Implicit interest rates and corporate balance sheets: an analysis using aggregate and disaggregated UK data

Working Paper no. 193

Andrew Benito and John Whitley

This paper examines evidence for effects on the cost of corporate debt finance from net worth. The central issue we address, confronted with data at both the aggregate and individual company levels, is whether implicit corporate interest rates reflect the strength of corporate balance sheets. In particular, such an effect is emphasised by the credit channel of the financial accelerator literature.

The analysis begins by exploring measures of implicit interest rates, deriving aggregate data from national accounts. Using a simple conceptual framework the paper estimates single time series models that relate implicit interest rates to risk-free rates and measures of corporate indebtedness. It finds evidence for a non-linear role for capital gearing, where gearing only changes the implicit interest rate when it is at relatively high levels. This is consistent with the prediction from the financial accelerator literature that balance sheet weakness should give rise to an increase in the external finance premium, although that does not depend on non-linearity.

The paper also uses company-level data to relate implicit rates to balance sheet measures of gearing and liquidity. Although non-linear effects are not found, the results confirm a significant positive relationship between implicit interest rates and balance sheet conditions.

Finally, the possible quantitative role that the finance premium may play in the propagation of shocks is considered. Simulations are presented using the time series equation in an aggregate macro model. The results show that implications for corporate liquidations can be quite sensitive to the presence of the non-linearity, although the sensitivity depends on the source of the shock.

A Merton-model approach to assessing the default risk of UK public companies

Working Paper no. 194

Merxe Tudela and Garry Young

The quantitative modelling of credit risk shows how the probability of company default can be inferred from the market valuation of companies under specific assumptions on how assets and liabilities evolve. This paper employs a Merton-style approach to estimate default risk for public non-financial UK companies and assesses the reliability of these estimates using a range of different techniques.

The original Merton model is based on some simplifying assumptions about the structure of the typical firm's finances. The event of default is determined by the market value of the firm's assets in conjunction with the liability structure of the firm. When the value of the assets falls below a certain threshold (the default point), the firm is considered to be in default.

To draw conclusions on financial stability and implement the right policy measures, the estimated probabilities of failure need to be both reliable and efficient. This paper assesses the reliability of the estimates by examining their success in predicting the failure or survival of both failed companies and survivors. The efficiency of the estimates is assessed by testing the extent to which the predictive power of the estimates could be improved by incorporating other information publicly available in company accounts. Models that combine a Merton approach with additional financial information are referred to in the literature as 'hybrid models'.

The probability of default derived from our Merton-model implementation provides a strong signal of failure one year in advance of its occurrence. For example, the mean value of the estimated one-year probabilities of default for our entire sample is 47.3% for those companies that went bankrupt, and 5.4% for those that did not.

Calculation of Type I and II errors (Type I errors are defined as the percentage of actual failures classified as non-failures, Type II errors are the percentage of non-failures classified as failures) suggests that the estimated probabilities of default are successful in discriminating between failing and non-failing firms. Classifying defaults as those firms with an estimated probability of default greater than or equal to 10%, the Type I error is relatively modest at 9.2% (with a Type II error of 15.0%).

Our implementation of the Merton approach clearly outperforms a reduced-form model based solely on company account data. But our analysis also shows that the type of hybrid models implemented here, ie those combining company account information and the Merton approach, outperform our implementation of the Merton approach, if only marginally.

Forecasting inflation using labour market indicators

Working Paper no. 195

Vincenzo Cassino and Michael Joyce

There are a large number of labour market indicators that could be used by monetary policy makers to assess the state of the labour market and the associated implications for inflationary pressure. A non-exhaustive list, taken from recent Bank of England *Inflation Reports*, would include the unemployment rate (measured from both claimant count and the Labour Force Survey), the employment rate, the non-employment rate, measures of skill shortages, and the ratio of vacancies to unemployment. This paper attempts to shed some light on how much weight should be attached to these and other labour market indicators by evaluating them against a simple criterion: their past performance in predicting price and wage inflation.

We compare the performance of 30 labour market indicators (derived from 16 underlying labour market variables) in forecasting three different price and wage inflation measures-based on the RPIX, the DGI-RPIX and the AEIover various sample periods from the mid-1970s to 2000. To model the relationship between inflation and each labour market indicator, we estimate a reduced-form inflation equation ('a backward-looking Phillips curve'), in which the change in inflation is specified as a data-determined function of past inflation, the labour market indicator itself and (in the case of RPIX and nominal earnings growth) real import price inflation. Where appropriate, we derive our indicator measures by first detrending the underlying labour market variable using a Hodrick-Prescott filter to form a 'gap' measure (ie an estimate of how far the variable is away from its trend) but, as a cross-check, we also separately examine the effect of using the first difference of the variable.

Two basic approaches are used to assess the

inflation-forecasting properties of each labour market indicator. We examine their *ex-post* forecast performance, by carrying out Granger causality tests based on data from the mid-1970s onwards, to see whether the indicators provide any information about movements in inflation not captured by the past history of inflation itself and (where appropriate) real import price inflation. Since they are backward looking, however, these tests do not tell us how useful particular labour market indicators would have been in genuine forecast situations. We therefore also consider the *ex-ante* forecast performance of the indicators, using simulated out-of-sample forecasting tests for the period 1985–2000. This procedure involves adding each of our selected labour market indicators to an inflation-forecasting equation that is estimated, either recursively or over a rolling sample, moving forward the end of the sample period one quarter at a time. The lag lengths of the variables in the equation are re-optimised over each period and the equation is used to forecast out of sample. By limiting our information set to data only available at the time of the forecast, this method should provide a better approximation to how the models would have predicted inflation in 'real time'. We then compare the out-of-sample forecasts of these indicator models with predictions from an autoregressive model of inflation and with other simple benchmark models.

The in-sample and out-of-sample criteria lead to rather different conclusions about the forecasting performance of the different indicators. According to the in-sample Granger causality analysis, most labour market indicators appear to be statistically significant in an inflation-forecasting equation. However, the out-of-sample forecasting analysis suggests that a much smaller number of labour market indicator models are better at forecasting changes in inflation than an autoregressive model, and that virtually none outperform this benchmark over the period since 1995. Moreover, the individual labour market indicator models that perform relatively well out of sample tend to be sensitive to the precise choice of inflation measure, sample period and estimation method. Interestingly, one seemingly robust result is that the unemployment rate gap, the most commonly used measure of labour market tightness, performs poorly across a range of specifications.

There are a number of possible reasons for the poor out-of-sample performance of most of the labour market indicator models examined. One contributory factor is that neither the Hodrick-Prescott or difference filters are likely to do a good job in capturing the time-varying trend of the underlying labour market variable. However, general model instability and overfitting in the estimation also contribute, probably reflecting the reduced-form nature of the analysis, which makes it vulnerable to structural and policy changes, as well as to changes in the pattern of shocks hitting the economy. Since no specific indicators are superior in all circumstances, we suggest that the best approach is to take into account a wide variety of information in forming an assessment of the labour market, in line with current practice. This conclusion is reinforced by the fact that simple combination forecasts, based on taking the median or trimmed mean of forecasts based on the individual indicator models, generally produce more reliable results.

UK business investment: long-run elasticities and short-run dynamics

Working Paper no. 196

Colin Ellis and Simon Price

Neoclassical theory tells us that a profit-maximising firm's desired capital/output ratio depends on the real user cost of capital: this is the long-run equilibrium relationship. On the steady-state growth path, firms remain at the optimal capital/output ratio by re-investing to offset depreciation and steady-state growth in the capital stock. With a stationary depreciation rate, this implies that in long-run equilibrium the investment/capital ratio is fixed. This is a second long-run equilibrium relationship.

In this paper we exploit a measure of the capital stock constructed at the Bank, and a real user cost of capital measure that explicitly incorporates relative prices. We relax the standard assumption of Cobb-Douglas technology that restricts the elasticity of substitution to unity, and instead use a constant elasticity of substitution (CES) production function that nests Cobb-Douglas as a special case. As described above, our theoretical framework implies two long-run equilibrium relationships: one between capital, output, and the real user cost; and the other between investment and capital. These theoretical long-run relationships imply restrictions on the model. They also imply a single reduced-form long-run relationship between investment, output and the real user cost.

We estimate this system as a vector error-correction mechanism (VECM) using the Johansen method. Our two long-run relationships form the basis for the two cointegrating vectors in the model. The model is statistically well specified and the overidentifying theoretical restrictions on the model are accepted. A key result is that the elasticity of substitution between labour and capital in production is significantly lower than unity at a little under 0.45. This estimate is obtained by a variety of measures and estimation techniques, and, as judged by external estimates, is plausible. This is a remarkable result, because most studies of aggregate investment have found it hard to find a significant relationship of the correct sign between investment and the user cost.

The model also tells us how investment and capital respond when the system is not in long-run equilibrium. Investment responds when the capital/output ratio is away from equilibrium, while capital responds when the investment/capital ratio is away from equilibrium. This last result is consistent with a log-linearisation of the capital accumulation identity. As with other aggregate investment models, the model takes a long time to reach the long-run equilibrium.

Despite the robust nature of our elasticity of substitution estimate, different estimation methods yield different results for the dynamics of investment. In particular, single-equation estimation results suggest that investment responds to disequilibrium in the investment/capital ratio, while our system estimation results suggest it does not respond to the investment/capital ratio.

We investigate this puzzle using simulations. We specify a model assuming the VECM results are correct, and use it to generate artificial data series for the four variables. Investment models are estimated on the artificial data using the single-equation and system techniques, and tested to see which technique correctly estimates the 'true' model. The system estimation is better at correctly estimating the dynamics than single-equation estimation, but rejects the restrictions from the theoretical long-run relationships too often. The single-equation results find the investment/capital ratio to be significant because they implicitly estimate the reduced-form long-run relationship, rather than the two separate long-run relationships.

E-barter versus fiat money: will central banks survive?

Working Paper no. 197

F H Capie, Dimitrios P Tsomocos and Geoffrey E Wood

Recent and extremely rapid development in computer technology has led to the emergence of what is called 'e-money'. This refers to technological developments which in effect give people much easier access to their bank accounts, and make the carrying of notes and coin unnecessary. Rather people carry 'electronic purses', cards which are loaded with monetary units in electronic form, and from which funds can be transferred directly, not intermediated through the banking system, onto another person's card or into a shop keeper's till. Those developments have been used so far on a geographically limited and essentially experimental basis, but this has not prevented conjecture that the development may go further.

This further stage is one where computer technology will replace money altogether. Goods will exchange directly for goods, and we shall return to barter, in electronic form, with computer technology lowering the costs of information storage and transmissions such that barter is a cheaper form of exchange than exchange using fiat money. In this paper the conditions under which the replacement could occur are analysed formally.

Key to our discussion is the medium of exchange function of money. We argue first that money evolved as a way of reducing the costs of transacting. It economises on information, by making all information about the buyer in a particular transaction irrelevant, and concentrating attention on what is being offered. As society evolves towards the use of a single money, so it evolves to a situation where the same information is needed for every transaction. One good will emerge as the money of a society, provided two conditions are satisfied. These are that not all goods are equally suitable as a medium of exchange and that the marginal costs of acquiring information about one good fall the more that good is used.

Having argued informally that the information-economising property of money is key to its evolution, a model of exchange based on that property is developed, and the cost of transacting in that system is compared with the cost of barter. The model is of a strategic market game, in which the stipulated means of exchange is fiat money and all transactions need cash in advance. (Note that constraint is not imposed arbitrarily, but emerges as a natural consequence of our prior argument that money exchange requires less information than does barter exchange.)

The model is then contrasted with one of 'electronic barter'. The costs of one method are compared with those of the other, and it is shown that unless inflation drives up the nominal interest rate substantially, fiat money exchange will continue to dominate electronic barter.

Further, it is shown that, if the government and the monetary authorities desire to do so, they can drive the costs of fiat money exchange towards zero. Accordingly, the paper concludes that in this model of an exchange economy fiat money will survive, and with it a meaningful and controllable price level. The paper also conjectures that these results would hold *a fortiori* in a model with production as well as exchange.

Non-interest income and total income stability

Working Paper no. 198

Rosie Smith, Christos Staikouras and Geoffrey Wood

It is generally believed that diversification by a firm reduces risk, just as diversification of investments by an individual does. In both cases, however, whether the desired risk reduction effect is achieved does of course depend on the correlation between the different activities or lines of business (in the case of the firms), and on the correlation between the prices of the different investments (in the case of the investing individual). Banks, like other firms, generally seek to reduce their risks by diversifying across various lines of business, although there is usually some degree of specialisation. In recent years, banks have started to move increasingly into areas that yield non-interest income-into activities that earn fees rather than interest. Some banks (traditionally, for example, the United Kingdom's merchant banks), have always concentrated on fee-earning activities, such as advising on how to raise capital. All banks have traditionally earned some fee income—the trustee business is a good example. The traditional fee-earning business was both a small part of most banks' earnings and indubitably little affected by such factors as fluctuations in the economy. But as the profitability of traditional banking activities has, for a wide variety of reasons, come under pressure in recent years, fee-earning activities have greatly increased their contribution to bank profits. It is therefore important to see whether these new activities offset fluctuations in other sources of earnings as successfully as did traditional fee-earning work. Addressing that question is the aim of this paper.

Having summarised the results of various earlier studies in the opening of the paper, the behaviour of interest and non-interest income is then discussed, first in broad outline for all EU banking systems, and then in some detail for the systems of Germany and the United Kingdom. The paper examines the variability of each source of income, as well as the correlation between these different sources. Where possible, the results are reported not just for the banking systems as a whole, in the countries of the EU but also by size and type of financial institution. By comparing the behaviour of fee and non-fee income the paper examines how the move into fee-earning activities changes the range of risk and return possibilities for banks. Other studies have considered whether increased fee income reduces or increases the riskiness of banks. But whether the change in risk is a result of banks' behaviour or of the nature of the two income streams is inevitably left obscure. This paper looks at the nature of non-interest income, concentrating in particular on the extent to which it represents not earnings from new activities, but earnings from performing the same business in a new way—for example earning a fee by arranging a loan for a customer rather than earning an interest spread by lending to the customer. Sources of non-interest income for depository institutions include securitisation and other major off balance sheet activities. Some of the different results obtained in this paper may be a product of the heterogeneity of fee-earning activity. The paper also considers why the changes have taken place; this may have implications both for the durability of the changes (were they the result of a passing fashion, or of some more durable change in conditions) and for future regulation or supervisory policy.

In broad outline, the paper finds that fee-earning income is less stabilising than seems generally believed; indeed, fee-earning income is actually more variable for most categories of banks than traditional interest income. It does, however, help in most cases to stabilise profit streams.

Credit risk diversification: evidence from the eurobond market

Working Paper no. 199

Simone Varotto

It is well known that portfolio risk can be reduced through diversification. Spreading portfolio holdings across countries and industrial sectors, for example, may help reduce portfolio volatility. It is less clear, however, whether these asset allocation strategies are effective in reducing return volatility from changes in credit spreads in a bond portfolio. While equity portfolio diversification has been widely investigated, diversification in portfolios of corporate bonds has only been analysed partially and sporadically.

This study looks at the effects of cross-country and industry diversification on credit risk. It also analyses other dimensions, namely maturity, seniority and credit rating diversification, because return uncertainty in bonds with different maturity, seniority and rating might be explained by different risk factors which are not perfectly correlated. For example, a firm's credit rating may determine the ease with which the firm can access financial markets for funding or decide the balance of power with customers and suppliers, when setting contractual obligations, such as terms of payments or speed of delivery. It follows that differences in credit standing may affect the firm's economic environment and the risk factors that influence the firm's profitability. This, in turn, allows the portfolio manager who invests in companies with varying credit quality to achieve diversification benefits.

The paper's analysis departs from the existing literature by introducing 'locally systematic' risk factors whose nature is systematic and idiosyncratic at the same time. Usually, diversification is defined as the reduction of idiosyncratic risk in the portfolio. The paper maintains the assumption that portfolio volatility is reduced through diversification of idiosyncratic risk as well as locally systematic risk. The latter is represented by country, industry, maturity, seniority and rating factors, estimated as deviations from the average market return. The average market return is truly systematic because it cannot be diversified away. Locally systematic risks, on the other hand, can be diversified away only if the portfolio is distributed across assets that are subject to different local factors. For example, to diversify the (locally systematic) German country effect in a portfolio of German bonds one needs to invest in other countries. Increasing the number of German securities would only reduce the idiosyncratic risk of the portfolio, narrowly defined as residual or unexplained bond volatility. Therefore, locally systematic risks are more persistent than idiosyncratic risk in that only a specific portfolio allocation strategy would cause their diversification. At the same time, they are not as persistent as the average market return since they too can be diversified away. This approach gives more structure to what was previously indistinctly described as idiosyncratic risk. It also provides a formal framework to describe phenomena that are already known and widely applied by portfolio managers.

The findings in the paper suggest that international diversification is most effective in reducing portfolio credit risk. Previous studies have shown that a similar conclusion also applies to equity risk. Surprisingly, diversification across maturity bands is found to be the second best strategy, superior to industry diversification.

Finally, the results may have a bearing on the ongoing debate on how to reform the current framework for setting banks' credit risk capital requirements. The capital adequacy rules in Pillar 1 of the New Basel Accord, as in the current Accord, do not take into account diversification effects on portfolio risk. Therefore, the results emphasise the potential importance of Pillar 2 of the new Accord within which supervisors are encouraged to take into account the extent of sectoral and geographical portfolio concentration when assessing the riskiness of banks relative to the capital they hold.