Estimating real interest rates for the United Kingdom

Working Paper no. 200

Jens Larsen, Ben May and James Talbot

The *ex-ante* real interest rate is a key variable in the transmission mechanism of monetary policy. Any change in the short-term nominal interest rate set by the monetary authority will—if prices are sluggish—lead to a change in real interest rates, which will affect demand, and subsequently inflation, via the consumption and savings or investment decisions of households and firms.

In general, there are few direct measures of the *ex-ante* real interest rate because almost all debt contracts are specified in nominal terms. So this paper explores a number of methods for calculating UK real interest rates. The pros and cons of each approach are evaluated carefully: after constructing a long and consistent time series of each measure, a rigorous sensitivity analysis is conducted and, where appropriate, error bands are constructed around the estimates in order to assess their accuracy.

The United Kingdom has a well-developed market for government bonds (gilts) that are indexed to the retail prices index (RPI), so the first approach considers real interest rates derived from these bond prices. But more recently, estimation has been complicated by the combined effect of limited supply and artificially price-inelastic demand on yields in this market.

A second approach uses yields on nominal gilts minus an appropriate measure of inflation expectations. But this method is also subject to several important problems. First, nominal bond yields will be subject to the same distortions identified above. Second, estimating inflation expectations is not an easy task. The paper adopts two approaches to devise inflation expectations: one is based on surveys and another one uses forecasts from a vector autoregressive model of inflation, unemployment and interest rates. Third, such estimates will include a measure of the inflation risk premium, and so are not directly comparable with those from index-linked gilts.

The third approach uses a 'consumption-based' measure—derived from manipulating the first-order condition of the standard household intertemporal optimisation problem. The basic (power utility) version of this model suffers from the standard problems outlined in the literature: the so-called 'risk-free rate' and 'equity premium' puzzles. So the basic framework is augmented to allow for (external) habit formation in consumption, and extended to estimate general *k*-period real interest rates.

Real interest rates at one, three and ten-year maturities derived using this approach look reasonably plausible: real interest rates peak during the recession of the early 1980s and fall during the economic expansions of the late 1980s and late 1990s. But because the model is based on a relatively simple process for consumption growth (a random walk), the term structure of interest rates contains less information, remaining relatively flat throughout the sample period.

Interestingly, although the different approaches outlined above can sometimes yield different estimates of real interest rates, all the measures move more closely together during the post-1992 inflation-targeting period than before. Before 1992, uncertainty about the monetary regime, coupled with persistent expectational errors, may have made it more difficult for agents to forecast real interest rates and inflation.

Another question is whether the fall in long real yields observed in the index-linked gilt market post 1997 is based on movements in real fundamentals? Evidence from the model with habit formation suggests that there has been some fall in the ten-year real interest rate since the mid-1990s. But it would appear that at least some of the decline observed in the index-linked market has been driven by the institutional factors described above, underlining the value of taking an eclectic approach when assessing movements in real interest rates.

Debt maturity structure with pre-emptive creditors

Working Paper no. 201

Prasanna Gai and Hyun Song Shin

Short-term liabilities play a central role in sovereign debt restructuring. Typically, the creditors of a debtor in distress must decide whether to extend further lines of short-term credit, or whether to cut their losses and refuse to lend. The greater the funding need that creditors must meet, the less likely it is that they will be persuaded to roll over their credit lines. This is because uncertainty about the assessments and actions of other creditors acts as a disincentive for an individual lender to extend credit. Thus, the greater the amount of short-term (immediate) debt outstanding, the more problematic the problem of coordinating creditors becomes.

In dealing with sovereign debt crises, policy-makers have proposed measures such as stays on creditor litigation, temporary payments suspensions, and concerted rollovers of credit lines, in an effort to target short-term debt. But following the use of concerted rollovers in Korea, creditors reacted pre-emptively to the crisis in Brazil—shortening maturities at the onset of crisis and cutting interbank lines sooner than might otherwise have been the case. This experience has led some to question the viability of rollovers and payments standstills as tools for crisis management. By encouraging creditors to 'rush for the exits', it is argued, such measures merely bring forward financial vulnerabilities by pushing debt maturities towards the shorter term.

This paper argues that such logic is not necessarily general. We model the 'rush for exits' as a pre-emption game among creditors. A debtor country undertakes an N-period project and creditors choose where, within the maturity spectrum, they prefer to extend credit. The fruits of the project, which are taken by long-term claimholders so long as premature liquidation is avoided, depend on the size of the funding gap and on the maturity structure of the debt—the shorter the maturity, the greater the probability of financial crisis.

Creditors face two conflicting incentives. First, there is the desire to be first in the queue (the shortest debt maturity) so as to be able to escape the losses associated with crisis. But if all creditors behave in this fashion, this maximises the chance of crisis. So some creditors choose longer maturities in the hope that funding problems do not arise. The balance of the two generates an equilibrium debt maturity profile for the project.

The analysis explores the effects of an orderly payments suspension on the creditor's choice of maturity and, hence, on the term structure of debt. We show that if such measures can boost recovery values in the event of crisis, then creditors may not seek short-lived claims. This is because there is a direct effect in increasing incentives to holding longer-term claims since the returns to holding these are now higher. And there is an indirect and reinforcing 'strategic' effect, as higher recovery rates brought about by such policy measures reduce the desire to engage in pre-emption in the first place.

Comparative static results suggest that the overall implications for the term structure of debt depend on the effectiveness of the crisis management framework as well as the length of time that the restructuring is expected to take. If payments suspensions are short-lived and have a positive effect on recovery values, they are unlikely to generate a move towards shorter maturity debt. Longer-lived debt workouts can push maturities towards the shorter term, however. Indeed, for suitably lengthy workouts, it is even possible that there can be a hollowing out of middle maturities as creditors move to either end of the maturity spectrum. It is not typically possible therefore to draw firm conclusions, *a priori*, about the shape of debt maturity profiles when measures such as payments standstills and concerted rollovers are used as part of crisis management.

Credit spreads on sterling corporate bonds and the term structure of UK interest rates

Working Paper no. 202

Jeremy Leake

This paper explores the relationship between credit spreads on sterling corporate bonds and the term structure of UK interest rates. In doing so, it addresses the extent to which credit spreads are reliable indicators of default risk.

Finance theory has suggested that there is a relationship between interest rates and default risk, and hence a relationship between interest rates and credit spreads. However, the theoretical models conflict as to the nature of this relationship.

On the one hand, 'structural' models based on option pricing suggest that higher interest rates might be associated with lower credit spreads. Such models view equity as a call option on the value of the firm, with the strike price equal to the face value of the debt. For example, in such models, the firm defaults on debt repayment if the value of the firm is less than the face value of the debt on the debt repayment date. A higher risk-free rate in this static model corresponds to a higher expected growth rate in the value of the firm (other things equal) and so a lower level of default probability over any given horizon.

A second view can be derived from 'reduced-form' models. Such models do not attempt to model why firms default on their debt but instead assume that some bonds will default on the balance of probability. There are numerous types of reduced-form models; in the one used in this paper investors demand compensation for default risk by grossing up the coupon paid on a default-free bond by the expected default probability. If interest rates rise by 1 basis point, the gross-up effect increases the coupon by more than 1 basis point. Thus, the differential between the coupon on the corporate bond and the coupon on the risk-free bond increases in absolute terms with the size of the default-free coupon and credit spreads rise when the default-free interest rate rises.

This paper examines the empirical relationship between credit spreads on single-A and Aa-rated sterling corporate bonds and the level and slope of the UK yield curve for the period 1990 to 1998. The corporate bond price data are quotes rather than actual trades. Corporate bonds can be much less liquid than government bonds and it is possible for some corporate bonds not to trade for long periods of time. As a result, price quotes may not reflect all current information and so calculated credit spreads might reflect delays in the arrival of information rather than economic factors. Any such bias is likely to be exacerbated at times when information arrives frequently—for example, when government bond prices move sharply over short periods of time.

The credit spread series calculated exhibits particularly high volatility during the second and third quarters of 1994, a period when prices in the gilt market fell sharply. This volatility could have been due to uncertainty in the corporate bond market or could have been a result of non-synchronous gilt and corporate bond data. The paper finds that, though the occurrence of unchanged prices was high in the data set, they were not particularly prevalent during the period of high credit spread volatility in 1994.

Due to the possibility of stale prices throughout the data set the paper runs two sets of regressions: one set using daily data with an adjustment for non-synchronous prices and a second set using weekly data with no such adjustment. The results of the two sets of regressions are similar: in both there is a negative correlation between credit spreads and the slope and level of the yield curve. However, most of the coefficients in the regressions using daily data are statistically insignificant, while the coefficients in the regressions using weekly data are more negative, and most are statistically significant. All results are economically small, in that a large change in the yield curve is required to produce measurable movements in spreads, suggesting the relationship is weak. The weak relationship between interest rates and credit spreads in this study gives us cause to doubt whether such credit spreads are reliable indicators of corporate bond credit risk. One potential explanation is that factors other than interest rates are more important in driving credit risk. More likely is the possibility that the credit risk component of credit spreads on investment-grade corporate bonds is small relative to factors such as liquidity and risk aversion. Studies using US data suggest that this latter explanation is correct. So an interesting extension to this work would be to undertake a similar study on sub investment-grade sterling corporate bonds, where one would expect the credit risk component of credit spreads to be greater. However, we will probably need to wait for more sterling sub investment-grade corporate bonds to be available and for a sufficiently long history of data before conclusions can be drawn from that market.

Analytics of sovereign debt restructuring

Working Paper no. 203

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

Recent years have seen an increasing incidence of sovereign debt crises, including Russia in 1998, Turkey in 2000 and 2001 and Argentina in 2001 and 2002. The costs of these crises imposed on debtors, creditors and the official sector have activated a heated debate on appropriate mechanisms for the restructuring of sovereign debt, particularly international bonds. Some commentators have argued in favour of market-based solutions, whereas others—most prominently the International Monetary Fund's First Deputy Managing Director Anne Krueger—have advocated statutory approaches akin to an international bankruptcy court.

This paper develops a simple theoretical model of sovereign debt restructuring to analyse the merits of some of these proposals. In the model, a debtor country with an unsustainable stock of debt makes a write-down offer to its creditors. Creditors vote whether to accept or reject the offer. Two market-based arrangements are analysed: under the first, contractual provisions require unanimous consent for the offer to go through; under the second, collective action clauses (CACs) bind the minority to the will of the majority. These arrangements are compared with a first-best welfare benchmark.

In recent years, coordination problems among creditors have intensified as a large part of sovereign credit is now held by a diverse set of bondholders and sellers of credit protection rather than in the loan books of a few internationally active banks. Our results suggest that, when intra-creditor coordination problems are severe, voluntary market-based debt exchanges with unanimity clauses lead to inefficiencies by providing incentives for certain types of creditor to hold out. As a result, resources are unnecessarily expended on legal proceedings and costly renegotiations preventing the debtor from exerting the socially desirable amount of adjustment effort. CACs can resolve these inefficiencies, provided that all parties have complete information about each other's preferences. In such a world, statutory mechanisms are unnecessary.

This conclusion, however, does not hold when the assumption that the debtor and its creditors have complete information about each other's preferences is dropped. In this case, the inefficiencies induced by strategic behaviour—the debtor-creditor bargaining problem—cannot be resolved by the parties themselves. One way of removing these inefficiencies would be the intervention of a third party—for example, an international analogue of a domestic bankruptcy court.

The dynamics of consumers' expenditure: the UK consumption ECM redux

Working Paper no. 204

Emilio Fernandez-Corugedo, Simon Price and Andrew Blake

Over the past 25 years, innumerable consumption error correction mechanisms (ECMs) have been estimated. In the United Kingdom in particular, research has concentrated on the variables explaining consumption in the long and short run. With single-equation consumption ECMs, the implication is that deviations from the common trend in consumption, income and wealth are corrected only through consumption. This is despite the fact that in the simplest intertemporal models of household consumption, there should be no consumption ECM. Instead, equilibration should operate via the income or wealth drivers. The former result does not hold with all extensions, for example to habit persistence, but the latter does. This issue, introduced by John Campbell in the 1980s, has been revived with a number of papers on US data by Sydney Ludvigson and her co-authors. In those papers, deviations from common trends tend to be corrected via changes in wealth. In this set-up, deviations from the long-run relationship appear to lead to changes in income or wealth. But the causality here is from expected future events to current consumption and saving decisions; it is not that (eg) higher consumption causes higher income growth through, say, some aggregate demand mechanism. In this paper we examine the evidence for the United Kingdom.

We pay some attention to the treatment of non-durable consumption. We construct a simple model of the consumption of both durable and non-durable goods. We construct appropriately defined data, and the short-run dynamics and long-run relationship between non-durable consumption, non-asset income, wealth and the relative price of durable goods are examined. One cointegrating relationship is found to exist. The relative price of durables to non-durables may play a role in this process. Estimating vector error correction mechanisms (VECMs), we find that adjustment towards the long-run common trend does indeed occur partly via changes in wealth. This is consistent with forward-looking behaviour by agents. It also means consumption can predict asset returns. This result is confirmed by a regression of excess returns to equities on the disequilibrium term from the long-run relationship.

We also perform a decomposition of shocks hitting the system into temporary and permanent components. Almost all of the variation in the consumption and income process can be ascribed to permanent shocks. Depending on the treatment of the relative price of durables, we find that between 30% and 90% of fluctuations in non-human wealth are transitory. Even if the lower figure applies, this means a substantial part of short-term fluctuations in wealth is decoupled from permanent consumption.

Our analysis implies that we can welcome the return of the UK consumption ECM, in the context of a complete VECM analysis of the system explaining the relationship between consumption and permanent income.

Empirical determinants of emerging market economies' sovereign bond spreads

Working Paper no. 205

Gianluigi Ferrucci

Yield spreads on emerging market economies' (EMEs') sovereign bonds are important indicators of financial fragility for country surveillance purposes. They are typically used as a measure of the markets' perception of the risk that a country might default and to assess EME external financing conditions. But EME spreads are influenced by a large number of determinants—credit risks, liquidity risks, and market risks—and inferring their exact information content is not straightforward.

This paper develops an empirical model relating secondary market sovereign spreads to a set of country-specific fundamentals, controlling for external factors, market risk and liquidity in bond markets. The aim is to explain the long-run determinants of EME bond spreads, together with some short-run dynamic behaviour. The estimated equation is reduced form, and posits that the fair-value spread is a function of the probability of default and the recovery rate in the event of default. In turn, the probability of default is linked to a set of macro-prudential indicators affecting the country's solvency and liquidity position. To underpin the selection of credit spread determinants (fundamentals), the paper discusses a simplified model of sovereign borrowing that formalises the consumption choices of an indebted small open economy. This model points to a set of variables that are important components of the internal and external constraints on government debt obligations. The data set for the estimation is a ragged-edge panel of secondary market spreads and a number of country-specific macro-prudential indicators obtained from a variety of sources. Estimates are obtained using the pooled mean group technique, which assumes a dynamic error correction equation with heterogeneous cross-sectional coefficients in the short-run equations and homogeneous coefficients in the long-run relationship.

We use this model to address three main questions. First, we ask what proportion of the change in market spreads is explained by changes in the underlying fundamentals, controlling for external factors, liquidity and market risk. Second, we provide a benchmark measure of sovereign risk against which to compare actual market spreads. Finally, we use the model to explain patterns in spreads, from an *ex-post* perspective. As a case study we analyse the generalised fall in secondary market EME bond spreads experienced between 1995 and 1997.

Data limitations highlighted in the paper mean that the results have to be interpreted with caution. Nevertheless, the model is informative and allows us to reach interesting conclusions. Our main finding is that market spreads broadly reflect fundamentals, but that non-fundamental factors also play an important role. Comparing market-based spreads against their fundamental-based counterparts we find that credit risk is typically priced fairly closely to a theoretical equilibrium level, based on the selected set of macro-prudential indicators. In the cases of large absolute misalignments, we identify whether the divergence is due to unmeasured fundamentals or is likely to depend on market imperfections. Finally, the model suggests that the fall in spreads between 1995 and 1997 cannot be explained solely in terms of improved fundamentals. Assuming that our model provides a fair picture of fundamental-based sovereign credit risk, the divergence must be due to capital market imperfections, such as higher investor risk appetite resulting from lower global interest rates.

The rise in US household debt: assessing its causes and sustainability

Working Paper no. 206

Sebastian Barnes and Garry Young

US households' debt relative to their income has increased to new highs in recent years, posing questions about the likely economic effects of this growth in indebtedness. This paper assesses possible causes of this rising indebtedness and considers how sustainable such borrowing behaviour might be.

The paper uses an overlapping generations model where differences between cohorts, ie households of different age, give rise to household sector borrowing and asset accumulation. Households borrow both because of a consumption-income motive, where young households with low current incomes borrow to raise their current consumption, and a housing-finance motive, where households borrow to fund owner-occupation of housing. Only the youngest households would choose to borrow due to the consumption-income motive but housing finance causes them to borrow more and later in their lives.

The model is calibrated to match a number of features of the US experience, both in aggregate and in the cross-section of the population. We also introduce an old-age borrowing constraint, which provides an alternative explanation for why older people choose not to borrow to finance owner-occupation towards the ends of their lives, even though this would allow them to consume more.

The debt to income ratio would have been stable if the economy were in steady state. So, we consider a number of shocks to the US economy that might possibly account for the rise in household debt over the past 30 years. Shocks to real interest rates and income growth expectations would affect the behaviour of individual households. Even with no change in household-level behaviour, demographic change such as the 'baby boom' might have affected total borrowing by altering the numbers of those most likely to borrow, ie the young, in the economy.

Combining observed shocks, we find that the rise in indebtedness during the 1990s is similar to that predicted by the calibrated model. However, the rise in debt during the 1980s is difficult to explain, as a number of factors suggest that it should have fallen during that time. This could reflect shortcomings in the model or the influence of other factors such as financial market liberalisation.

What does this imply for the sustainability of US household debt? The model suggests that household borrowing would be expected to increase further over coming years, reflecting the gradual adjustment to shocks during the 1990s, albeit at a slower rate. However, the sustainability of current behaviour depends critically on the realisation of the expectations on which households have made their borrowing decisions.

A quantitative framework for commercial property and its relationship to the analysis of the financial stability of the corporate sector

Working Paper no. 207

John Whitley and Richard Windram

In the past property-related lending has been a significant cause of losses for UK financial institutions and the property cycle has been assigned a role in accentuating the general business cycle. Whereas the role of residential property has been well documented in terms of the transmission of shocks and its relationship to the overall macroeconomy, this has not been the case for commercial property.

This paper sets out a quantitative framework for considering the implications of commercial property developments for the financial stability of the corporate sector. It builds on previous work extending the Bank of England's macroeconometric model to the household and corporate sector's balance sheets by constructing a model of real estate companies and linking this model into the aggregate corporate sector through the role of property as collateral. Previous modelling work on commercial property has focused on market studies or used single-equation relationships. Few studies have attempted to link the commercial property market to financial markets, and even fewer to the rest of the macroeconomy. Lack of suitable data has been a major constraint on attempts to model the commercial property sector.

In this paper we attempt to fill this gap. Data for more than 1,000 real estate companies are used to calibrate the financial accounts of real estate companies. We then combine various rules of thumb that are consistent with these accounts with econometric analysis to build an overall simple model of real estate companies' behaviour, related to macroeconomic factors. One of the reasons why previous modelling attempts have not been developed for practical forecasting and projections is that the models have required projections of other variables that are either related to the property market itself, or are difficult to project. The principal objective is to ensure that the model is capable of being used for both forecasting and simulation, either in isolation, or in combination with a wider macroeconomic model.

Econometric analysis of rental income and bank lending is the main behavioural element in the real estate model. The bank lending equation does not find a consistent role for borrowing costs relative to returns from property over the whole sample period. However, the gap between property yields and the base rate appears to explain a large proportion of the growth of bank lending after 1999. Without allowing for this influence, the equation appears to break down. Sale and lease-back deals between non real estate and real estate companies also appear to have boosted bank lending after 1999. This model can be used to derive an estimated probability of default for real estate companies, drawing upon other research being developed within the Bank of England on corporate failure. Further work on data and on possible supply influences might help to resolve the puzzles highlighted in this paper.

A dynamic simulation of the property model starting in 1990 illustrates the overall performance of the model. The model fails to capture all the cyclicality of capital values and bank lending since 1990. This can be traced back to a failure to predict changes in the discount rate applied to rental income during this period, and may reflect a (temporary) change in risk premia. The model itself does not include an explicit treatment of expectations or risk premia. Future work might usefully examine the role of alternative expectations mechanisms.

An important feature of the work is that it enables the specification of the role of commercial property in influencing the financial health of the overall corporate sector. The key link between the real estate model and the rest of the corporate sector is through the capital value of commercial property. Capital values are derived from rental flows using a simple discount model. Changes in capital values are found to affect aggregate corporate liquidations, because they alter the collateral security that backs corporate loans. Simulations of changes in macro variables using the Bank of England's macroeconometric model illustrate the importance of this property link and the potential for adverse developments in the commercial property sector to amplify the sensitivity of corporate default to macroeconomic shocks. So, although there are no feedbacks to broader macroeconomic aggregates, the model enables further quantification of financial stability risks.

A matching model of non-employment and wage pressure

Working Paper no. 208

Andrew Brigden and Jonathan Thomas

In contrast to previous cyclical upswings where both the unemployment and inactivity rates have declined in tandem, the fall in the total non-employment since the mid-1990s has been almost completely accounted for by a decrease in the unemployment rate, while the inactivity rate has remained broadly flat.

The monetary policy implications of these developments are unclear. It is possible that the relatively stable inactivity rate has helped to moderate any extra wage pressure arising from the decline in unemployment. However, it is equally plausible to argue that the inactive are so detached from the labour market that they have no impact on wage bargaining.

In order to address these issues, this paper develops a model of the labour market that explicitly distinguishes between the unemployed and the inactive, rather than treating all those who are out of work as unemployed. The key difference between the groups is the value that they place on non-work related activities such as leisure. It is assumed that unemployed people have a relatively low valuation on such uses of their time. Consequently, they search harder for jobs, are prepared to accept lower pay, and therefore enter employment more readily.

We then use the model to examine the behaviour of inactivity, unemployment and wage growth over 1994–2000. Specifically, we attempt to identify the underlying shocks that can explain the observed trends in unemployment and inactivity over this period. We consider shocks to the benefits received by the unemployed and the inactive, the costs incurred by the firm when hiring and firing workers, and the share of individuals with low search effort in the working-age population. The most plausible impulses involve a rise in the fraction of individuals with low search effort, and a reduction in benefits to the unemployed. The rise in the proportion of students in the working-age population over the 1990s could have raised the share of individuals with low search effort correspondingly, while the stricter benefit regime since the mid-1990s could have increased the attractiveness of working compared with being unemployed. Both these shocks imply movements in unemployment and inactivity that would not be accompanied by a rise in wage pressure.

Settlement bank behaviour and throughput rules in an RTGS payment system with collateralised intraday credit

Working Paper no. 209

Simon Buckle and Erin Campbell

High-value payment systems are critical elements of the economy and typically take one of two forms: deferred net settlement (DNS) and real-time gross settlement (RTGS). In a DNS system, banks make payments to each other during a specified period (usually one day) and then settle the net amounts at the end of that period. Until settlement is completed, banks are effectively extending unsecured and possibly unmonitored loans to each other. The amount of credit risk in such systems was one of the main drivers for the introduction of RTGS systems in Europe and elsewhere.

RTGS systems eliminate the counterparty credit risk present in DNS systems by requiring participants to settle payments on a gross basis in real time. But this credit risk reduction comes at the cost of a requirement for potentially expensive intraday liquidity. Central banks have sought to reduce liquidity costs for settlement banks, for example by providing collateralised intraday liquidity and good system design. Even so, intraday liquidity in RTGS systems is not in general free and unlimited.

An important determinant of the liquidity efficiency of an RTGS payment system is the extent to which the system design gives settlement banks an incentive to manage their payments in a socially efficient way. In an RTGS system, one bank's payments are a source of intraday liquidity for the recipient bank, which it may then subsequently use to make its own payments. If banks recycle liquidity sufficiently quickly, the aggregate requirement for intraday liquidity can be significantly reduced.

This paper provides a simple analytical model with which to study RTGS system design. In the context of this two-bank model, we show that banks will delay payments when they care about payment imbalances between them in the first period, leading to an inefficient degree of liquidity recycling. When banks do not care about first-period payments imbalances, there is no unique equilibrium outcome but one possible symmetric outcome is efficient—when each bank posts the same amount of collateral, equal to half of the value of payments each wants to make, and uses all its available liquidity to make payments in the first period. This results in the maximum possible degree of liquidity recycling and the lowest aggregate collateral requirement.

In practice, banks do care about payment imbalances between them during the day because of competitive and/or liquidity risk concerns. While some degree of liquidity recycling is likely to emerge even in these circumstances, in particular due to the repeated nature of the interaction between settlement banks, we argue that full efficiency is not guaranteed, largely because of imperfect information and the competitive dynamics of the payment industry. Using the model, we show how regulation—in this case a throughput rule—can be used to achieve the efficient outcome even in this situation. Throughput rules, which stipulate the proportion of each settlement bank's usual daily payments that must be made by a certain cut-off time, substantially reduce the overall requirement for intraday liquidity in an RTGS system and may also increase the contestability of the payments market, encouraging a higher degree of direct access to payment systems. Consequently, throughput rules could have risk-reduction benefits if they help to reduce the level of tiering in the financial system.

We also address the question of how to design throughput rules in practice. Our model suggests that increasing the number of throughput rules would enhance efficiency indefinitely, although at an ever diminishing rate. It seems likely, however, that there is some upper limit to the efficient number of throughput rules, for at least two reasons. First, real payments have a finite size and are sometimes very large and urgent. If the value of payments required to be made in a given period was less than the size of a large, urgent payment, banks needing to make such payments would be forced to use more liquidity than they receive back from other banks—the original problem that throughput rules were designed to solve. The second reason why the feasible number of throughput rules will be bounded is that real payments between banks are stochastic and, at least in part, unknown at the start of the day. Assuming the throughput rules are based on the average value of payments, as the number of rules increases, eventually a point will be reached where, on a day when the demand for payments is low, one or more of the banks will just not have sufficient customer payments to meet the final throughput requirement.

A further potentially important design issue highlighted in this paper is that aggregate throughput rules may not be adequate in a payment system with more than two settlement banks, since they could not prevent banks from forming cartels to disadvantage other banks or new potential entrants. While we have no evidence of such behaviour within the UK high-value payments system, there may be merit in considering the feasibility of applying throughput rules on a bilateral basis or putting other equivalent incentive mechanisms in place.

Company accounts based modelling of business failures and the implications for financial stability

Working Paper no. 210

Philip Bunn and Victoria Redwood

Corporate failure poses a threat to financial stability if firms who fail default on their debt. Although the failure of an individual firm is unlikely to have systemic implications, if a number of firms with large amounts of outstanding debt fail simultaneously there may be systemic implications. Previous work in the Bank of England, which aimed at monitoring these risks to financial stability arising from corporate failure, has been relatively qualitative. The aim of this paper is to supplement that work with a more quantitative approach. We use firm-level data to develop a model of corporate failure, which we then use to analyse both the aggregate risks and the distribution of those risks.

The early literature used balance-sheet information and a discriminate-analysis method to try and predict firm failure. More recent articles have favoured probit models, and this is the approach we take. Most of the papers in the existing literature use a relatively small number of firms or a relatively short timescale. We attempt to address this problem by using a sample of over 100,000 observations from 29,361 public and private UK non-financial firms between 1991 and 2001.

We estimate a probit model for individual company failure using firm-level balance-sheet information and aggregate data on macroeconomic conditions. We find that there is a negative relationship between profitability and corporate failure, but this relationship is non-linear, with negative profitability being associated with the largest marginal effect on the probability of failure. There is a positive association between the debt to assets ratio and the likelihood of failure, and there is an additional positive impact on the probability of failure if above-average capital gearing coincides with the firm making a loss. The probability of a firm failing is found to be negatively related to its interest cover and liquidity. If a firm is large and a subsidiary, it is less likely to fail, holding all other factors constant. Our model controls for industry: firms in the service sector are less likely to fail than those in manufacturing, primary industries and utilities. We incorporate macroeconomic effects into the model by including GDP growth and find a negative correlation between GDP growth and failure, even after controlling for all of the firm-level characteristics.

We use the firm-level probabilities of failure generated by the model and apply these to the analysis of risks to financial stability arising from the UK corporate sector. We do this by defining debt at risk: the probability of failure of an individual firm multiplied by its debt. To derive an aggregate measure of financial risk we sum debt at risk across all firms in each year. We find that this micro-based measure of financial risk performs better in predicting debt at risk of default than a macro-based approach, which involves multiplying the average probability of failure by the total stock of debt and therefore does not fully exploit the firm-level dimension of the data. Debt at risk, as a proportion of total debt, was at its highest in the early 1990s, and it has been relatively stable since 1993, although the stock of debt has risen over this period.

As well as analysing aggregate measures of debt at risk, the paper also looks at the distribution. The distribution appears heavily skewed, with debt at risk being concentrated among a small number of firms. The implication of this is that we should particularly focus on these firms in order to monitor what is happening to the aggregate measure. While debt at risk is concentrated among a relatively small number of firms, in general these are not the firms with the highest probabilities of failure. The firms with the highest probabilities of failure tend to be small firms, which do not hold large amounts of debt in absolute terms.