The determinants of household debt and balance sheets in the United Kingdom

Working Paper no. 266

Merxe Tudela and Garry Young

The outstanding debt of the UK household sector moved above £1,000 billion in 2004, equivalent to around 140% of household income (compared with around 105% ten years earlier). The rapid accumulation of debt has raised questions about the ability of people to repay what they owe, especially in the event of a sudden change in economic circumstances. This could have implications for both monetary policy, if the combination of high debt levels and a worsening economic outlook were to cause a slowdown in spending by households, and financial stability, if an increasing number of households were to default on their debts. It is therefore important to understand what lies behind the increase in debt and to assess its future sustainability.

Debt sustainability cannot satisfactorily be addressed by looking at the aggregate balance sheet of the household sector alone. There are substantial differences across households and shocks to the household sector are likely to affect different households in different ways. This paper proposes a framework for understanding aggregate indebtedness in terms of individual optimising decisions and adopts a model to explain the rise in borrowing. The model is set up to be consistent with the aggregate, cross-sectional and cohort experience of British households using information from the British Household Panel Survey. This process of calibrating the model reveals some inconsistencies between the basic life-cycle model of household behaviour used here and what is observed in practice. In particular, the level of debt is lower than expected at both extremes of the age spectrum. We therefore modify the basic model so that it can account for the observed cross-sectional balance sheet position of British households.

The model may be used to look at how balance sheets might develop in the future, on the assumption that it adequately captures current and future household behaviour and dependent on future trends in its determining factors such as real interest rates, house prices and incomes. This can be used as means of assessing the 'sustainability' of recent high debt levels. Sustainability of debt can be judged in two ways: whether debt will remain at or above current levels; and whether it is affordable. On the first test, this depends critically on the expected path for key determining variables. The paper shows that different future paths for the real interest rates could lead to a higher or lower debt-income ratio, suggesting that sustainability can only be assessed conditional on a view of how these determining factors are likely to develop. In neither case, however, do recent debt levels look unaffordable to the typical individual. Even if real interest rates were to revert to the higher levels seen in the late 1990s, the future consumption of even the most indebted cohorts would exceed that enjoyed by older cohorts today, reflecting the impact of past and future economic growth. Of course, the emergence of unexpected shocks would have an adverse impact on households. We have illustrated the effect of higher interest rates, lower house prices and lower pension incomes. All would cause a contraction in household spending and change the equilibrium debt-income ratio. The more severe the shock the more likely that the sustainability of debt would become an issue. While we are unable to assess the likelihood of such shocks with the current model. it is nevertheless a useful tool for assessing the severity of their impact.

Bank loans versus bond finance: implications for sovereign debtors

Working Paper no. 267

Misa Tanaka

Since the 1990s, syndicated bank lending to emerging market sovereigns has declined steadily, while eurobond issuance has increased. This paper tries to explain why these countries have recently shifted towards bond finance and considers the implications.

In this model, sovereigns' incentive to repay their debt arises from their desire to avoid a financial crisis which could be triggered by a default. Sovereigns have different risk characteristics, and the information about their creditworthiness can only be obtained through costly monitoring. Whereas banks can monitor their borrowers directly, the cost of monitoring is too high for small individual bondholders. But sovereigns wishing to issue bonds can hire a credit rating agency to monitor them and publish its assessment. Therefore, the critical difference between bank lending and bond finance is that banks act as private monitors and keep their assessment of the borrower private, whereas rating agencies act as public monitors and disseminate this information not only to the existing bondholders but also to third parties — ie potential future creditors. Consequently, bank loans are non-transferable whereas public monitoring makes bonds transferable by eliminating the information asymmetry between the existing creditors and potential future creditors.

When the timing of cash flow is uncertain, borrowers prefer long-term financing because short-term credit entails a risk of interim debt restructuring and crisis. Transferability makes bonds cheaper for long-term financing compared to bank loans, given that it is costly for banks to commit to holding a claim for multiple periods. Thus, when the cost of information dissemination is low and crisis costs are large, borrowers issue long-term bonds for financing projects with uncertain timing of cash flows, and use bank loans only for financing strictly short-term projects.

Our analysis shows that there are two inefficiencies in the current international financial system which is dominated by long-term bond financing. First, although the possibility of a financial crisis is necessary to prevent strategic defaults, it is *ex post* a deadweight cost if a default is unavoidable. Second, long-term bond issuers are subject to moral hazard, because the fear of a financial crisis prevents them from restructuring their unsustainable debt at an early stage. We demonstrate that state-contingent debt and IMF intervention to prevent a crisis conditional on the restructuring of an unsustainable debt are both welfare improving.

Forecasting using Bayesian and information theoretic model averaging: an application to UK inflation

Working Paper no. 268

George Kapetanios, Vincent Labhard and Simon Price

Recently, there has been increasing interest in forecasting methods that utilise large data sets. There is a huge quantity of information available in the economic arena which might be useful for forecasting, but standard econometric techniques are not well suited to extract this. In an effort to assist in this task, econometricians began assembling large macroeconomic data sets and devising ways of forecasting with them. Standard regression techniques cannot be used in this context, as the number of variables is far too large. Instead, broadly speaking there are two methodologies that can be applied: factor modelling and forecast combination. In the former, a factor structure is imposed on the data and then techniques such as principal components are used to extract the factors that are subsequently used in forecasts. This approach has been widely used in macroeconomic forecasting in recent years.

The alternative methodology is forecast combining, often of simple and probably misspecified models. This grew out of the observation by forecast practitioners in the 1960s that combining forecasts (initially by simple averaging) produced a forecast superior to any single forecast. If it were possible to identify the correctly specified model and the data generating process (DGP) is unchanging, then this approach would not be sensible. However, models may be incomplete, in different ways; they employ different information sets. Forecasts might be biased, and biases can offset each other. Even if forecasts are unbiased, there will be covariances between forecasts which should be taken into account. Thus combining misspecified models may, and often will, improve the forecast.

Despite this, combining forecasts will not in general deliver the optimal forecast, while combining information will. Nevertheless, it may not be practicable to estimate the fully encompassing model, not least because the set of variables is vast. Thus we have a justification for combining forecasts. One could call this the frequentist misspecification case. It should be clear that in this context forecast combining is viewed as mainly a stop-gap measure that works in practice but would be surpassed by an appropriate model that addressed the underlying misspecification. A further practical problem is that with standard combining methods the forecast weights can only be reliably constructed for a relatively small number of models. Nevertheless, given that the true DGP may involve a vast number of variables, it is clear that forecast combination is a route into the combining of information, and this is how it is interpreted in the literature relating to large data sets.

Forecast combining can be also be interpreted in a Bayesian framework. Here it is assumed that there is a distribution of models. The basic problem, that a chosen model is not necessarily the correct one, can then be addressed in a variety of ways, one of which is Bayesian model averaging. A chosen model is simply the one with the best posterior odds, but posterior odds can be formed for all models under consideration and offer weights for forecast combinations.

There is an analagous frequentist information theoretic approach, on which we focus in this paper. Given we have a set of models, we can define relative model likelihood. Model weights within this framework have been suggested by Akaike in a series of papers. In practical terms such weights are easy to construct using standard information criteria. Our purpose, then, is to consider this way of model averaging as an alternative to Bayesian model averaging.

We address this in two ways. We first assess the performance of information theoretic and other model averaging techniques by means of a Monte Carlo study. We then examine how various schemes can perform in forecasting UK inflation. For this, we use a UK data set which emulates a well-known data set constructed by Stock and Watson for the United States. We find that model averaging techniques can be beneficial with the information theoretic weights performing very well. Our findings partly confirm that Bayesian model averaging can provide good inflation forecasts, but we find that the frequentist approach also works well, and dominates in a large subset of the cases we examine for UK data. It is unlikely that a single technique would be more useful than all others in all settings. Nevertheless, our work indicates that information theoretic model averaging provides a useful addition to the forecasting toolbox of macroeconomists. Indeed, we find that the information theoretic method is the most robust of those we examine.

This paper does not describe the way in which the Bank of England generates its forecasts. The findings in this paper pertain to a specific type of forecasting model which is only part of a much broader approach to forecasting applied at the Bank. The Bank does not use a single model to forecast inflation or other variables; instead it uses a 'suite' of many models ranging from purely theoretical through purely data driven to the Bank's macroeconometric model, the Bank of England Quarterly Model (BEQM). All these models are useful in a particular context: in no case will any one model provide a uniquely best forecast.

Accounting for the source of exchange rate movements: new evidence

Working Paper no. 269

Katie Farrant and Gert Peersman

Considerable research has previously been carried out to try to explain past movements in exchange rates. We examine this issue by estimating a structural vector autoregression, with sign restrictions, for the United Kingdom, the euro area, Japan and Canada versus the United States. The structural vector autoregression identifies not only demand, supply and monetary policy shocks, which may be important in explaining exchange rate movements, but also specific exchange rate shocks. These exchange rate shocks can be thought of primarily as movements in the exchange rate which are not explained by fundamentals. As far as we are aware, this is the first time that specific exchange rate shocks have been identified using sign restrictions, which is a much more general and less stringent approach than traditional identifying procedures.

We find that, while fundamentals have been important in explaining movements in exchange rates, there are also specific exchange rate shocks that have had a significant influence in determining exchange rate paths over time. This is in contrast to a number of other studies, which suggest that exchange rate movements can primarily be explained by demand shocks. Applying the traditional identifying strategy based on long-run restrictions to our data set, however, supports the findings of these other studies, suggesting that the identification strategy is important in determining the results.

A model of bank capital, lending and the macroeconomy: Basel I versus Basel II

Working Paper no. 270

Lea Zicchino

The process of reforming the 1988 Basel Accord, that started in 1999, has been motivated by the goal of more closely matching regulatory capital to the risk profile of banks' asset portfolios. The rationale for minimum capital requirements is that they mitigate financial institutions' moral hazard. Regulators are imposing a cost on bank owners to 'encourage' them to avoid costly default. However, the limited number of risk categories in the current framework has created opportunities for banks to increase the risk to which they are exposed without increasing the amount of regulatory capital.

The new Basel Accord is widely recognised as a much needed effort to deal with the shortcomings of the current system. By realigning capital adequacy rules with banks' incentives it aims at restoring the link between risk and capital holding. Nonetheless, a number of questions have been raised by central bankers, regulators and practitioners regarding the impact of a more risk-sensitive regulatory framework on macroeconomic stability. Among them, there is the issue of the potential procyclical effects of the new capital adequacy requirements, ie the possibility that during periods of weak economic growth, a fall in capital ratios and an increase in regulatory requirements implied by a deterioration in the risk profile of banks' assets might increase the likelihood of credit contraction and, therefore, a further weakening of growth.

This paper analyses the relationship between banks' capital holdings, banks' loans and macroeconomic activity under risk-sensitive capital adequacy requirements. In particular, it compares the impact of macroeconomic shocks on banks' choices of capital structure and loan supply under the old and new capital adequacy regimes. It does so by extending a model that investigates the impact of monetary policy on lending in an economy where banks operate in an oligopolistic market and are subject to minimum capital requirements. In order to analyse banks' reaction to changes in macroeconomic conditions under the new capital adequacy regime, I extend the model by assuming a link between loan risk-weights and borrowers' creditworthiness. In particular, I introduce asset risk-weights that vary with macroeconomic performance, which is a major determinant of credit risk.

The first result of the paper is that the response of banks to shocks that affect loan demand differs when the minimum capital requirements are calculated with asset risk-weights that are sensitive to macroeconomic conditions. In particular, bank capital is less volatile than under capital requirements with constant risk-weights. The intuition behind this result can be understood by considering, for example, a positive shock to macroeconomic conditions that increases both current and future loan demand. If the capital constraint is binding, banks may not be able to expand loan supply in the current period and they may need to raise capital to increase supply in the future. Therefore, if capital requirements do not change with borrowers' risk, capital increases in response to positive macroeconomic shocks and decreases after negative shocks. But when asset risk-weights depend on macroeconomic conditions, bank capital might not need to increase for banks to be able to expand their credit supply. In fact, following a positive macroeconomic shock the risk-weights decrease and the capital constraint thus become looser. This insight has an important policy implication. On the one hand banks will tend to operate above the minimum regulatory capital to avoid the capital constraint becoming binding in future periods. On the other hand banks may not voluntarily accumulate capital in times of good macroeconomic conditions because it is during these times that the capital constraint becomes looser. This means that if banks are affected by an adverse shock during a period of credit expansion, they might be forced to raise capital at a time when market conditions are unfavourable. A second and related result of the paper concerns the effect of macroeconomic shocks on loan supply. Since capital is more difficult to accumulate in a recession, and easier to accumulate when the economy experiences a positive shock, bank credit is likely to be more procyclical under the new Accord than under the current one.

Consumption, house prices and expectations

Working Paper no. 271

Orazio Attanasio, Laura Blow, Robert Hamilton and Andrew Leicester

Over much of the past 25 years, the cycles of house price and consumption growth have been closely synchronised. Three main hypotheses for this co-movement have been proposed in the literature. First, that an increase in house prices raises households' wealth, which increases their desired level of expenditure. Second, that house price growth increases the collateral available to homeowners, reducing credit constraints and thereby facilitating higher consumption. And third, that house prices and consumption have tended to be influenced by common factors (eg productivity growth or tax changes), which cause revisions to households' expected lifetime income. This paper uses individual household level data to assess the importance of these different hypotheses. Revisiting this link seems particularly timely, as the housing market has cooled since the end of 2004, generating widespread press speculation about the outlook for prices. In addition, there is the puzzle, discussed in a box in the Bank of England's Inflation Report in November 2004, about the recent decline in the correlation between house price and consumption growth, and hence the likely impact of house prices on consumption in the future.

Many previous related studies have focused on the late 1980s consumption and house price booms. Attanasio and Weber recognised that microeconomic data on individual households' expenditure provides a way to distinguish between the competing wealth and common causality hypotheses. If wealth effects were important, older homeowners - who are less likely to demand more housing services in the future — should be the primary beneficiaries of a house price boom and should increase their consumption the most. In contrast, if house prices and consumption are both influenced by common expectations of income growth, younger consumers, with a greater remaining lifespan to realise the gain, should be the ones to raise consumption the most. Their paper argued that common causality was the more likely explanation for the late 1980s correlation. But since then, many other studies, mainly relying on aggregate data, have argued that there is a direct wealth effect.

This paper extends and updates Attanasio and Weber's results, covering data spanning the consumption and house price weakness of the early 1990s, and developments up to and including 2001. We estimate various specifications for individual households' consumption using pseudo-cohorts drawn from 24 years of the Family Expenditure Survey between 1978 and 2001/02. In our baseline specification, the consumption of a household in a given year depends on the cohort to which it belongs, the age of the head of household and various other demographic and household characteristics. We then assess the extent to which adding various house price terms to our baseline model can help explain the consumption patterns over time. By analysing the results for households in different age groups, we determine whether house price movements appear to be a more important determinant of the consumption of younger or older households, or of renters or homeowners, using similar identifying assumptions to those previously used by Attanasio and Weber.

We find several pieces of evidence which suggest that common causality has been the most significant explanation for the co-movements between house price and consumption growth. First, younger cohorts had the largest swings in expenditure during the consumption and housing cycles. Second, the effect of regional house price growth on consumption is found to be stronger for these younger households. Third, the coefficient on the regional level of house prices is as large for younger as for older households, while they had a greater response to the effect of 'unexpected' house price movements. And fourth, the consumption of both homeowners and renters are equally aligned with the house price cycle. Of course, it remains likely that the wealth and collateral channels are important for some households at some points in time. But the evidence in this paper suggests that the main reason for house prices and consumption being correlated in the past is changes in common driving factors — like income expectations.

What caused the early millennium slowdown? Evidence based on vector autoregressions

Working Paper no. 272

Gert Peersman

This paper analyses the underlying causes of the recent slowdown and preceding expansion for the industrialised world (proxied by an aggregate of 17 countries), the United States and the euro area. In order to do the analysis, vector autoregressions (VARs) are estimated for the sample period 1980 Q1–2002 Q2 containing output, inflation, interest rates and oil prices. The impact of aggregate supply, aggregate demand, monetary policy and oil price shocks is estimated.

A crucial problem when using VARs is the identification of the structural shocks. We compare the results of two identification strategies. The first one is based on conventional zero contemporaneous and long-run restrictions. Specifically, a number of restrictions are imposed on the immediate impact of a shock on certain variables (for instance, allowing no immediate effect of monetary policy on output) or on the long-run effects of specific shocks (for instance, ensuring the long-run neutrality of monetary policy). These restrictions are, however, very stringent in many cases. Short-run restrictions are typically not based on theoretical considerations, and long-run restrictions can be highly misleading. We therefore propose an identification strategy based on more recent sign restrictions as an alternative (for example, after a restrictive monetary policy shock, the sign of the output reaction is not positive). Hitherto, this type of restriction has only been used to identify monetary policy shocks. We extend this method to our larger set of structural shocks. The advantage of this procedure is that we do not have to impose strong and perhaps implausible constraints. By contrast, our alternative approach only makes explicit

use of restrictions that researchers often use implicitly. Often, researchers experiment with the model specification until the results look reasonable; for example, a restrictive monetary policy shock is expected to have a negative impact on prices and a temporary effect on output. This *a priori* theorising is made more explicit with sign restrictions, and at the same time, no additional short and long-run conditions are necessary. As a result, this approach is much more general.

We show that the identification strategy is indeed important, in particular for oil prices and monetary policy shocks. The difference between both approaches is statistically and economically very important. After a restrictive monetary policy shock, the maximum impact on output is -0.3% with conventional restrictions, whilst the impact is estimated to be between -0.4% and -1.0% with sign constraints.

When applying both methods on recent output fluctuations, we find that the recent slowdown was caused by a combination of several shocks. Across both methodologies, we find an important role for negative aggregate spending shocks. In addition, there were negative aggregate supply shocks, negative effects of restrictive monetary policy in 2000 and a negative impact of oil price increases in 1999. The magnitude of the latter two is significantly different between both approaches. We find an important role for oil price shocks with conventional restrictions and for monetary policy shocks using sign conditions. The shocks are also more pronounced in the United States than in the euro area.

'Real-world' mortgages, consumption volatility and the low inflation environment

Working Paper no. 273

Sebastian Barnes and Gregory Thwaites

This paper considers the interaction between the microeconomic decisions facing households and the macroeconomic environment in a setting where households have 'real-world' mortgage contracts. In particular, we consider the possible consequences of the important changes in the framework for setting monetary policy in the United Kingdom in recent decades.

The change in monetary policy regime from the 1980s to the 1990s has been associated with greater stability of the macroeconomic environment. 'Real-world' mortgages may provide an explanation of how more stable economic conditions have contributed to reducing the volatility of aggregate consumption through effects not captured in elementary textbook models of consumption with debt. In these models, it is typically assumed that household borrowing takes the form of successive one-period debt contracts, denominated in units of consumption. Actual mortgage contracts — the biggest financial commitments that most households ever make — look very different to this: they are denominated in nominal terms with repayments over many periods, sometimes with fixed nominal interest rates. This paper is concerned with the role of such real-world mortgage contracts in consumption volatility.

We use a model of real-world mortgages to show the effects at household level of the change in monetary policy regime under adjustable-rate and fixed-rate mortgages. We use this to model aggregate consumption uncertainty in a partial equilibrium overlapping generations framework. At household level, we find that non-housing consumption would be smoother over the life cycle in the more stable 1990s regime. The change of regime generates substantial welfare gains for mortgage holders. Even though households now have more mortgage debt than in the past, we find that households could still enjoy similar levels of utility from non-housing consumption in the 1990s as in the 1980s regime. This suggests that households may have increased their demand for housing in response to the lower cost and greater certainty of mortgage borrowing in the 1990s.

The main parameterisation of the model suggests, counterintuitively, that aggregate consumption volatility under the 1990s regime would actually be higher in the steady state than in the 1980s regime, other things being equal. Although macroeconomic shocks have become less pronounced, this result suggests that households' responses to shocks have become more synchronised.

Furthermore, higher indebtedness in the 1990s has also tended to make aggregate consumption less stable. This result shows how the more stable economic environment associated with the 1990s regime would not necessarily translate into greater stability of aggregate consumption given real-world mortgages. If the assumptions necessary for this result hold, the observed fall in aggregate consumption volatility in the 1990s would either have to be explained by other offsetting factors or because the economy was in a period of transition between two regimes rather than the new steady state; households in the 1990s actually benefited simultaneously from the more stable macroeconomic environment, and the lower levels of indebtedness inherited from the past.

The substitution of bank for non-bank corporate finance: evidence for the United Kingdom

Working Paper no. 274

Ursel Baumann, Glenn Hoggarth and Darren Pain

The aim of this paper is to investigate empirically the links between alternative forms of corporate debt finance using data on the UK economy. Based on a small panel data set of UK-owned banks for the 1986 Q3-2001 Q3 period, we estimate equations for the quantity of bank credit to the corporate sector. In particular, we investigate the extent to which changes in non-bank finance — either from (bond and other debt securities) markets or from non-bank financial institutions — affect the growth in corporate loans of UK-owned banks. In doing so, we aim to investigate the degree of substitutability or complementarity between bank and non-bank finance. Moreover, we examine whether these relationships are different in periods when non-bank finance falls sharply to assess whether bank credit acts as a back-up source of funding when other forms of finance are not readily available.

In order to understand the potential interaction between bank and non-bank markets, an important distinction relates to the separate influences of supply and demand factors. But there is an identification issue: observed changes in corporate bank and non-bank finance will reflect movements in both the supply of and demand for external funds and it is difficult to disentangle the two. To address this issue, we exploit information on the average interest rates banks charge on their corporate loan portfolios. By considering how these loan rates respond to developments in non-bank finance markets in conjunction with the changes in the amount of credit extended, we hope to throw light on whether supply or demand influences are more important, particularly during periods of stress in non-bank finance.

Our results suggest that there is substitutability for companies between bond finance and bank loans from

the large UK-owned banks. In particular, the growth in bank lending of the major UK-owned banks increases around some periods of bond market stress as well as during more tranquil periods when bond spreads widen. In general, the loan rates of the large UK banks are not found to be sensitive to changes in non-bank finance. This could reflect a relatively flat loan supply curve whereby banks increase the amount of credit extended when, for example, bond spreads rise substantially without increasing their loan rates. This would be consistent with firms using their arranged loan facilities with banks to absorb shocks in the availability of other forms of external finance. In this way, banks may passively accommodate shifts in the demand for bank loans that are associated with disturbances in non-bank finance.

However, there are some variations in the results for different forms of non-bank finance. This suggests that banks' responses may depend on the nature of the shock. In periods when bond spreads widen sharply, bank loans would seem to provide alternative finance for corporates, at largely unchanged interest rates. This would be indicative of companies switching their demand for external finance away from capital market financing to bank loans, and is consistent with the notion of substitutability between alternative forms of finance. However, disruptions to the amount of corporate bond and commercial paper issuance seem to be associated with an increase in loan rates and either a fall or unchanged bank lending growth. This appears to be consistent with higher corporate demand for bank finance being choked off by a decline in loan supply by banks.