Defined benefit company pensions and corporate valuations: simulation and empirical evidence from the United Kingdom

Working Paper no. 289

Kamakshya Trivedi and Garry Young

Shareholders of sponsoring companies are primarily responsible for ensuring the solvency of the defined benefit (DB) pension schemes that firms offer their workers. Hence, even though the assets and liabilities of such pension schemes are distinct from the company's balance sheet, corporate sponsors are clearly the residual claimants or guarantors, and hence they should be analysed together. This paper investigates whether this feature of UK company pensions affects how company stock prices respond to common shocks. We consider two channels through which common shocks to companies' real business values can be amplified. First, to the extent that defined benefit pension liabilities are debt-like, they add to the overall leverage or indebtedness of companies. For given asset risk, we should expect that more highly levered stocks are more volatile. Second, in the United Kingdom pension scheme assets are largely invested in equities of other UK companies. These cross-holdings of equity mean that common shocks to company valuations are transmitted among each other via their defined benefit pension schemes, and the response of stock prices to such a shock can be amplified.

If it does exist, this kind of amplification is clearly of relevance to systemic financial stability, since it can rapidly push corporate valuations upwards or downwards, and there could be corresponding knock-on effects on the wider macroeconomy. For example, if capital investment is sensitive to corporate valuations, through either cost of capital or Tobin's Q effects, then this could exacerbate the real economic cycle. In addition, stock return volatility can also be costly for individual companies and their shareholders. Higher volatility can increase a company's perceived riskiness, and therefore its cost of external capital. Alternately, as a company's stock price becomes a less informative signal of 'true' value, stock-based compensation becomes less effective at providing appropriate incentives to managers.

To investigate these issues we start with a stylised model of a company's balance sheet — in which pension fund assets and liabilities are treated in exactly the same way as a company's ordinary, or on balance sheet, liabilities. Using the model we demonstrate how common shocks can be amplified on account of 'economic leverage' and equity cross-holdings. We then calibrate this model for about 90 of the FTSE 100 companies and simulate it to illustrate the possible size of such amplification effects. We perform two simulations where the company's business value is reduced by 5%. In the first simulation the total effect of the shock is the sum of the effect from the cross-holdings channel and the leverage channel. We compare these effects with a second simulation where we switch off any effects from the cross-holdings channel (consistent with the company's pension fund equity assets being invested abroad). The comparison allows us to break down the total impact of the 5% shock into the part that comes from additional leverage and the part that comes from cross-holdings. Our main result is that on average, the shock causes a 10.5% reduction in market value. Of the additional 5.5% reduction, 1.4% was due to companies holding other companies' equity in their pension funds. The remainder is due to the higher leverage induced by pension liabilities.

We also examine whether such effects, in fact, exist in data within the framework of a standard Capital Asset Pricing Model (CAPM). Empirical analysis using matched balance sheet data (from Datastream) and pension scheme data (collected by hand from individual FRS 17 disclosures) suggests that stock return volatility is systematically related to proxies for the two channels of amplification discussed above. These effects are statistically significant and robust to the inclusion of control variables and the exclusion of outliers.

UK monetary regimes and macroeconomic stylised facts

Working Paper no. 290

Luca Benati

The UK historical experience, with the remarkable variety of its monetary arrangements over the course of the past few centuries — from the *de facto* silver standard prevailing until 1717, up to the post-October 1992 inflation-targeting regime — and the high quality of its historical data, provides a unique 'macroeconomic laboratory' for the applied monetary economist. This paper exploits the marked changes in UK monetary arrangements since the metallic standards era to investigate continuity and changes across monetary regimes in key macroeconomic stylised facts in the United Kingdom. Our main findings may be summarised as follows.

First, the post-1992 inflation-targeting regime appears to have been characterised, to date, by the most stable macroeconomic environment in recorded UK history. Since 1992, the volatilities of the business-cycle components of real GDP, national accounts aggregates, and inflation measures have been, post-1992, systematically lower than for any of the pre-1992 monetary regimes or historical periods, often markedly so, as in the case of inflation and real GDP. The comparison with the period between the floating of the pound *vis-à-vis* the US dollar (June 1972) and the introduction of inflation targeting (October 1992) is especially striking, with the standard deviations of the business-cycle components of real GDP and inflation having fallen by about 50% and 70%, respectively.

Second, the so-called Phillips correlation between unemployment and inflation at business-cycle frequencies appears to have been weakest under the gold standard, and strongest between 1972 and 1992. Under inflation targeting the correlation has exhibited, so far, the greatest extent of stability in recorded history. In line with Ball, Mankiw and Romer, evidence points, overall, towards a positive correlation between average inflation and the strength of the Phillips correlation, both across monetary regimes and over time (especially over the post-WWII era).

Third, historically inflation persistence — broadly speaking, the tendency for inflation to be comparatively high (low) in one period, having been comparatively high (low) in previous periods — appears to have been the exception, rather than the rule. Inflation is only found to have been very highly persistent only during the period between the floating of the pound and the introduction of inflation targeting. Under inflation targeting, inflation exhibits little or no persistence based on all the price indices we consider. In line with a recent, and growing, literature, in particular the recent work of Cogley and Sargent, and in contrast with the 'traditional' position of, eg, Fuhrer and Moore or Blanchard and Gali, our results provide compelling evidence that high inflation persistence is not an intrinsic, structural feature of the economy. Instead, the extent of inflation persistence may crucially depend on the monetary regime in place over the sample period.

Fourth, we document a remarkable stability across regimes in the correlation between inflation and the rates of growth of both narrow and broad monetary aggregates at the very low frequencies. The exception is base money growth under the current inflation-targeting regime, for which the correlation clearly appears to have been, so far, negative. Our results, in particular, suggest that a key finding in Rolnick and Weber, a stronger correlation between inflation and the rates of growth of monetary aggregates under fiat standards than under commodity standards, may find its origin in their exclusive focus on the raw data (in other words, in their failure to distinguish between the different frequency components of the data).

Finally, we show how Keynes, in his dispute with Dunlop and Tarshis on real wage cyclicality, was entirely right: during the inter-war period, real wages were strikingly countercyclical. By contrast, under inflation targeting they have been, so far, strongly procyclical.

Affine term structure models for the foreign exchange risk premium

Working Paper no. 291

Luca Benati

The ability to produce reliable estimates of foreign exchange risk premia would be of potentially paramount importance for policymakers. For example, a given appreciation of the currency bears markedly different implications for monetary policy when it originates from a movement in the risk premium, as opposed to (say) a change in the equilibrium exchange rate. Four decades ago, Fama first called the attention of the economic profession to the so-called 'forward discount anomaly', a puzzling violation of the uncovered interest parity (UIP) hypothesis according to which future foreign exchange rate depreciation should exactly reflect the current spread between foreign and domestic interest rates. Given that the presence of a time-varying foreign exchange risk premium represents a possible explanation for the failure of UIP to hold, in the intervening years economists have been trying to estimate risk premia within several different econometric frameworks. A first strand of literature has tried to estimate models based on strong theoretical restrictions, encountering, as of today, near-universal lack of success. Typical problems found within this approach include implausible estimates of the degree of risk aversion and, almost always, the empirical rejection of key theoretical implications of the underlying model.

A second group of studies has reacted to the rejection of models based on strong theoretical restrictions by pursuing a radically alternative strategy, namely by adopting a pure time-series approach that imposes a minimal theoretical structure on the data. While studies in this vein are capable of identifying a predictable component in the foreign exchange excess return, they typically suffer from the drawback that, by not imposing enough structure on the data, they cannot guarantee that such an estimated predictable component truly is a risk premium.

In this paper we adopt an intermediate approach, based on semi-structural models imposing minimal restrictions on the two countries' so-called pricing kernels — the processes on which all of the assets within the two countries, and the nominal exchange rate between them, can be priced. Such models should be considered as a 'bridge' between the two previously discussed groups of studies, imposing on a time-series structure a set of restrictions just sufficient to identify a foreign exchange risk premium with a reasonable degree of confidence, but otherwise leaving the model largely unconstrained. Although, on strictly logical grounds, it is clearly suboptimal — ideally, we would like to be able to impose a solid theoretical structure capable of generating a time-varying risk premium — at the moment such an approach is probably the most promising.

We extract historical estimates of foreign exchange risk premia for the pound with respect to the US dollar based on two affine (ie linear) term structure models. The term structures of interest rates for the two countries are estimated jointly, together with the dynamics of the nominal exchange rates between them, via maximum likelihood. The likelihood function is computed via the Kalman filter, and is maximised with respect to unknown parameters. Particular attention is paid to the robustness of the results across models; to the overall (filter plus parameter) econometric uncertainty associated with risk premia estimates; and to the ability of estimated structures to replicate Fama's 'forward discount anomaly', the key conditional stylised fact pertaining to the foreign exchange market.

The paper's main results may be summarised as follows. First, the risk premia estimates generated by the two models, although exhibiting a qualitatively similar time profile, are numerically quite different, to the point of casting doubts about the possibility of using them within a policy context. Second, both models fail to replicate the forward discount anomaly. Third — and not surprisingly, given the well-known difficulty of forecasting exchange rates — the estimated models exhibit virtually no forecasting power for foreign exchange rate depreciation.

Switching costs in the market for personal current accounts: some evidence for the United Kingdom

Working Paper no. 292

Céline Gondat-Larralde and Erlend Nier

Bank current accounts play a pivotal role in the relationship between a bank and its customers and may serve as a gateway through which banks can cross-sell other products. This paper analyses the competition in the market for personal current accounts in the United Kingdom. Using the Financial Research Survey (FRS) data collected by National Opinion Poll (NOP), we first describe some stylised facts on market shares and prices associated with the current account, such as the interest rate offered on positive balances and the rate charged on overdraft. While the level of concentration has remained high in this market, the market appears to have become gradually more competitive, with building societies and direct banks making some significant inroads during the 1996–2001 period. Against this, we find a marked dispersion in price, which appears to persist through time.

To assess the level of competition in the current account market more formally, we derive the elasticity — that is the sensitivity — of bank market shares with respect to the set of prices that relate to the current account product. This analysis controls for differences in current account characteristics (such as the extent of the branch network) in order to isolate the effect of price differentials on changes in market share. We find a moderate sensitivity of changes in market share to differences in the current account rate across banks. The elasticity of market share with respect to the overdraft rate is considerably lower. Overall our findings are consistent with a moderate degree of imperfect competition in the market for personal current accounts.

We proceed to investigate further the type of friction in this market that best characterises the data. We find a positive relationship between levels of market share and price — again controlling for non-price characteristics. This finding points to the importance of the cost of changing banks and is consistent with dynamic models of competition with switching costs developed recently. The basic intuition is that each bank faces a trade-off: raising the price increases the profit the bank achieves on its existing customer base, but also implies that the bank might lose some of its present customers and is less likely to attract new customers. The bank's current market share determines how this trade-off is resolved. A bank's incentive to raise its price is more pronounced, the larger is the bank's current market share. The model also predicts that the relationship between market share and price should be stronger, the lower the elasticity of demand with respect to price. Consistent with this prediction, we find that the relationship between market share and price is strongest for the overdraft rate, for which the elasticity of demand is lowest.

Since the end of our sample period, there have been several initiatives to facilitate switching. In response to the Cruickshank report in 2000, the Government asked a group led by DeAnne Julius to review the Banking Code. One set of recommendations in the report that has since been implemented specifically focuses on ways to facilitate switching accounts. Moreover, the banks have implemented improvements to the logistics of the switching process — eg as regards the exchange of information between the switchers' old and new banks — to improve the speed and the accuracy of the account transfer. In addition to initiatives to reduce the cost of switching, steps have also been taken to increase consumer awareness of the potential benefits of changing banks. Even though it may be too early to assess the impact of these initiatives empirically, the results of this study appear broadly supportive of such initiatives, in that they document empirically the presence of switching costs in the UK market for personal current accounts.

Resolving banking crises — an analysis of policy options

Working Paper no. 293

Misa Tanaka and Glenn Hoggarth

This paper develops a simple but general framework which can be used to analyse alternative policies to restructure failed banks when the authorities cannot observe banks' balance sheets. We demonstrate that without regulatory intervention, weak banks have the incentives to hold on to the non-performing loans (NPLs) and gamble for the small chance of recovering these loans ('gamble for resurrection'). But if the authorities cannot force weak banks to liquidate their NPLs because they cannot observe their balance sheets, they may have to rely on financial incentives to induce banks to liquidate their bad assets. Our paper considers the optimal design of such financial incentives, taking into account their impact both on managerial moral hazard and fiscal cost of resolution.

We first examine actual policies used in recent banking crises to clarify why certain choices have been made. Subsequently, we use a model to consider five different policy options for resolving banking failures when the authorities cannot observe the level of non-performing loans held by each bank. When faced with this asymmetric information, the first-best outcome is achievable when the authorities can close all banks that fail to raise a minimum level of new capital. But when the authorities cannot close banks and must rely instead on financial incentives to induce banks to liquidate their NPLs, equity (Tier 1 capital) injection would be the second-best policy, whereas subordinated debt (Tier 2 capital) injection is suboptimal. If the authorities do not wish to hold an equity stake in a bank, they should subsidise the liquidation of non-performing loans rather than inject subordinated debt. We also show that the cost of this subsidy can be reduced if it is offered in a menu that includes equity injection. Thus, our analysis clarifies the conditions under which each policy should be used, and provides a practical guidance to policymakers in resolving bank failures when they cannot immediately assess the problems at each bank.

How does the down-payment constraint affect the UK housing market?

Working Paper no. 294

Andrew Benito

Buying a home usually requires a significant amount of cash. Lenders typically require that a home-buyer has some equity in the home. There are good reasons for why this should be the case. This paper considers the implications of this borrowing constraint for the UK housing market.

For the aggregate housing market, the paper shows that several features can be explained by the model which attaches an important role to the down-payment constraint: first, a positive correlation between the rate of change of house prices and transactions; second, the greater volatility in the rate of change of house prices among former owner-occupiers' properties than for first-time buyers; third, the presence of more former owner-occupiers relative to first-time buyers in the market when the rate of change of house prices is high; and fourth, house prices are more sensitive to the incomes of the young than to aggregate income.

An important feature of the model highlighted in this paper is that it is based on the economic fundamentals of the housing market. This contrasts with some discussions of the housing market which draw on the idea of housing market 'bubbles' to attempt to rationalise outcomes, in particular significant swings in activity and prices. Any model based on bubbles is difficult to test. Moreover, used in this paper also suggests that there can be episodes of price 'overshooting' in the housing market, as prices increase beyond their new equilibrium in response to an increase in income and then decline. Traditional models find this difficult to explain. This may be why, by default, some commentators have attempted to explain house price fluctuations by appealing to notions of bubbles instead.

Much commentary on the housing market appeals to ratios such as the ratio of house prices to incomes or earnings as being a key attractor to which house prices should return in the long run. Yet basic economic theory suggests that prices are not determined by averages, but instead, are set at the margin. If the marginal buyer is a young first-time buyer then this suggests that the prices should be more sensitive to the incomes of the young than to average income. This paper demonstrates that in the early 1990s, when house prices declined significantly, there was a notable decline in incomes among young, potential first-time buyers relative to the wider population, suggesting a greater sensitivity to their income than to the wider population. More generally, higher volatility in the incomes of the young than for the population as a whole suggests that house prices will be more volatile than if they were related to average incomes.

The paper also explores variation across districts. Despite some remarkable movements witnessed in house prices in recent years, there is much more variation across districts than over time in the rate of change of house prices. Examining these differences across districts can also shed light on the behaviour of the housing market. Market professionals themselves argue that different districts should be thought of as quite distinct housing markets: so using aggregate data to examine changes in house prices could be misleading. But there are few, if any, studies of local housing markets in the United Kingdom that can be said to cover a large part of the country.

By focusing on variation in house price inflation across districts, the paper examines another key implication of these down-payment models, namely the role for leverage (loan to value ratios) in influencing the response of local house prices to incomes. The paper finds that a large incidence of households with relatively high loan to value ratios in an area increases the response of prices in that area to local incomes and financial shocks. This justifies many commentators' focus on loan to value ratios in their discussion of the housing market. In recent years loan to value ratios have been declining in the United Kingdom among first-time buyers, suggesting a lower sensitivity of house prices to shocks in future.

Productivity growth, adjustment costs and variable factor utilisation: the UK case

Working Paper no. 295

Charlotta Groth, Soledad Nuñez and Sylaja Srinivasan

The aim of monetary policy is to keep inflation low and stable. A key influence on inflationary pressure is the balance between the demand for and the economy's capacity to supply goods and services. This capacity depends both on the quantities and qualities of the inputs into the production process (capital and labour), and on the efficiency with which they are combined. The latter concept is often referred to as total factor productivity (TFP). A good understanding of past and current TFP growth is thus important for understanding aggregate supply capacity, and so is relevant for the conduct of monetary policy.

During the 1990s, productivity growth did not increase in the United Kingdom while it rose sharply in the United States. This diverging performance looks puzzling, especially when considering that, following the 1990–92 recession, the macroeconomic environment in the two countries was similar. This research tries to estimate underlying productivity growth by accounting for a number of factors that may bias the standard estimate of productivity growth, and thereby give us a distorted picture of underlying technological progress. By doing so, it tries to assess and account for the lack of a pickup in UK productivity growth during the 1990s.

The starting point of the analysis is a standard measure of aggregate TFP growth, or the so-called Solow residual. This is calculated as that part of aggregate output growth that cannot be accounted for by the primary factors of production, under the assumptions of perfect competition, constant returns to scale, no costs to adjusting the factors of production and therefore full utilisation of available factors.

When any of these assumptions is violated, the Solow residual may not correctly measure underlying technological progress. For example, increasing returns to scale in the production of output may cause this measure of TFP growth to rise whenever input growth rises. And if firms face adjustment costs when hiring and firing workers or changing the level of capital, they could respond to short-run fluctuations in demand by adjusting the intensity with which they use labour and capital. This would cause larger fluctuations in output than in capital and labour, and hence procyclical movements in measured TFP growth. In addition, if firms face costs to adjusting capital and labour, marketable output (which matters for the Solow residual) may be low during periods of rapid investment or hiring growth. This is because firms may spend resources internally to install capital or labour, rather than producing marketable output. In this paper, we try to control for these types of non-technological factors, to see whether this affects our conclusions about the United Kingdom's productivity performance during the 1990s.

It is not possible to observe how hard companies are working capital and labour — or their utilisation levels — directly. But by assuming that firms maximise profits, we can derive links between variables such as hours worked and the amount of intermediate inputs used, and changes in the rate of utilisation of capital and labour. The paper also tries to account for the amount of resources that is used by firms to install new capital and hire new labour, instead of producing marketable output.

The results suggest that the aggregate Solow residual underestimates underlying UK total factor productivity growth through the 1990s, since it does not account for falling utilisation rates and high capital adjustment costs. We find, however, that these non-technological factors had a similar impact on the Solow residual during the first and the second half of the 1990s. The broad movement in the aggregate Solow residual through the 1990s is therefore similar to that of our estimate of underlying productivity growth. Thus the puzzle of the apparent lack of a pickup in UK productivity growth during the 1990s remains.

In a comparison with the United States, the paper notes that the US experience of a rise in TFP growth between the first and the second half of the 1990s was, to a large extent, driven by strong growth in ICT-producing industries, the distribution sector and financial services. A broadly similar pattern is found for the United Kingdom. One difference, however, is that whereas the US durables manufacturing sector as a whole contributed to rising rates of TFP growth, UK estimates suggest that most durables industries did not see an increase in TFP growth over the same periods. So the results suggest that the rise in TFP growth appears to have been more broadly based in the United States than in the United Kingdom, and this may partly explain the difference in the aggregate data.

Sterling implications of a US current account reversal

Working Paper no. 296

Morten Spange and Pawel Zabczyk

The US current account deficit reached a new high of 6.3% of GDP in 2004 Q4. The deficit is large in comparison with the current account balances of other countries and this has led a number of commentators to question its sustainability. This paper explores the potential implications for sterling of a restoration of the US current account deficit to balance. The analysis is based on a model calibrated to represent the United Kingdom, the United States and a third region covering the rest of the world. Different triggers that might bring about a realignment of the US current account deficit are considered. We begin by analysing the implications of a negative shock to US consumers' demand. In addition, we study a scenario in which such a demand shock is supplemented by a positive productivity shock in the US tradable sector — helping the United States bridge its trade deficit and so improving the current account. Finally, we also assess the impact of revaluation effects on international investment positions and how this affects the results.

Our analysis suggests that the magnitude of sterling adjustment depends heavily on (a) the cause of the US current account adjustment, ie the type of shock that brings it about; (b) the assumptions made about the associated adjustments of the United Kingdom and rest of the world current account deficits, ie how the adjustment to the US unwinding is split geographically; and (c) assumptions about key judgements such as the degree of substitutability between different types of goods (tradable and non-tradable) and goods produced in different regions.

Assuming that the UK current account deficit deteriorates in proportion to sterling's share in the dollar effective exchange rate index (ERI), we can derive estimates for movements in the sterling real effective ERI ranging from a depreciation of 1.4% to an appreciation of 4.2%, depending on different judgements about substitutability and the cause of the adjustment. If we assume that the dollar pegs maintained by a number of Asian economies result in a larger proportion of the adjustment falling on the United Kingdom, then the model generates estimates ranging from a depreciation of the sterling real ERI of 0.7% to an appreciation of 4.9%. However, in the event that all current accounts were to move to balance (implying a UK current account improvement) the model predicts a real ERI sterling depreciation in the range of 0.6% to 7.8%. It is important to note that the exchange rate movements presented in this paper are a symptom of rebalancing global demand, and they are not associated with unemployment or recessions.

Optimal monetary policy in a regime-switching economy: the response to abrupt shifts in exchange rate dynamics

Working Paper no. 297

Fabrizio Zampolli

A common concern among central bankers is that the true or perceived existence of financial imbalances or asset price misalignments could at some point in time lead to sudden and large adjustments in asset prices, with potentially adverse consequences for inflation and output. For instance, one of the major risks that has worried some members of the Bank of England's Monetary Policy Committee (MPC) in the past has been the possibility that sterling could suddenly fall by a material amount. Other risks routinely debated by actual policymakers, including oil price hikes or abrupt changes in key econometric relationships, may also be asymmetric — that is, a given change may be more likely to occur in one direction than in the opposite. Nevertheless, modelling of asymmetric risks is not very common in the monetary policy literature, possibly because of the lack of readily-applicable technical tools.

In this paper we examine the trade-offs that the policymaker faces when the exchange rate can experience sustained deviations from its fundamental value (ie the value implied by interest rates absent any economic shock) and occasionally collapse. To do so we use a simple method which has rarely been applied in the economics literature. The method allows us to solve for the optimal monetary policy in an economy subject to regime shifts, while retaining the flexibility and simplicity of more commonly applied methods. The method could be applied in other ways that are not considered in this paper and can be considered as a general tool for studying uncertainty in monetary policy. In particular, it provides an example of how policymakers can incorporate judgemental information about a potential misalignment (and the uncertainties associated with it) into their macroeconomic model, and work out the best policy response based on that judgement.

Our analysis is based on a small open economy model, comprising a demand equation, a Phillips curve which determines prices, and an equation linking the real exchange rate to the domestic real interest rate. We modify this model to incorporate regime switching in the exchange rate. In one regime, which we call the bubble regime, any shock can lead the exchange rate to increasingly deviate from its fundamental value. Depending on the sign of the shock, the exchange rate can continue to rise above its fundamental value, or it can continue to fall below it. In the other regime, which we call the no-bubble regime, the exchange rate displays transitory fluctuations around its fundamental value. The times at which the bubble begins and ends are uncertain to the policymaker. Moreover, the size of the correction in the exchange rate, which occurs when the economy switches from the bubble to the no-bubble regime, will vary over time as it depends on the past behaviour of the exchange rate as well as the interest rate.

Analysis of the optimal regime-switching policy rule shows the existence of an intuitive link in the bubble regime between the optimal response of the interest rate to the exchange rate and the expected duration of a bubble. When the bubble is expected to last for at least two years, the optimal interest rate is negatively related to movements in the real exchange rate and becomes more responsive as the expected duration of the bubble lengthens (an increase in the exchange rate being an appreciation). Similarly, in the no-bubble regime there is an intuitive link between the response to the exchange rate and the probability of the bubble emerging: for lower probabilities of bubbles the interest rate is positively correlated with exchange rate fluctuations (reflecting the likely transitory nature of exchange rate movements) but becomes less responsive as the probability of a bubble increases. For high probabilities of the bubble the interest rate responds negatively and becomes more reactive to exchange rate fluctuations as the probability rises further (reflecting the likely onset of a bubble). Another characteristic of the optimal regime-switching interest rate rule is that in both regimes the interest rate is for the most part less responsive to inflation and output fluctuations than in the absence of regime uncertainty, with the degree of caution increasing as both transition probabilities approach a half.

A key result of the paper concerns the assumptions that the policymaker makes about the (unknown) probabilities of moving between bubble and no-bubble regimes. These probabilities could be highly uncertain since historical experience might provide little or no help in quantifying them. We find that there are 'robust' values of the probabilities corresponding to more muted policy responses, where by 'robust' we mean values of the probabilities which can be assumed by the policymaker without fear of causing unnecessary volatility in output and inflation, were they to prove wrong in hindsight. This result is interesting as in the robust control literature uncertainty is often found to lead to more reactive policy responses than in the absence of uncertainty.

Optimal monetary policy in Markov-switching models with rational expectations agents

Working Paper no. 298

Andrew P Blake and Fabrizio Zampolli

Uncertainty is one of the major problems faced by policymakers. Economic models are simple representations of how the economy works, and might turn out to be wrong. For example, the way the economy works might change over time in an unanticipated manner which would not be captured by normal economic models. This paper focuses particularly on this type of uncertainty. As interest rates normally affect output and inflation with a lag, rates must therefore be set while bearing in mind how the economy might change by the time that the interest rates exert influence on inflation and aggregate output. Unfortunately, the normal way of modelling the economy is to assume that it does not change over time and that the only uncertainty faced by the policymaker is about the type and duration of the shocks that hit the economy — for example, changes in foreign demand. To put it differently, the normal way of modelling the economy is to assume that the policymaker knows how economic shocks affect inflation and output (ie the transmission mechanism), and also to assume that this mechanism will not change. In this paper, instead, we consider an economy in which the transmission mechanism can change over time in an uncertain manner. For example, aggregate demand may become more sensitive to changes in interest rates, or the degree to which the exchange rate affects consumer prices can become larger. This implies that the shocks hitting the economy might not have always the same impact on the variables targeted by policymakers. By ignoring these potential changes, policymakers might be in danger of missing the inflation target more often than otherwise, or to cause inflation and output to be more volatile than is really necessary.

The main contribution of this paper is to develop simple methods for working out the best interest rate response to shocks in such an evolving economy. More specifically, the economy is modelled as a so-called Markov-switching framework. That is, the economy is assumed to alternate over time between a number of regimes (eg high and low exchange rate pass-through regimes) according to some given probabilities. It is also assumed that in this economy the private sector forms so-called rational expectations. That is, in forming their views about the future they understand what the transmission mechanism is in the different regimes and they also understand how policymakers set the interest rate in response to shocks. The paper also shows how the methods for calculating the best interest response can be applied to the case in which policymakers and the private sector differ in their views as to the probability of the regime change. Another important feature we consider in this paper is the possibility of assuming that uncertainty is asymmetric — that is, a given change is more likely to occur in one direction than in the opposite (eg an increase in the sensitivity of aggregate demand to interest rates is more likely than a fall of the same size).

We apply our procedure to a small open economy model in which some of its key features can suddenly change. In this application we are considering so-called time-consistent policies, ie policies which continue to be the best possible as time passes. With such policies the monetary authority is unable to affect the private sector's expectations. In our results, which should be thought of as first steps, we find that for the most part interest rates are set more cautiously when uncertainty about changes in the economy is symmetric. That is, in response to shocks the interest rate is varied by less than when such uncertainty is absent or ignored. Being less cautious would make the economy more volatile without the benefit of an improved trade-off between output and inflation, which would result from the ability of policymakers to affect the private sector's expectations. We also find that the optimal policy can be significantly affected by differences between the policymaker and the private sector in their views about the probabilities of parameter changes. When changes in the economy are asymmetric, the findings about the optimal policy response cannot be easily generalised.

Optimal discretionary policy in rational expectations models with regime switching

Working Paper no. 299

Richhild Moessner

Structural change is an important feature of economies. One aspect of such change is that features of the macroeconomy may vary over time — for example, intrinsic inflation and output persistence, the interest elasticity of demand, or the persistence of shocks. Moreover, uncertainty is an important issue facing policymakers, including uncertainty about structural change, about the best model of the economy, as well as about shocks hitting it. It is therefore interesting to study the implications for policymakers of structural changes that are not known with certainty. This paper considers policy design in the presence of structural change which is not known with certainty, and which may take the form of time variation in the parameters of an economic model. We handle this time variation by assuming there are Markov processes underlying the parameters, so that they can take on several different values and switch between them according to given probabilities. Moreover, structural change may take many different forms, and in particular it may be abrupt, transitory and asymmetric in nature; modelling structural change as Markov processes also enables us to capture these features. By contrast, other work on optimal monetary policy with parameter uncertainty, which assume that policymakers have symmetric uncertainty about parameters, do not capture all of these features.

Optimal policy with Markov switching in model parameters has previously been considered for backward-looking models. This paper extends the analysis to forward-looking models of the economy for the case of discretionary policy, when both the central bank and the private sector face uncertainty about model parameters. Deriving the solution for the case of forward-looking models with rational expectations is useful, since in contrast to purely backward-looking models, such models include forward-looking private sector expectations. This makes the treatment of private sector expectations consistent with the forward-looking behaviour of the policymaker. The macroeconomic models currently used for economic policy analysis mainly incorporate rational expectations, to ensure consistency, and to be able to base them — at least in part — on optimising microeconomic behaviour. In related work at the Bank, Fabrizio Zampolli derives optimal policy for the case of Markov switching of model parameters in backward-looking models, while Andrew Blake and Fabrizio Zampolli consider time-consistent optimal policy in forward-looking models within a semi-structural model representation. In related academic work, Lars Svensson and Noah Williams derive optimal policy with Markov switching in forward-looking models under both commitment and discretion.

As an illustration, we apply our method to study optimal monetary policy in the presence of structural changes in output persistence, within a forward-looking model estimated for the euro area. The main reason for adding this output persistence to the basic forward-looking model is to improve the fit with the data. Output persistence may change, for example, because of changes in the degree to which firms' investment decisions are constrained by cash flow, rather than being purely forward-looking. We assume there is a Markov process driving these changes. We find that the coefficients of the optimal policy rule depend on the state of the economy characterised by different values of output persistence, and the coefficients depend on the transition probabilities of the Markov process governing the structural change. For uncertainty about output persistence, the optimal policy rule is non-linearly related to the transition probabilities. We find that if the probability of moving from a state with low output persistence to a state with high output persistence is high, it is optimal for monetary policy in the former state to respond more aggressively to the lagged output gap, lagged inflation and the two shocks (to output and inflation) we consider, than in the absence of uncertainty about changes in output persistence.