The real exchange rate and quality improvements

Summary of Working Paper no. 320 Karen Dury and Özlem Oomen

Much of the literature on the New Open Economy Macroeconomics (NOEM) focuses on technological progress that manifests itself through improvements in productivity, ie increased efficiency in the production of a given range of goods, which is also known as *process innovation*. A common finding in this literature is that a positive productivity shock in the home country tends to depreciate the real exchange rate. This is because a positive domestic productivity shock increases the supply of home relative to foreign goods, which then reduces the relative price of home goods causing a real depreciation. But technological progress can also come about via improvements in the quality of a given range of products, known as *product innovation*. We might envisage that this type of technological progress in the home country will cause the real exchange rate to appreciate. This is because higher quality goods, in general, can command higher prices, which then tend to increase the relative price of home goods leading to a real appreciation. The simple model presented in this paper aims to take the NOEM literature a step towards modelling this type of technological progress. In particular, we are interested in understanding the theoretical link between quality improvements and real exchange rates.

The real exchange rate is defined as the ratio of the two countries' price indices, expressed in a common currency. But in a world where goods become obsolete and are replaced due to quality improvements, the relevant real exchange rate is the real exchange rate measured in terms of quality-adjusted prices. In practice, price indices may not (fully) capture the quality improvements in goods, and the real exchange rate may consequently be miscalculated. Here, we examine the impact of quality improvements on two measures of the real exchange rate: the *quality-adjusted* and the *quality-unadjusted* measure. The former measure is calculated using price indices that aggregate prices per quality unit, and hence, by construction, it fully accounts for product quality. The latter measure, on the contrary, is calculated using price indices that aggregate unit prices only, and hence by construction, it fails to account for product quality.

Our analysis shows that a quality improvement can lead to either a depreciation or an appreciation of either measure of the real exchange rate depending on how costs of production are affected by the quality improvement. We also find that the real exchange rate defined in terms of unit prices does not always move in the same direction as the real exchange rate defined in terms of prices per quality unit, illustrating the importance of measuring quality correctly.

Comparing the pre-settlement risk implications of alternative clearing arrangements

Summary of Working Paper no. 321 John P Jackson and Mark J Manning

This paper analyses the risk implications of different arrangements for clearing securities and derivatives markets. In this context clearing refers to the set of procedures in place for calculating the net exposures arising from a set of financial market trades and managing the credit risks arising from these trades in the period prior to their final settlement.

This is a topic of considerable policy interest. For instance, there is a live debate underway in policy and industry circles regarding the potential risk-reduction benefits of centralised clearing arrangements for a broader range of over-the-counter (OTC) derivative products. Another topical issue, particularly in an EU context, is whether significant efficiency gains could be realised by merging several domestic central counterparty clearing houses (CCPs) into a single cross-market entity.

This paper provides an analytical framework for evaluating quantitatively the relative cost and risk implications of a range of clearing methods, covering different constellations of products, trader profiles and market structures. This is done by simulating agents' pre-settlement costs and risks under a range of bilateral and multilateral clearing arrangements. Two metrics for pre-settlement risk are analysed: the magnitude of replacement cost losses; and the distribution of such losses.

Replacement cost risk arises during the period between trade and settlement and reflects the cost to a trader of replacing a trade on which a counterparty has defaulted. Agents can mitigate replacement cost risk by collecting collateral (known as margin) from their trading counterparties during the pre-settlement period; hence a trader (or CCP) will only incur a replacement cost loss if there is a coincidence of events: an adverse change in the underlying contract price in excess of the per-unit value of margin collected from a counterparty, combined with a default by that counterparty. However, the requirement to post margin may impose a significant cost on agents, which in our analysis is quantified and compared across arrangements.

We analyse three distinct clearing and settlement arrangements for futures markets: (i) bilateral clearing; (ii) ring clearing; and (iii) CCP clearing. These may be defined as follows:

– In bilateral clearing, trading agents post margin on the basis of their net bilateral obligations. This remains the typical clearing arrangement for off-exchange and OTC trading, particularly in less standardised products.

– The second approach, ring clearing, is a way of achieving multilateral netting of exposures without requiring a CCP to become the legal counterparty to all trades. Rather, the original bilateral exposures are extinguished and multilateral net exposures reallocated, according to some pre-determined algorithm, among members of the ring. A ringing arrangement reduces collateral costs, but agents retain some counterparty credit exposure to one another. There are, to our knowledge, no formal ringing arrangements in operation at present, although services for multilateral contract terminations can achieve something similar.

– The final approach analysed, CCP clearing, takes ringing a step further by introducing novation of all trades to a central counterparty; novation refers to the process by which the CCP interposes itself as legal counterparty to both the buy and sell-side of all the trades it clears. In the absence of counterparty default, the CCP has a balanced book and does not, therefore, face any market risk. At the same time, agents are no longer exposed to their original counterparties, instead having a single net exposure in each asset with the CCP. By providing centralised risk management and facilitating anonymous trade, CCP clearing is particularly beneficial in the case of exchange-traded assets, particularly those with long settlement periods, such as derivatives.

We identify two basic sources of replacement cost risk differentials across the arrangements under consideration: netting ratios and margin pooling. We show that replacement cost losses and the opportunity costs from posting collateral under CCP or ring clearing decline as the number of bilateral trading counterparties increases.

In the context of multi-asset clearing, we find that 'margin pooling' is an important effect. This is the benefit derived when an agent's margin payments in respect of multiple positions can be pooled, such that, in the event that the agent defaults, the margin-taker can draw upon any residual margin in the pool (either from profitable, or only modestly loss-making, positions) to cover a margin shortfall arising on any individual position(s). Our simulations show how this effect can vary according to the degree of price and position correlation across assets. Our results imply that a merger of CCPs has the potential to significantly reduce the risks and costs faced by traders.

Finally we allow trader credit quality to vary in order to analyse agents' individual incentives to adopt particular clearing arrangements. We show that restricted access or tiered clearing arrangements, where risky traders are not able to become a member of the CCP but must clear their trades through a more creditworthy agent who is a member, may then emerge naturally.

An affine macro-factor model of the UK yield curve

Summary of Working Paper no. 322 Peter Lildholdt, Nikolaos Panigirtzoglou and Chris Peacock

Understanding which factors drive movements in the term structure of interest rates is of potential interest to policymakers for a number of reasons. For example, the extent to which changes in the short-term policy rate feed through to longer-term yields is important since it represents a key part of the transmission mechanism of monetary policy by affecting the spending, saving and investment behaviour of individuals and firms in the economy. Moreover, the yield curve has been found to be a good predictor of future real activity and inflation. The term structure also contains information about market participants' expectations of the future path of interest rates. But there is strong empirical evidence to suggest that time-varying risk premia drive forward rates away from these expectations. The decomposition of forward rates into expectations of future interest rates and risk premia is one of the key contributions of this paper.

In this paper, we estimate various models of the term structure of interest rates for the United Kingdom, where the underlying factors that drive movements in the term structure have a macroeconomic interpretation. The first factor is an unobserved inflation target, the second factor is annual inflation, and the third factor reflects, among other things, the output gap and monetary policy shocks. We find that the long end of the yield curve is primarily driven by changes in the unobserved inflation target. At shorter maturities, yield curve movements reflect mainly the other two factors.

Our preferred model implies that agents require compensation (ie a risk premium) for risks associated with output gap and inflation shocks but do not require compensation for shocks to the inflation target. This result seems consistent with simple asset pricing models with an assumed representative (homogenous) agent. Our yield curve models can be used to back out a path for an unobserved time-varying inflation target. This path is shown to be closely linked to other measures of long-run inflation expectations, such as those from market-based ten year ahead breakeven inflation rates and long-run Consensus forecasts of inflation.

Time series of risk premia on long forward rates from the preferred yield curve model have declined since the 1970s, which is consistent with perceptions of declining macroeconomic uncertainty or perhaps more efficient macroeconomic stabilisation policies. Model-derived risk premia at short maturities are shown to be highly correlated with survey-based risk premia, which indicates that the model could be useful for the purpose of extracting market-based interest rate expectations. This is comforting because we have not used survey data for estimation or even model selection. As such, it provides support for the estimated models.

Forecast combination and the Bank of England's suite of statistical forecasting models

Summary of Working Paper no. 323 George Kapetanios, Vincent Labhard and Simon Price

Monetary policy at the Bank of England and at many other central banks is forward looking. So it is essential to be able to forecast accurately the future evolution of the economy. Consequently, the Bank of England maintains a large number of models, ranging from the purely statistical to data-free theoretical models, which we call upon to answer not only forecast but also other questions. As part of this general philosophy, the Bank has developed a range of purely statistical forecasting models (referred to hereafter as the 'Suite') which can be used to construct judgement-free statistical forecasts of inflation and output growth and which form one of many inputs into the Monetary Policy Committee's (MPC's) forecast process. This process culminates in the forecast fan charts reported in the *Inflation Report* which show a range of possible outcomes. These encapsulate the MPC's collective judgement of the prospects for inflation and growth, and are conditioned on specific assumptions, including interest and exchange rates and some exogenous variables, as well as on general views about the future.

We describe the Suite as it stood when it was first created in May 2005. Naturally, this is merely a snapshot, as the Suite continues to evolve; models or model combinations may be added or dropped, and the data continually change. On the evidence of the data and models that we examine in this paper, combinations of statistical forecasts generate good forecasts of the key macroeconomic variables, which can serve as judgement-free benchmarks to compare with the policymaker's projections. Moreover, changes in forecasts as new data arrive provide a summary measure of the relevant news in the data, giving a natural indicator of changing inflationary pressure over the horizons of policy interest.

We use two broad types of models. The first uses only univariate models (using only the variable to be forecast), which capture information solely in the forecast variable's history. Within this broad class we include linear and non-linear models of various types, including ones which may be more robust to some types of structural change. The second comprises multivariate models (including more than one variable), which capture a wider range of information. The data sets here vary in size, the largest using over 60 variables. Here too we include models which may be robust to structural change.

One important issue is the 'attractor', the value to which the forecast tends in the long run. If models fit the data well they will tend to produce a long-run forecast close to the average of the past. In the case of inflation, the monetary regime has changed over the sample period: the recent average inflation rate is substantially lower than over the whole sample period, reflecting the success in meeting the inflation targets in place since 1992. We test for structural breaks in the mean, and then forecast the inflation rate less this mean.

Individual forecasts are then combined to produce a single forecast. Forecast combination has a good track-record of improving forecasts. The combinations we use are a simple average of all the forecasts in the Suite, where all individual forecasts have an equal weight, which has been shown to work well in practice; and our preferred method based on goodness-of-fit, which we have shown may have a superior forecast ability.

This exercise is essentially practical, and success is measured by improved forecasts. Data typically has some obvious short-run cyclical variation that has to be accounted for, but it is often possible to capture this with a simple autoregressive (AR) process (where the model is a combination of past values of the variable being forecast). So we assess the forecasts since Bank independence in 1997 Q2 to 2005 Q1 relative to a benchmark AR forecast. Over our sample the AR forecasts are hard to beat, especially for inflation, with most of the models doing worse for most periods, although two non-linear models do better at most horizons. However, the benchmark combinations can beat the AR at many horizons for both growth and inflation. Thus the Suite appears to be fit for its intended purpose, as a statistical benchmark forming one of many inputs into the MPC's forecast process.

Housing equity as a buffer: evidence from UK households

Summary of Working Paper no. 324 Andrew Benito

The links between consumer spending and the housing market have been the focus of much debate. On several occasions in the past, swings in consumption and house prices have coincided. Precisely how to interpret that is by no means clear. One view has it that house prices are an asset price for an essential commodity, shelter, and that they largely reflect macroeconomic conditions with no special role of their own. Another view is that there is an important causal effect of housing in providing collateral for households' borrowing and spending decisions. Previous work by the Bank of England in this area has emphasised both views. Much of the comovement of house prices and consumption is driven by common movements in other variables. But there is likely to have also been a causal effect on consumption from house prices that results from the collateral channel. That reflects the fact that after a rise in house prices, homeowners enjoy capital gains that improve the terms on which they can obtain credit. For those wishing to borrow, that may have implications for their spending. Both of these views downplay the notion that house prices have an aggregate 'wealth effect' on consumption.

In this paper, one stage in that collateral channel is examined in finer detail: the withdrawal of home equity by households that enjoy gains in home equity (often referred to as 'mortgage equity withdrawal'). To look at this, the paper uses data on UK households over the period 1992 to 2003. Most previous analysis of mortgage equity withdrawal has used aggregate data. But aggregate mortgage equity withdrawal data conflate together rather different decisions by distinct groups of households. For example, the decisions by homeowners to actively borrow against the value of their homes are quite different from decisions by last-time sellers. Employing data for large numbers of homeowners and focusing solely on whether they actively borrow against their home equity or not avoids that drawback. In the period examined, the UK housing market went from a depressed state to one of its most remarkable booms. That varied experience, also reflected in the use of housing equity withdrawal, makes it a particularly useful period to examine.

Modelling the decision to withdraw versus retaining equity in the home helps us build up a picture of what influences this decision. This leads to a view that, among other things, housing equity plays the role of a financial buffer, being retained in normal times but drawn upon (and withdrawn) when a temporary, adverse shock has been experienced.

As well as the experience of some adverse financial shock, the findings also indicate that households are more likely to withdraw home equity if they are liquidity constrained, if they hold relatively large amounts of home equity and if they have higher incomes. There is also some evidence that they are more likely to withdraw equity if they face less house price uncertainty. Changes in marital status are particularly important. It is much more common for withdrawn equity to flow into housing investment than into consumption. That suggests the 'collateral channel' should be stronger for housing investment than consumer spending.

An emphasis on housing equity as a buffer is consistent with the view that many households seem to look at their housing equity as an asset that they would be prepared to draw on in an emergency to support their consumption plans. This is important since most homeowners have relatively little cash but relatively large amounts of home equity. The precautionary savings literature has, however, generally emphasised the use of liquid assets as a buffer.

Inter-industry contagion between UK life insurers and UK banks: an event study

Summary of Working Paper no. 325 Marco Stringa and Allan Monks

One of the Bank's core purposes is to detect and reduce threats to the financial system as a whole. The UK banking sector is a cornerstone of the UK financial system. Hence, contagion from one financial sector to the UK banking system may potentially have relevant implications for financial stability.

Over the past decade, correlations between equity price movements of UK banks and life insurers have increased markedly, most likely due to banks' increased involvement in the life insurance market. During the equity market decline between 2001 and 2003, UK life insurers were adversely affected. Consequently, the potential for contagion from the insurance sector to the UK banking sector became an important and much debated issue. This paper uses that period to assess the extent to which events in the life insurance sector have the potential to spillover to the banking system in times of stress.

Previous work at the Bank has identified potential channels by which shocks may be transmitted between sectors. Such interlinkages do not only originate from direct channels — ie counterparty exposures — but also from indirect channels via the impact of adverse and unexpected news on financial markets and consumers' confidence. Although accounting data provide a means to obtain a first estimate of counterparty exposures, they are less useful in measuring the magnitude of indirect channels. This paper aims to capture all three possible channels of contagion by using unexpected changes in equity price movements. In other words, we use equity prices as a tool to gauge the degree of inter-industry contagion from the UK life insurance sector to the UK banking sector.

The paper also uses information on equity trading volumes, in order to detect any significant reactions not captured by equity prices. For example, when investigating the presence of interlinkages, a mix of positive and negative reactions may lead to misleading conclusions since opposite interpretations of news can offset each other resulting in non-significant changes in equity prices. Therefore, we originally introduce the use of trading volumes to detect any significant reaction not captured by equity prices.

After undertaking a rigorous selection process to identify suitable events that originated in the life insurance sector between 2001 and 2003, we split them into two categories: events that impacted on specific life insurance companies and those that affected the life insurance sector as a whole. The results show that none of the firm-specific disturbances spilled over to the UK banking sector. There was, however, some evidence that elements of the banking system responded to events that affected the life insurance sector as a whole — but these reactions were not uniformly pervasive. On closer inspection of the banking sector, the results show that bancassurers, defined as those banks that have large holdings of life insurance assets, were the only group whose equity prices were significantly affected by disruptions in the UK life insurance sector. These results suggest that the most significant channel for spillover to the banking sector is via UK banks' ownership of life insurers, while indirect channels were not found to be materially significant.

Our study is based upon a relatively recent period, as changes in banks' business models as well as structural changes to the economy may alter the magnitude and nature of interlinkages. Consequently, our analysis employs a relatively small sample. Further research could investigate whether the results presented in this paper can be replicated for other countries.