on each other.

ANNEX
by Niki Anderson

This annex provides a brief overview of the main issues in the estimation of yield curves and implied probability distributions from asset prices. The first step is to offer a precise definition of the information that can be extracted, and to discuss how this relates to the observable data. A parametric approach to extracting this information is then described and some problems that have been encountered are highlighted. Finally, there is a discussion of how non-parametric methods may be used to overcome some of these problems.

What information can be extracted?

Focusing first on the yield curve, the main text has described a number of ways in which data on debt instruments can be used to inform monetary policy. Some of the most useful pieces of information that can be derived from these data are implied forward rates of interest. These are important in their own right as they reflect—albeit imperfectly—market beliefs about the future path of interest rates. But they also provide the building blocks for other types of information including the synthetic bond prices that can be created in order to derive credit spreads from observed corporate bond yields.

Implied forward rates are the marginal rates of return investors require in order to hold bonds of different maturities. Suppose, for example, that it is possible to observe two zero-coupon bond prices with terms to maturity of t and $t+s$. The implied s-period forward rate in t period's time measures the marginal difference in price between these two bonds i.e.:[19]

$$f(t, t+s) = \ln B(t) - \ln B(t+s) \qquad \textbf{(1)}$$

where f is the forward rate and B is the price of a zero-coupon bond. More generally, it might be desirable to observe forward rates over a range of different periods, s. This can be achieved by first defining what are known as instantaneous forward rates. These correspond to the case for which s tends to zero. Equation (1) implies that these are related to bond prices as follows:

$$B(t) = \exp\left[-\int_0^t f(t)dt \right] \quad \textbf{(2)}$$

These rates are the building blocks of the yield curve and can be used to price bonds of any maturity and hence calculate the implied forward rate of interest over any horizon in the future.[20] Thus these are the rates which would ideally be measured.

Turning to information from options, these can be used to extract an implied probability distribution around potential future outcomes of the underlying asset. To see how it is possible to extract this information requires only the definition of a call price:[21]

$$C(S, X, t) = \exp(-rt)\hat{E}[\max(0, X - S_t)] \quad \textbf{(3)}$$

where C is the price of the call, X is the strike price S_τ is the terminal underlying asset price, r is the riskfree rate of interest and t is the maturity of the option. The '^' denotes that the expectation is taken under the risk-neutral probability measure, which basically means that it does not necessarily represent the true beliefs of the investor, rather those adjusted for risk.

This expectation term will be determined by the set of (risk-neutral) probabilities the investor attaches to the likelihood of each future outcome for the price of the underlying asset. Equation (3) implies that the price of the call is related to this probability distribution as follows:

$$C(S, X, t) = \exp\{-rt\}\int_X^\infty (S_t - X)g(S_t)dS_t \quad \textbf{(4)}$$

---

[19] Notice that this assumes the forward rate is continuously compounded.

[20] By definition, the s-period implied forward rate in equation (1), $f(t, t+s)$, is equal to the sum $\int_t^{t+s} f(t)dt$.

[21] Strictly speaking, the expectation term should be shown to be conditional on information about the current value of the underlying asset, S. Similarly, in later equations, the implied probability distribution should be conditional on this information. These points are excluded for the sake of brevity.

where $g(S_\tau)$ is the risk-neutral probability that the terminal value of the asset will be S_τ. Abstracting from any issues of interpretation, these probabilities are clearly useful in describing market perceptions about the value of the underlying asset.

Why is this information difficult to measure?

To measure the set of implied instantaneous forward rates directly from the market requires a set of observable zero-coupon bond prices across a continuum of maturities, where this is referred to as the discount function. By definition:

$$f(t) = -\frac{\partial B(t)}{\partial t}\frac{1}{B(t)} \tag{5}$$

In practice, however, only the prices of coupon-bearing bonds can be observed.[22] As a result, the discount function embodied in their prices cannot be measured directly. All that can be done is to express the price of each observable bond as follows:

$$P(c,t) = \sum_{i=1}^{n} cB(t-i) + 100B(t) \tag{6}$$

where t denotes the maturity of the bond, c is the coupon payment made in each period and n refers to the number of such payments outstanding. A more fundamental problem is that these bonds are issued only across a finite set of maturities. To define a continuous discount function, and hence an instantaneous forward rate curve, effectively requires a method for 'filling in the gaps'.

A similar problem is encountered when trying to extract an implied probability distribution from option prices. In this case, it can be shown that the risk-neutral probability, $g(S_\tau)$, is given by:[23]

---

[22] In the gilts market, since the introduction of strips in December 1997, it has been possible to observe these prices directly. However, it is not clear that trading in the strips market is sufficiently active to produce reliable prices for yield estimation purposes.

[23] See, for example, Breeden and Litzenberger (1978).

$$g(S_t) = \exp\{rt\}\frac{\partial^2 C}{\partial X^2} \qquad \textbf{(7)}$$

To measure this further requires a continuum of option prices across different strike prices. Only then would it be possible to directly observe the second derivative of the call pricing function with respect to X. In practice, however, there are usually only a handful of contracts at different strike prices, which are assumed to be priced using equation (4).

Parametric Approaches

The simplest approach to solving each of these problems is to define the two variables of interest, $f(t)$ and $g(S_\tau)$ as functions of a set of unknown parameters. In the forward rate case, for example, the following functional form developed by Svensson (1994, 1995) is often used:

$$f(t) = b_0 + b_1 \exp\left(-\frac{t}{a_1}\right) + b_2\left[\left(\frac{t}{a_1}\right)\exp\left(-\frac{t}{a_1}\right)\right] + b_3\left[\left(\frac{t}{a_2}\right)\exp\left(-\frac{t}{a_2}\right)\right] \qquad \textbf{(8)}$$

where b denotes the vector of parameters, $\{b_0, b_1, b_2, a_1, a_2\}$. Given this function, equations (2) and (6) can be used to express the price of any coupon bond as a function, $P(c,t,b)$, of the underlying parameters together with the known variables, c and t. The problem of deriving the implied forward rate curve then reduces to one of estimating the parameters using an objective criterion to compare observed and fitted prices. Having estimated these parameters, an estimated implied forward rate curve can be derived using equation (8).

Similarly, a parametric expression for the implied probability distribution, $g(S_\tau)$ can be derived. A common assumption is that this is defined by a mixture lognormal distribution which is characterised by a set of parameters, $b=\{b_1, b_2, \alpha_1, \alpha_2, \theta\}$, i.e.[24]:

$$g(S_t) = qL(a_1, b_1) + (1-q)L(a_2, b_2) \qquad \textbf{(9)}$$

---

[24] See, for example, Bahra (1996, 1997) and Melick and Thomas (1997).

where L is the lognormal distribution function evaluated at S_τ. Again this means that the observed data can be expressed as a function, $C(S, X, t)$ of the fixed variables, X, r and t, and the set of unknown parameters, b. All that remains to do is to specify an objective function which compares fitted and observed prices so that the parameters can be estimated in such a way that the value of this function is minimised over all available observations. The implied probability distribution function is then given by equation (9).

Shortcomings

These two methods have both been used extensively at the Bank over the past few years. Experience of doing so has highlighted a number of shortcomings with this parametric approach. These are described in more detail in forthcoming working papers.[25] The main problem is that the derived yield curves and implied probability distributions tend to be fairly unstable. By this it is meant that small changes in the underlying data can produce fairly large differences in the estimated curves. This can produce undesirable properties in the two types of curve.

In our work, these observations have been attributed to the parametric nature of the functional forms used to fit the curves. A small change in one of the datapoints will have a direct effect on at least one of the parameter estimates. Via the functional form, this feeds through to the fitted values of each of the remaining datapoints so that, via the objective function, other parameters are inevitably affected by the change. This can lead to a significantly different curve. Since these idiosyncratic price movements are fairly common, the result is that the estimated yield curves and implied probability distributions can fluctuate significantly across time, even if there has been little change overall in the observed prices upon which they are based.

Non-parametric Methods

The approach taken at the Bank to deal with these problems is to examine instead non-parametric estimation of the two types of curve. In the case of the yield curve, this takes the form of a spline-based specification of the implied instantaneous forward rate curve. Instead of defining a

---

[25] These papers will be available on the Bank of England’s website at http://www.bankofengland.co.uk

continuous curve—so that forward rates at all maturities are a function of the same underlying set of parameters—forward rates are described by a number of piecewise cubic polynomials which are joined at knotpoints. This effectively localises the influence of any movement in the price of an individual bond to one portion of the curve.

Given the relative lack of data points, a slightly different approach is used in the case of implied probability distributions. As mentioned previously, the distribution could be calculated directly if there were a continuous set of call prices. An alternative approach, therefore, is to first use a non-parametric curve-fitting technique to 'fill in the gaps' between the observations that are available.[26] This produces a continuous call pricing function. From this, the implied probability distribution is derived directly using equation (7). Notice that if it were possible to observe zero-coupon bond prices, a similar exercise could be performed to derive a continuous discount function, from which to derive instantaneous forward rates using equation (5).

Of course, non-parametric techniques have their own problems. Experience in the case of the yield curve, for example, suggests that they can be too flexible, fitting too closely to individual bonds rather than capturing the underlying shape of the term structure. Methods are available, however, to deal with these problems. For example, Fisher, Nychka and Zervos (1996) estimate an additional parameter that penalises excess curvature in the forward rate curve. This is modified slightly by Waggoner (1997): he allows the penalty parameter to vary across maturity in order to account for the fact that the shape of the term structure is likely to be more well-defined at the short end (so that less smoothing is required) than at longer maturities.

[26] In practice, these observations are dealt with on the basis of their implied volatilities. Having fitted a non-parametric function through these volatilities, we then transform this curve to derive a continuous call pricing function.

REFERENCES

Alchian, A and Klein, B (1973), ‘On a Correct Measure of Inflation’, *Journal of Money, Credit, and Banking*, Vol 5, pages 173-191.

Bahra, B (1996), ‘Probability Distributions of Future Asset Prices Implied by Option Prices’, *Bank of England Quarterly Bulletin*, Vol 36(3), pages 299-311.

Bahra, B (1997), ‘Implied risk-neutral Probability Density Functions from Option Prices: Theory and Application’, Bank of England Working Paper No 66.

Bank of England (1999), *Economic Models at the Bank of England*.

Bernanke, B and Gertler, M (1999), ‘Monetary Policy and Asset Price Volatility’, paper presented at the Federal Reserve Bank of Kansas City conference on ‘New Challenges for Monetary Policy’, Jackson Hole, Wyoming, August.

Breeden, D and Litzenberger, R (1978), ‘Prices of State-contingent Claims Implicit in Option Prices’, *Journal of Business*, Vol 51, pages 621-651.

Campbell, J (1995), ‘Some Lessons from the Yield Curve’, *Journal of Economic Perspectives*, Vol 9, pages 129-52.

Fisher, M, Nychka, D and Zervos, D (1995), ‘Fitting the Term Structure of Interest Rates with Smoothing Splines’, Finance and Economics Discussion Series, 95-1, Federal Reserve Board.

Goodhart, C (1999), ‘Time, Inflation and Asset Prices’, paper presented at conference on ‘The Measurement of Inflation’, Cardiff Business School, August.

Kent, C and Lowe, P (1997), ‘Asset-Price Bubbles and Monetary Policy’, Reserve Bank of Australia Research Discussion Paper 9709.

King, M (1995), 'Credibility and Monetary Policy: Theory and Evidence', *Bank of England Quarterly Bulletin*, Vol 35(1), pages 84-91.

Melick, W and Thomas, C (1997), 'Recovering an Asset's Implied PDF from Options Prices: An Application to Crude Oil During the Gulf Crisis', *Journal of Financial and Quantitative Analysis*, Vol 32, pages 91-115.

Monetary Policy Committee (1999), *The Transmission Mechanism of Monetary Policy*, Bank of England.

Söderlind, P and Svensson, L (1997), 'New Techniques to Extract Market Expectations from Financial Instruments', *Journal of Monetary Economics*, Vol 40, pages 383-429.

Svensson, L (1994), 'Estimating and Interpreting Forward Interest Rates: Sweden 1992-4', CEPR Discussion Paper No 1051.

Svensson, L (1995), 'Estimating Forward Interest Rates with the Extended Nelson and Siegel Method', *Sveriges Riksbank Quarterly Review*, Vol 95, pages 13-26.

Wadhwani, S (1999), 'The US Stock Market and the Global Economic Crisis', *National Institute Economic Review*, No 167, pages 86-105.

Waggoner, D., (1997), 'Spline Methods for Extracting Interest Rate Curves from Coupon Bond Prices', Federal Reserve Bank of Atlanta Working Paper 97-10.