

Corporate Defaults and Large Macroeconomic Shocks*

Mathias Drehmann[†]
Bank of England

Andrew J. Patton[‡]
London School of Economics
and Bank of England

Steffen Sorensen[§]
Bank of England

First draft, October 2005

Abstract

We explore the impact of possible non-linearities on credit risk in a VAR set-up. We look at two measures of credit risk: quarterly aggregate liquidation rates and quarterly firm specific probability of defaults, which are derived by a new method from annual default data. We show three important results. First, non-linearities matter for the level and shape of impulse response functions of credit risk following small as well as large shocks to systematic risk factors. Second, in the non-linear model the impact of a shock depends significantly on the starting level of exogenous variables. Depending on actual conditions and the forecast horizon, the level as well as the sign of the impact can change. Third, we account for uncertainty in our estimates and show that ignoring this can lead to a substantial underestimation of credit risk, especially in extreme conditions.

Keywords: credit risk, impulse response functions, stress testing, nonlinear time series, VAR models
J.E.L. Codes: G33, C32.

*The authors would like to thank Phil Bunn for generously providing the extended dataset of corporate defaults. This paper represents the views of the authors and do not necessarily reflect those of the Bank of England or the Monetary Policy Committee members. All errors remain our responsibility.

[†]Corresponding author. Contact address: Financial Stability, Bank of England, Threadneedle Street, London EC2R 8AH. Email: mathias.drehmann@bankofengland.co.uk.

[‡]Contact address: Financial Markets Group, London School of Economics, Houghton Street, London WC2A 2AE. Email: a.patton@lse.ac.uk.

[§]Contact address: Financial Stability, Bank of England, Threadneedle Street, London EC2R 8AH. Email: stef-fen.sorensen@bankofengland.co.uk.

1 Introduction

Traded credit risk products have been one of the biggest growth areas in recent years. An accurate forecast of the probability of default (PD) of companies is one of the key variables for pricing these products. PDs have received a lot of attention from a capital setting perspective as well as banks can use their own PD estimates as input for regulatory capital under the new Basel II rules. Basel II also requires banks to stress test their credit portfolios. Essentially, a stress test is a forecast of how a shock to a systematic risk factor will impact on future losses. Conceptually, stress tests are similar to out-of-sample forecasting exercises. But, there is one key difference: the “out-of-sample” nature of a stress test lies in the consideration of large shocks to the predictor variables. That is, rather than trying to predict some variable for a time period beyond the sample period, a stress test tries to predict a variable conditioning on a large shock that sometimes even did not occur in the sample period. Stress tests generally study the impact of shocks that are ‘severe but plausible’. For normally distributed random variables this would consist of studying shocks of size between three and five standard deviations, rather than the one standard deviation shocks usually studied in macroeconomic analyses. If the data generating process (DGP) is linear, then this difference is not important and standard econometric models, such as vector autoregressions, are adequate. If interest lies in studying the impact of small shocks around the equilibrium of the process, then a standard linear model may produce adequate forecasts even if the true DGP is non-linear. In such a case the linear model may be interpreted as a first-order Taylor series approximation to the true DGP. However, stress tests do not consider small shocks. And, nor is it certain whether the data generating process is truly log-linear. We explore, therefore, the generally made assumption of log-linearity. We show that allowing for non-linearities leads to substantially different predictions of portfolio losses in scenarios typically considered for stress testing. But not only are predictions of non-linear models different when looking at large shocks. We also show that linear models tend to overestimate credit risk for small shocks. Hence, accounting for non-linearities is important for the setting of capital as well as the pricing of credit risk.

In the first part of the paper we look at a reduced form credit risk model. First we estimate a non-linear vector auto regression model (VAR) of the underlying macroeconomic drivers of credit risk where we concentrate on the three key macroeconomic factors: GDP growth, inflation and the interest rate. We then investigate how macroeconomic shocks feed through to the aggregate liquidation rate, used as a first proxy for aggregate credit risk.

Nonlinear models for macroeconomic variables have been studied by Koop, *et al.* (1996) and

Potter (2000). But in this paper we employ the methodology of Jorda (2005). Jorda's approach builds on the fact that a standard VAR can be interpreted as a first-order approximation to the true unknown DGP. Thus, a more flexible approximation may be obtained by considering, for example, a quadratic or cubic approximation. An important implication of considering a standard linear VAR as a linear approximation to the true DGP is that it is no longer clear that forecasts or stress tests of horizons greater than one period should be obtained by iterating the one-period model forward which is the standard practice when deriving impulse response functions for VAR models. As Jorda (2005) notes, if the one-period model is mis-specified, then iterating it forward may well lead to a compounding of mis-specification error. He suggests an alternative approach, namely to estimate a different approximation model for each horizon of interest. If the DGP is truly a VAR then this approach is consistent but not efficient, while if the DGP is not a VAR then this approach offers the best approximation at each horizon, rather than just at the one-quarter horizon. This modelling approach has its roots in the direct multi-step versus iterated forecasting approaches (see for example Stock and Watson, 1999). A benefit of piece wise regressions is that simple ordinary least square (OLS) techniques can be used to estimate the non-linear VAR.

Whereas aggregate liquidation rates may be useful as a proxy for aggregate credit risk, they could severely underestimate the impact of large shocks to PDs of individual companies and, thus, the risk of an overall portfolio. Starting with Beaver (1966) and Altman (1968) there is a large literature on the drivers of bankruptcy. Early papers such as Altman's (1968) Z-score or Ohlson's (1980) O-score use accounting variables such as leverage as explanatory variables. Amongst others Wilson (1997a and 1997b)¹ identified not only firm specific but also macroeconomic factors as systematic risk drivers of PDs². More recent models, such as Shumway (2001) or Hillegeist et al (2004), use hazard rate models with time dependent covariates to predict PDs. Most of the latest models find that accounting data, macroeconomic factors and market data, for example the Merton model (1976) implied distance to default, have a significant impact on PDs. Generally, the literature focuses on predicting PDs over a one year horizon. But it has been recognized recently (see for example Duffie et al, 2005) that the forecast horizon matters. For example, the conditional probability of default during the second year must not be the same as for the first year, even if the predictor variables remain constant. Therefore, Duffie et al (2005) estimate the term structure of

¹Several papers apply this model in a stress testing context (see e.g. Boss, 2002, or Virolainen, 2004).

²While both idiosyncratic risk and systematic risk factors affect PDs, the difference is that idiosyncratic factors are uncorrelated across firms. What matters for stress testing and portfolio management are systematic risk factors which impact on all obligors in a portfolio and are correlated between each other.

hazard rates based on a mean reverting time series process for macro and firm specific variables. Alternatively, Campbell et al. (2005) undertake a more reduced form approach and directly estimate probabilities of defaults at particular horizons in the future, conditional on survival up to that period.

In line with Campbell et al (2005) we follow a reduced form approach in the second part of the paper and estimate quarterly PDs for 2 years, conditional on survival until that quarter. But we introduce several innovations. First, we extend Jorda's methodology to estimate the potential non-linear impacts of macroeconomic variables on PDs. Note that the estimated probit specification is already a non-linear transformation from the independent macroeconomic and firm specific variables into default forecasts. But Jorda's methodology goes further as it allows the inclusion of squares and cubes in the specification, which should lead to a better fit especially in the tails of the distribution.

Second, we explicitly consider estimation uncertainty in our forecasts by constructing confidence intervals around our PD estimates. While this is routinely done for VAR models, this has so far been neglected when looking at credit risk models. We show that ignoring estimation uncertainty can lead to a significant underestimation of credit risk.

The third innovation of this paper is that we show a new method to derive consistent quarterly PDs from annual data. The main drawback of the available dataset of defaulting and non-defaulting UK companies is that defaults and company account data are only reported annually. Ideally one would have quarterly recordings of defaults as this would allow us to link the credit risk to quarterly macroeconomic developments. By assuming that the quarterly liquidation rate truly reflects the number of defaults in a given quarter, we show that it is possible to construct a quarterly proxy series from annual data and generate unbiased estimates of the impact of quarterly macroeconomic variables on quarterly PDs.

All our PD estimates are based on a large dataset of over 30,000 UK companies over the period 1991 to 2004. The dataset covers all UK industries and firms of various sizes. Such a rich dataset should reduce the overall estimation error of our specification. But the main benefit of it is that the data cover publicly traded as well as private companies, and, therefore, should reflect the exposure of banks much better when used for a stress test in contrast to a model based on a Merton type credit risk models. The latter requires equity data as input, which are only available for a relatively small number of companies in the UK³.

We show that the results of the non-linear VAR are significantly different to results using standard linear models, especially when considering large shocks. This can be seen in the simple

³See Drehmann (2005) for a credit risk stress test based on a Merton model.

three variable macro model of inflation, GDP and a short term interest rate. More importantly, we show that accounting for non-linearities in the underlying macroeconomic environment leads to substantially different conclusions for credit risk projections in stressed conditions. For the firm specific model, we show that for small shocks linear models seem to overestimate credit risk, whereas for large ones they tend to underestimate it. We also show that when using the non-linear model the levels of macroeconomic variables before the shock occurs matter. Not only do they impact on the level of projected credit risk, but they also influence the shape of the impulse response function.

However, we would caution against a literal interpretation of our results at the moment as we have not undertaken sufficient robustness checks yet. Preliminary robustness checks indicated that in some cases the shape of the impulse responses can depend on sample of the estimation. However, overall conclusions seem robust. If future research confirms our findings, our results would have serious implication not only from a regulatory perspective concerned about capital levels at 99.9% confidence level but also from a pricing perspective. Given the rapid increase in trade credit risk products and the introduction of Basel II, this should be of great interest to market participants.

The remainder of the paper is as follows. In Section 2 we present a more formal motivation for the consideration of non-linear multivariate models when studying the impact of large shocks. In Section 3 we discuss the estimation of the macro model and the resulting impulse response functions. In Section 4 we introduce our model for liquidation rates as well as corporate defaults, and present the results of our analysis of large macroeconomic shocks on default probabilities. Section 5 concludes. Technical details and estimation results are presented in the Appendix.

2 Why non-linearities matter

Suppose we are interested in a scalar variable y_t , which is follows the following general process:

$$y_t = h(y_{t-1}, \varepsilon_t; \theta), \quad t = 1, 2, \dots \quad (1)$$

where the residual ε_t is independent of y_{t-1} , h is some (possibly non-linear) function, and θ is a parameter vector. In the standard linear setting we would have

$$h(y_{t-1}, \varepsilon_t; \theta) = \phi_0 + \phi_1 y_{t-1} + \varepsilon_t \quad (2)$$

and so y_t would follow a first-order autoregressive process. In this case the conditional mean function,

$$\mu(y) \equiv E[y_t | y_{t-1} = y] = \phi_0 + \phi_1 y \quad (3)$$

is affine in y_{t-1} and so a first-order Taylor series approximation of μ corresponds exactly to μ . For other data generating processes the conditional mean function need not be affine in y_{t-1} . For example, in Appendix A we describe a simple non-linear specification for h , under which y_t is unconditionally Normally distributed, the first-order autocorrelation coefficient is 0.5, and the conditional mean function is nonlinear. The conditional mean function is plotted in Figure 1.




Figure 1: The conditional mean of Y_t given Y_{t-1} assuming $Y_t \sim N(0, 1)$ and using a Clayton copula implying first-order autocorrelation of 0.5.

The figure illustrates why linear approximations may be inadequate for stress testing: such approximations may work satisfactorily in the middle of the distribution, but can perform poorly in the tails. In this example the linear approximation matches the true conditional mean function reasonably well for $|Y_{t-1}| \leq 1.5$, but deviates outside this region⁴. For the study of “small” shocks to this variable the linear approximation may be acceptable, but for the study of large shocks (two or more standard deviations) it is not. For example, if $Y_{t-1} = -3$, corresponding to a three

⁴The linear approximation to the true conditional mean function is obtained by noting that the optimal mean squared error approximation is a line through zero with slope equal to one-half. This follows from the fact that $E[Y_t] = 0$, $V[Y_t] = 1$ and $Cov[Y_t, Y_{t-1}] = 0.5$. The quadratic approximation can similarly be derived analytically from the properties of the joint distribution of (Y_t, Y_{t-1}) .

standard deviation shock in this setting, the linear approximation would predict $Y_t = -1.5$, while the true conditional expected value of Y_t is -2.7 . Consider now a quadratic approximation to the conditional mean function, also plotted in Figure 1. This approximation is very close to the true conditional mean function for $|Y_{t-1}| \leq 2$. More importantly, for our interest in studying large shocks, the quadratic approximation also does a lot better in the tails. Continuing our previous example, if $Y_{t-1} = -3$ the quadratic approximation predicts $Y_t = -2.8$, close to the true value of -2.7 . Thus this simple example highlights the potential for more flexible models to provide better estimates of the impact of “large” shocks.

2.1 A nonlinear VAR(p) model

Consider a $(k \times 1)$ vector of macroeconomic variables, \mathbf{Y}_t . By far the simplest and most widely used model for the dynamics of macroeconomic time series is the vector autoregression.

$$\mathbf{Y}_t = B_0 + B_1 \mathbf{Y}_{t-1} + \dots + B_p \mathbf{Y}_{t-p} + \mathbf{e}_t$$

More generally, we can think about the mapping between $\mathbf{Y}_{t-1}^p \equiv [\mathbf{Y}'_{t-1}, \dots, \mathbf{Y}'_{t-p}]'$ and \mathbf{Y}_t as some general unknown function, $\mathbf{g} : \mathbb{R}^{pk} \rightarrow \mathbb{R}^k$

$$\mathbf{Y}_t = \mathbf{g}(\mathbf{Y}_{t-1}^p) + \mathbf{e}_t \quad (4)$$

and interpret the standard VAR above as a simple first-order Taylor series approximation to the unknown function \mathbf{g} . For notational simplicity let us assume that all variables have mean zero.

$$\begin{aligned} \mathbf{g}(\mathbf{Y}_{t-1}) &= \mathbf{g}(\mathbf{0}) + \nabla \mathbf{g}(\mathbf{0}) \mathbf{Y}_{t-1}^p + R_1 \\ &\approx \mathbf{g}(\mathbf{0}) + \nabla \mathbf{g}(\mathbf{0}) \mathbf{Y}_{t-1}^p \\ &\equiv B_0 + B_1^p \mathbf{Y}_{t-1}^p \\ &\equiv B_0 + B_1 \mathbf{Y}_{t-1} + \dots + B_p \mathbf{Y}_{t-p} \end{aligned} \quad (5)$$

If we are primarily interested in studying the dynamics of \mathbf{Y}_t “near” its unconditional mean, then the first-order approximation of \mathbf{g} provided by a standard VAR may be sufficiently accurate. By convention, VAR studies show impulse response functions to one standard deviation shocks. But it is well known (see for example Koop, *et al.*, 1996) that for standard linear VAR models the magnitude of the shock has no impact on the shape of the impulse response function; it merely affects the scale. As discussed above, in stress testing studies interest lies not in small- or medium-sized shocks, but extreme shocks. Considering three or five standard deviation shocks means

considering the dynamics of the variables “far” from their unconditional mean, which may lead us to question the quality of a first-order Taylor series approximation to the true, unknown dynamics.

An obvious extension is to expand to a second or third-order Taylor series approximation of \mathbf{g} . First, we look at the second order expansion for the first element of \mathbf{g} , denoted g_1 , where $g_1 : \mathbb{R}^{pk} \rightarrow \mathbb{R}$

$$g_1(\mathbf{Y}_{t-1}^p) \approx g_1(\mathbf{0}) + \nabla g_1(\mathbf{0}) \mathbf{Y}_{t-1}^p + \frac{1}{2} \mathbf{Y}_{t-1}^{p'} \nabla^2 g_1(\mathbf{0}) \mathbf{Y}_{t-1}^p \quad (6)$$

We can re-write this expression in a more convenient form by making use of the *vec* operator and the Kronecker product (denoted \otimes):

$$\begin{aligned} g_1(\mathbf{Y}_{t-1}) &\approx g_1(\mathbf{0}) + \nabla g_1(\mathbf{0}) \mathbf{Y}_{t-1}^p + \frac{1}{2} \underbrace{\text{vec}(\nabla^2 g_1(\mathbf{0}))'}_{(pk \times pk)} \underbrace{(\mathbf{Y}_{t-1}^p \otimes \mathbf{Y}_{t-1}^p)}_{(pk \times 1) \quad (pk \times 1)} \\ &= g_1(\mathbf{0}) + \nabla g_1(\mathbf{0}) \mathbf{Y}_{t-1}^p + \frac{1}{2} \text{vec}(\nabla^2 g_1(\mathbf{0}))' \text{vec}(\mathbf{Y}_{t-1}^p \mathbf{Y}_{t-1}^{p'}) \end{aligned}$$

We can stack the equations to obtain:

$$\begin{aligned} \mathbf{g}(\mathbf{Y}_{t-1}) &\approx \mathbf{g}(\mathbf{0}) + \nabla \mathbf{g}(\mathbf{0}) \mathbf{Y}_{t-1}^p + \frac{1}{2} \begin{bmatrix} \text{vec}(\nabla^2 g_1(\mathbf{0}))' \\ \text{vec}(\nabla^2 g_2(\mathbf{0}))' \\ \vdots \\ \text{vec}(\nabla^2 g_k(\mathbf{0}))' \end{bmatrix} \text{vec}(\mathbf{Y}_{t-1}^p \mathbf{Y}_{t-1}^{p'}) \\ \text{Let } \nabla^2 \mathbf{g}(\mathbf{0}) &\equiv \begin{bmatrix} \text{vec}(\nabla^2 g_1(\mathbf{0}))' \\ \text{vec}(\nabla^2 g_2(\mathbf{0}))' \\ \vdots \\ \text{vec}(\nabla^2 g_k(\mathbf{0}))' \end{bmatrix} \\ \text{Then } \mathbf{g}(\mathbf{Y}_{t-1}^p) &= \mathbf{g}(\mathbf{0}) + \nabla \mathbf{g}(\mathbf{0}) \mathbf{Y}_{t-1}^p + \frac{1}{2} \nabla^2 \mathbf{g}(\mathbf{0}) \underbrace{(\mathbf{Y}_{t-1}^p \otimes \mathbf{Y}_{t-1}^p)}_{(p^2 k^2 \times 1)} \\ &\equiv B_0 + B_1^p \mathbf{Y}_{t-1}^p + \underbrace{B_2^p \text{vech}(\mathbf{Y}_{t-1}^p \mathbf{Y}_{t-1}^{p'})}_{(pk(pk+1)/2 \times 1)}, \text{ collecting terms} \\ &\equiv B_0 + \sum_{m=1}^p B_{1m} \mathbf{Y}_{t-m} + \sum_{i=1}^p \sum_{j=i}^p B_{2ij} \underbrace{\text{vech}(\mathbf{Y}_{t-i} \mathbf{Y}_{t-j}')}_{(k(k+1)/2 \times 1)}, \text{ collecting terms} \end{aligned} \quad (7)$$

where *vech*(X) stacks only the lower triangle of the matrix X . We use the *vech* function rather than the *vec* function as $\mathbf{Y}_{t-1}^p \mathbf{Y}_{t-1}^{p'}$ includes both $Y_{1,t-1} Y_{2,t-1}$ and $Y_{2,t-1} Y_{1,t-1}$, for example, and we can collect such terms. Moving from the penultimate to the final line above also follows from a collection of terms, further reducing the number of free parameters.

Note that the number of unknown parameters is larger for the second-order Taylor series approximation than the first-order: the first-order approximation has $k + pk^2$ free parameters, while the more flexible model has $k + pk^2 + pk^2(p + 1)(k + 1)/4$ free parameters. We can consider numerous methods for reducing the number of free parameters: one possibility is to restrict all second-order effects in equation i to include $Y_{t-m,i}$. Alternatively, we could restrict the second-order terms to only include lagged squared terms. The same analysis can easily be repeated for the third-order Taylor series approximation, adding greater flexibility at the cost of more free parameters.

3 Estimation results for the non-linear macro VAR

3.1 Estimation of flexible non-linear approximations

We employ a third-order approximation in our models of the relationships between the macroeconomic variables and the measures of corporate default. In the interests of parsimony we drop all cross-product terms from this approximation, and consider only one lag of the higher-order terms. Thus, our model for the macroeconomic variables at the one-quarter horizon is:

$$Y_{jt} = \beta_{0j} + \sum_{m=1}^p \beta'_{1jm} \mathbf{Y}_{t-m} + \gamma'_{2j} \mathbf{Y}_{t-1} \odot \mathbf{Y}_{t-1} + \gamma'_{3j} \mathbf{Y}_{t-1} \odot \mathbf{Y}_{t-1} \odot \mathbf{Y}_{t-1} + e_{jt} \quad (8)$$

for $j = 1, 2, 3$, where \odot is the Hadamard product. This model is estimable via OLS, and thus is very simple to implement.

As discussed in the introduction, an implication of considering a standard linear VAR as a linear approximation to the true DGP is that it might be problematic that forecasts or stress tests of horizon greater than one period should be obtained by iterating the one-period model forward. Jorda (2005) proposes an alternative approach, namely to estimate a different approximation model for each horizon of interest. Following this argument, we, therefore, have a set of models for the three macroeconomic variables and eight horizons:

$$Y_{j,t+h-1} = \beta_{0j}^h + \sum_{m=1}^p \beta_{1jm}^{h'} \mathbf{Y}_{t-m} + \gamma_{2j}^{h'} \mathbf{Y}_{t-1} \odot \mathbf{Y}_{t-1} + \gamma_{3j}^{h'} \mathbf{Y}_{t-1} \odot \mathbf{Y}_{t-1} \odot \mathbf{Y}_{t-1} + e_{jt}^h \quad (9)$$

for $j = 1, 2, 3$ and $h = 1, 2, \dots, 8$. For each variable and each horizon this model is estimable via OLS.

It should be noted that estimating the model for each horizon of interest has the additional benefit of making the obtaining of confidence intervals on the forecast or stress test very simple:

they come directly from the covariance matrix of the parameters estimated for each horizon. This is in contrast with the traditionally used linear VAR approach, where the confidence intervals for horizons greater than one period must be obtained either via the “delta” rule, or a bootstrap procedure.

3.2 Data

We use data on three key macroeconomic variables, GDP growth, the three-month Treasury bill rate, and the inflation rate, to summarise the state of the macro economy. Our macro model is small relative to some of the macroeconomic models used in the analysis of credit risk, Pesaran, *et al.* (2005) being a prominent example. But it is large enough to convey the main ideas of this paper.

Our sample period is 1992Q4 to 2004Q3. These series are available for a much longer period, but we focus on data after 1992Q4. At this point the UK adopted an inflation targeting regime and it has been recognized that inflation targeting in the UK and other countries lead to a significant reduction in the volatility of macroeconomic series (see for example Kuttner and Posen, 1999 or Benati 2004). It is, therefore, reasonable to assume that the introduction of inflation targeting induces a structural break in the macroeconomic time series of the UK in 1992Q4. In Figure 13 in the Appendix we plot the macroeconomic variables.

3.3 The estimated macroeconomic non-linear VAR

It has been long understood (for example see Koop *et al.*, 1996) that for standard linear VARs the size and the sign of the shock do not change the shape of the impulse response function (IRF). Furthermore, starting values are not material. However, for the non-linear VAR the size, the sign and starting values are important. In all cases we evaluate the IRF holding all non-shocked variables at their unconditional averages⁵. Consistent with the extant macroeconomic literature, we order the variables as GDP growth, inflation, interest rate.

In line with Jorda (2005) and much of the VAR literature we use a Cholesky factorisation of the Covariance matrix of errors to obtain our scenarios. What we, therefore, look at are unexpected one and three standard deviation shocks to GDP, inflation and the interest rate. Even though the Cholesky factorisation takes account off the correlation between the unexpected shocks to GDP and unexpected shocks to inflation and the interest rate (and, respectively unexpected shocks to

⁵For details on the computation of IRFs in this setting see Jordà (2005).

inflation and unexpected shocks to the interest rate) we label these scenarios as GDP, inflation and interest rate shocks. With the simple Cholesky factorisation we are unable to identify for example whether GDP shocks are driven by demand or supply shocks, which would be of interest from a monetary policy perspective. But for this paper a simple scenario selection is sufficient to illustrate our results. And the methodology is general enough to consider any scenario in the future.

In Figures 14 - 17 in the Appendix we plot the impulse response functions (IRFs) of the three-variable macroeconomic VAR to shocks of various sizes and signs. Figure 14 reveals that in most one-standard deviation IRFs the cubic and the linear models yield similar results. However the response of interest rates to GDP growth shocks and interest rate shocks do differ substantially: the response of interest rates to a GDP growth shock and an interest rate shock is significantly greater, for horizons 1 through 4 quarters, if cubic terms are considered than if these are ignored.

Figure 15 shows the IRFs for a -1 standard deviation shock. For the linear model this figure is just a sign change of the plots in Figure , whereas this is not necessarily so for the cubic model. For example, the response of interest rates to a positive GDP shock was significantly greater using the cubic model than the linear model, whereas this difference between the two models essentially disappears for a negative GDP growth shock.

In Figure 16 we present the IRFs for a positive 3 standard deviation shock. For the linear model these IRFs are just 3 times the IRFs from Figure 14, while this is not so for the cubic model. Some interesting differences appear comparing Figures 14 and 16. For example, the response of interest rates to a 1 standard deviation inflation shock was small and positive (negative) for the linear (cubic) model, slowly increasing as the horizon approached eight quarters. However, for a three standard deviation shock the cubic model suggests a large positive response of interest rates for the first 5 quarters, followed by a loosening of interest rates at the 8th quarter. This indicates a difference between small shocks to inflation, which lead to modest changes in interest rates, and very large shocks to inflation, which lead to much different interest rate reactions.

4 Macroeconomic risks and corporate default probabilities

In addition to non-linearities between the macroeconomic risk factors, the linkage between macroeconomic shocks and corporate default probabilities may be non-linear. Hence, it is not obvious that, for example, a positive three standard deviation shock to interest rates should affect default probabilities by exactly three times the impact of a one standard deviation shock. Of course, since the liquidation rate is bounded between zero and one, and a corporate default indicator is either 0

or 1, the models used to analyse defaults are usually non-linear in the first instance, such as Logit or Probit models for example. But as discussed in the Introduction, we go further than this and apply Jorda's (2005) methodology to allow for non-linearities in the underlying process.

In our first analysis of macroeconomic shocks and corporate default, we employ an aggregate liquidation rate to summarise the economy-wide probability of corporate default. This series has the benefit of being easy to model and interpret, but is by its nature just a summary variable. Our second analysis employs a large panel of corporate default indicators, across over 30,000 firms and 14 years. We use standard Probit analysis to relate accounting information and macroeconomic shocks to probabilities of default. By carefully considering how current and lagged macroeconomic variables affect corporate defaults we are able to use the impulse response functions from the non-linear VAR to trace out the impact of large macroeconomic shocks up to two years into the future.

4.1 The estimated impact of macroeconomic shocks on corporate liquidation rates

To estimate the impact of macroeconomic shocks on the aggregate liquidation rate we estimate the following Logit model:

$$\Lambda^{-1}(P_{t+h-1}) = \beta_0^h + \alpha_1^h \Lambda^{-1}(P_{t-1}) + \beta_{1j}^{h'} \mathbf{Y}_{t-1} + \gamma_{2j}^{h'} \mathbf{Y}_{t-1} \odot \mathbf{Y}_{t-1} + \gamma_{3j}^{h'} \mathbf{Y}_{t-1} \odot \mathbf{Y}_{t-1} \odot \mathbf{Y}_{t-1} + e_{t+h-1} \quad (10)$$

for $h = 1, 2, \dots, 8$, where $\Lambda(x) = 1/(1 + e^{-x})$ is the standard logistic function and $\mathbf{Y}_t = [Y_{1t}, Y_{2t}, Y_{3t}]'$ is the vector of the three macroeconomic variables. Since the model above involves a nonlinear transformation of the liquidation rate, to obtain the results from a stress test, or impulse response function, we use a simulation-based method: The stress test results and confidence intervals are obtained by estimating the above model and then simulating 10,000 draws from the asymptotic distribution of the parameter estimates, and a Normal distribution for the regression residual, to obtain an estimated impact of a shock above what would be observed when all variables are held at their unconditional values. That is, if $\tilde{\mathbf{Y}}_{t-1}$ represents the vector of stressed values of the the

macroeconomic variables, then:

$$\begin{aligned}
\tilde{P}_{t+h-1}^{(j)} \left(\tilde{\mathbf{Y}}_{t-1} \right) &= \Lambda \left(\hat{\beta}_0^{(j)h} + \hat{\alpha}_1^{(j)h} \Lambda^{-1} (\bar{P}) + \hat{\beta}_{1j}^{(j)h'} \tilde{\mathbf{Y}}_{t-1} \right. \\
&\quad \left. + \hat{\gamma}_{2j}^{(j)h'} \tilde{\mathbf{Y}}_{t-1} \odot \tilde{\mathbf{Y}}_{t-1} + \hat{\gamma}_{3j}^{(j)h'} \tilde{\mathbf{Y}}_{t-1} \odot \tilde{\mathbf{Y}}_{t-1} \odot \tilde{\mathbf{Y}}_{t-1} + e_{t+h-1}^{(j)} \right) \\
\bar{P}_{t+h-1}^{(j)} &= \Lambda \left(\hat{\beta}_0^{(j)h} + \hat{\alpha}_1^{(j)h} \Lambda^{-1} (\bar{P}) + \hat{\beta}_{1j}^{(j)h'} \bar{\mathbf{Y}} \right. \\
&\quad \left. + \hat{\gamma}_{2j}^{(j)h'} \bar{\mathbf{Y}} \odot \bar{\mathbf{Y}} + \hat{\gamma}_{3j}^{(j)h'} \bar{\mathbf{Y}} \odot \bar{\mathbf{Y}} \odot \bar{\mathbf{Y}} + e_{t+h-1}^{(j)} \right) \\
IRF^{(j)} \left(h, \tilde{\mathbf{Y}}_{t-1} \right) &\equiv \tilde{P}_{t+h-1}^{(j)} \left(\tilde{\mathbf{Y}}_{t-1} \right) / \bar{P}_{t+h-1}^{(j)} \\
\text{where } \left\{ \left[\begin{array}{c} \hat{\boldsymbol{\theta}}^{(j)h} \\ e_{t+h,t}^{(j)} \end{array} \right] \right\}_{j=1}^J &\sim iid N \left(\left[\begin{array}{c} \hat{\boldsymbol{\theta}}^h \\ \mathbf{0} \end{array} \right], \left[\begin{array}{cc} \hat{V}_\theta^h & \mathbf{0} \\ \mathbf{0} & \hat{\sigma}_{e,h}^2 \end{array} \right] \right) \\
\text{and } \hat{\boldsymbol{\theta}}^{(j)h} &\equiv \left[\hat{\beta}_0^{(j)h}, \hat{\alpha}_1^{(j)h}, \hat{\beta}_{1j}^{(j)h'}, \hat{\gamma}_{2j}^{(j)h'}, \hat{\gamma}_{3j}^{(j)h'} \right]'
\end{aligned}$$

where \hat{V}_θ^h is the covariance matrix of the parameter estimates, and $\hat{\sigma}_{e,h}^2$ is the variance of the residual from the h -horizon regression. We set $J = 10,000$ and present the mean, 0.025 and 0.975 quantiles of the simulated distribution of $IRF^{(j)} \left(h, \tilde{\mathbf{Y}}_{t-1} \right)$. As for the macroeconomic nonlinear VAR in the previous section, we consider stressing each of the three macroeconomic variables via the Cholesky factorisation of their unconditional covariance matrix.

4.2 Creating quarterly firm default data from annual firm default data and quarterly liquidation rates

A problem with the panel dataset is that we have information on the default status of the company only within a specific year. Ideally we would want to know in which quarter a failing firm defaults to get a better understanding of the shape of the IRF. But by constructing a proxy series from the quarterly observable liquidation rate and the annual company accounts data, we can estimate unbiased quarterly PD models. For our method to hold we only require that the distribution of defaulting companies across the quarters in each fiscal year the same as that for the entire economy (which is summarised by the liquidation rate series). This assumption is not verifiable with the annual data we have available, but given the large number of companies covered by our data set it seems a reasonable assumption.

$$\text{Let } X_{it} = \begin{cases} 1, & \text{if firm } i \text{ defaulted in quarter } t \\ 0, & \text{else} \end{cases} = \text{quarterly default indicator.}$$

$$Y_{i,4t} = \sum_{j=0}^3 X_{i[4t]-j} = \text{annual default indicator, only observed every 4 quarters.}$$

where $[a]$ = smallest integer greater than a .

$$p_t = \frac{1}{K_t} \sum_{i=1}^{K_t} X_{it} = \text{liquidation rate for quarter } t.$$

$$\tilde{p}_{4t} = \sum_{j=0}^3 p_{[4t]-j} = \text{annual liquidation rate, only observed every 4 quarters.}$$

$$Z_{it} = \text{macro variables and firm specific variables.}$$

If we could observe X_{it} we could run the following Logit regression:

$$X_{it} = \Lambda(Z_{it-1}\beta_0) + \varepsilon_{it}$$

and so $E[X_{it}|Z_{it-1}] = \Lambda(Z_{it-1}\beta_0)$

But we only observe liquidation rates and annual default indicators. Therefore, we consider the following proxy:

$$\hat{X}_{it} = \begin{cases} 0, & \text{if } Y_{i,[t/4]} = 0 \\ p_t/\tilde{p}_{[t/4]} & \text{else} \end{cases}$$

That is, \hat{X}_{it} is equal to zero if the annual default indicator for that year equals zero. In this case we know $\hat{X}_{it} = X_{it}$, and this proxy is without error. When the annual default indicator for that year equals 1 we set \hat{X}_{it} equal to the probability that this firm defaulted in this quarter. Given the observed liquidation rate this probability equals $p_t/\tilde{p}_{[t/4]}$. Notice that

$$\hat{X}_{it} = \Pr[X_{it} = 1 | \ddot{Y}, \ddot{p}] = E[X_{it} | \ddot{Y}, \ddot{p}]$$

where \ddot{Y} and \ddot{p} are the whole samples of $Y_{i,t}$ and $p_{i,t}$ (including past, present and future values).

Our key assumption is

$$E[X_{it} | \ddot{Y}, \ddot{p}] = E[X_{it} | \ddot{Y}, \ddot{p}, Z_{i,t-1}] \text{ for all } (\ddot{Y}, \ddot{p}, Z_{i,t-1}) \quad (11)$$

This equality is known to hold for $Y_{i,[t/4]} = 0$. When $Y_{i,[t/4]} = 1$ the equality may not hold, but as \ddot{Y} and \ddot{p} are both already functions of $Z_{i,-1}$, it is reasonable to assume that the expectation of the

quarterly default indicator, conditional on \ddot{Y} and \ddot{p} , is independent of Z_{it-1} . In other words, $Z_{i,t-1}$ carries no additional information beyond that which is conveyed through \ddot{Y} and \ddot{p} .

If the equality in 11 holds, then

$$\begin{aligned} E \left[\hat{X}_{it} | Z_{it-1} \right] &= E \left[E \left[X_{it} | \ddot{Y}, \ddot{p} \right] | Z_{it-1} \right] \\ &= E \left[E \left[X_{it} | \ddot{Y}, \ddot{p}, Z_{it-1} \right] | Z_{it-1} \right], \text{ by assumption} \\ &= E \left[X_{it} | Z_{it-1} \right], \text{ by the law of iterated expectations} \end{aligned}$$

which implies that we can estimate

$$\hat{X}_{it} = \Lambda (Z_{it-1} \beta_0) + \varepsilon_{it}$$

and obtain unbiased parameter estimates of the infeasible regression using the true X_{it} variable. The parameter estimates we obtain will, in general, be less precise than the case where X_{it} is observable, but will nevertheless be unbiased. This reasoning allows us to combine the information in our annual default indicator series and our quarterly liquidation rate series into an proxy quarterly default indicator series. Using standard maximum likelihood techniques (see for example Davidson and MacKinnon, 1993) we are, thus, able to obtain quarterly Probit parameter estimates that are unbiased. This allows us to generate impulse response functions.

We estimate the following Probit models

$$\Lambda^{-1} (P_{t+h-1}) = \beta_0^h + \alpha_1^h \Lambda^{-1} (P_{t-1}) + \beta_{1j}^{h'} \mathbf{Y}_{t-1} + \gamma_{2j}^{h'} \mathbf{Y}_{t-1} \odot \mathbf{Y}_{t-1} + \gamma_{3j}^{h'} \mathbf{Y}_{t-1} \odot \mathbf{Y}_{t-1} \odot \mathbf{Y}_{t-1} + e_{t+h-1} \quad (12)$$

$$\begin{aligned} \hat{X}_{i,t+h-1} &= \Phi \left(\beta_0^h + \alpha_1^h \mathbf{W}_{i,t-1} + \beta_{1j}^{h'} \mathbf{Y}_{t-1} \right. \\ &\quad \left. + \gamma_{2j}^{h'} \mathbf{Y}_{t-1} \odot \mathbf{Y}_{t-1} + \gamma_{3j}^{h'} \mathbf{Y}_{t-1} \odot \mathbf{Y}_{t-1} \odot \mathbf{Y}_{t-1} \right) + e_{t+h-1} \end{aligned}$$

for $i = 1, 2, \dots, K_{t+h-1}$, the number of firms in the data base at time $t+h-1$, and for $h = 1, 2, \dots, 8$. The firm-specific balance sheet variables are included in $\mathbf{W}_{i,t-1}$ ⁶.

4.3 Data

Our first measure of aggregate credit risk is the aggregate liquidation rate for the UK. This is defined as the number of companies going into liquidation as a percentage of the stock of active

⁶The above model could have been estimated using any other distribution, such as the widely used Logit model. We used the Logit model as a robustness check and found the difference in results indistinguishable.

companies. While firms may suffer financial distress without formally entering liquidation, this measure should be highly correlated with actual aggregate credit risk. The liquidation rate from 1992 onwards is shown in Figure 13 in the Appendix.

The second measure of credit risk is based on company accounts level data for all public as well as private companies registered in the UK from 1991 to 2004⁷. This dataset is an up-dated version of the dataset used in Bunn and Redwood (2003) and is described in Table B1 and B2 in the Appendix. Due to incomplete accounting information we exclude all companies with less than 100 employees. Following Bunn and Redwood (2003), we also delete any observations if it has or more missing values for any explanatory variable used. A default is recorded for a company in a specific year, if we observe accounts in the previous year and if the company is either in receivership, liquidation or dissolved. Take-overs are not considered a failure. But this definition includes voluntary liquidation and dissolution as we are unable to distinguish between voluntary and compulsory failures. However, we do not consider this to be a large distortion of our dataset. Individual firm can appear as a separate observation in each sample year. The data cover over 30,000 UK companies between 1991 to 2004. With the exception of financial services, the dataset includes companies of all UK industries and firms of various sizes. The mean default rate in the whole sample is 1.78%.

As our focus is on systematic risk factors for PDs, we do not experiment with different firm specific explanatory variables as found in Bunn and Redwood (2003) for the same data. The firm specific factors are: the interest cover, the current ratio, the debt to asset ratio, the number of employees, the profit margin and industry dummies. However, in line with our macro variables we also look at higher powers of the firm specific explanatory variables. We find that the square and cube of interest cover and the number of employees are significant in explaining the firm specific PDs.

4.4 Impulse response function for the liquidation rate

In Figure 2 and 3 and Figures 18 and 19 in the Appendix we present the results of the stress tests for the liquidation rate. Again, in all cases we evaluate the IRF holding all non-shocked variables at their unconditional averages. Since the variable of interest is a probability, it is more natural to present the results as a proportion of the base case, i.e. as $E \left[\tilde{P}_{t+h-1}^{(j)} \left(\tilde{\mathbf{Y}}_{t-1} \right) / \bar{P}_{t+h-1}^{(j)} \right]$, rather

⁷All accounts data are taken from the Bureau van Dijk FAME database. By using accounting information we assume that there is no systematic distortion in the data from either management action to smooth profits or off-balance sheet financing.

than as a difference from the base case, $E \left[\tilde{P}_{t+h-1}^{(j)} \left(\tilde{\mathbf{Y}}_{t-1} \right) - \bar{P}_{t+h-1}^{(j)} \right]$.

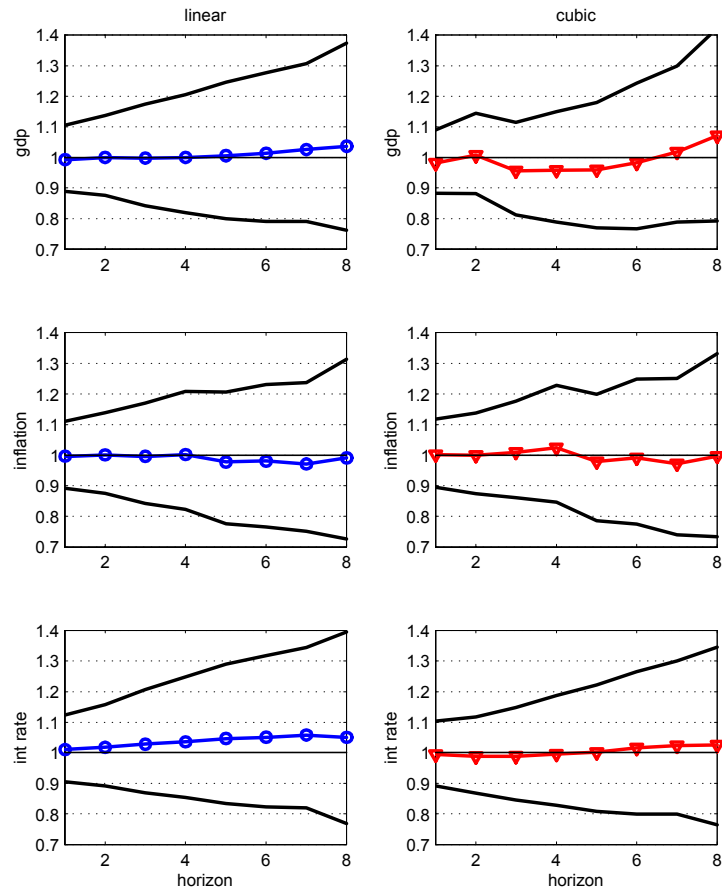


Figure 2: These figures show the response of the liquidation ratio to a 1 standard deviation shock, relative to the baseline liquidation ratio of 1.32% per year. The rows indicate the shocked variables; the columns show the model used, either a linear projection or a cubic projection. 95% confidence intervals are denoted with a thick line.

The baseline average default rate was 1.32% over our sample period, so a stress test value of 2 indicates a doubling of the probability of default. Figure 2 presents the results of a 1 standard deviation shock to each of the three macroeconomic variables, using either the linear or the cubic model. This figure reveals that the mean impact on the corporate liquidation rate from a macroeconomic shock is relatively small: the largest impact on the liquidation rate occurs 8-quarters after a shock to GDP (as estimated from the cubic model) where the liquidation rate increases by almost 10%

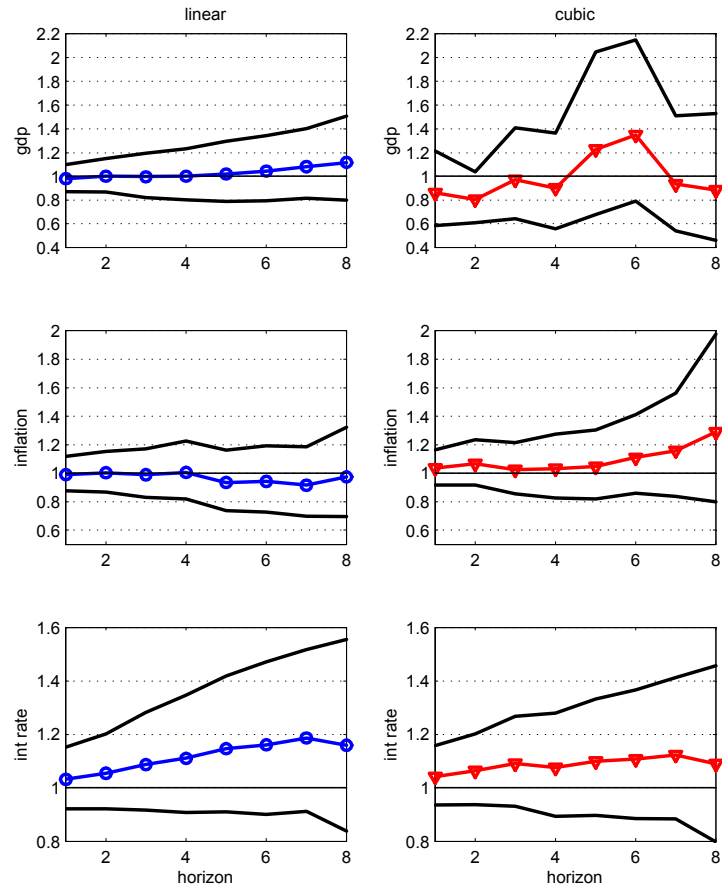


Figure 3: These figures show the response of the liquidation ratio to a 3 standard deviation shock, relative to the baseline liquidation ratio of 1.32% per year. The rows indicate the shocked variables; the columns show the model used, either a linear projection or a cubic projection. 95% confidence intervals are denoted with a thick line.

from its sample average. In Figure 3 we show the results of the 3 standard deviation macroeconomic shocks. In comparison to Figure 2 it is clear that large shocks have a significantly different impact on the liquidation rate than small shocks. For a positive 3 standard deviation shock to GDP the non-linear model predicts a significant fall in the liquidation rate 2 quarters ahead - based on the linear VAR one would falsely conclude that the liquidation rate does not change in the quarters following large GDP shocks.

In all cases, the confidence bounds on the impulse responses are wide. In some applications the fact that the confidence intervals always include 1, the base case, would be taken as evidence that corporate liquidations are independent of business cycle shocks. But the estimation uncertainty is important for regulators as well as banks to set capital at a sufficiently conservative level. Therefore, we should pay more attention to the upper confidence interval bound which is the upper bound on the mean impact of a shock to each of these variables at the 95% confidence level in our figures. Following a 1 standard deviation shock the upper bound is approximately 1.1 for all three shocks at the one-quarter horizon, and is between 1.3 and 1.4 at the eight-quarter horizon. Thus it is plausible that the mean impact of a one standard deviation shock is a 30% to 40% increase in the probability of default, meaning that the liquidation rate could move from 1.32% to as high as 1.85%, a substantial increase. This effect is even more significant if the shocks are large, where the impact from macro shocks can lead to a more than a 100% increase of PDs at a 95% confidence level.

Of course, the upper confidence interval bound could be made arbitrarily close to 100% by including more and more, potentially irrelevant, variables in the model for the liquidation rate. Doing so would increase the estimation error, increasing the uncertainty surrounding the estimated impact of a shock, and thus increase the upper confidence bound. For this reason it is important to carefully consider which variables to include in the model and the degree of flexibility to allow. By including numerous irrelevant variables we will likely obtain an upper confidence bound that is too conservative; by excluding the possibility of non-linearities and other effects we may instead obtain an upper confidence bound that is too low that does not truly reflect the uncertainty faced.

Our conclusions are confirmed when looking at the impact of small and large negative shocks (Figure 18 and 19 in the Appendix). Following a negative 3 standard deviation shock to GDP the liquidation rate increases significantly (borderline) in the second quarter by roughly 50% relative to the sample average. Independent of the forecast horizon we find that the liquidation rate falls strongly, and significantly, in the quarters after large negative shocks to the interest rate. The maximum fall occurs after 8 quarters when the liquidation rate is roughly 30% lower than on

average. Overall, our results suggest that large positive as well as negative unexpected changes in interest rates have a significant impact on the corporate liquidation rate both in the very short and in the more intermediate term (up to 2 years).

Looking at the write-off ratio of UK banks, Hoggarth et al (2005) find that the sample period is important for their conclusions on the link between credit risk and macroeconomic shocks. In particular, they find that the impact of a shock to output, relative to potential, is stronger in the years after the UK adopted an inflation targeting regime. As discussed in Section 3.2, this may be due to a structural break in the relation between the macroeconomic variables. Obviously, if a structural break is present in 1992 then it would be better to focus on the post 1992 estimations as we have done so far. However, we extend our sample back to 1985Q1 to investigate whether our conclusions hold once a recession is included in the sample. The IRFs from the estimated VAR based on the longer sample shocks can be found in Figure 20 and 21 in the Appendix. We only show the impact on the liquidation rate following large macroeconomic shocks as small shocks are have an insignificant impact on the liquidation rate at all horizons.

This robustness check leads to some interesting conclusions. First, we find that large positive GDP shocks imply a significant fall in the corporate liquidation rate up to 1 year following the shock. Although negative GDP shocks are found to increase the corporate liquidation rate in the short run the impact is insignificant - the maximum impact from a large negative GDP shock is after 3 quarters, rather than 2 quarters using the 1992Q4-2004Q3 sample. Second, we find a substantially different impact from interest rate shocks on the corporate liquidation rate in comparison with the VAR estimated on the 1992Q2-2004Q3 sample. Large positive shocks to the interest rate increases the corporate liquidation rate significantly at all horizons with a maximum impact after 8 quarters, which is 8 times as high as its sample average. Large negative interest rate shocks, on the other hand, decrease the corporate liquidation rate significantly, in particular 2-6 quarters following the interest rate shock. The largest fall in the liquidation rate after large negative interest rate shocks is in the fifth quarter when the level of the liquidation rate is around 90% lower than its sample average.

The robustness checks tentatively suggest that the corporate liquidation rate is much more strongly related to the interest rate once the recession of the early 1990s is included in the sample. Large negative shocks to GDP would not lead to a significant increase in the corporate liquidation rate whereas both large positive GDP and large negative interest rate shocks would imply a fall in the corporate liquidation rate. Overall, these findings imply that one should be cautious with our results before we have not undertaken more robustness checks.

4.5 Impulse responses functions using the Probit model

In this section we continue to look at the same shocks as before. As before we investigate the IRFs of a plus/minus one/three standard deviation shocks to GDP, inflation and the interest rate. However, instead of using aggregate liquidation rates we use the estimated quarterly probit model to capture the impact of our scenarios on credit risk. In line with our analysis for the liquidation rate we use the estimated parameter vector and the associated covariance matrix to compute the IRFs. To account for estimation uncertainty we take 10,000 random draws from the multivariate normal distribution of the parameter vector. Conditioning on the average of the firm specific variables and the shocked macroeconomic variables we, thus, obtain 10,000 corporate default probabilities. In all graphs below we, show IRF and associated error bands at a 95% confidence level. Again, we plot all IRFs as a proportion of the benchmark PDs of no shocks. Although we obtain error bands in both the linear and cubic model we only show error bands for the cubic model⁸.

In Figure 4 and Figure 5 we show the IRFs for positive and negative shocks. As for the liquidation rate, it is apparent that especially for large shocks the predictions of the cubic model are substantially different from the linear model. In several cases this difference is even statistically significant. One of the most interesting differences is for a +3 standard deviation shock to GDP. Whereas the linear model does not predict a positive impact on credit risk beyond the 4th quarter, the cubic model shows that the PD is 40-50% lower than the base case for the 3rd to 4th quarter. The figures also confirm that large shocks to the interest rate are the most significant driver for PDs. In the most extreme case quarterly PDs rise by over 25% following a 3 standard deviation shock to the interest rate.

One of the surprising insights is that especially for small shocks the point predictions of the linear model seem to be more conservative than the point predictions of the cubic model. For large shock this is often reversed. But when looking at the estimated Taylor-series expansion and plotting the underlying polynomials this is not so surprising anymore, as for small shocks the linear model lies above the cubic model⁹.

Some of the results in Figure 4 and Figure 5 seem counter-intuitive. Take the example of the estimated impact of a positive one standard deviation shock to the interest rate. Figure 4 indicates a marginally positive impact in the short run which becomes significant for year two of the forecast. The latter can be explain when looking at the IRF of the macro VAR, which show that for a +1

⁸We refer from showing the estimated impulse response function from the quadratic model as the model represents a special case between the linear and cubic model and does not reveal any interesting insights.

⁹The graphs of the polynomials can be provided by the authors on request.

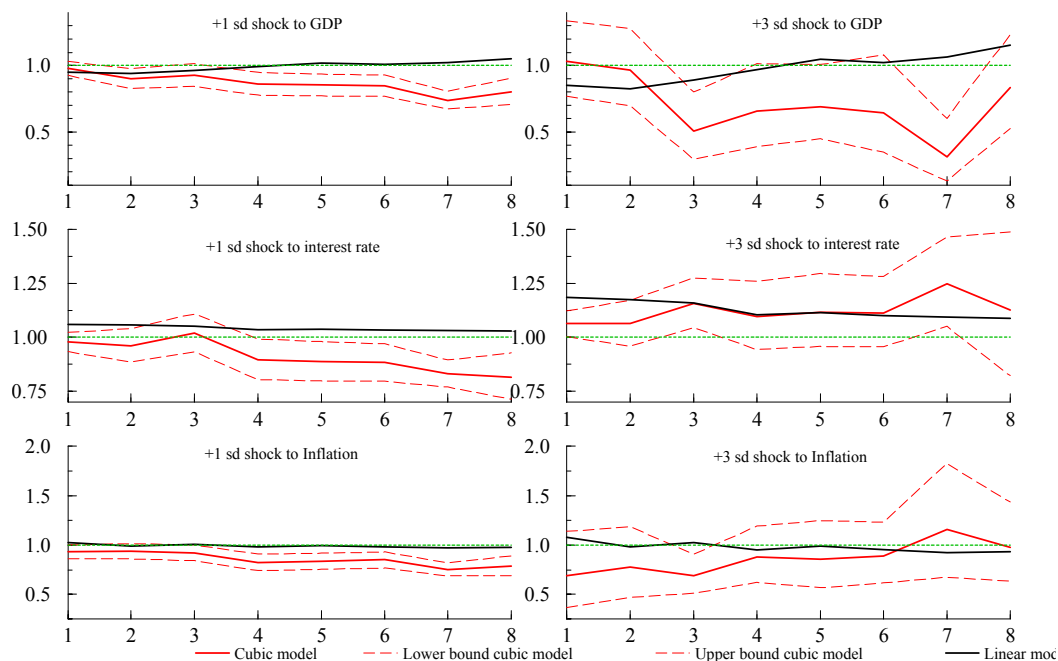


Figure 4: Impact of positive macroeconomic shocks on the PD. Impulse responses and error bands obtained using 10000 Monte Carlo simulations. Confidence bounds at a 95% level. Impulse responses are relative to base.

standard deviation shock GDP growth increases significantly in the second year (see Figure ?? in the Appendix). Some of the short term impact can be explained by starting values. Whilst starting values do not matter for the linear model, they do for a non-linear model. So far we have looked at IRF, where we set base conditions of the explanatory variables to the sample average for each variable. Figure 6 compares the impact of positive interest rate shocks on credit risk when taking 2003 values as the starting point (left panel) as well as when taking 2003 as starting values but setting the interest rate to 2% (right panel)¹⁰. Just comparing the left and right panel it is clear that starting values matter significantly. It is also apparent that the shape of the IRF changes as well. Whereas in average conditions the impact of a +3 standard deviation shock is very large in the long run (see Figure 4) it is insignificant in the second year for the hypothetical world of 2003 with a 2% interest rate (see Figure 6). But the impact is much more significant in the hypothetical world over the first three quarters.

¹⁰For firm specific variables we take the sample average for 2003.

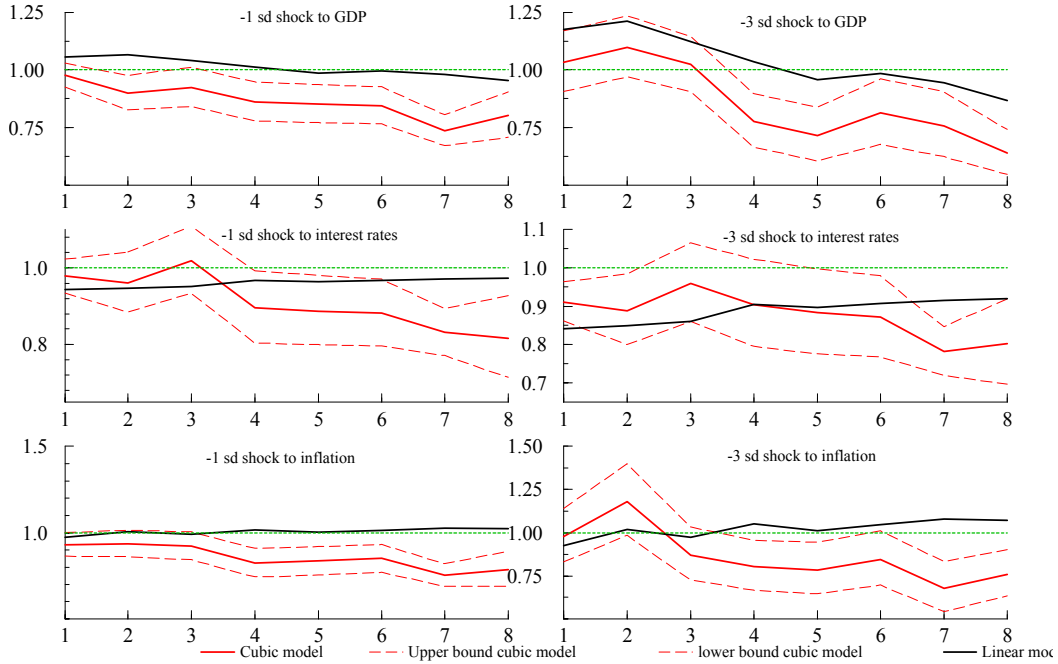


Figure 5: Impact of negative macroeconomic shocks on the PD. Impulse responses and error bands obtained using 10000 Monte Carlo simulations. Confidence bounds at a 95% level. Impulse responses are relative to base.

4.6 Stressed PDs for different firm characteristics

For a portfolio of credit risks, it is not only important to look at the starting conditions but to take the full distribution of firm specific characteristics into account. To have a better understanding of the distribution of PDs in the UK economy we now look at the impact of our stress scenarios taking account of different firm characteristics in 2003.

Instead of 10,000 random draws for the coefficients, we take for this exercise 3,000 random draws of the parameter vector from a multivariate normal distribution. For each of these draws and the specific characteristics of a each company in the 2003 sample we compute 3,000 PDs with the macro variables shocked relative to their average value.

In Figure 8 we plot the distributions of corporate PDs for different small (1 standard deviation) and large (3 standard deviation) adverse (positive shocks to the interest rate and negative macroeconomic shocks to GDP) for two different forecast horizons, 1 and 2 quarters. First, the graph shows that most companies in the UK have a very low PD but there is a fat tail with relatively high PDs. Second, the graph reveals that a small adverse macroeconomic shock hardly

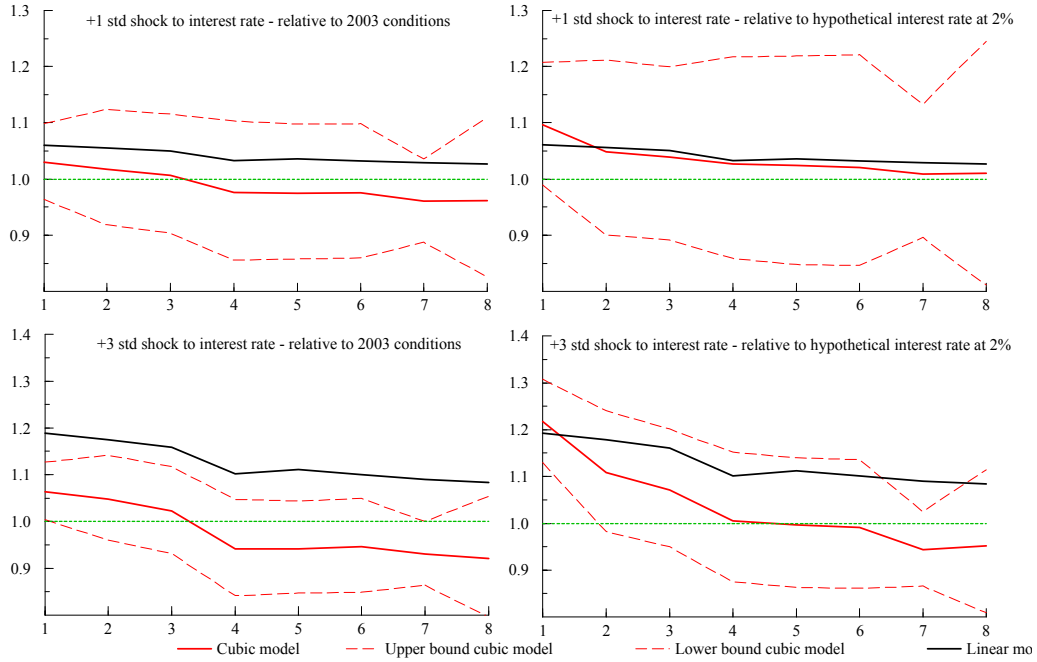


Figure 6: Impact of positive interest rate shocks on the PD using the average level of firm characteristics in 2003 and shocks to either interest rates relative to 2003 average or shocks relative to a hypothetical 2% interest rate level.

shifts the distribution of the probability of default (2003 sample). It looks rather different when the adverse macroeconomic shocks are large. Following large adverse shocks the distribution shifts to the right from the base scenario with more corporate default probabilities in the tail. The mean of the distribution, following a small adverse shock to the interest rate, is 1.85% with a one quarter forecast horizon (1.85% for two quarter forecast horizon as well); the mean of the distribution at a one quarter forecast horizon is 2.0% had the adverse interest rate shock been large (2.02% with a 2 quarter forecast horizon). The mean of the distribution following adverse GDP shocks is 1.82% when the shock is small and the forecast horizon is 1 quarter (1.90% with a 2 quarter forecast horizon) and 1.94% if the adverse shock is large (2.07% if forecast horizon is 2 quarters). Hence, consistent with our findings in the previous section using the average 1990-2004 quality distribution of companies we find that small adverse macroeconomic shocks hardly changes the distribution of corporate PDs in 2003 relative to had the macroeconomic variables remained at their mean level whereas the large adverse shocks have a much higher impact. We note also that the impact of a large adverse interest rate shocks would have been even stronger if forecasting at longer horizons (consistent with what was found from the impulse response functions in Figure 4).

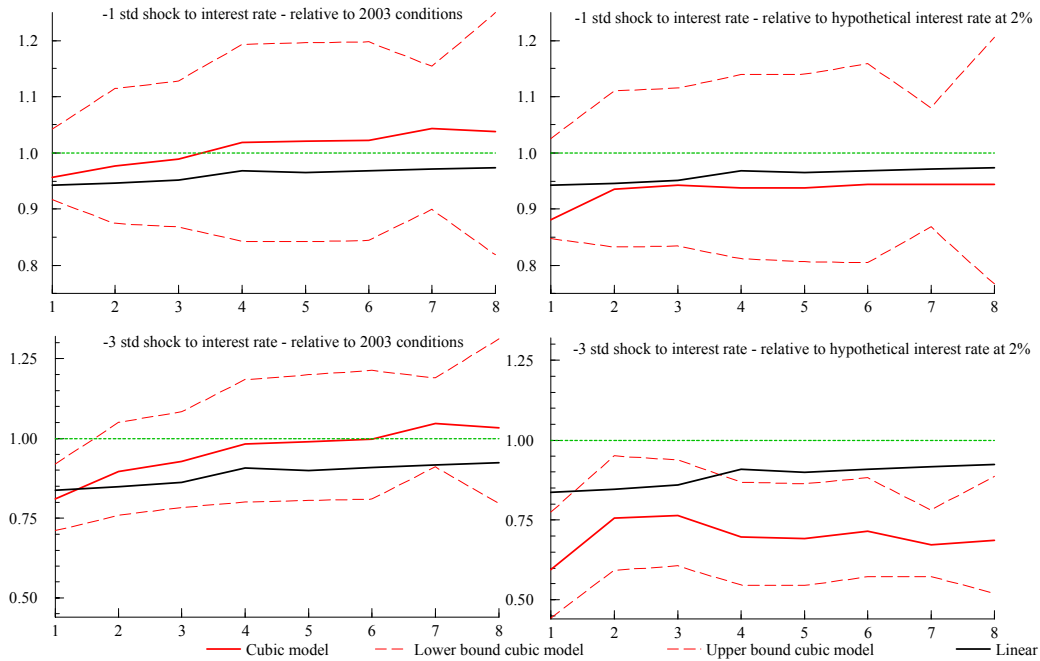


Figure 7: Impact of negative interest rate shocks on the PD using the average level of firm characteristics in 2003 and shocks to either interest rates relative to 2003 average or shocks relative to a hypothetical 2% interest rate level.

In Figure 9 a similar experiment is performed but this time comparing the shift in the distribution of the corporate default probability depending on whether using a linear or a cubic model. We see again that the linear model is a more conservative estimate for small shocks.

In Figure 10 we investigate how the distribution of the corporate default probability vary across different forecast horizons during periods of large adverse shocks. We do not report the distribution at all forecasting horizons in order not to blur the picture - it appears that the tail of the distribution is much thicker at short lags (mostly so when forecasting horizon is 2 lags) than at longer lags such as 8 quarters following large adverse GDP shocks - the opposite is the case for interest rates - more default probabilities are in the tails when the forecasting horizon is 7 or 8 quarters ahead following a large positive shocks to interest rates. Hence we see less dispersion in the distribution of PDs at shorter future horizons after large negative GDP shocks whereas this dispersion is much larger at longer horizons following a large positive interest rate shocks.

The error bands on the impulse responses presented in the previous section reflects the parameter uncertainty underlying the estimated impulse responses. We can obtain similar confidence




Figure 8: The distribution of corporate PDs in 2003 for different adverse macroeconomic shocks (based on cubic model). Base is the base distribution where macro variables remain at sample average. The label indicates the macroeconomic variable that is shocked (lag is the forecast horizon).

bounds when conditioning on firm specific data in 2003. From each of the 3000 Monte Carlo simulations we can compute the maximum value of the PD and the 97.5% (for instance) highest PD for each company. In Figure 11 we plot the distribution of the 97.5% highest PD and maximum PD for all the companies in the 2003 sample. The distributions are plotted against the stressed base distribution (i.e. average distribution). We note that the maximum uncertainty to the estimated stressed distributions is much higher with a forecast horizon of 2 quarters rather than 1 quarter and that the distribution of the maximum PD shifts much more to the right relative to the 97.5% critical value also when the forecast horizon is 2 quarters (relative to 1 quarter). Following a large negative shock to GDP and a forecasting horizon of 1 quarter, the mean of the stressed distribution is 1.94%, with a mean of the 97.5% distribution of 2.21% and the max distribution with a mean of 2.48% (the equivalent means of the distributions with a 2 quarter forecast horizon is 2.07%, 2.47% and 2.84%). Following large adverse interest shocks and a forecasting horizon of 1 quarter the means are 1.99%, 2.19% and 2.36% and finally with a forecasting horizon of 2 quarters the means are 2.02%, 2.39% and 2.74% respectively.




Figure 9: The distribution of corporate PDs in 2003 following different adverse shocks - linear vs cubic model. Note the label indicates which model is used for the estimation (linear vs cubic) and which of the macroeconomic variable that is shocked. Lag indicates the forecast horizon.

5 Conclusion

In this paper we investigate the impact of possible non-linearities on credit risk in a VAR setup. As standard VAR models are unable to deal with non-linearities we use the method proposed by Jorda (2004). The key insight of Jorda was to interpret a general VAR as a first order Taylor series approximation of the unknown data generating process, and to instead estimate more flexible approximations, which capture possible non-linearities in the data. We apply Jorda's method to a small model of the macro economy and extend it to analyse liquidation rates as well as a quarterly default model based on company accounts data to capture possible non-linear impacts of macroeconomic factors on credit risk. We show that the results of the non-linear VAR are significantly different to results using standard linear models, especially when considering large shocks. This can be seen in the simple three variable macro model. More importantly, we show that accounting for non-linearities in the underlying macroeconomic environment leads to substantially different conclusions for credit risk projects. We show that for small shocks linear models seem to

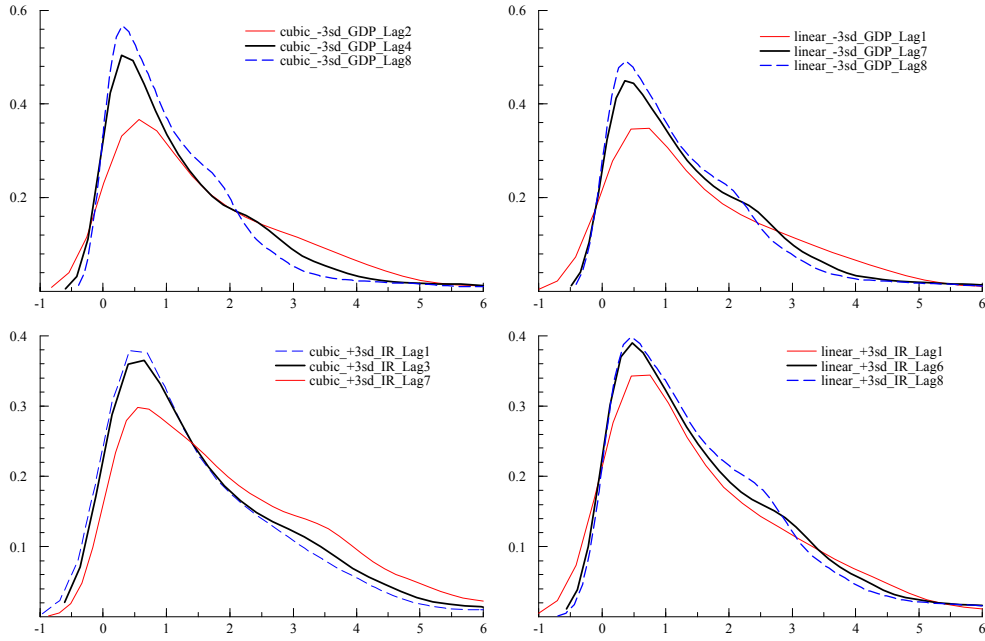


Figure 10: The distribution of corporate PDs forecast K periods ahead, with the firm characteristics in 2003, following large adverse macroeconomic shocks and comparison of the linear vs the cubic model.

overestimate credit risk, whereas for large ones they tend to underestimate it. We also show that in the non-linear model starting values impact not only the level of projected credit risk, but they also influence the shape of the impulse response function.

In contrast to other papers we explicitly take account of the underlying estimation uncertainty of the models. We show that this can have significant implications for the estimated level of credit risk, especially when looking at the tails of the credit risk distribution.

The third innovation of the paper is to propose a method to construct a proxy series from annual default data to generate IRFs of quarterly macroeconomic risk factors on quarterly firm specific PDs. We show that our estimates are unbiased, as long as the liquidation rate truly reflects the aggregate number of corporate defaults in our dataset. Given that we have data of more than 30.000 UK companies we are convinced that this assumption holds for our dataset. This method enables us to not only generate measures of quarterly aggregate credit risk but also to look at the distribution of credit risk across UK companies on a quarterly basis. We show that the tails of this distribution widen under stressed conditions. But given the non-linear impact and the dependence on starting values it is hard to derive unambiguous results with respect to the level as well as shape

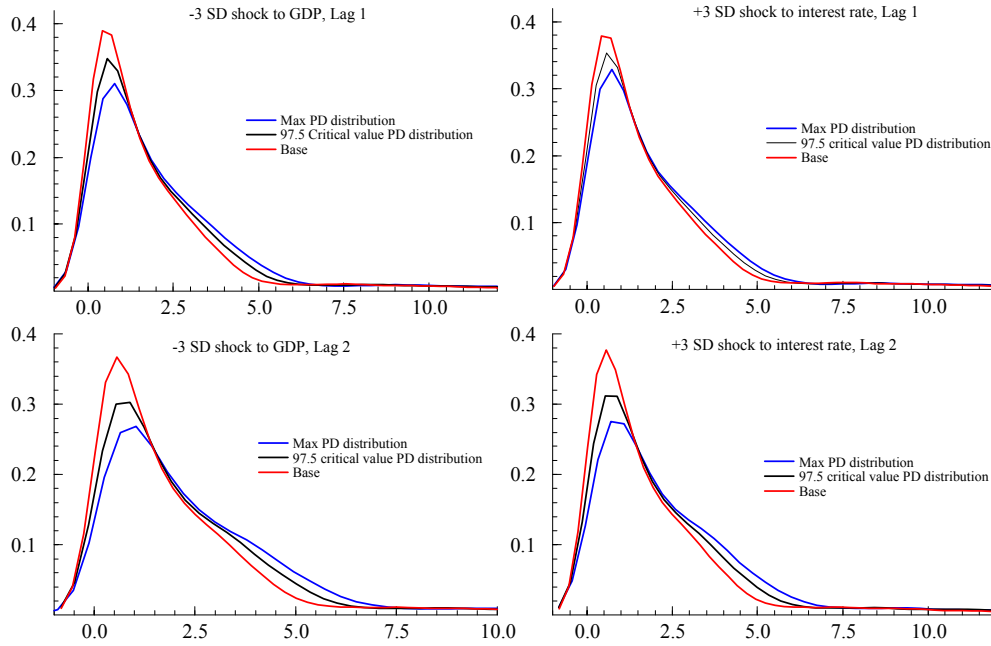


Figure 11: The distribution of the estimated maximum PD, distribution of the 97.5% critical value PD given severe adverse stress. The two distributions are plotted against the base distribution of the 2003 PDs (i.e. the distribution of the PD given the macrostress).

of the IRFs of different shocks to the interest rate, GDP or inflation. Overall, our analysis confirms previous papers (see for example Benito et al 2001) that especially a large increase in interest rates is a key driver of credit risk and that large positive shocks to GDP tend to reduce risk significantly.

However, we would caution against a literal interpretation of our results for the moment. So far we have not undertaken sufficient robustness checks, especially as our robustness checks for liquidation rates indicated that the shape of the IRF can depend on the sample, the model was estimated on. But even in this case, the overall conclusion that large standard deviation shocks seem to have a substantially different impact than small ones remains. In the future, we will undertake more research into the stability of our results. As said in the introduction, a confirmation of the findings of this paper would have serious implication not only from a regulatory perspective concerned about capital levels at 99.9% confidence level but also from a pricing perspective. Given the rapid increase in trade credit risk products and the new Basel II rules, this should be of great interest to market participants.

6 Appendix

6.1 Appendix A: A simple non-linear model

Assume that (y_t, y_{t-1}) have the following joint distribution

$$(y_t, y_{t-1}) \sim F = C_C(\Phi, \Phi; \kappa)$$

where F is some bivariate distribution with standard Normal marginal distributions (denoted Φ) connected with Clayton's copula, C_C , with dependence parameter κ . This type of time series process was first studied in economics by Chen and Fan (2002?). This implies that we can write

$$\begin{aligned} y_t &= h(y_{t-1}, \varepsilon_t; \kappa) \\ \text{where } \varepsilon_t | y_{t-1} &\sim N(0, 1) \\ C_C(u|v; \kappa) &\equiv \frac{\partial C_C(u, v; \kappa)}{\partial v} \\ h(y, \varepsilon; \kappa) &\equiv C_C^{-1}(\Phi(\varepsilon) | \Phi(y); \kappa) \end{aligned}$$

which is a general, stationary, non-linear data generating process that is simple to simulate. In Figure 1 we show one simulated sample path from this process for $k = 1.1$, which yields $\text{Corr}[y_t, y_{t-1}] = 0.5$. In Figure A1 below we simulate a path for y_t .




Figure 12: One sample path for Y_t using a nonlinear DGP.

6.2 Appendix B: Tables

Table B1

| Variable definitions | |
|----------------------|--|
| Def | Proportion of companies defaulting |
| pmn | Dummy with value of 1 if negative profit margin, 0 otherwise |
| pml | Dummy with value 1 if $0 \leq \text{Profit margin} < 0.03$, 0 otherwise |
| pmm | Dummy with value 1 if $0.03 \leq \text{Profit margin} < 0.06$, 0 otherwise |
| IC | Interest cover |
| DA | Debt to asset ratio |
| PD35 | Dummy with value 1 if negative profit margin and a debt to asset ratio >0.35 |
| CR | Current ratio |
| IE | Log of number employed |
| S | Dummy with value 1 if subsidiary, 0 otherwise |
| PNS | Dummy with value 1 if negative profit margin and subsidiary, 0 otherwise |
| I4 | Dummy with value 1 if construction industry, 0 otherwise |
| I5 | Dummy with value 1 if wholesale and retail industry, 0 otherwise |
| I6 | Dummy with value 1 if hotels and restaurants industry, 0 otherwise |
| I7 | Dummy with value 1 if transport, storage and communication industry, 0 otherwise |
| I9 | Dummy with value 1 if real estate, renting and business activities industry, 0 otherwise |
| I10 | Dummy with value 1 if other services industry, 0 otherwise |
| gdp | Quarter on quarter GDP growth |
| int | The nominal base rate |
| inf | Quarter on quarter growth in producer prices |

138155 observations of which 10981 are in 2003

Table B2

| Descriptive statistics | average | average | Standard dev | Standard dev |
|---------------------------|-----------|---------|--------------|--------------|
| | 1991-2004 | 2003 | 1991-2004 | 2003 |
| Def | 0.0178 | 0.01878 | 0.1324 | 0.1357 |
| pmn | 0.1879 | 0.2196 | 0.3906 | 0.4139 |
| pm | 0.2342 | 0.2543 | 0.4235 | 0.4355 |
| pmm | 0.2083 | 0.1991 | 0.4061 | 0.3993 |
| IC | 7.0276 | 7.3263 | 9.5175 | 9.9194 |
| DA | 0.3802 | 0.4038 | 0.3160 | 0.3644 |
| PD35 | 0.1191 | 0.1410 | 0.3239 | 0.3480 |
| CR | 1.2243 | 1.2182 | 0.6481 | 0.6858 |
| IE | 5.8712 | 5.8868 | 1.1756 | 1.1847 |
| S | 0.6043 | 0.5567 | 0.4890 | 0.4968 |
| PNS | 0.1263 | 0.1417 | 0.3321 | 0.3488 |
| I4 | 0.0623 | 0.0710 | 0.2425 | 0.2569 |
| I5 | 0.1642 | 0.1822 | 0.3705 | 0.3860 |
| I6 | 0.0393 | 0.0477 | 0.1944 | 0.2132 |
| I7 | 0.0670 | 0.0738 | 0.2501 | 0.2615 |
| I9 | 0.1764 | 0.1816 | 0.3812 | 0.3855 |
| I10 | 0.0766 | 0.0791 | 0.2660 | 0.2700 |
| gdp | 0.0063 | 0.0041 | 0.0066 | |
| int | 0.0145 | 0.0047 | 0.0088 | |
| inf | 0.0069 | 0.0052 | 0.0071 | |

138155 observations of which 10981 are in 2003

6.3 Appendix C: Charts




Figure 13: The Corporate liquidation rate and macroeconomic variables included in the VAR.

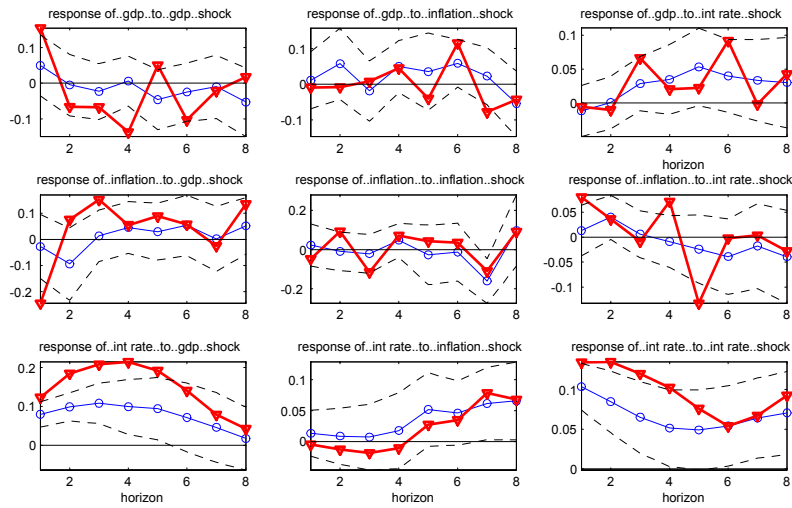


Figure 14: The thick line marked with triangles is the impulse response for a 1 standard deviation shock from the cubic projection; the thin line marked with circles is the impulse response from the linear projection; the dashed lines are the 95% confidence bounds on the impulse response from the linear projection.

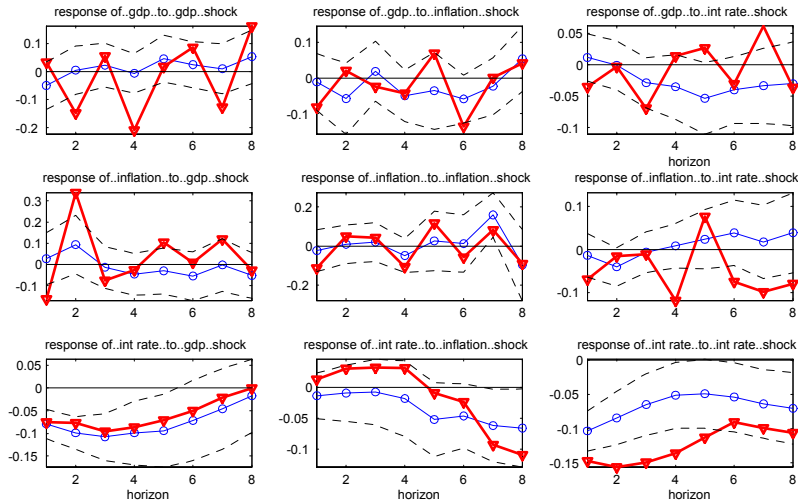


Figure 15: The thick line marked with triangles is the impulse response for a -1 standard deviation shock from the cubic projection; the thin line marked with circles is the impulse response from the linear projection; the dashed lines are the 95% confidence bounds on the impulse response from the linear projection.

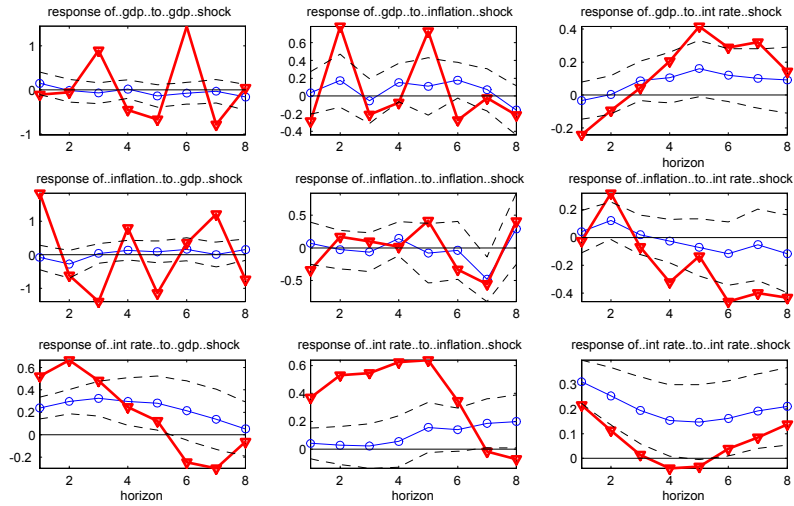


Figure 16: The thick line marked with triangles is the impulse response for a 3 standard deviation shock from the cubic projection; the thin line marked with circles is the impulse response from the linear projection; the dashed lines are the 95% confidence bounds on the impulse response from the linear projection.

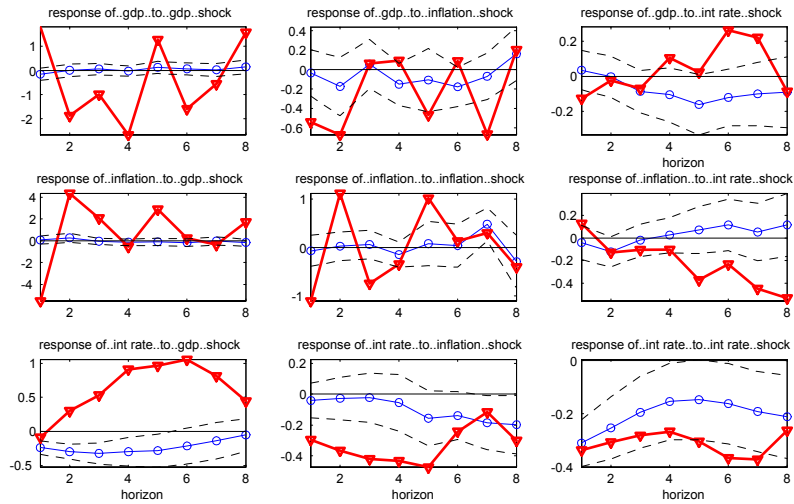


Figure 17: The thick line marked with triangles is the impulse response for a -3 standard deviation shock from the cubic projection; the thin line marked with circles is the impulse response from the linear projection; the dashed lines are the 95% confidence bounds on the impulse response from the linear projection.

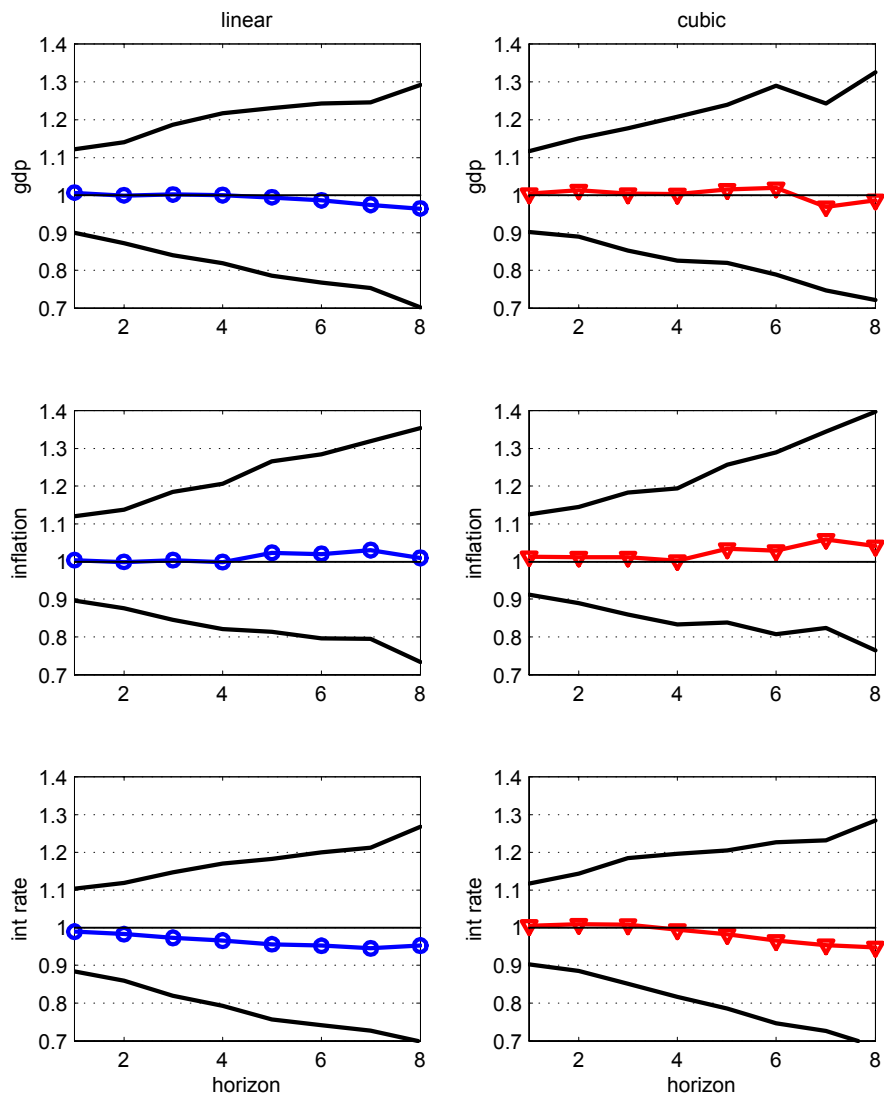


Figure 18: These figures show the response of the liquidation ratio to a -1 standard deviation shock, relative to the baseline liquidation ratio of 1.32% per year. The rows indicate the shocked variables; the columns show the model used, either a linear projection or a cubic projection. 95% confidence intervals are denoted with a thick line.

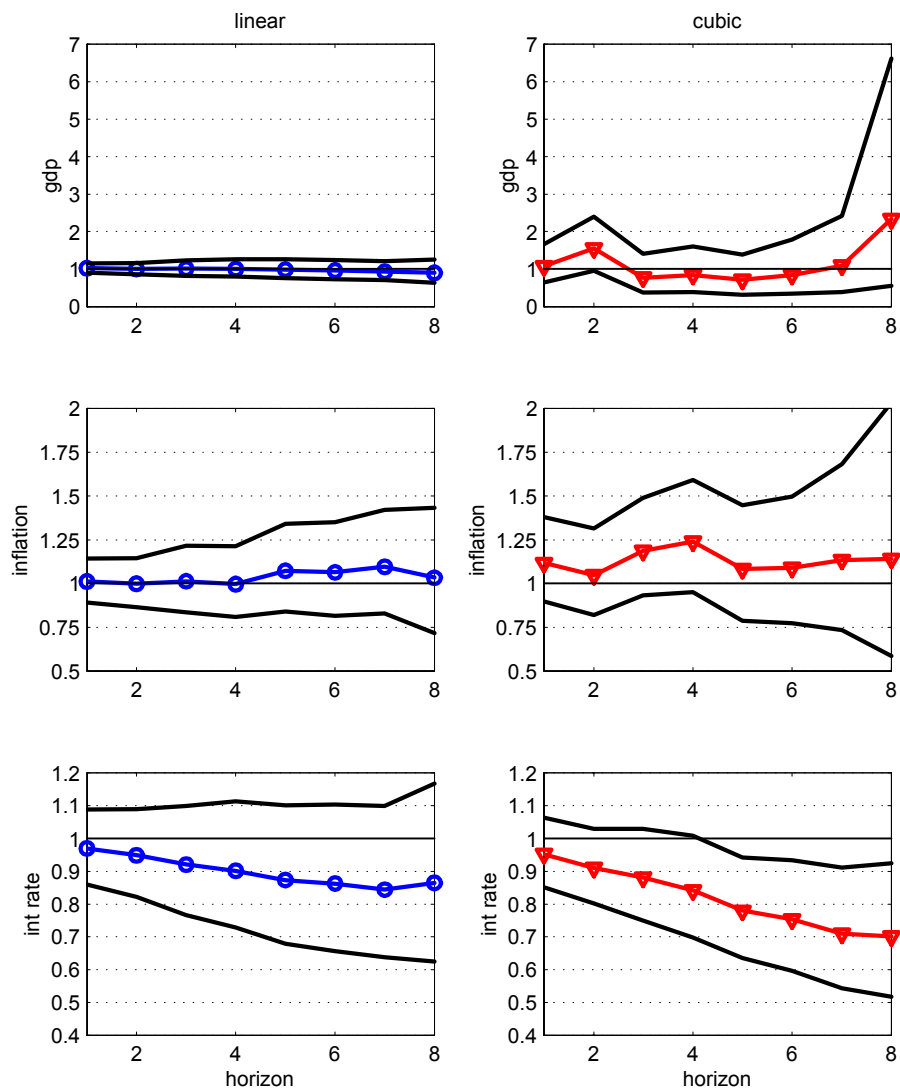


Figure 19: These figures show the response of the liquidation ratio to a -3 standard deviation shock, relative to the baseline liquidation ratio of 1.32% per year. The rows indicate the shocked variables; the columns show the model used, either a linear projection or a cubic projection. 95% confidence intervals are denoted with a thick line.

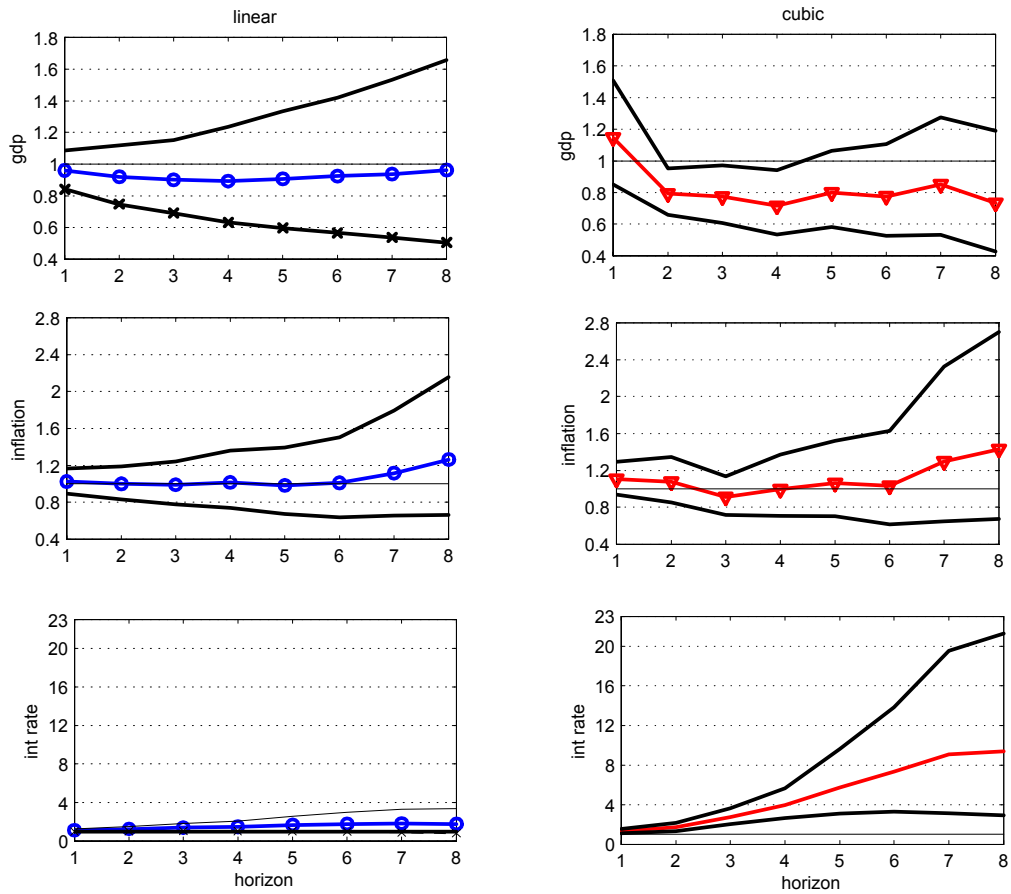


Figure 20: These figures show the response of the liquidation ratio to a +3 standard deviation shock, relative to the baseline liquidation ratio of 1.32% per year. The rows indicate the shocked variables; the columns show the model used, either a linear projection or a cubic projection. 95% confidence intervals are denoted with a thick line. Using sample from 1985Q1-2004Q3.

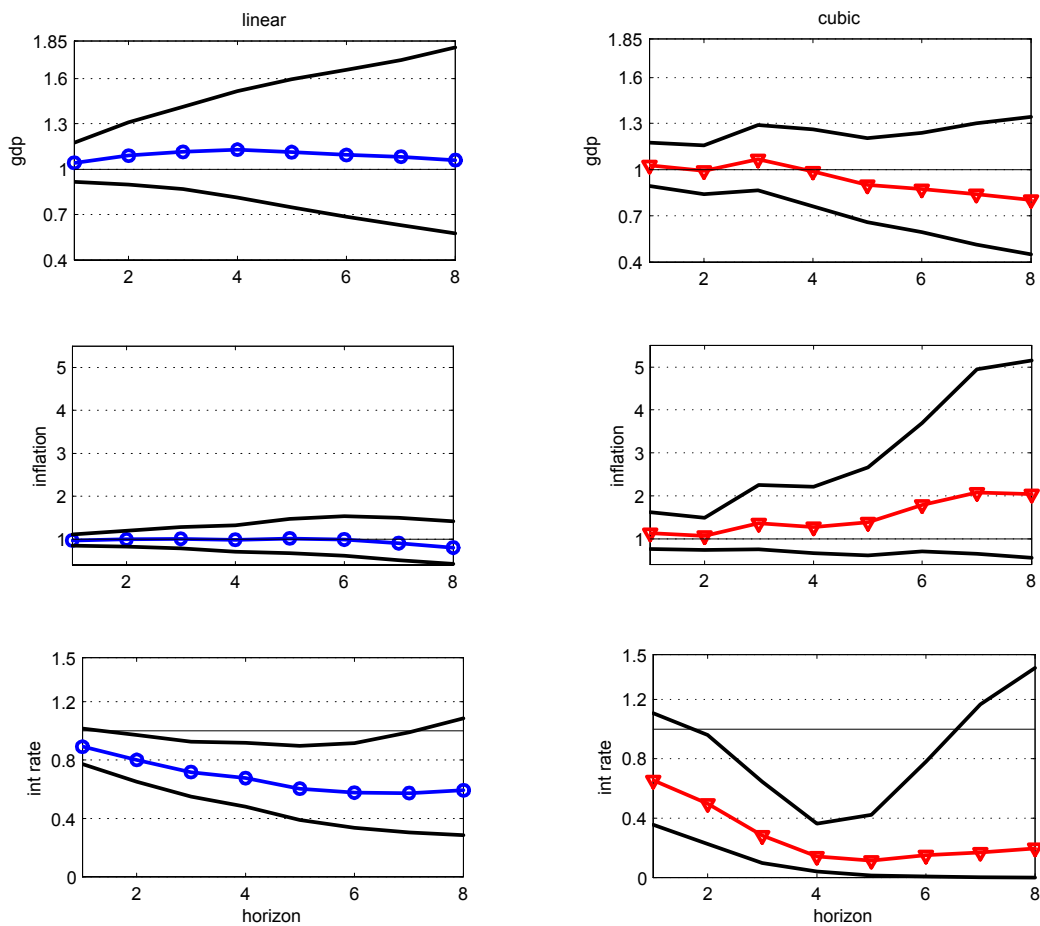


Figure 21: These figures show the response of the liquidation ratio to a -3 standard deviation shock, relative to the baseline liquidation ratio of 1.32% per year. The rows indicate the shocked variables; the columns show the model used, either a linear projection or a cubic projection. 95% confidence intervals are denoted with a thick line. Using sample from 1985Q1-2004Q3.

References

- [1] Altman, E.I., 1968, Financial Ratios, Discriminant Analysis, and the Prediction of Corporate Bankruptcy, *Journal of Finance* 23
- [2] Beaver, B., 1966, Financial Ratios as Predictors of Failure, *Empirical Research in Accounting: Selected Studies, Supplement to Journal of Accounting Research Autumn*,
- [3] Benati, L., 2004, Evolving post-World War II UK economic performance, *Bank of England Working Paper 232*
- [4] Benito, A, Whitley, J and Young, G, 2001, Analysing Corporate and Household Sector Balance Sheets, *Bank of England Financial Stability Review*, December
- [5] Boss, M., 2002, A Macroeconomic Credit Risk Model for Stress Testing the Austrian Credit Portfolio, *Financial Stability Review*, Oesterreichische Nationalbank, 4.
- [6] Bunn, P. and V. Redwood, 2003, Company Accounts Based Modelling of Business Failures and the Implications for Financial Stability, *Bank of England Working Paper 210*.
- [7] Campbell, J. R., J. Hilscher and J. Szilagyi, 2003, In Search of Distress Risk, *Harvard Institute of Economic Research Discussion Paper 2081*.
- [8] Davidson, R. and J. G. MacKinnon, 1993, *Estimation and Inference in Econometrics*, Oxford University Press
- [9] Drehmann, M., 2005, A Market Based Macro Stress Test for the Corporate Credit Exposures of UK Banks, *mimeo, Bank of England*
- [10] Duffie, D., L. Saita, and K. Wang, 2005, Multi-Period Corporate Default Predictions With Stochastic Covariates, *mimeo*
- [11] Hillegeist, S.A., E.A. Keating, D.P.Cram and K.G. Lundstedt, 2003, Assessing the Probability of Bankruptcy, *Review of Accounting Studies* 9
- [12] Hoggarth, G., Sorensen, S. and L. Zicchino, 2005, Stress tests of UK banks using a VAR approach, *Forthcoming Bank of England Working Paper*
- [13] Jordà, O., 2004, Estimation and Inference of Impulse Responses by Local Projections, *American Economic Review*, forthcoming.
- [14] Koop, G., M. H. Pesaran, and S. M. Potter, 1996, Impulse Response Analysis in Nonlinear Multivariate Models, *Journal of Econometrics*, 74, 119-147.
- [15] Kuttner, K. and A. Posen, 1999, Does Talk Matter After All? Inflation Targeting and Central Bank Behavior, *Institute for International Economics Working Paper 99-10*.
- [16] Merton, R. C., 1974, On the Pricing of Corporate Debt: The Risk Structure of Interest Rates, *Journal of Finance*, 29(2).
- [17] Ohlson, J., 1980, Financial Ratios and the Probabilistic Prediction of Bankruptcy, *Journal of Accounting Research* 19

- [18] Pesaran, M. H., T. Schuermann, B.-J. Truetler and S. M. Weiner, 2005, Macroeconomic Dynamics and Credit Risk: A Global Perspective, *Journal of Money, Credit, and Banking*, forthcoming.
- [19] Shumway, T., 2001, Forecasting Bankruptcy More Accurately: A Simple Hazard Rate Model, *Journal of Business* 74
- [20] Sorge, M., 2004, Stress Testing Financial Systems: An Overview of Current Methodologies, *BIS Working Paper 165*.
- [21] Wilson, T. C., 1997a, Portfolio Credit Risk (I), *RISK*.
- [22] Wilson, T. C., 1997a, Portfolio Credit Risk (II), *RISK*.
- [23] Virolainen, K., 2004, Macro Stress Testing with a Macroeconomic Credit Risk Model for Finland, *Bank of Finland Discussion Papers*, 18.