Competitiveness, inflation, and monetary policy

Working Paper no. 246

Hashmat Khan and Richhild Moessner

Differences in the degree of competition among firms — 'competitiveness' — may affect the rate of inflation in the short term and the monetary transmission mechanism. In addition, cyclical variations in competitiveness may affect inflation dynamics and the transmission mechanism. It follows that the examination of both types of changes is potentially important for monetary policy makers.

We examine how differences in the level of steady-state competitiveness and the trend rate of inflation might affect inflation responses to monetary policy shocks, using a standard New Keynesian model. We extend the model to allow for positive trend inflation and cyclical variations in competitiveness. This allows us to quantify separately the impact of differences in steady-state levels of and cyclical changes in competitiveness on inflation dynamics, in high and low inflation environments. We apply this model to scenarios chosen to capture broadly the conditions in the UK economy in the early 1990s and more recently.

We show that in a low inflation/high competitiveness environment, the higher degree of price stickiness implied by the low inflation environment, and the higher degree of steady-state competitiveness both have the effect of dampening the inflation response to monetary policy shocks, compared with the high inflation/low competitiveness scenario. By contrast, in the low inflation/low competitiveness environment, we find that the effect of lower steady-state competitiveness partially offsets the effect of the higher degree of price stickiness in the low inflation environment, so that the inflation responses in the high and low inflation environments are similar to each other. Moreover, we quantify the extent to which procyclical changes in competitiveness dampen the impulse response of inflation to a given monetary policy shock, and the extent to which countercyclical changes amplify it.

The exposure of international bank loans to third-country risk: an empirical analysis of overdue claims

Working Paper no. 247

Drew Dahl and Andrew Logan

The paper investigates whether the credit quality of UK-owned banks' international lending is sensitive to conditions in borrower countries' largest trading partners. Borrowers may be dependent on export earnings or other income generated by economic activity in the third country to repay the loan. A deterioration in economic conditions in the third country could impair borrowers' ability to meet their loan obligations.

The existence of trade-based interdependencies has implications for banks' risk management and the authorities. Interlinkages limit banks' ability to diversify away credit risk by lending to different countries. Moreover, banks' risk management techniques will need to address the cross-country correlations in borrowers' ability to repay. Central banks and banking regulators with responsibility for financial stability or prudential supervision also need to take account of the impact of trade-based spillovers. In the Bank of England's case, any judgment of the likely impact of an adverse shock to a particular country on UK financial institutions and markets will need to factor in the knock-on effects on borrowers in third countries.

The paper has some similarities with previous empirical studies on interdependencies in international lending. One branch of the literature focuses on how the quantity of credit supplied by international banks in a country varies with financial conditions in other countries. Another investigates whether risk in international bank lending is systematic (global) or diversifiable (local). A problem in applying these studies to the issue of interdependencies, however, is that they view risk from a global perspective. They do not explore the direction or strength of interdependencies between countries and do not use a direct measure of performance.

The measure of credit quality used in this study is the proportion of the principal of cross-border and overseas operations' non-local currency loans that is in arrears. The data are annual, bank specific and disaggregated by country. This means that we have information on the credit performance of individual banks in a particular country in a given year. We are unaware of any prior empirical study that identifies performance outcomes in different countries for individual creditors. The data are confidential and collected as part of the Bank of England's suite of monetary and banking returns.

The study focuses on the credit quality of 28 of the largest UK-owned banks' international loans to 17 countries between 1991 and 2000. The 17 countries selected are those to which UK-owned banks had the greatest exposure in 2002 and for which the relevant macroeconomic data were available. They are predominately large industrialised countries. The banks included held virtually all of the foreign assets and 96% of total assets owned by UK-owned banks. The banks are important lenders within Europe — which, in turn, accounts for the majority of worldwide international lending — and have country-exposure rankings that are highly correlated with other BIS-area banks. To this extent, the results for our sample of UK-owned banks are applicable to banks in other countries.

We model overdue credit as a function of credit composition, bank characteristics and situational factors and macroeconomic conditions in the largest export market country. Two indicators of financial condition in the linked country are used. The first measures the percentage change in merchandise exports from the country of the borrower to the linked country. The second measures the percentage change in output in the linked country. This effect could encompass merchandise exports, but also trade in services or other international transactions. It may also capture effects unrelated to trade (such as collateral).

We find that economic conditions in a country are transmitted to another country whose borrowers have obtained credit from international banks. As exports to a linked country increase, or gross domestic product in the linked country increases, repayment performance in the country of the borrower improves. We find that this relationship is pronounced in countries, such as Ireland and Mexico, that have close ties to a larger economy and during the later years of our sample period (1997 to 2000).

Concepts of equilibrium exchange rates

Working Paper no. 248

Rebecca L Driver and Peter F Westaway

The term equilibrium exchange rate has been used to mean many different things by many different people. For some, the concept is clearly a long-run one. For others, even short-run movements in exchange rates may represent equilibrium behaviour. So pinning down exactly what people mean when they use the term equilibrium exchange rate may therefore be important for understanding how to interpret the information it provides.

The aim of this paper is to discuss why a range of views on the nature of the equilibrium may be valid. It does this by highlighting the distinction between short, medium and long-run concepts of equilibrium and arguing how, at any point, all these will be relevant to understanding the economy. It emphasises how the choice of equilibrium will depend on the question of interest.

The paper also emphasises that, because real exchange rates are a measure of relative prices, several different definitions can be used, where again the choice will also depend on the question of interest. But for any short, medium or long-term equilibrium concept, there will be an equilibrium configuration of relative prices associated with these different measures.

The paper briefly discusses some of the ways in which different approaches can be assessed, where again the metric employed will depend on the question of interest. For example, assessing a measure of long-run equilibrium using short-term forecast performance is inappropriate. If, however, the concept of equilibrium to be measured is short run, distinguishing between different measures using forecast performance may well be useful.

Finally, the paper provides a taxonomy of the different approaches that researchers have used to analyse equilibrium exchange rates and attempts to highlight the similarities and differences between them.

Optimal collective action clause thresholds

Working Paper no. 249

Andrew G Haldane, Adrian Penalver, Victoria Saporta and Hyun Song Shin

The potential advantages of collective action clauses (CACs) to facilitate the restructuring of debts have long been recognised and have been standard in English law bonds since the 19th century. Collective action clauses allow a contractually specified super-majority of bondholders to agree on a revision to the payment terms that is binding on all, even those who voted against. But until recently it has been the market convention in New York not to include collective action clauses. Achieving a comprehensive debt restructuring has required the unanimous agreement of all bondholders. This is generally felt to be suboptimal. Unanimous agreement means that debt restructurings are potentially held hostage to the actions of recalcitrant or rogue creditors who hope to receive better terms in subsequent offers. This can delay restructuring deals that are to the benefit of the majority of creditors and the debtor and can leave the debtor vulnerable to opportunistic legal action. In response to these concerns, a push to change the market convention in New York to include collective action clauses was first made by the official sector after the Mexican crisis in 1994–95. Little action followed. A second push was made by the official sector in 2002, following crises in Turkey, Brazil and most prominently Argentina. Again under the auspices of the Group of Ten, a working group was set up to draft model CACs. The aims of the Group of Ten working group were twofold. First, to examine a range of potential contractual clauses that could be included in sovereign bonds and to recommend which ones to include. Second, to set a new market standard. These clauses specified a majority voting threshold of 75% for changes in a bond's financial terms.

In February 2003, Mexico made a policy decision to include collective action clauses in its sovereign bonds issued under New York law, contrary to market convention. The bonds issued by Mexico followed closely the G10 model clauses, including a 75% threshold. But some subsequent issues by other countries — Brazil, Belize, Guatemala and Venezuela — opted for higher, 85%, thresholds. Some within the official sector have taken a dim view of these developments. First, because these higher thresholds take us closer to a 100% unanimity bond, thereby increasing the risk of holdouts. And second, because different voting thresholds risk a splintering of the market standard.

One contribution of this paper is to use a theoretical model of financial crisis to examine what factors might determine the choice of voting threshold — is lower always better? — and whether there are valid reasons why different issuers may want to set different, but country-specific, thresholds. We find that individual countries may wish to set different thresholds because of differing risk preferences and creditworthiness. Strongly risk-averse debtors put much greater weight on pay-offs during crisis periods than during non-crisis periods and are therefore more likely to choose lower CAC thresholds than less risk-averse debtors. The worse the creditworthiness of risk-averse debtors, however, the more likely they will want to issue bonds with high collective action clauses.

A second contribution of this paper is to develop a model that nests both liquidity runs and debt restructuring following a solvency crisis. Typically, the two are treated separately. In practice, however, it is rarely straightforward to partition crises in this way. Liquidity crises affect prospects for solvency; and expected recovery rates for creditors following a debt restructuring will in turn affect short-term decisions on liquidity. These interactions mean that most crises lie in the 'grey zone' between pure liquidity and pure insolvency. The model presented here is one such 'grey-zone' model, which allows behavioural interactions between short-term liquidity and debt restructuring following a solvency crisis.

Asset price based estimates of sterling exchange rate risk premia

Working Paper no. 250

Jan J J Groen and Ravi Balakrishnan

Many structural exchange rate models, as well as open-economy policy models, use uncovered interest rate parity (UIP) as a building block, despite the fact that UIP is strongly violated for floating currencies. Several explanations for this phenomenon have been put forward, including the presence of time-varying risk premia. Existing empirical models of the foreign exchange rate risk premium, however, are not able to generate risk premium estimates that are sufficiently variable to explain the variability in deviations from UIP.

In this paper we attempt to estimate the risk premium for several bilateral sterling exchange rates as well as the sterling effective exchange rate index (ERI). Within intertemporal utility optimisation models, the foreign exchange rate risk premium equals the conditional covariance between the future exchange rate change and the future marginal rate of substitution of the representative investor. In conventional models the marginal rate of substitution equals a linear function of future consumption growth, which is often proxied by the future real return on a stock market portfolio. This motivates the use of an unconditional (or otherwise known as static) linear factor model for this marginal rate of substitution, with either consumption growth or the real stock return as a factor. In this paper we allow for habit persistence in the consumption behaviour of a representative international investor when we derive our measure of the marginal rate of substitution. This derivation can be used to motivate the use of a conditional linear factor model for the marginal rate of substitution in which it still is related to the future real return on the agent's stock portfolio, but the model parameters are time-varying and this time variation is related to movements in the slope of the term structure of interest rates. The slope of the term structure is used, as this variable has predictive power for future turning points in the real return on the stock market portfolio. Another novel feature relative to the existing literature is that our risk premium measures are related to a representative investor who operates on a global level instead of a representative investor from a particular country.

Our estimates of unconditional and conditional factor models for the global representative investor show that, in contrast to the unconditional factor model, the conditional factor model is accepted on a monthly 1987–2001 sample of nine major sterling exchange rates. We combine the resulting conditional estimates of the marginal rate of substitution for the global investor with the covariance between the relative change in a particular sterling rate and the real return on a 'world' stock portfolio to proxy the risk premium in the effective sterling exchange rate, the sterling/DM rate and the sterling/dollar rate. The resulting sterling risk premia exhibit large swings and seem especially important for the sterling/DM rate. A graphical analysis of the estimated sterling risk premia shows, however, that the impact of the risk premium movements on sterling exchange rates seems to be limited to the short to medium run.

The foreign exchange risk premium is unobservable, and it therefore is difficult to assess whether our estimates of the foreign exchange risk premium are accurate. However, our estimates of both the marginal rate of substitution and exchange rate risk premia indicate that our approach has some empirical validity. Risk-adjusted UIP test regressions indeed indicate that relative to the major European currencies the usage of our estimated sterling exchange rate risk premia improves the parameter estimates slightly in favour of UIP, albeit not significantly so.

The stock market and capital accumulation: an application to UK data

Working Paper no. 251

Demetrios Eliades and Olaf Weeken

Estimates of the capital stock derived from National Accounts investment data suffer from a number of potential shortcomings. These are related to the difficulty in measuring investment in intangible assets and frequent data revisions. Provided that they are fairly valued, financial markets measure the value of firms' productive assets, ie their capital stock. Being less prone to revisions and arguably better suited to measure intangible assets, such market-based estimates address some of the shortcomings of National Accounts estimates. In his influential work, Robert Hall provides such market-based estimates for the US capital stock and shows that they differ substantially from National Accounts based estimates. His model is based on the well-known result that, under the assumptions of constant returns to scale in technology and in the adjustment cost function and the firm being a price taker, *marginal q*, as derived from the first-order condition of the market value maximising firm, equals *average q*. In this framework, and under certain assumptions about adjustment costs, the volume of the capital stock can be derived by equating *marginal q* to *average q*. This paper applies Hall's model to the United Kingdom to provide a market-based estimate of the UK business sector capital stock. Qualitatively, the results for the United Kingdom mirror those of Hall for the United States, with substantial discrepancies between the market-based and National Accounts based estimates. In particular, market-based estimates of the UK business capital stock were higher in the late 1990s than National Accounts based estimates. These results are robust across a range of different depreciation rates and starting values, and for all but the largest adjustment costs. These differences could reflect financial markets better capturing intangible assets than the National Accounts. However, they could also reflect an asset price bubble or economic rents that the model would mistakenly interpret as intangible assets. The results differ from Hall, in that they show a prolonged period of 'negative intangibles' for the United Kingdom. The sensitivity analysis suggests that this result is qualitatively robust throughout a wide range of adjustment costs, depreciation rates and starting values. In spite of the possible explanations for periods of 'negative intangibles', the length and magnitude of 'negative intangibles' in the United Kingdom are puzzling.

Real-Time Gross Settlement and hybrid payment systems: a comparison

Working Paper no. 252

Matthew Willison

This paper considers the issue of payment system design. Real-Time Gross Settlement (RTGS) has become the foremost system for the settlement of high-value payments in developed economies. The rationale behind the trend towards RTGS has been the perceived need to reduce the risk potentially found in deferred net settlement (DNS), the predominant system for settling high-value payments previously. RTGS entails payments being settled on a gross basis in real time. As a consequence, credit risk and settlement risk between settlement banks are eliminated. But RTGS does not dominate DNS in all respects. With payments settled on a gross basis, settlement banks' liquidity needs under RTGS are greater than those under DNS. This could induce settlement banks to delay payment activity in order to wait for incoming payments that can then be used as liquidity. Central banks have mitigated these problems to an extent by providing intraday liquidity. Of course, lending liquidity generates credit risk for the central bank. Thus, this lending is collateralised to remove this risk. Yet, the benefits from reducing the risk associated with DNS systems are considered to exceed the costs of greater liquidity needs; hence, the number of RTGS systems has grown.

The debate surrounding the optimal basis for settling payments has shifted of late with the advent of so-called hybrid payment systems. A hybrid, as the name suggests, combines features of both RTGS and DNS. More precisely, a hybrid typically takes one of the designs and augments it with features associated with the other design. Given that RTGS is the prime payment system design, recent debate has mainly concentrated on the benefits of complementing RTGS with a liquidity-saving feature called payment offset. A payment is offset when it is settled simultaneously with a set of other payments rather than being settled individually like in RTGS. When payments are settled simultaneously, the payments are self-collateralising, to the extent that their values are alike. Settlement banks only need liquidity equal to the net value of their payments in the set to settle these payments. An important design feature of any hybrid is that payments can be placed in a central queue. While payments are in this queue, the system operator searches for offsetting payments. Otherwise, payments can still be settled by RTGS without necessarily entering the central queue. The benefit of complementing RTGS with a payment offset facility is that liquidity needs can fall and the incentive to delay placing payments into the system is reduced, relative to RTGS; but offset does not necessarily reintroduce the risk present under DNS.

In this paper, we examine the issue of optimal payment system design. We compare the performance of an RTGS system against six hybrid systems based on payment offset. We assume that, when payments are offset, they are considered legally to be final and irrevocable. So the hybrid systems introduce no credit risk relative to the RTGS benchmark. We compare the system designs based on two criteria: their liquidity demands and the speed with which payments are settled. The second criterion captures the potential impact of operational risk, since any operational event will have a larger effect the more payments still remain to be settled when the event occurs. There is a trade-off between liquidity efficiency and exposure to operational risk.

We assume that some payments are time-critical. Hence, if a settlement bank delays payment settlement it faces a cost. Each settlement bank faces a trade-off between the costs of obtaining liquidity from the central bank and the costs of delaying payments when choosing how many payments to settle in certain periods of the day. We assume that the existence of payments is common knowledge but the time-criticality of payments is not publicly known. This particular kind of information asymmetry enables us to focus on the problems settlement banks face in co-ordinating their usage of the central queue and how certain features of hybrid systems could potentially affect their ability to overcome these difficulties. We analyse the effects that the frequency at which payments are offset can have. We find that, when payments can only be offset late in the day, a hybrid will not offer improvements on RTGS according to either criterion. However, when offset occurs early or all day the first-best outcome is obtained.

Decomposing credit spreads

Working Paper no. 253

Rohan Churm and Nikolaos Panigirtzoglou

Corporate credit spreads are important indicators for both monetary policy and financial stability purposes. The Bank of England therefore regularly monitors movements in such spreads, both domestically and internationally. Credit spreads contribute to the cost of external debt financing for the corporate sector, which forms part of the cost of capital that affects firms' investment decisions. Spreads also reflect perceptions about the financial health of the corporate issuers, and can thus indicate potential stress in specific sectors in the economy.

This paper addresses the factors behind credit spread movements. We know that compensation for expected default is only one component of credit spreads. Another component can also be related to credit risk, that is, compensation for the uncertainty about the probability of default. The final component is due to non-credit risk factors, which are driven by differences between government and corporate bonds and the markets in which they are traded, such as liquidity, regulation and tax. The implications for policy of an increase in spreads driven by higher expected default are different from those due to an increase driven by changes in liquidity.

The technical issue raised in this paper is the quantification of the above components. In particular, we perform two exercises. The first is to calibrate a structural model of credit risk to firms' historical default frequencies, both investment-grade and high-yield. We choose the Leland and Toft model developed in 1996 because of its simplicity and intuition, and use US data, as default frequency data for UK companies are insufficient for this purpose. UK data are available for a much shorter period and a smaller sample of companies. The purpose of this exercise is to assess the ability of the model to match firms' historical default behaviour by not only examining the fit of the model to historical default frequencies, but also the plausibility of the derived estimates of asset volatility and risk premium. In addition, this exercise allows us to calculate an average historical compensation for credit risk and compare it with the average observed spread.

The second exercise involves the use of contemporaneous forward-looking information for the equity risk premium and equity volatility, along with the Leland and Toft model, to generate time-series decompositions for the observed credit spreads of UK and US investment-grade companies, as well as US high-yield companies.

The results from these two exercises are as follows:

- The historical estimates generated for US investment-grade companies are around 20.6% for asset volatility and an asset risk premium just above 4%. This is equivalent to an equity volatility of 35% and an equity risk premium of 6.3%. The estimates for high-yield firms are 25.4% for asset volatility and 7.5% for the asset risk premium. This is equivalent to an equity volatility of 78% and an equity risk premium of 11.5%.

- These parameters imply that the average compensation for credit risk factors is 72 basis points for investment-grade firms, 55% of the average observed spread of 136 basis points. We conclude that a large part of the investment-grade credit spread is due to non-credit risk factors. The corresponding numbers for high-yield firms are 430 and 523 basis points. We therefore find that a higher proportion, 82%, of the spread is explained by credit risk for riskier (high-yield) debt.

- The contemporaneous decomposition shows that, on average, a significant proportion of the observed credit spread is due to non-credit risk factors. This is consistent with the historical decomposition. The actual spreads and the compensation for credit risk we calculate are highly correlated. The component that compensates investors for expected default, which is the only credit risk compensation risk-neutral investors would require, is significantly more stable than the spreads we observe.

- The non-credit risk component is closely related quantitatively to swap spreads for our investment-grade decomposition. Previous studies have found that a small proportion of variation in swap spreads is due to credit risk. This provides support for identifying the residual of our decomposition with the non-credit risk component. For high-yield companies, the non-credit risk component is significantly higher than swap spreads, although they follow similar patterns. This may reflect higher liquidity premia required in the high-yield corporate bond markets.

The above results imply that the information content of credit spreads as a macroeconomic indicator or predictor of corporate sector default rates would depend on the source of the shock. However, data availability restricts the available history of the decomposition. As more data become available we would be able to test the predictive ability with respect to both default rates and future growth. Another avenue for future research would be to quantify the possible determinants of the non-credit risk component of credit spreads by closely examining the structure of the different markets.

On the consumption-real exchange rate anomaly

Working Paper no. 254

Gianluca Benigno and Christoph Thoenissen

One of the well known puzzles in international finance is the so-called consumption-real exchange rate anomaly. Most international business cycle models predict that, under the assumption of perfect financial markets along with supply disturbances, consumption should be higher in the country where its price, converted into a common currency, is lower. This feature of the models is in sharp contrast with the empirical evidence, which suggests that the consumption differential across countries does not comove in any systematic pattern with its relative price (ie the real exchange rate). The removal of the assumption of perfect financial markets is not sufficient in replicating the observed evidence: indeed, recent research shows that the same anomaly in the behaviour of consumption and the real exchange rate does continue to hold. This paper explores the extent to which the introduction of non-traded goods, along with a limited international financial market structure, might account for the aforementioned anomaly. Our results suggest that the combination of these two factors is a promising avenue for understanding the behaviour of consumption across countries as well as the real exchange rate. Indeed, in our model, the calibrated moments are close to reproducing the observed behaviour of the data for a wide range of plausible parameter values. Two key features are important in accounting for our results. By assuming that international asset trade is limited to a riskless bond we break the link between the real exchange rate and relative consumption that would arise under perfect financial markets. Whereas by introducing non-traded goods, we allow for the possibility that, depending on the origin of the shock (ie traded versus non-traded), the real exchange rate and relative consumption across countries can move in opposite directions. In particular, following a positive shock to the traded goods sector in the home economy, home consumption increases with respect to consumption abroad. On the other hand, the real exchange rate appreciates if the effect coming from the relative price of non-traded to traded goods (the so-called Balassa-Samuelson effect) outweighs the terms of trade effect that would imply a depreciation of the real exchange rate. The first effect will be stronger the more dominant the shocks to the traded goods sector relative to the non-traded goods sector, while the second effect will be stronger the higher the degree of home bias in preferences. More generally, the structure of the disturbance and the specification of preferences determine the overall cross-correlation between real exchange rate and relative consumption.