Health, disability insurance and labour force participation

Working Paper no. 218

Brian Bell and James Smith

The 1990s witnessed a rapid expansion of employment in the United Kingdom and an associated decline in unemployment to levels last seen in the 1970s. Yet over the same period, the aggregate participation rate was flat. This aggregate picture masks diverging trends in the activity rates of the sexes: female participation continued to increase, but over half a million men of working age left the labour market. Moreover, the decline in labour force participation was most pronounced among prime-age men (aged between 25 and 54), with early-retirement trends explaining very little of the change.

In this paper we focus on two important features in the data on rising male inactivity. First, the overall rise for men was accompanied by a rise in the numbers saying that they were too ill to work. The participation rate of prime-age males fell by a mere 0.7 percentage points between 1971 and 1989, but fell by 2.9 percentage points over the course of the 1990s. A feature of this fall was an increase in the number of those who cited health reasons for their inactivity. Many of these men also claimed disability benefits. This suggests that any explanation for declining male participation needs to address the rise in the inactivity among the long-term ill and the higher incidence of those claiming disability benefits. Second, the decline in male labour force participation was more pronounced among those with little or no formal qualifications. For those males who left school with no qualifications, the participation rate dropped almost 13 percentage points over the course of the 1990s.

One explanation that has been suggested for these trends is that a deterioration in the labour market opportunities for the low skilled coincided with increasing generosity of disability benefits, producing incentives for these workers to drop out of the labour market. However, testing such a hypothesis is not straightforward. The incentive for workers to drop out of the labour market and claim disability benefits depends upon the relative pay-off from looking for work. This makes it difficult to estimate the effect of benefits on individuals' labour supply decisions, as variation in benefits is driven primarily by differences in earnings. Since workers' earnings are likely to be highly correlated with taste for work, it is difficult to isolate the behavioural effects of disability benefits from these taste differences.

One way to get around this problem is to use a 'natural experiment', exploiting variation in benefit levels that is unrelated to tastes for work. Such variation in benefits occurred in 1995, when the UK disability benefits program was reformed. Prior to 1995, those claiming disability benefits received an Additional Pension (AP) based on earnings history: people becoming sick were entitled to higher amounts of AP depending upon earnings. After 1995, new cohorts lost entitlement to AP. This reduced the value of the disability program to new cohorts of older men, but left younger men—with only a short earnings history—largely unaffected. We exploit the resultant variation in benefit levels to estimate labour supply elasticities.

Using this approach we obtain significant positive effects from benefits on labour supply. The elasticities are particularly large for the least educated males. These results support the hypothesis that relatively generous disability benefits encouraged the early accommodation of health problems for those males who were most at risk of job loss.

The participation rate is a key determinant of sustainable supply capacity. Therefore future inflation will depend on whether or not the trends seen over the 1990s continue. So what does our analysis suggest? As entry into the disability benefits system tends to be a decision that results in permanent exit from the labour market, it seems unlikely that future demand shocks will generate similar-size flows out of the labour market. There are two reasons for this. First, future shocks may not have the same skill characteristics as those observed over the previous two decades. Second, the generosity of disability benefits has fallen significantly since the recession of the early 1990s. Hence the pull-factor of disability benefits has been reduced and workers are more likely to remain within the labour market following job loss.

Female labour force participation in the United Kingdom: evolving characteristics or changing behaviour?

Working Paper no. 221

Maria Gutiérrez-Domènech and Brian Bell

The period 1984–2002 was characterised by a substantial increase in female participation rates in the United Kingdom, whereas the opposite trend was observed for males. Understanding participation trends in the United Kingdom is important for monetary policy, since participation affects the supply capacity of the economy. The balance between that capacity and aggregate demand in turn affects inflationary outturns.

Because males and females have experienced such different participation trends, it is necessary to analyse them separately. This paper focuses mainly on females and uses an accounting framework to quantify how much of the rise in female participation was related to changes in the characteristics of the female population, and how much was linked to changes in behaviour and other uncontrolled factors. This exercise suggests that two thirds of the growth in female participation over 1984–2002 was associated with changes in the socio-demographics of the female population, especially education and fertility. As these two variables may be endogenous to participation, we cannot say anything about causality. The remaining one third of the rise in female participation was linked to changes in behaviour, such as women with the same observable characteristics responding differently over the period, and/or driven by changes in other variables not accounted for in the model.

Most of the increase in female participation between 1984 and 2002 took place in 1984–92, when both characteristics and 'behaviour' contributed positively to participation growth. In the 1980s, changes in behaviour contributed significantly to participation growth. The majority of the increase over the 1990s was driven by changes in the characteristics of the female population, whereas the slowdown in participation growth was mainly accounted for by a reversal of the behavioural effects.

The paper uses the same method to analyse briefly the evolution of male participation. The data reveal that the decline in male participation was mainly driven by changes in behaviour, especially after 1993.

The roles of expected profitability, Tobin's Q and cash flow in econometric models of company investment

Working Paper no. 222

Stephen Bond, Alexander Klemm, Rain Newton-Smith, Murtaza Syed and Gertjan Vlieghe

Econometric models of company investment face the problem that current investment decisions depend on expectations of future conditions, but these expectations are generally not observed. This makes it difficult to know whether significant coefficients on financial variables, such as cash flow, in empirical investment equations indicate the importance of financing constraints, or whether these variables simply provide additional relevant information about current expectations of future profitability. In this paper we construct explicit measures of expectations of future profitability for UK firms to address this question.

The *Q* model of investment relates investment to the firm's stock market valuation, which is meant to reflect the present discounted value of expected future profits. Under certain assumptions about the firm's technology and competitive environment, the ratio of the stock market value of the firm to its replacement cost (Tobin's *Q*) should be a sufficient statistic for investment. Significant coefficients on cash-flow variables after controlling for Tobin's *Q* can then not be attributed to additional information about current expectations. However, if the above conditions are not satisfied, or if stock market valuations are influenced by 'bubbles' or any factors other than the present discounted value of expected future profits, then Tobin's *Q* would not capture all relevant information about the expected future profitability of current investment. In this case additional explanatory variables, like current or lagged sales or cash-flow terms, could proxy for the missing information about expected future conditions.

This problem is particularly important in the literature that tests for an impact of financing constraints or capital market imperfections on corporate investment. Many empirical studies have added cash-flow variables to empirical models that relate investment rates to Tobin's *Q*, and interpreted significant coefficients on these cash-flow terms as evidence of 'excess sensitivity' of investment to the availability of internal funds. Although these findings are consistent with the presence of a cost premium for external sources of investment finance, they may also be explained, in the absence of financing constraints, by observed cash-flow or profits variables containing additional relevant information about expected future profitability not captured by Tobin's *Q*.

Recent findings for US data suggest that much, if not all, of the significance of cash-flow variables in conventional estimates of Tobin's *Q* investment equations can be attributed to the failure of Tobin's *Q* to capture all relevant information about the expected profitability of current investment. Previous studies using UK company data have reported significant coefficients on cash-flow variables, both in the context of models that relate investment to Tobin's *Q*, and in the context of reduced-form empirical models without explicitly forward-looking controls for expected profitability. The aim of the present study is to consider the robustness of these findings to alternative controls for expected future profitability. We obtain data on earnings forecasts from IBES International for around 700 publicly traded UK companies between 1987 and 2000. We match this information with stock market valuations and company accounts data on investment, cash flow and other financial variables obtained from Datastream International. Our main finding is that, whereas lagged cash flow is highly significant conditional on a standard measure of Tobin's *Q*, the coefficient on this cash-flow variable becomes insignificantly different from zero when we include our direct measures of expected future profitability. This parallels the results found for US data by other researchers. We also examine subsamples of firms, and find that the results are robust across subsamples of smaller firms and low-dividend firms.

Although cash-flow variables become insignificant when we control for expected profitability in this way, we find positive coefficients on both sales growth and cash-stock variables that remain statistically significant after conditioning on our measures of expected profits. These additional variables could be capturing expectations of profitability in the longer term that are not captured by our explicit measure of expectations. These longer-term expectations would be relevant for explaining investment rates under the maintained structure of the *Q* model. Alternatively, our findings could reflect misspecifications of the basic *Q* model, such as market power, decreasing returns to scale, or non-convex components of adjustment costs. In principle, the significance of these additional variables could also be due to the presence of financing constraints, although our results for subsamples do not suggest that this is a likely explanation. The coefficients on the additional sales growth and cash-stock terms appear to be broadly similar between subsamples of firms that have been considered elsewhere to be more or less likely to be subject to significant financing constraints. So the additional information these variables provide appears more likely to be explained by more general features of the investment behaviour of UK firms.

Real exchange rates and the relative prices of non-traded and traded goods: an empirical analysis

Working Paper no. 223

Jan J J Groen and Clare Lombardelli

Empirical real exchange rate studies mainly reflect one of two views of real exchange rate behaviour. Purchasing power parity (PPP) assumes that any measure of the real exchange rate is mean-reverting in nature and therefore constant in the long run. An alternative hypothesis makes a distinction between the empirical behaviour of the tradable and non-tradable components of the real exchange rate. This approach assumes that cross-country differences in the prices of tradable goods expressed in the same currency should eventually be eliminated, that is the Law of One Price (LOOP) across tradable goods between countries holds. In this case, the long-run movements in real exchange rates are related to movements in the ratio of the relative price of non-tradable and tradable goods between countries.

Based on evidence in the literature it seems sensible to assume that the real exchange rate contains a unit root. We carry out unit root tests on the data, which show this assumption is appropriate. Although this phenomenon is not consistent with PPP, it can be reconciled with the second approach; that national price indices have non-tradable components, which in turn affects real exchange rate behaviour. In this context, short to medium-run deviations between the real exchange rate and the ratio of the relative price of non-tradable and tradable components are possible. These occur as a consequence of temporary deviations from the LOOP. Hence LOOP deviations can only dominate the variability of the real exchange rate in the short to medium run.

In this paper we test this hypothesis for movements in UK real exchange rates relative to a sample of six main OECD partners. The identification of a long-run relationship between the real exchange rate and the ratio of the relative price of the non-tradable and tradable components requires us to choose a method for constructing these components. Determining precise indices that accurately capture the price of traded and non-traded goods is virtually impossible. Given these inevitable constraints we use two different methods to construct indices to capture movements in the prices of traded and non-traded goods in each country in our sample. One method decomposes the consumer prices index into its tradable and non-tradable components; the other uses the producer prices index as a proxy for tradable goods prices.

The analysis presented examines the existence of a long-run relationship between bilateral UK real exchange rates and the corresponding relative prices of non-traded to traded goods. Consistent with the findings elsewhere in the literature, using cointegrated vector autoregressive (VAR) models for these series, otherwise known as vector error correction (VEC) models, we find little support for the LOOP; there is only limited evidence for a cointegrating relationship in the dollar and euro bilateral rates. Using an autoregressive model for the relative price of tradable goods, we quantify the severity of the deviations from the law of one price. This provides evidence that such deviations are persistent relative to the time span of our data set. This finding motivates the use of a multi-country panel cointegration-testing framework. It provides evidence for a cointegrating relationship between the real exchange rate and the relative price of non-tradable goods for the United Kingdom, using both the CPI and the PPI-based decompositions.

Out-of-sample evaluation shows that the estimated time series based cointegrating VAR models are inferior to a naive random-walk model. But we find evidence that a novel panel VEC approach can, for most bilaterals, provide a significantly more accurate prediction of movements in the real exchange rate than a random-walk model. Our results show that by using a panel-data framework we are able to identify a long-run relationship between bilateral UK real exchange rates and the corresponding relative prices of non-traded to traded goods.

The informational content of empirical measures of real interest rate and output gaps for the United Kingdom

Working Paper no. 224

Jens D J Larsen and Jack McKeown

Comparing short-term nominal interest rates with some benchmark level can provide a measure of policy stance and hence may provide an indication of whether inflation will rise or fall in the future. One such benchmark is the natural rate of interest. In this paper, the unobservable natural real rate is estimated and the leading indicator properties of the real rate gap—the difference between the estimated actual real rate and the estimated natural rate—for inflation over the past 40 years are assessed. The estimates of the natural rate of interest in this paper can be interpreted as being like an intercept in a policy rule: so a real rate gap of zero—that is setting actual real rates equal to the natural rate—is consistent with an output gap of zero and with stable inflation in the medium term, while a negative (positive) real rate gap is consistent with a positive (negative) output gap and with rising (falling) inflation.

Because the natural rate of interest is unobservable there are a variety of possible approaches to obtain estimates of it. At one end of the spectrum, one could use a fully specified dynamic general equilibrium model. The main advantage of this approach is that the estimates of the natural rate—and other elements in the model—could be given full structural interpretation. However, where such models have been log-linearised around a non-stochastic steady state they cannot be used to make inferences about low-frequency movements in the natural rate of interest, because the long-run natural rate, by construction, is constant. At the opposite end of the spectrum of possible approaches one could use long-maturity index-linked bond yields or simple filtering. This approach has the advantage that it does not require any estimation or modelling. However, such an approach would not allow a structural interpretation of the estimates, and they could not be construed as a direct guide to monetary policy. The approach taken in this paper lies somewhere in between these two extremes.

Here, the natural rate of interest is estimated using Kalman filtering techniques in a small semi-structural model of the UK economy. Because these estimates are obtained using a semi-structural approach, they can be interpreted as economically meaningful, so they are preferable to estimates obtained from bond yields or simple filtering. On the other hand, the combination of a relatively sparse theoretical structure with a data-driven filtering approach allows for low-frequency movements in the natural rate of interest and means that the estimated levels of the natural rate are not tied to some calibrated long-run value.

The paper provides estimates of expected inflation along with estimates of real rate gaps, output gaps and unemployment gaps, which all appear broadly plausible. The estimates of the real rate gap are found to have had leading indicator properties for both the estimated output gap and inflation over the sample as a whole. However, the paper also finds evidence of substantial variation in the indicator properties over time. Breaking the sample into four subsamples it appears that the leading indicator properties for both the output and real rate gap were substantially stronger for the subsample that covers most of the 1980s. After the introduction of the inflation target, post 1992, the relationship between the real interest rate gap and the output gap strengthens, but the leading indicator properties of the estimated gaps for inflation diminish, as might be expected under an inflation-targeting regime.

Exploring the relationship between credit spreads and default probabilities

Working Paper no. 225

Mark J Manning

Theoretically, changes in the yield spread between risky and risk-free bonds should reflect changing expectations about the likelihood of loss from default, which will itself be determined by variability in the probability of default and expected recovery. Our principal interest in this paper is to explore the extent to which variability in sterling corporate credit spreads corresponds to the theory, drawing, in particular, upon the predictions of a structural (Merton-style) model of corporate failure. Although credit spreads are often cited as indicators of such expectations, the empirical literature has found little evidence that idiosyncratic default risk is the principal driver of variability in credit spreads.

The recent empirical literature has generally adopted one of two approaches to examine the validity of structural models of default. Researchers have either compared actual credit spreads with those implied by a fully calibrated structural model or else they have regressed changes in spreads upon a reduced form of the model. In this study we take a different approach. First, we adopt an error-correction method in order to capture both the long-run relationship between spreads and default probabilities, and short-run deviations from trend. Second, while analysis of a reduced form of the structural model allows the key relationships to be identified, the non-linear interactions between the model inputs are not exploited. Hence, some of the power of the model is lost. In this paper, therefore, we apply the structural approach more directly, employing a Merton-style model, developed at the Bank of England, to generate a panel of implied default probabilities. Finally, much of the previous work in this area has drawn upon data from the US non-government bond market. Our work, by contrast, employs a sample of 78 sterling bond issues by 42 UK industrial companies. For each, we have up to 83 monthly observations for both asset swap credit spreads and Merton-generated implied probabilities of default, thereby creating a diverse data set, covering a segment of the market that, to our knowledge, has not previously been studied in this way.

The application of this approach is revealing. In a pooled regression, we find that variability in the implied probability of default can explain just 8% of the probability of default in the highest quality credit spreads (AAA/AA), and 11% of that in A-rated credits. With the probability of default for these issuers generally low, and often lacking variability, the relative importance of systematic factors tends to increase. Indeed, we find that the addition of time dummies to the specification increases explanatory power considerably, perhaps reflecting the influence of common factors such as liquidity conditions not explicitly included in the specification. Our results for lower investment-grade issues, those rated BBB, are more supportive of the structural model. Here we find that the probability of default explains around a third of the variation in credit spreads in a pooled regression, which is higher than previous empirical studies have found.

Comparison with a broadly equivalent specification to that employed elsewhere suggests that this is a reflection of the more direct application of the Merton approach; in particular, capturing the non-linearity inherent in the structural model, which is most important for companies that are closer to the default point.

In a further round of tests, we allow for heterogeneity in responses across individual issues in the ratings subgroups. Heterogeneity does indeed appear to be an important feature of the data set, with explanatory power increasing to 28% for high-quality issues, and almost 50% for BBB issues. This argues in favour of not only applying the Merton model directly, but also allowing for potential idiosyncrasies in factors such as liquidity and recovery rates.

Finally, we consider whether we are losing valuable information in the annualisation process for our implied default probabilities. If investors have short horizons, they may place greater weight on near-term default probabilities, and this will perhaps be more important for lower-grade bond issues. This hypothesis is supported by the data. Returning to a pooled specification with common coefficients, but retaining differences across ratings, we find once again that almost half of the variability in BBB credit spreads is explained by the regression specification. Explanatory power remains at just 12% for high-quality issues.

Many of these results would appear to have an intuitive interpretation. Previous research has established that the theoretical relationship between credit spreads and default expectations does not hold fully in practice, and this paper concurs with that finding. Spreads would appear to be influenced by market factors, such as liquidity premia, and these are likely to be time varying. Thus, it is intuitive that, for high-quality issuers, where both the level and variability of the probability of default is likely to be lower, the relative contribution of default expectations is likely to be much smaller.

Corporate capital structure in the United Kingdom: determinants and adjustment

Working Paper no. 226

Philip Bunn and Garry Young

The balance sheet position of non-financial companies goes through phases of strength and weakness. At the end of 2003 the amount of debt on corporate balance sheets was at a historically high level in relation to the market value of the capital that ultimately provides the means by which the debt is serviced. Past patterns would suggest that such high gearing situations do not persist and that companies act to bring down their indebtedness. This paper addresses the factors that determine the level of gearing that companies appear to aim for over time and what actions companies take to adjust when their debt gets out of line with their desired level.

Our analysis provides an empirical test of the 'trade-off' theory of corporate capital structure, which suggests that firms have an equilibrium level of capital gearing that is determined by trading off the advantages of holding debt against the expected costs of financial distress, which becomes more likely at high debt levels. We consider only the tax benefits of holding debt, since the other factors that make debt an attractive form of business finance are difficult to quantify. The tax advantage of debt arises from the deductibility of interest payments against corporation tax payments, but the magnitude of the benefit depends in a complex way on the personal tax rates faced by shareholders. This paper uses a theoretical model of corporate behaviour to derive an expression for the tax gains to corporate gearing, which we construct for the United Kingdom from 1970 onwards and use as a basis for the empirical part of the paper. The tax gains to gearing were high in the second half of the 1970s and the early 1980s when corporation tax rates were high, but they have fallen since and are currently at a historically low level.

We find evidence that firms in the United Kingdom have target levels of capital gearing at the aggregate level, which in the long run depend on the tax advantages of debt and on the probability of bankruptcy (which will be related to the expected costs of financial distress). This finding provides empirical support for the 'trade-off' theory of corporate capital structure, and it reinforces the results of previous firm-level work at the Bank which also found that firms have target levels of gearing. The current level of long-run equilibrium capital gearing at market value for the UK PNFC sector implied by our model is approximately 16%.

The paper then goes on to test how firms adjust their balance sheets to eliminate deviations in actual gearing from the implied equilibrium level. We find that most of the adjustment in response to above-equilibrium gearing takes place through reduced dividend payments and increased equity issuance. There is only weak evidence that firms adjust through more restrained capital investment. This is consistent with the 'new view' of corporate behaviour which suggests that real adjustment will only take place once dividends cannot be reduced any further. These findings are also consistent with firm-level work for the United Kingdom which has found evidence of adjustment in dividends and new equity issuance, with the proviso that investment appears to be more responsive to a flow measure of financial pressure than a stock measure of balance sheet disequilibrium.

Illustrative simulations show how firms may adjust their balance sheets in response to shocks that move gearing further away from its implied equilibrium. Although firms appear to respond quickly and make relatively large adjustments to the flows, the actual adjustment process is likely to be protracted because the flows of dividends, equity issuance and investment are all small in relation to the stock of debt.

The Phillips curve under state-dependent pricing

Working Paper no. 227

Hasan Bakhshi, Hashmat Khan and Barbara Rudolf

The Phillips curve has long served as a useful description of monetary policy effects on inflation. In modern New Keynesian models, it is explicitly derived from the pricing decisions of firms. One advantage of this new approach is that, because the relationship has a structural interpretation, we can, for example, infer implications for the transmission of inflation following a shock; the Phillips curve is no longer a 'black box'. But if there are structural changes in the economy, such as the move to a low-inflation environment witnessed since the 1990s in the United Kingdom and several other countries, the price-setting behaviour of firms may change and affect inflation dynamics. From a policy perspective, therefore, two important issues arise. First, how sensitive are short-term inflation dynamics to such shifts in the economic environment? Second, how well does a Phillips curve based on the assumption of unchanged price-setting behaviour of firms describe inflation dynamics of an economy where this assumption does not hold?

One approach to modelling firms' price-setting behaviour is to assume that firms choose their prices optimally, while the timing of their price changes is exogenous (time-dependent pricing). This approach underlies the New Keynesian Phillips curve (NKPC), which suggests that current inflation is determined by the expectation of next period's inflation and a measure of current economic activity. The time-dependent pricing assumption implies that firms may not adjust the time pattern of their price adjustments in response to changes in macroeconomic conditions. This is hardly plausible if we think of an environment with shifts in trend inflation, for example, and therefore it may limit the value of these models for monetary policy analysis. In response to this problem, approaches with an endogenous timing of price changes have been developed. These approaches allow the firms' time pattern of price changes to respond to the state of the economy (state-dependent pricing).

This paper derives a closed-form solution for short-term inflation using a state-dependent pricing model. The resulting equation is more general than the NKPC and it nests the latter as a special case. It relates inflation to lagged inflation, expected future inflation, and current and expected future real marginal costs. The number of leads and the coefficients are endogenous and depend on the level of steady-state inflation and on firms' beliefs about future adjustment costs associated with price changes. This structural equation is referred to in this paper as the state-dependent Phillips curve (SDPC).

In contrast to the NKPC, the SDPC allows lagged inflation terms to affect current inflation. This is an interesting feature since recent empirical evidence suggests that the NKPC extended by a lagged inflation term provides a better description of inflation dynamics than the purely forward-looking NKPC for several countries. In fact, specifications with lagged inflation terms have been derived before by several authors. But all these studies were based on the assumption of an exogenous timing of price changes. The SDPC, therefore, has the advantage that it explicitly captures the aggregate effects of state-dependent pricing behaviour on current inflation.

The paper uses the SDPC framework to examine whether a hybrid NKPC (NKPC extended by a lagged inflation term) can adequately describe inflation dynamics of a realistically calibrated state-dependent pricing economy. To explore this issue, artificial data sets for a state-dependent pricing economy are generated based on various calibrations of price adjustment costs under both low and high trend inflation environments. We use these data to estimate the hybrid NKPC and to assess the specification by examining both the estimated coefficients and the correlations between the simulated inflation and the inflation predicted by the hybrid NKPC. The findings suggest that the hybrid NKPC provides a good reduced-form description of inflation dynamics for a wide range of state-dependent pricing behaviours, particularly in the low-inflation environment. The fit of the hybrid NKPC is similar to that reported in the literature for estimations using real-world data. An interpretation of this finding is that the hybrid NKPC may be a good proxy for inflation dynamics implied by more realistic models of price-setting. Consequently, structural interpretation of its parameters may not be straightforward.

The UK labour force participation rate: business cycle and trend influences

Working Paper no. 228

Mark Schweitzer and David Tinsley

Policymakers will frequently be interested in how 'tight' the labour market is currently and how tight it can be expected to be in the future. This assessment will in turn depend upon a view of how the demand for labour compares with its availability. Looking at the unemployment rate alone might not be a sufficient statistic for gauging labour availability, since the inactive population represents a large potential source of labour supply. And the distinctions between some forms of inactivity and unemployment can be fairly weak, so that certain types of inactive people are as likely to fill jobs as the unemployed.

The decision whether to participate in the labour market is subject to numerous long-term 'trend' influences. In the United Kingdom, these long-term influences have included an increase in the number of students, as well as in the number of individuals who report themselves as long-term sick. But alongside these trend influences some aggregate business cycle effects are also likely to operate.

This paper investigates the extent to which the participation rate is influenced by structural trends and by the business cycle. We propose a modelling strategy that pools the available micro and macro-level data to produce a mutually consistent model of the trend and cyclical components of participation.

We find a significant procyclical pattern to participation in the available time-series data. However, we also identify some distinct trend influences on the participation rate, using longitudinal microdata. Together, these factors help to explain some of the movements seen in overall participation over the 1990s.

Our approach also allows us to construct forecasts for the participation rate, which would be a useful input into the sort of macroeconometric models used by policymakers. We assess our approach by conducting out-of-sample forecasts and find that it outperforms some conventional macroeconometric forecasts.

On the resolution of banking crises: theory and evidence

Working Paper no. 229

Glenn Hoggarth, Jack Reidhill and Peter Sinclair

Over the past quarter of a century, unlike in the preceding 25 years, there have been many large bank failures around the world. Moreover, cross-country estimates suggest that output losses during banking crises have been, on average, large—over 10% of annual GDP.

This paper reviews the merits of the various techniques used by authorities when resolving individual or widespread bank failures in developed and emerging market economies. In particular, the various banking crisis resolution techniques available to the authorities are classified and then compared with the techniques that have been used in practice, drawing on both the available evidence and our own analysis.

There is a range of options for dealing with insolvent banks. At one extreme, a bank can be kept open through an injection of capital. At the other extreme, a bank can be closed with its assets sold and depositors and possibly other creditors paid off. Between these extremes, a bank's licence may be removed but the bank may be sold off to another bank, in full or part, to preserve its activities. The extent of involvement of the authorities may also vary. It may be limited to encouraging or organising private sector support, or extended to official financial support, in the limit through government takeover.

Faced with a banking crisis the authorities often face a trade-off between maintaining financial stability today through intervention and jeopardising future financial stability through increasing moral hazard later on. To the extent that the public sector becomes involved in crisis resolution, moral hazard and the resolution costs can be limited by ensuring that bank 'stakeholders'—shareholders, managers, depositors and other creditors—share at least some of the losses. Clarity and transparency over restructuring programmes may also speed up the resolution process and reduce both present costs and future risks.

In practice, faced with individual bank failures the authorities have usually first sought a private sector solution. Any losses have been passed on to existing shareholders, managers and sometimes uninsured creditors, and not to taxpayers. Most recent systemic crises have typically been caused by an adverse macroeconomic shock weakening the whole financial system, rather than resulting from the impact of contagion following the failure of just one individual bank. In these cases policy options have often been limited. Finding a domestic private sector solution has often been difficult, so there has been more reliance on foreign takeovers and government intervention. Also, the authorities have been faced with the dilemma that imposing losses on to the banks' stakeholders could exacerbate rather than ameliorate the liquidity crisis.

In practice, in most recent systemic crises:

- early on central banks have provided liquidity to failing banks and governments have given blanket guarantees to depositors. In nearly all cases investor panics have been quelled but at a cost to the budget and increasing the risk of future moral hazard;
- open-ended central bank liquidity support seems to have prolonged crises, thus increasing rather than reducing the output costs to the economy;
- bank restructuring has usually occurred through mergers, often government assisted, and some government capital injection or increase in control. Bank liquidations have been used only occasionally, and typically for smaller institutions. Shareholders have usually lost their capital and senior managers their jobs, but creditors, including uninsured ones, have rarely made losses; and
- resolution measures have been more successful in improving banks' balance sheet positions than in restoring their profits or credit to the private sector. In many cases, bank lending and profitability have remained subdued for years after a banking crisis. However banking crises are handled, the adverse effects on the economy are likely to be large. This suggests that ensuring that the financial system is robust in the face of even substantial shocks should be a key objective of financial stability policy.

Financial interlinkages in the United Kingdom's interbank market and the risk of contagion

Working Paper no. 230

Simon Wells

A well functioning interbank market is essential for efficient financial intermediation. But interbank credit exposures imply the possibility of direct contagion: the sudden insolvency of a single institution may trigger multiple bank failures due to direct credit exposures. This paper aims to examine the potential for direct contagion in the UK interbank market.

Economic theory suggests that the potential for direct contagion depends, to some extent, on the exact structure of the interbank market. The problem is that the precise network of interbank exposures is unobservable. So this paper uses available data to estimate bilateral exposures between UK-resident banks. The estimates are used to assess the potential for direct contagion by tracing the path of assumed insolvency shocks through the banking system. We simulate the failure of each individual bank in the model and estimate the losses suffered by other banks as a result of the initial shock. We assume that contagion occurs (ie a bank fails outright) if a bank suffers a loss that exceeds its Tier 1 capital holdings.

Analysis is performed on three alternative estimates of the UK interbank structure. In each case, data on the total borrowing and lending positions of each UK-resident bank with the entire UK system are used to estimate the complete map of bilateral exposures. The first model (the benchmark case) assumes interbank borrowing and lending is as widely dispersed as possible, given each bank's observed total interbank assets and liabilities. This estimate is not conditional on market structure and so may be a poor representation of reality. The second estimate incorporates information from a database of bilateral exposures reported by banks. These data do not provide a complete map of interbank exposures: they include only exposures exceeding a certain threshold for a sample of banks. Nevertheless, incorporating information from this database into the model may mean that it better reflects concentrations in the UK interbank market. The final model is a restricted version of the benchmark case, where smaller banks and foreign banks are assumed to transact only with large UK-owned banks. In this case, the large banks can be thought of as a money centre for all banks in the UK system.

Data constraints mean that it is difficult to draw reliable conclusions about the potential for contagion. First, only banks that are resident in the United Kingdom are modelled. This means that the estimates capture only the exposure of UK-owned banks to the UK branches and subsidiaries of foreign banks and not to entire foreign banking groups. Given London's position as an international financial centre, failing to capture the full extent of exposures to foreign banks rules out a potentially important channel of contagion. Second, suitable data are only available for interbank money market loans and deposits. Although these capture a large part of unsecured interbank activity, exposures arising from other instruments (such as interbank holdings of Certificates of Deposit and financial derivatives) are not included. Third, each model of the interbank market is derived from partial information and we show that the results depend on the assumed distribution of lending across banks.

Despite these caveats, our results give some useful information on the general potential for contagion in the UK interbank market. We explore the effect of one type of extreme event—the sudden and unexpected insolvency of a single bank. Our results show that an insolvency shock, idiosyncratic to a given bank, can lead to a substantial weakening in the capital holdings of other banks, but in most cases does not result in additional (or knock-on) bank failures. But assuming complete loss given default, our stylised model suggests that, in extreme cases, a single bank insolvency could trigger knock-on effects leading in the worst case to the failure of up to one quarter of the UK banking system. At the same time, a further quarter of the banking system would suffer losses amounting to more than 10% of their Tier 1 capital. For loss given default levels of less than 50%, contagion affects, at worst, less than 1% of total banking system assets. However, even with low loss given default, a narrow shock can considerably reduce the capital reserves of many banks. And, if the initial shocks hit during a period where the banking system is already weakened (say during a period of large macroeconomic fluctuations), the effect of contagion can be larger.