

# Dynamics of the term structure of UK interest rates

## Summary of Working Paper no. 363 Francesco Bianchi, Haroon Mumtaz and Paolo Surico

---

A number of recent papers have analysed the evolving dynamics of output and inflation using systems of equations known as vector autoregressions (VARs): a set of equations where the explanatory variables in each equation are the complete set of lagged variables in the system. GDP growth, inflation and the nominal interest rate are the typical variables included in VARs that describe the transmission mechanism of monetary policy. These empirical models are subject to the criticism that they include a limited amount of information. If, in reality, the central bank examines a wider set of variables when setting policy, estimates of the monetary policy shock derived from these small empirical models may be biased — ie not completely disentangled from non-policy shocks. As a consequence an accurate assessment of structural shifts may be hampered.

The aim of this paper is to investigate the evolution of UK macroeconomic dynamics using a VAR model that is less susceptible to this criticism. In particular, we augment the standard three-variable VAR with variables that describe the level, slope and curvature of the yield curve, which shows the pattern of interest rates at different maturities. These yield curve variables contain information about private sector expectations. This additional information may alleviate the biases referred to above by ensuring that the forward-looking aspect of monetary policy is accounted for in our empirical model. In addition, we allow the relationship between the yield curve and the macroeconomy (embodied in our VAR) to change over time. We use this model to investigate how the

dynamics of UK macroeconomic variables have changed over time and how these changes are related to changing properties of the yield curve.

The main results can be summarised as follows. First, the level, slope and curvature factors display substantial time variation, with the level factor moving closely with measures of inflation expectations. Second, our estimates indicate a large decline in the volatility of both yield curve and macroeconomic variables around 1992, when the United Kingdom adopted inflation targeting. Third, and more important, during the inflation-targeting regime, monetary policy shocks have been more muted and inflation expectations have been lower than in the pre-1992 era. Fourth, the link between the macroeconomy and the yield curve has also changed over time, with fluctuations in the level factor becoming less important for inflation after Bank of England independence in 1997. In particular, policy rates appear to have responded more systematically to inflation and unemployment in the current regime. Finally, we use our time-varying macro-finance model to revisit the evidence on the expectations hypothesis (ie the hypothesis that in any given period the yield on a long-maturity bond is equal to the discounted sum of the expected yields on short-maturity bonds over the lifetime of the long-maturity bond). Our results suggest that time-varying dynamics in both the yield curve and the structure of the economy may explain part of the deviations from the expectation hypothesis found in fixed-coefficient models.

# What lies beneath: what can disaggregated data tell us about the behaviour of prices?

## Summary of Working Paper no. 364 Haroon Mumtaz, Pawel Zabczyk and Colin Ellis

---

How do prices respond to changes in interest rates? Most previous work has tried to answer this question by looking at aggregate price measures, such as the consumer prices index (CPI) or the National Accounts consumption deflator. This paper takes a different approach. Following recent work on US data, we examine the behaviour of both aggregate and disaggregated prices in the United Kingdom using a large volume of data covering prices, volumes, money and asset prices.

In this paper, we summarise these data by using 'principal components', or 'factors'. Factor analysis uses linear transformations of data series to identify common components that underlie those series. The 'factors' are calculated by creating combinations of the underlying data series to make new series that in turn capture the largest possible amount of variation in the data set as a whole, while remaining statistically independent of each other. We then use these factors to estimate a simple model (known as a vector autoregression, or VAR), which in this case relates these factors to their previous values and the interest rate. The resulting model is known as a 'factor-augmented vector autoregression', or FAVAR for short.

The advantages of a FAVAR are that it encompasses a large number of data series but, at the same time, is relatively simple to estimate. By estimating a FAVAR on disaggregated data, we are able to examine how individual disaggregated prices respond to monetary policy and other macroeconomic shocks. The model also tells us how important these macroeconomic factors are, compared to sector-specific factors that affect the individual disaggregated series.

Our benchmark results match those of previous studies and suggest that aggregate demand falls before aggregate inflation when interest rates rise. However, our disaggregated results offer a number of insights that are not captured by aggregate models.

- First, while macroeconomic factors are very important for aggregate data such as CPI inflation, they are much less important for disaggregated inflation measures. Sector-specific factors are at least as important for disaggregated prices.
- Second, we find evidence of significant aggregation bias — aggregate inflation is far more closely related to its previous values than disaggregated inflation measures. This suggests that aggregate inflation measures do not offer a good guide to the behaviour of underlying prices. In other words, trying to infer the statistical properties of individual prices from those of aggregate price indices is likely to be misleading.
- Third, different disaggregated prices respond differently to changes in interest rates, suggesting that monetary policy can affect relative prices in the economy.
- Fourth, there is some evidence that competition within industries plays a role in determining how companies set prices — in particular, companies in less competitive industries may be more able to pass on changes in prices to customers.

# Foreign exchange rate risk in a small open economy

## Summary of Working Paper no. 365 Bianca De Paoli and Jens Søndergaard

Investors require compensation (or a 'premium') to hold risky financial asset. So if some currencies are perceived to be riskier than others, investors may demand a foreign exchange (FX) premium to invest in those currencies. This paper presents a small open economy model that can explain why FX premia arise in currency markets. We use this model to examine how well it resolves the so-called uncovered interest rate parity (UIP) puzzle. UIP is simply a condition that follows from financial market arbitrage. It ensures that the interest rate return on a domestic currency asset should equal the interest rate on each foreign currency assets, less the expected appreciation of the domestic currency. The puzzle stems from the empirical observation that high interest rate currencies tend to appreciate — contrary to what UIP would predict.

A key feature of our model is that households are assumed to have consumption habits, ie households get used to a 'habit' level of consumption, and only attain higher utility if actual consumption rises relative to that level.

We demonstrate that our model will only resolve the UIP puzzle if it produces significant precautionary savings effects, where savings rise in response to increased uncertainty. And these savings effects will only occur if we assume quite persistent productivity shocks combined with very slow-moving consumption habits.

In our model, changes in precautionary savings are a result of changes in households' attitude towards risk, and changes in economic prospects. In the face of bad shocks, for example, households increase their precautionary savings if they expect consumption to be low relative to their habits level. Thus, the slower is the adjustment of habits to the shock, the larger will be the revisions in precautionary savings. These revisions are also larger when the shocks are more persistent.

To understand the combined role of slow-moving consumption habits and persistent shocks in resolving the UIP puzzle, consider how a temporary fall in productivity in the rest of the world works its way through our model. The drop in

foreign productivity causes an *ex ante* excess demand for foreign goods which is eliminated by a rise in the relative price of foreign goods, ie a domestic currency depreciation. But since this is ultimately a temporary shock, the domestic currency is expected to appreciate back towards its initial steady state.

However, the same negative foreign shock also triggers a large increase in foreign precautionary savings, putting downward pressure on foreign interest rates and hence causing domestic interest rates to exceed foreign rates at the same time as the domestic currency is expected to appreciate, thus potentially resolving the puzzle. But at the same time the increase in foreign households' borrowing to smooth their consumption (known as 'intertemporal substitution') will tend to put upward pressure on foreign interest rates and hence cause domestic rates to lie below foreign rates at the same time as the domestic currency is expected to appreciate (ie in line with the predictions of UIP). So we can only account for the tendency for high interest rate currencies to appreciate if the precautionary savings effects outweigh the intertemporal substitution effects. This would be the case if the shock is very persistent and consumption habits are very slow-moving.

We initially show our result at work in a model with fixed labour supply. We then examine how our result changes when we allow domestic households in our small open economy to vary their hours worked. In our model, this extension makes domestic consumption less synchronised with foreign consumption. To ensure that risk is efficiently shared across countries, the real exchange rate would have to fluctuate more. We find that a more volatile real exchange rate combined with a stronger precautionary savings effect actually improves the model's ability to address the UIP puzzle. But when we allow both domestic and foreign households to vary their hours worked, consumption is both smooth and synchronised across countries. This dampens the FX premium volatility and impedes the model's ability to resolve the UIP puzzle.

# Common determinants of currency crises: role of external balance sheet variables

## Summary of Working Paper no. 366 Mirko Licchetta

This paper investigates the role of external balance sheet variables as determinants of episodes of currency crises in both advanced and emerging market economies (EMEs). There is a relatively well-established literature on the determinants of currency crises but only recently has some attention been made to the role of a country's external capital structure as a potential source of vulnerability. Since the Asian crisis, many economists have focused on the destabilising role of short-term debt flows, suggesting that their liberalisation between the late 1980s and the early 1990s was a major cause of episodes of crises in EMEs. More recently, the development of the balance sheet approach to financial crises has emphasised the role of external assets and liabilities in affecting a country's financial strength, and some empirical studies have provided support to the idea that debt flows are particularly prone to sudden stops in times of stress.

This paper uses a model to investigate the role of external balance sheet variables as determinants of currency crises in emerging and advanced economies over the January 1980 to December 2004 period. Using a new database on external assets and liabilities, this paper investigates the role of the size and the composition of the stock of gross external liabilities as possible determinants of a country's degree of vulnerability to crises. Our central finding is that the probability of a crisis is found to increase with the size of total liabilities (relative to GDP) and, particularly in EMEs, to decrease with the share of foreign direct investment (FDI) in total liabilities.

There are reasons in support of the idea that a country's vulnerability to crises increases with the stock of external debt. A large stock of external debt implies a large dependence on foreign sources of finance. Therefore, the larger the stock of external liabilities, the larger is the amount of capital that can potentially be withdrawn in a sudden stop. Then, from an empirical perspective, there is evidence to suggest that international capital flows are determined by external factors as well as domestic ones. Therefore, the larger the inflow, the more sensitive a country's external financing is likely to become to external conditions. However, it is still debated within the empirical literature whether a high level of debt necessarily increases the likelihood of a currency crisis.

As for the role of the composition of external liabilities on the determination of currency crises, there are reasons to suggest that a higher (lower) share of external debt (FDI) liabilities increases (decreases) the susceptibility of a crisis. First, a lot of external debt is short term, whereas FDI is less fungible and, thus, more

difficult to withdraw in a crisis. Second, contractual obligations for debt financing — unlike for equity — are unrelated to the performance of the economy, so an adverse shock may cause EMEs debt repayment difficulties and forward-looking investors may withdraw in anticipation of these problems. The empirical evidence supports the view that a high FDI (debt) share is likely to reduce (increase) the vulnerability of an economy to crises. Short-term debt flows are usually found more sensitive to shocks to other capital flows and more volatile than FDI. Moreover, both bank loan and bonds debt flows are largely reversed during periods of stress whereas portfolio equities are found to be less sensitive and FDI stable.

Our results also suggest that the composition of external liabilities has a more important impact on the degree of vulnerability of emerging rather than advanced economies. This might be due to the shorter maturity of debt that EMEs traditionally experience. In the presence of mismatches between short-term liabilities and long-term assets, a country is likely to be particularly vulnerable to crises. Another explanation may be related to the so-called 'debt intolerance' of emerging market economies, which suggests that in most emerging markets external debt to GNP ratio needs to be lower than 35% (and even lower if a country has a long history of crises or defaults) to be regarded as 'safe'. This is because emerging market economies tend to have a weaker fiscal structure, less-developed financial systems and a worse record of macroeconomic management and inflation than more advanced economies. Therefore, they are felt as less able to tolerate higher levels of indebtedness.

Countries with a fixed exchange rate regime are found to be more sensitive to external balance sheet variables than economies with more flexible regime. Under a flexible exchange rate, banks and firms may be more likely to be sensitive to currency risks. Indeed, they have a stronger incentive to match foreign currency liabilities with dollar assets than in the presence of a fixed exchange rate. On the other hand, for a given external liability structure, fixed exchange rate regimes are more likely to lead to currency mismatches because economic agents believe the government commitment to the peg will immunise them from exchange rate fluctuations.

This paper also provides further support to standard leading indicators of currency crises and it reinforces the view that crises during the 1990s were likely to be less 'fundamentally' driven than those in the 1980s.

# Labour market flows: facts from the United Kingdom

## Summary of Working Paper no. 367 Pedro Gomes

---

Macroeconomic policy makers need to have a good understanding of the state of the supply side of the economy in order to set monetary policy appropriately, and an important part of supply is in the labour market. Close attention is paid to the stocks of employed, inactive (ie, those not working or looking for work) and unemployed people, as well as to the balance between demand and supply often referred to as 'tightness'.

But the labour market stocks and aggregate indicators are fundamentally driven by the behaviour of flows between employment, unemployment and inactivity. These flows are very large. On average, between 1996 and 2007 nearly a million people moved into new jobs every three months, with a slightly smaller number leaving. A smaller but comparably large number of people shifted jobs each quarter as well. So these gross flows are massive. For example, at 60,000 per quarter, the average increase in employment over this period (the net flow) was less than a tenth the size of either of those two gross employment flows. It is clear therefore that an understanding of all the relevant flows is essential to our understanding of labour market dynamics and business cycle fluctuations. Moreover, from an academic point of view, they lie at the heart of many recent theories of unemployment.

Thus the simple objective of this paper is to describe the main developments in, and establish a number of key facts about, the recent history of these important UK labour market flows. For policy makers, knowledge of those facts can help improve the monitoring of business cycles, the detection of inflection (turning) points and the assessment of labour market tightness. For macroeconomists, they provide a summary of the empirical features that theoretical models should ideally have.

It is possible to draw out some broad features of the data from the analysis of the Labour Force Survey over the period 1996 to 2007. On average, in each quarter 7% of the working-age population change status between inactivity, employment and

unemployment and 2% of the working-age population change their employer. In expansions, although jobs become easier to find, as the labour market becomes tighter there are fewer movements between the three pools. The cyclical behaviour of flows between inactivity and employment seem to have changed in recent years. They were not related to the business cycle until 2001, but became positively related (procyclical) thereafter.

Every quarter 7% of all employees search for a different job, and they are seven times more likely to change jobs than those who are not searching. In booms, there are less people searching for a different job, but they are more likely to change employer. In booms, more people resign their jobs, but there are less people being fired. Involuntary separations dominate the employment-to-unemployment flows, while 70% of all employment-to-inactivity flows occur because of personal reasons. Inactive people who want a job are twice as likely to move into the labour force, and four times more likely to move into unemployment, than those inactive people who do not want a job.

Some of the structural changes in the UK labour market seem to be due to changes in the education level of the working-age population, particularly due to the increasing share of the highly educated. There are substantial differences in the employment, unemployment and inactivity rates of different education categories, as well as in the transition probabilities (the chances of moving between different labour market states). The less-educated individuals face unemployment and inactivity rates that are three times greater than those with higher education, as well as double the separation and half the job-finding rate.

Job-finding and job-separation rates are equally important determinants of unemployment fluctuations. The job-finding rate has been more important over the past ten years, but further analysis of claimant count data has revealed that the job-separation rate was particularly relevant in the period between 1989 and 1996.

# The real exchange rate in sticky-price models: does investment matter?

## Summary of Working Paper no. 368 Enrique Martínez-García and Jens Søndergaard

Explaining exchange rate movements remains a challenging area of research for academics and policymakers alike. This is so partly because exchange rates are volatile; the standard deviations are roughly five times larger than GDP. It can also take three to five years for exchange rates to return halfway to their long-run values. So previous researchers have therefore attempted to build macro models where exchange rates are both excessively volatile and very slow to revert. In one strand, it is argued that high exchange rate volatility is due to the combination of sticky prices and nominal (eg monetary) shocks. But nominal shocks tend not to produce sufficiently persistent exchange rates. A more recent strand has argued that real shocks can potentially account for both the volatility and persistence of real exchange rates. However, the literature has so far not paid attention to the link between real exchange rate dynamics and what the model assumes about physical capital. Given how volatile investment flows are relative to output over the business cycle, we think this omission is not inconsequential.

Dynamic stochastic general equilibrium (DSGE) models take into account the evolution over time of interrelationships between agents in the economy, where there are random ('stochastic') shocks hitting the economy. We build a two-country open economy model that features optimising households and firms, as well as sticky output and import prices. It is symmetric with two equally sized economies (ie the United States and the euro area). As is common in the academic literature, we compare how the model matches the main features of US/euro-area data including the \$/€ real exchange rate (the exchange rate accounting for price differences). We are particularly interested in whether our results hinge on what the model assumes about capital.

A key assumption is that households and firms have access to complete international financial markets (ie they can buy a set of securities that ensure risk is effectively shared across countries). We find that assuming this imposes a tight link between real exchange rate and consumption volatility. Real exchange rate volatility is approximately determined by how correlated consumption is across countries as well as how volatile consumption is. So whether the model produces volatile real exchange rates or not depends on how well households can smooth consumption over time ('intertemporally') in response to country-specific shocks. Capital formation (investment) represents an additional intertemporal smoothing channel, and so tends to lead to smoother consumption and, hence, less

volatile exchange rates. But this ability to smooth consumption hinges on how costly it is for households to adjust their capital stock. Capital adjustment costs regulate the volatility of investment and indirectly control the degree of consumption smoothing in the model. So ultimately real exchange rate volatility depends on the degree of capital adjustment costs imposed in the model.

We show that our sticky-price DSGE model with capital adjustment costs can produce volatile exchange rates as they are in the data when business cycles are exclusively driven by nominal (monetary) shocks. But adding variable capital utilisation reduces the volatility of real exchange rates. Capital utilisation offers a way around the investment constraints imposed by capital adjustment costs and hence facilitates consumption smoothing in the model with monetary shocks. This, in turn, reduces real exchange rate volatility in our model by 50%. We then focus on the role of real shocks, and show that these shocks produce too little real exchange rate volatility, especially when we allow for capital accumulation. For instance, our model with no capital accumulation and real shocks generates 60% of the observed real exchange rate volatility. But when we allow households to alter their capital stock — subject to investment adjustment costs — that proportion falls to 30%.

Our model's ability to generate real exchange rate persistence depends on the type of shock driving business cycles. Monetary shocks alone cannot generate real exchange rates that are as persistent as in the data. The key to high real exchange rate persistence depends on whether the particular shock causes long-lived real interest rate differentials. But this is partly determined by how monetary policy responds to the particular shock. For instance, an expansionary monetary shock increases inflation and output which requires the central bank to subsequently raise nominal interest rates sharply. Since this endogenous policy response quickly brings inflation and output back to target, the initial monetary shock dissipates very quickly and produces little real exchange rate persistence. In contrast, real shocks do produce more persistence. Positive real shocks increase output but push down on inflation. Monetary policy responds by lowering interest rates, further supporting demand but also increasing inflation. This allows the shock to propagate for a longer time leading to longer-lasting real interest rate differentials and greater real exchange rate persistence. Including or excluding capital in the model turns out to be inconsequential in most instances for persistence.



# Multivariate methods for monitoring structural change

## Summary of Working Paper no. 369 Jan J J Groen, George Kapetanios and Simon Price

Monetary policy makers need to know what is happening now in the economy, and also to have some idea what will happen in the future. To do the latter, they need to forecast. But a major practical problem is that one of the main causes of forecast failure is structural change. Often, as David Hendry and Mike Clements have emphasised, this manifests itself as a 'mean shift'; for example, a step change in a growth rate. In some cases, policymakers might have a good idea when such changes take place. For example, the shift to inflation targeting in the United Kingdom in 1992 and the move to Bank independence in 1997 were clear structural changes, with likely consequences for inflation. But in other cases, such as the period after a large rise in energy prices, the case may not be so obvious. It would be helpful to have statistical techniques that help us look for evidence of such changes in the data.

However, such a continuous 'monitoring' of series for structural changes, period after period, raises well-known econometric issues. Statistical tests are designed so that accepting a false hypothesis happens only a small proportion of times, often set at 5%. The idea is that if we do such a test only once, then there is only one chance in 20 of making this type of mistake. In this way we can be quite confident that results are unlikely to have been generated by chance; it is a cautious approach. But it is easy to see that if such a test is repeated many times then eventually it will accept a false hypothesis (in this case, that a break has happened) purely by chance. This has led to the development of techniques looking at single time series accounting for this problem. But in practice such tests are not very successful in detecting breaks, as they must be inherently conservative.

The simple insight explored in this paper is that if several time series of data have a structural break at roughly the same time ('co-break'), as may often be plausible, then it is possible that simultaneous examination of a set of such variables helps identify changes with higher probability or more rapidly than when each is examined on a case-by-case basis. Naturally, this need not necessarily imply that there is a break in some aggregate series of interest, although it may do so. Some statistical theory is developed for such a method, which cumulates forecast errors from many series and picks the maximum at each point to construct a 'CUSUM' (cumulated

sum) detection test, or 'detector'. Breaks leading to forecast failure often manifest themselves by mean shifts, even if the shock to the variable is temporary. We therefore focus on this type. Monte Carlo experiments (simulating data generating processes thousands of times) suggest that there is an improvement in detection relative to a single variable test over a wide range of experimental parameters, given a sufficiently large number of co-breaking series. It should be clear, however, that this method only has the potential to detect the existence of a general break; it is not informative about the precise nature of that object.

One very natural application is UK retail prices index (RPI) inflation in the period after 2001. This is partly because many subcomponent series are published (about 80 on a consistent basis over the relevant period). But there are also several reasons to suppose that at least some of these may have experienced breaks. Although inflation is determined by monetary policy in the long run, in the short run large fluctuations in important prices may lead to the breakdown of the type that we consider in empirical relationships. From 1992 to 2001 there was a high degree of stability in aggregate RPI inflation. But thereafter house price inflation fluctuated fairly widely (peaking at over 25% per year) and energy prices rose dramatically after 2004. As policymakers, we are mainly interested in the aggregate, and not all of the subcomponents need to have co-broken for there to be an impact on total RPI, so this method could be useful. On the other hand, we should recognise that breaks may be offsetting, so that there may be no effect on the aggregate series.

It turns out that univariate methods would not have detected any breaks in the aggregate series, but the multivariate method would have indicated a potential break in 2005, which would then suggest further examination using other methods. And it appears from an examination of the data over the whole period, and therefore with the benefit of hindsight, that such a break may have occurred in the aggregate RPI series. It should be noted, however, that the evidence is not overwhelming, and there is no sign of a break in the inflation series that was targeted for much of this period (RPI excluding mortgage interest payments). Nevertheless, this new method may be a useful addition to the toolkits of policymakers and other forecasters.

# Banks' intraday liquidity management during operational outages: theory and evidence from the UK payment system

## Summary of Working Paper no. 370 Ouarda Merrouche and Jochen Schanz

We investigate how banks that are direct members of the United Kingdom's large-value payment system, CHAPS, react to operational problems that prevent their counterparties from making payments in CHAPS. It handles nearly all large-value same-day sterling payments between banks, other than those relating specifically to the settlement of securities transactions. Every day, about £270 billion worth of payments are settled using the system. In such real-time gross settlement systems, these direct members — also referred to as settlement banks — rely to some extent on incoming payments to fund their own payments: in 2006, five settlement banks settled £5–£10 worth of payments for each pound of liquidity they had available at the start of the day, and five other banks settled even more than £10. (There were fifteen direct members at that time: the Bank of England and CLS were excluded in these calculations, and the Royal Bank of Scotland and NatWest treated as one entity.)

Occasionally, a settlement bank experiences operational problems which prevent it from sending payment instructions to CHAPS (an 'outage'). Frequently, such a bank — we refer to it as a 'stricken bank' — remains able to receive payments on its account with the Bank of England. If its counterparties continue to make payments, a stricken bank involuntarily absorbs liquidity: it becomes a 'liquidity sink'. To the extent that healthy banks relied on incoming payments to fund their own payments, they find themselves short of liquidity. Thus, if banks do not sufficiently monitor their outgoing payments, operational risk at one bank can be a source of liquidity risk to the payments systems as a whole.

In this paper, we investigate by how much, and when, banks on average reduce their outgoing payments to a stricken bank that is able to receive but unable to send payments.

We first present a game-theoretic model to understand under which circumstances we would expect banks to withhold payments to a stricken bank. The distinctive feature of such models is that the 'players' take into account the likely response of each other; in other words, they play 'strategically'. The model covers payments behaviour on a single day. Two banks decide at the start of the day how much liquidity to borrow from the central bank. In the subsequent periods, they decide whether to delay the execution of their payment instruction(s). Whether delay is attractive depends

on how much liquidity each bank has available, the other bank's strategy, and whether operational shocks have hit one or both banks.

We show that, under reasonable conditions, banks ensure they retain sufficient liquidity to be able to execute unexpected urgent payment instructions immediately — if necessary by temporarily withholding payments to a stricken counterparty. Because these urgent payments are more likely early in the day, a healthy bank is more likely to withhold payments temporarily if a shock occurs in the morning rather than in the afternoon.

These results are supported by our empirical estimates. We focus on the activity of the five major CHAPS settlement banks. (They execute 80% of all payments in value terms.) Our data set includes eight days at various points in 2007 when at least one of these banks was unable to send any payment during a certain time interval. We find that during the outage, healthy banks on average reduce their payment outflows to the stricken bank by 40%. The trough is reached at 50% about 40 minutes into the outage; afterwards, payment flows pick up. As our model predicts, the reduction is substantially stronger (by a factor of between two and four) when the outage starts in the morning rather than in the afternoon.

Importantly, by selectively reducing their liquidity outflows to the stricken bank, healthy banks successfully prevented any significant spillover from the outage; the average decrease in payment flows between healthy banks is not statistically different from zero.

We execute a number of robustness checks for our empirical results. Our outages differ in several aspects: the bank that experiences the outage, the time of day at which they occur and the date (some occur during the liquidity shortage in the second half of 2007), and, finally, their length. With just eight outages, we cannot independently identify each of these influences. We therefore group the outages into smaller groups which are more homogeneous: for example, in one group, we exclude outages that occurred during the crisis; in another, we exclude the two longest and the two shortest outages. Our results prove robust to these variations.



# Payment systems, inside money and financial intermediation

## Summary of Working Paper no. 371 Ouarda Merrouche and Erlend Nier

Reforms to improve the efficiency of interbank payment systems have the potential to improve welfare by increasing the demand for inside money (demand deposits created 'inside' the banking system) and reducing the need for outside money (currency and reserves, created by the central bank and therefore 'outside' the banking system). This is because inside money can be put to productive use: banks use deposits to finance lending to the corporate sector. An economy which increases its reliance on inside money relative to outside money may therefore be able to support a higher level of capital and increased growth as that capital is accumulated.

In this paper, we study the effects of payment systems reform, using as a laboratory the transition from paper-based to modern, automated payment systems in Eastern European countries during the period 1995 to 2005. After 1989, these countries undertook major reforms of their financial and banking systems. As part of that effort they introduced modern, secure, automated interbank payment systems. These reforms were introduced to improve the reliability and security and to increase the efficiency of accounts-based payments. At the same time there have been sizable shifts in the amount of funds intermediated by the banking system. Indeed, many commentators have referred to a credit boom in a number of our sample countries. This paper investigates whether payment system reform and credit creation are causally linked. We also study the channels that might effect such a link.

Two channels are investigated. First, innovations in payments technology enhance the speed and security of inside money as a payment medium for customers and therefore affect the split between holdings of cash (outside money) and holdings of deposits (inside money). Second, innovations in payment systems help establish well-functioning interbank markets for end-of-day funds. This reduces the need for banks to hold excess reserves (outside money) to self-insure against end-of-day outflows and thus helps credit creation.

We find that upon the introduction of efficient payment systems there is a marked increase in the trend growth of financial intermediation. This finding exploits differences in the timing of reform across our sample countries, and it is robust to the inclusion of control variables that take account of macroeconomic determinants of credit growth, such as general economic development and the extent of foreign capital inflows, as well as other dimensions of structural change in our sample, such as variation in government ownership of banking assets.

Further investigation suggests that the trend increase in credit supplied to the private sector is associated with a trend decrease in the use of currency, relative to demand deposits. By contrast, the evidence in favour of a reduction in banks' holding of excess reserves is less strong.

To establish whether any of these two channels might have caused the observed increase in the trend growth of credit around payment reform, we investigate whether there is more generally a positive response of credit to either a shock to the demand for deposit (relative to currency) or the ratio of reserves to deposits. Here we find a corroborating, more general relationship between credit and deposits, but not between credit and reserves. This suggests that while a shift away from cash and towards demand deposits around the reform dates can be argued plausibly to have caused an increase in credit, a shift away from reserves cannot.

Last, while our main results rely on a reform variable that is 'zero-one' we also employ alternative measures of payments systems development, such as the number of credit transfers effected across the payment system. These alternatives provide a continuous and more direct measure of the system's use. Our empirical tests using these alternatives confirm that payment system development is an important contributing factor in accounting for the observed increases in credit creation in some of our sample countries.