# Changes in the transmission of monetary policy: evidence from a time-varying factor-augmented VAR

**Summary of Working Paper no. 401** Christiane Baumeister, Philip Liu and Haroon Mumtaz

Several recent studies have documented that the volatility of output and inflation in the United States showed a remarkable decline after the mid-1980s in common with the experience in many countries. In addition, there is evidence to suggest that the persistence of inflation also fell after this date. A growing empirical literature has examined this apparent change in the dynamics of the US economy. These papers usually employ empirical models that contain a limited amount of macroeconomic variables — typically 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 simple VARs that describe the transmission mechanism of monetary policy. 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.

This paper therefore explores the dynamics of the US macroeconomy using a VAR model that incorporates a larger amount of economic information than a tri-variate model. In particular, we use an extended version of the 'factor-augmented VAR' (FAVAR) model recently proposed in the literature. The idea behind the FAVAR model is that the bias created by the difference in the information set of the researcher and the agents described in the model can be alleviated by augmenting the standard VAR with *common factors* that are extracted from a large set of macroeconomic indicators. These common factors summarise the relevant information in the macroeconomic indicators and therefore provide a proxy for the information set of agents in the model.

Our FAVAR model for the United States contains common factors extracted from data on real activity, inflation, money and credit and asset prices in addition to a short-term nominal interest rate. The innovation in our work is that we also allow the coefficients of the model and the variances of the shocks to vary over time. When this model is estimated on artificial data, it provides robust inference on changes in impulse response functions suggesting that the model is well suited to the task at hand.

The model is estimated over the period 1960 Q1 to 2008 Q3 (largely predating the recent recession). Our main results suggest that time variation is indeed a pervasive feature of key macroeconomic variables like output measures, price indices, money aggregates and asset prices. In this respect, we find important differences in the responses obtained from our FAVAR specification compared to low-dimensional systems. More specifically, in our data-rich environment we find that economic activity declines by less in more recent times after a restrictive monetary policy shock, whereas no time variation is detected in small-scale VARs.

We find no evidence of a 'price puzzle' (the common and counterintuitive finding that prices rise after a monetary contraction) for any of the aggregate price measures throughout the sample period. This may suggest that the extra information captured by the factors leads to more robust structural estimates in that it mimics the central bank's practice of examining and reacting to a wide variety of data series. However at the disaggregate level, a considerable portion of sectoral price responses displays a significant price puzzle at short horizons during the 1970s which ameliorates from the early 1980s onwards. Our evidence therefore provides a case for the price puzzle not being a puzzle at disaggregate level, but rather a distinctive feature of sectoral dynamics. This should allow us to infer something about the price-setting behaviour of firms in reaction to monetary surprises.

Our results suggest that durable goods are most sensitive to interest rate innovations and show a considerable fall in consumption volumes and a decline in the price level. Durable goods also contribute the least to the dispersion of sectoral prices since individual impulse responses are closely aligned. Instead, non-durable goods and to some extent services account for a large share of cross-sectional heterogeneity, with price responses widely dispersed, covering a broad range of positive and negative values.

# DSGE model restrictions for structural VAR identification

**Summary of Working Paper no. 402** Philip Liu and Konstantinos Theodoridis

Monetary policy making in central banks requires a profound understanding of the way the economy reacts to the shocks that continually bombard it. So banks call upon a wide range of economic models to help them in this undertaking. Since the pioneering work of Sims, vector autoregressive (VAR) models have been used extensively by applied researchers, forecasters and policymakers to address a range of economic issues. These models comprise equations explaining a small number of key macroeconomic variables where each equation includes the same set of explanatory variables, lagged values of all the variables in the system. The basic VAR is therefore unable to tell us about the detailed structure of the relationship or shocks, which is what the policymaker really wants to know, as it is a 'reduced-form' model. To unpack the shocks hitting the system and their effects on the economy, we need to 'identify' the model with extra assumptions.

Although VARs have been very successful in capturing the dynamic properties of macroeconomic time-series data, the decomposition of these statistical relationships back to coherent economic stories is still subject to a vigorous debate. However, the outcomes of the VAR analysis depend crucially on these assumptions and the various competing identification restrictions cannot be easily tested against the data. Even though several procedures have been proposed in the literature, shock identification remains a highly controversial issue.

A type of model that is not susceptible to this problem is the dynamic stochastic general equilibrium (DSGE) model. In this case, economic theory is used to define all the linkages between variables. The tight economic structure solves the identification problem, but at a cost. As theory is never able to fully explain the data, an agnostic VAR will almost certainly 'fit' the data better.

This paper proposes an identification strategy for VARs that extends an idea introduced by Harald Uhlig, a 'penalty function' that effectively weights various restrictions suggested by theory — in his case, the signs of various effects. So we construct a penalty function that is based on *quantitative* restrictions implied by a DSGE model. To assess the usefulness of the proposed identification strategy, we present a series of Monte Carlo experiments (where many experiments are carried out on an artificial model, randomly differing in the shocks hitting the system). First, we investigate the ability of the method to recover the true set of structural shocks; second, we examine the source of bias in the identified VAR responses relative to the true data generating process; and third, we assess how the proposed identification strategy performs using restrictions from a misspecified model. We also present an application using a seven-variable VAR model estimated on US data. The structural shocks are identified using restrictions from a classic medium-scale DSGE model developed by Frank Smets and Raf Wouters.

A number of interesting results emerge from the analysis. First, by using the correct model restrictions, the identification procedure is successful in recovering the initial impact of the shocks from the data. Second, despite using restrictions from misspecified models, the data tend to push the VAR responses away from the misspecified model and closer to that of the true data generating process. Third, the proposed identification strategy systematically gives smaller bias compared with other popular identification schemes.

# Monetary policy rules and foreign currency positions

**Summary of Working Paper no. 403** Bianca De Paoli, Hande Küçük-Tuğer and Jens Søndergaard

Over the past decade, international financial markets have become increasingly integrated. This process of financial globalisation is reflected in the rapid expansion of the external balance sheets of countries which record cross-border ownership of assets and liabilities. In this world of interlinked balance sheets, exchange rate movements can lead to significant shifts in a country's external position. This 'valuation effect' depends crucially on the size as well as currency composition of a country's external position. For example, if a country's foreign assets are predominantly denominated in foreign currency, a weakening in the domestic currency will increase the domestic currency value of its net foreign asset position.

The empirical evidence suggests that an indirect link exists between the currency composition of a country's external position and its monetary policy. In particular, inflation-targeting countries appear to hold relatively more foreign debt liabilities denominated in foreign currency than non inflation targeting countries.

This paper formalises this empirical link between monetary policy and foreign asset holdings. It uses a model of endogenous portfolio choice explaining why agents hold particular assets, under the assumption of incomplete markets (that is, in the absence of complete insurance against risk). A framework is developed where optimal foreign currency portfolios are directly linked to exchange rate dynamics. Whether the domestic currency depreciates or appreciates in periods of relatively low consumption determines whether investors take a long or short position in the foreign currency (in other words, whether their portfolio is overweight or underweight in foreign bonds).

The key insight of this analysis is that different monetary regimes change the cyclical properties of the exchange rate and hence alter agents' hedging incentives (whereby agents take positions that protect themselves against adverse movements in their consumption). For instance, if the central bank is assumed to target money growth — or follow an interest rate setting 'Taylor rule', ie a rule that has interest rates responding not only to movements in inflation but also some measure of output growth or the output gap — agents would choose a portfolio that is underweight (short) in domestic bonds and overweight (long) in foreign bonds. Intuitively, any adverse real country-specific shocks will — with these particular monetary policy rules — be associated with a nominal depreciation of the domestic currency. Being overweight in domestic currency denominated assets is therefore a bad hedge.

On the other hand, when the central bank conducts policy through an inflation-targeting Taylor-type rule, the same adverse shock will trigger a nominal domestic currency appreciation. So holding domestic currency denominated assets is a good hedge and agents will choose an optimal portfolio that is overweight in domestic currency denominated bonds.

The paper also illustrates how the endogenous portfolio choice determines the cross-border transmission of monetary policy shocks via a valuation channel. In the case of money-growth rules, agents are overweight in foreign bonds. So monetary policy shocks that cause a domestic currency depreciation generate an increase in the domestic country's net external wealth position. Thus the valuation effects of monetary policy shocks are beggar-thy-neighbour. By contrast, monetary policy shocks with an inflation-targeting Taylor rule cause international valuation effects that are beggar-thy-self. Since agents are holding a portfolio short in foreign bonds, a domestic nominal depreciation will imply a decline in the country's net external wealth.

# The impact of payment splitting on liquidity requirements in RTGS

**Summary of Working Paper no. 404** Edward Denbee and Ben Norman

This paper examines the impact that payment splitting could have upon the liquidity requirements and efficiency of a large-value payment system, such as the United Kingdom's CHAPS. Under payment splitting, a threshold value for payments is defined. Any payments larger than this threshold are split into equal pieces, each smaller than the threshold, and are then settled. In this study we use real UK payments data and the Bank of Finland Payment and Settlement Simulator to test two hypotheses: that (i) payment splitting can reduce the length and impact of payment queues prior to settlement; and, equivalently, (ii) payment splitting can reduce the liquidity requirements of the system.

A number of systems worldwide already adopt payment splitting, either as a formal mechanism or through informal guidance and practice, as a means of being more liquidity efficient. In CLS, a foreign exchange cash settlement system, a currency threshold is set for each currency that it processes. Any eligible transaction above this threshold, in either currency, is split into smaller, equally sized transactions. The Swiss SIC payment system, the Japanese large-value payment system, BoJ-Net and the Canadian securities settlement system, CDSX, all have guidelines or rules that encourage participants to split the largest payments into smaller pieces to aid payment co-ordination and liquidity efficiency.

Our results suggest that if banks were liquidity constrained and, hence, payments were queued prior to settlement, payment splitting could significantly reduce the length of these queues. Splitting allows partial settlement of payments where otherwise none would have been possible. This directly reduces the value of payments queued. Beyond this the recipient bank may be able to use this liquidity to settle queued payments of its own resulting in a favourable 'payment cascade' effect. Reducing the splitting threshold generally results in greater reductions in payment queues.

We also find that, equivalently, payment splitting can reduce banks' liquidity requirements. Splitting payments into smaller pieces and releasing them piecemeal can help banks to co-ordinate their incoming and outgoing payments resulting in less demand for liquidity. By spreading the largest-value payments over time, banks are able to use incoming payments to fund the remaining pieces of an outgoing split payment.

Given the potential benefits from payment splitting, it is worth asking why it is not more widespread. We identify two issues that may discourage systems from adopting payment splitting. First, the liquidity savings that result from this approach are not uniformly distributed. In our simulations, most banks made savings, whereas a few saw an increase in their liquidity needs. The latter tended to be those banks whose payment flows are most dependent on the arrival of incoming payments. In practice we expect that these banks would change their behaviour following the introduction of payment splitting.

Second, we recognise that some legal questions could be raised by payment splitting: above all, if a bank goes into administration after having only partially completed a payment, what is the status of that payment? Whether, and (if so) to what degree, this introduces risk depends upon the type of transaction (if any) that is underlying the payment. While a risk for some underlying transaction types, we conclude that in some cases splitting may actually reduce credit risk. We do not attempt to address the legal questions in detail but merely highlight the issues that a system operator would need to consider if it were to implement payment splitting functionality.

This paper does not seek to propose the adoption of payment splitting functionality in the United Kingdom but rather contributes to the growing literature on mechanisms for making real-time gross settlement payment systems more liquidity efficient.

# Monetary policy, capital inflows and the housing boom

## Summary of Working Paper no. 405 Filipa Sá and Tomasz Wieladek

A range of hypotheses have been put forward to explain the boom in house prices that occurred in the United States from the mid-1990s to 2007. This paper considers the relative importance of two of these hypotheses. First, global imbalances increased liquidity in the US financial system, driving down long-term real interest rates. Second, the Federal Reserve kept interest rates low in the first half of the 2000s. Both factors reduced the cost of borrowing and may have encouraged the boom in house prices. We develop an empirical framework to separate the relative contributions of these two factors to the evolution of residential investment and real house prices. Two types of shocks are identified: an increase in capital flows to the United States and an expansionary monetary policy shock.

The results suggest that capital flows shocks played a much larger role in increasing house prices than monetary policy shocks. We find that compared to monetary policy, the effect of a capital inflows shock on US house prices and residential investment is about twice as large and substantially more persistent. This finding is confirmed by the results of variance decompositions which show that, at a forecast horizon of 20 quarters, capital flows shocks explain 15% of the variation in real house prices, while monetary policy shocks explain only 5%.

A simple counterfactual exercise suggests that if the Federal Reserve had kept policy rates constant since the end of 1998, house prices might have been 8% lower by the end of 2007. Similarly, if policy rates had been set according to the Taylor rule, house prices might have been 5.5% lower. House prices would have been considerably lower (13%) if the ratio of the current account deficit to GDP had remained constant since the end of 1998.

The evidence suggests that global imbalances played an important role in generating the housing boom that characterised the run-up to the current crisis. This result would lend support to calls for the development of policies to prevent the build-up of large current account imbalances in the future, making the international monetary system more resilient to crises like the one we recently experienced.

# Forecasting in the presence of recent structural change

**Summary of Working Paper no. 406** Jana Eklund, George Kapetanios and Simon Price

Forecasting is a central activity for central banks, not least because policy takes effect with a lag. Inevitably, policy is forward looking. Thus in many central banks, including the Bank of England, the published forecast is a key tool in communicating judgements about monetary policy and the economy. The Bank's forecast, published in the *Inflation Report*, represents the judgements of the Monetary Policy Committee and is not mechanically produced by a single model. However, many forecasting models — a 'suite' of models — help the Committee determine its judgement, including simple largely atheoretical models of the type considered in this paper.

One common cause of forecast failure is that structural changes or 'breaks' keep on occurring in the underlying relationships in the economy, and this paper addresses that problem. Dealing with this has two aspects. First, detection; and subsequently the right forecasting strategy. Consequently, there are many papers on the identification of breaks, and forecasting methods that are robust to them. But these are mainly in the context of fairly distant events. The fact that in practice forecasters have to forecast after recent changes has received remarkably little attention. Yet this is a pervasive and profound problem.

Furthermore, in practice we may be continually 'monitoring' for breaks, and this raises a subtle issue. In that case the forecaster inevitably carries out repeated tests. This matters, because if statistical tests are repeated enough times, then even if one never occurs in reality by pure chance they must eventually flag a break. Luckily, there are methods to take care of this. But the subsequent problem of how to then adapt the forecasting strategy has hardly been discussed. We therefore address two important issues. First, we ask whether the forecaster should attempt to detect and react to breaks each period, or instead adopt robust forecasting strategies. Second, we consider two quite different environments. In one case, breaks are unique events (or are rare enough to be treated as such), and in the other they recur.

The monitoring strategy we examine is to look for evidence of breaks and then combine forecasts from models that do and do not use data before the change. And the alternative is simply to use methods that are robust to breaks. We examine several commonly used methods of this type, all of which work by in one way or another giving more weight to recent observations (less likely to be affected by breaks).

We first derive some analytical results for the forecast performance of the robust methods relative to a benchmark using the full sample. For random breaks in a simple model we obtain rankings, but not under deterministic breaks. Clearly, it is hard to draw theoretical conclusions. So we experiment with 'Monte Carlo' simulations (creating many randomly drawn artificial data sets) for a variety of cases. The best methods can vary widely according to the particular break and choice of parameters. With the monitoring method we find the gains are small, although equally the costs (in cases where there are small breaks) are also small. Other methods can do much better where there are large breaks. The results make it hard to recommend a single method. But a method based on averaging over many different samples often improves on the full-sample benchmark and rarely comes with a large penalty where there are frequent or small breaks.

Finally, we take the methods to real data. We examine simple forecasting models using about 200 US and UK time series. For the United Kingdom, where there are relatively many breaks identified in the full sample, the best-performing method is forecast averaging, consistent with the Monte Carlo results.

We conclude that monitoring for breaks will not lead to a deterioration in forecast performance relative to using the full sample, but not much benefit either. Instead methods that discount past data in various ways are to be preferred. The averaging method we explore seems to be a useful default choice.

# Extracting information from structured credit markets

## Summary of Working Paper no. 407 Joseph Noss

Assessing the stability of an economy frequently involves assessing the risk of bad states of the world materialising. It is often necessary to judge how many firms are likely to default on their debt obligations over a certain time horizon. The likelihood of a large number of firms defaulting is of particular interest to policymakers, particularly if this is caused by some 'systemic shock' that presents a particular threat to financial stability.

Structured credit instruments are created by collecting defaultable assets, such as mortgages or corporate bonds, into portfolios and issuing claims of different seniority against these portfolios. Claims' seniorities determine the order in which they receive cash flows from the underlying assets, with more senior claims being paid first. Their prices therefore reflect market perceptions about the chance of these cash flows materialising, or equivalently, the likely extent of defaults of the underlying credit instruments. While the values of standard credit instruments, such as corporate bonds, offer an insight into the market-perceived probability of a given firm defaulting, the values of structured credit instruments provide a richer view of the likely extent of corporate defaults away from this central case. Claims of different seniorities incur loss only if defaults reach different magnitudes; their relative value therefore affords an insight into the likelihood of losses being of different severities.

Information can be recovered from the prices of structured credit by modelling the default of the different underlying credit instruments and then fitting the resulting modelled prices to those observed in the market. Correctly modelling the distribution of defaults, and in particular their codependence, is crucial in order to find a model whose tranche premia fit those traded in the market. For example, only if a large number of firms default together will senior claims incur loss. Previous attempts to model this interdependence have used a 'Gaussian copula model', based on the Gaussian or normal distribution, to capture the correlation between firms' defaults. However, this gives insufficient weight to the 'tail event' of multiple firms defaulting together.

The framework presented here instead uses a gamma distribution that is more able to capture the possibility of extreme dependence between defaults. It is therefore more successful in matching the traded prices of structured credit products. The model is also extended to include 'catastrophe' and 'becalmed' states that represent the possibility of very high degrees of systemic risk in credit markets, and its reduction perhaps due to government intervention; it therefore offers an intuitive explanation for the large fluctuations in codependence witnessed during the recent credit crisis.

This work offers three key outputs. First, it allows the market-implied probability distribution of firms' defaults to be inferred from the traded value of structured credit instruments. These distributions may be of use to policymakers, particularly because they offer an insight into the risk of 'tail outcomes' involving the default of large numbers of firms. This is likely to be of particular interest to policymakers seeking to measure and mitigate systemic risk. Second, the model offers an insight into the nature and magnitude of the risks firms face. It allows the average probability of a firm defaulting to be decomposed into components relating to default events of different severities. For example, it can estimate how the probability of a particular firm defaulting depends on the likelihood of a very severe event such as widespread financial crisis. Finally, in common with other models of structured credit that go beyond the Gaussian copula, this work is of potential use to those who trade structured credit products. It gives rise to a set of parameters that determine the structure of the codependence of default between credits, which could form the basis of an investor's 'hedging strategy' that allows positions in different tranches to be hedged against each other. This has the potential to protect them from changes in the nature of default codependence that reduce the value of their portfolio.