Model Uncertainty and Aggregated Default Probabilities: New Evidence from Austria

Paul Hofmarcher, Stefan Kerbl, Bettina Grün, Michael Sigmund, Kurt Hornik

Research Report Series
Report 116, December 2011

Institute for Statistics and Mathematics
http://statmath.wu.ac.at/

Model Uncertainty and Aggregated Default Probabilities: New Evidence from Austria

Paul Hofmarcher[a], Stefan Kerbl[b], Bettina Grün[c], Michael Sigmund[b], Kurt Hornik[a]

January 5, 2012

[a] Department of Finance, Accounting and Statistics, WU Wirtschaftsuniversität Wien
[b] Oesterreichische Nationalbank
[c] Department of Applied Statistics, Johannes Kepler University Linz

Abstract

Understanding the determinants of aggregated default probabilities (PDs) has attracted substantial research over the past decades. This study addresses two major difficulties in understanding the determinants of aggregate PDs: Model uncertainty and multicollinearity among the regressors. We present Bayesian Model Averaging (BMA) as a powerful tool that overcomes model uncertainty. Furthermore, we supplement BMA with ridge regression to mitigate multicollinearity. We apply our approach to an Austrian dataset. Our findings suggest that factor prices like short term interest rates and energy prices constitute major drivers of default rates, while firms' profits reduce the expected number of failures. Finally, we show that the results of our baseline model are fairly robust to the choice of the prior model size.

JEL Classification: E44, C52, E37.

1 Introduction

Understanding the driving factors of aggregated probabilities of corporate defaults is an important topic both for financial institutions and supervisors. For example, conducting meaningful stress tests requires the translation of macroeconomic scenarios into portfolio losses. The same applies when financial institutions and supervisors are interested in forecasting the credit quality of portfolios on an aggregated level. Both in the field of macro prudential supervision and strategic risk management a knowledge of the determinants of aggregated defaults is crucial.

Consequently, estimating the link between macroeconomic variables and probabilities of defaults has been a long-standing topic in research, as numerous papers testify (see below). However, the classical approach of regression faces a major challenge: Due to the sparse theoretical framework of how firm defaults are linked to specific macroeconomic variables, researchers are compelled to draw on their intuition which macro variables to include or not. Such a procedure neglects the uncertainty in the model choice and might end up with wrong conclusions. This challenge, commonly known as model uncertainty, is a problem shared with many other empirical fields of research. In what follows we present a state-of-the-art statistical approach of dealing with model uncertainty, a combination of Bayesian Model Averaging and ridge regression which we then apply to Austrian data.

Motivated by the high interest in the topic from industry and supervisors, there is a growing body of literature examining the relationship between firm defaults and economic conditions. Altman (1983) uses augmented distributed lags to demonstrate the effect of GNP, money supply and corporate profits on firms' ability to survive. Altman (1984) presents a survey discussing different business failure models that have been tested and developed outside the United States. Liu and Wilson (2002) use a time-series model to construct measures showing that interest rate and insolvency legislation are important variables in explaining firm bankruptcy. Similarly, Virolainen (2004) regresses Finnish sector-specific default rates on macroeconomic indicators like

GDP, interest rates and levels of corporate indebtedness. Liu (2004) uses an error-correction model to investigate the macroeconomic determinants of UK corporate failure rates. Liu (2009) extends this research by implementing a vector error-correction model specifically accounting for policy-induced changes in the macroeconomy, concluding that macro variables like the interest rate and inflation impact firm failures. Simons and Rolwes (2009) use macroeconomic-based models for estimating default probabilities using a Dutch dataset. Additionally, they compare their results with Austrian data. They conclude that for both countries their model delivers different results, deducing that their provided model is country specific. Further contributions are Koopman and Lucas (2005) who analyse the co-movement of credit and macro cycles in the US and Foglia et al. (2009) who examine Italian default frequencies per sector.

Screening the literature reveals that authors have to rely on expert knowledge when deciding upon the inclusion or non-inclusion of macro variables. To the best of our knowledge, uncertainty about the correct model specification for aggregated probabilities of default has not explicitly been addressed yet. The approach we present here refrains from assuming that there is one "true" model but instead averages over a huge number of potential models.

This approach is known as Bayesian Model Averaging (BMA) (see Hoeting et al., 1999). Thereby, the researcher controls the model size via a prior model inclusion probability for each variable[1]. Sampling from the set of regressors BMA then computes a huge number of models, which are weighted by their marginal likelihood and subsequently averaged. This simple procedure reveals important determinants of the dependent variable and their respective coefficients.

As noted above, BMA is becoming a central tool applied in dealing with model uncertainty, or in general settings with large numbers of potential regressors and relatively limited numbers of observations (see Ley and Steel, 2009). In the literature on growth determinants

[1]The approach we follow here attaches the same prior inclusion probability to each variable (see Section 2). However, in general the researcher could attach a higher probability to variables deemed to be of special relevance.

Fernandez et al. (2001) and Sala-I-Martin et al. (2004) propose BMA to identify robust drivers of countries' average growth. Wright (2008) and Avramov (2002) use BMA to forecast exchange rates and stock returns respectively. Empirical results have shown that BMA might outperform single model in prediction (see Hoeting et al., 1999).

However, at least in our case highly correlated candidate variables (multicollinearity) constitute an issue to be accounted for. To some extent this fact arises due to the inclusion of lagged explanatory variables, which display a particularly high correlation. To explicitly deal with this correlation structure we supplement BMA with a shrinkage method, *ridge regression* (Hoerl and Kennard, 1970a,b). Ridge regression aims at avoiding the commonly observed characteristic upon inclusion of highly correlated variables: coefficients display high absolute magnitudes which are cancelled out by coefficients of correlated cousins of comparable magnitude with reversed sign. By adding a penalty term dependent on the size of coefficients ridge regression indeed overcomes this issue.

The remainder of the paper proceeds as follows. Section 2 presents the methodological approach outlined above, i.e. BMA and ridge regression. We then apply this approach in Section 3 and Section 4, whereby the former presents the dataset and the latter the results. Finally, Section 5 concludes and provides discussion on further research.

2 Model Specification and Estimation

In the following subsections we give a brief overview of the methods we apply. First, we highlight the advantage of ridge regression. Second, we refine our methodology by introducing the spike and slab approach, a specific BMA technique to account explicitly for model uncertainty.

In order to introduce the methodological approach presented in this paper, we start with the familiar framework of linear regression. Here, we assume that the relationship between the *logit transformed* aggregated default rates, as response variable, $\mathbf{y}$ $(N \times 1)$ and the design matrix of the explanatory variables (here the macro variables) $\mathbf{X}$ $(N \times$

K) is given by the linear regression

$$\mathbf{y} = \mathbf{X}\boldsymbol{\beta} + \epsilon, \tag{1}$$

where $\epsilon \sim N(0, \sigma^2 \mathbf{I}_N)$. The vector $\boldsymbol{\beta}$ denotes the parameter vector of interest. Assuming that the explanatories X are highly correlated the standard OLS estimator $\hat{\boldsymbol{\beta}} = (\boldsymbol{X'X})^{-1}\boldsymbol{X'y}$ might be ill-conditioned (multicollinearity). In particular, at least one of the eigenvalues η_k of $\boldsymbol{X'X}$ will move towards zero, inflating the variance of the OLS estimator $E((\hat{\boldsymbol{\beta}} - \boldsymbol{\beta})'(\hat{\boldsymbol{\beta}} - \boldsymbol{\beta})) = \sigma^2 \sum_K \eta_k^{-1}$.

2.1 Ridge Regression and Bayesian Ridge Regression

Ridge regression (see Hoerl and Kennard, 1970a) belongs to the class of shrinkage methods in the context of linear regression models. In contrast to well known subset selection algorithms (e.g. Forward Stepwise Selection) it does not retain a subset of predictors and discard the rest but shrinks the size of predictors proportionally in accordance with their importance (Friedman et al. 2009). To see why this is so valuable imagine the usual setup of highly correlated variables in the design matrix leading to large positive and negative coefficients and thus to unreliable results. Indeed, multicollinearity may result in poorly determined parameters. One way to deal with multicollinearity is the use of ridge regression. From a frequentist point of view, ridge regression solves the optimization problem

$$\hat{\boldsymbol{\beta}}^{\text{ridge}} = \underset{\boldsymbol{\beta}}{\operatorname{argmin}} \left\{ (\mathbf{y} - \mathbf{X}\boldsymbol{\beta})'(\mathbf{y} - \mathbf{X}\boldsymbol{\beta}) + \lambda \sum_{j=1}^{K} |\beta_j|^2 \right\}. \tag{2}$$

The Lagrangian parameter λ defines how much the classical OLS–βs are shrunk. If λ moves towards 0 then the constraint is not binding and one arrives at the OLS solution.

As for OLS, it is possible to give a closed solution of the ridge regression problem, which is given by

$$\hat{\boldsymbol{\beta}}^{\text{ridge}} = (\mathbf{X}'\mathbf{X} + \lambda\mathbf{I})^{-1}\mathbf{X}'\mathbf{y}. \tag{3}$$

The ridge regression solution is very similar to the OLS solution (except for the term $\lambda\mathbf{I}$) and is linear in the response variable $\mathbf{y}$. The proportional shrinkage of the ridge parameters via the L_2 norm in Equation 2 provides the ability to cope with correlated variables as large coefficients are penalized. Clearly, a precondition of ridge regression is the standardisation of regressors in order to treat variables measured on different scales equally. An analogous approach to ridge regression is available in a Bayesian setting. *Bayesian ridge regression* was first introduced by Hsiang (1975). Keeping the assumptions of linear regression and setting $\lambda = \sigma^2/\tau^2$ one implements the following hierarchical Bayesian model:

$$\mathbf{y}|\boldsymbol{\beta}, \sigma^2 \sim N(\mathbf{X}\boldsymbol{\beta}, \sigma^2\mathbf{I}_n), \tag{4}$$

where the prior specifications of the coefficients $\boldsymbol{\beta}$

is given by

$$\boldsymbol{\beta}|\tau^2 \sim \prod_{j=1}^{P} N(0, \tau^2), \tag{5}$$

with proper priors[2] for the variances σ^2 and τ^2.

The prior on $\boldsymbol{\beta}$ conditional on τ and the fact that $\boldsymbol{y} \sim N(\boldsymbol{X}\boldsymbol{\beta}, \sigma^2\boldsymbol{I}_n)$ allows for the use of Markov Chain Monte Carlo (MCMC) sampling to estimate the posterior distribution of interest.

2.2 Model Uncertainty

As outlined in the introduction, an important task in statistical modeling is the choice of an optimal model from the set of all possible models. With K potential explanatory variables, one faces 2^K possible combinations of regressors. Selecting the best model out of 2^K linear models is a challenging task. In addition, several models with similar performance might arise which does not allow for an unambiguous single best choice. Thus, the uncertainty associated with a selected model is an important aspect, especially when it comes to forecasting

[2]Any inverted gamma prior for σ^2 and τ^2 would maintain conjugacy. Here we use the limiting improper priors $\frac{1}{\sigma^2}$, respectively $\frac{1}{\tau^2}$.

(see Steel, 2011). One natural way to deal with model uncertainty is to pool over the considered models — as BMA does. Thereby, weights of the single models depend on how much the data support each model via the posterior distribution. An excellent review of BMA is given in Hoeting et al. (1999). Using BMA, one obtains the distribution of some quantity of interest $\boldsymbol{\beta}$, e.g., the effect of a macro-variable, by averaging inference over all models M_k

$$P(\boldsymbol{\beta}|\boldsymbol{Z}) = \sum_{l=1}^{2^K} P(\boldsymbol{\beta}|M_l, \boldsymbol{Z}) P(M_l|\boldsymbol{Z}), \tag{6}$$

where $P(M_k|\boldsymbol{Z})$ is the posterior probability of model M_k given the whole dataset $\boldsymbol{Z}$ ($\boldsymbol{X}$ and $\boldsymbol{y}$ combined) and is derived by

$$P(M_k|\boldsymbol{Z}) = \frac{P(\boldsymbol{Z}|M_k)P(M_k)}{\sum_l P(\boldsymbol{Z}|M_l)P(M_l)}, \tag{7}$$

where $P(M_k)$ is the prior probability of model M_k and $P(\boldsymbol{Z}|M_k)$ is the marginal or integrated likelihood of model M_k obtained by integrating over the parameters (see Hoeting et al., 1999). Suitable choices of prior inclusion probabilities $P(M_k)$ allow to control the expected model size, i.e., the number of included parameters. In order to sample different models M_k of varying size and average across them, we make use of *spike and slab* priors (Mitchell and Beauchamp, 1988; George and McCulloch, 1993, 1997).

2.2.1 Model Uncertainty via the Spike and Slab Approach

The central point in using spike and slab priors is to assign each coefficient a prior which is a mixture of a point mass at zero and a specified "slab" distribution. This allows to exclude variables from the regression. In this sense spike and slab constitutes an optimal supplement to ridge regression which alone does not provide variable selection.

Formally, we modify the prior defined in Equation 5 and use for all considered regressions discussed in this work a *coefficient prior* of the form

$$P(\beta_j|c_j, \tau, \sigma^2) \sim (1 - c_j)I_0 + c_j\pi(\tau), \tag{8}$$

where c_j is a binary random variable with success probability $\gamma = P(c_j = 1)$ (which we set to the same value for all candidate regressors j). $\pi(\tau)$ is the prior distribution of β_j defined by Equation 5.

The Posterior Inclusion Probability ($P(c_j = 1|\mathbf{Z})$ or PIP[3]) of each variable j contains valuable insights about the importance of variable j. In particular, the PIP is of high value as it displays the fraction of models visited in which variable j was selected, $P(c_j = 1|\mathbf{Z})$.[4] PIP can thus be understood as a measure of "posterior importance" of a given variable and is a widely used measure in Bayesian Model Averaging (see Sala-I-Martin et al., 2004).

2.2.2 Model size

We have not yet discussed in detail the specification of the prior variable inclusion probabilities used by the spike and slab approach in Equation 8.

One possible approach would be to assign each variable β_j an uninformative inclusion probability of $\gamma = 0.5$, i.e. c_j is drawn from a Bernoulli distribution $Be(0.5)$. This has the odd and troubling implication that we assume the number of included variables to be very large (see Sala-I-Martin et al., 2004). In particular the expected model size, $E[M_j]$, equals $K \times 0.5$, where K is the number of candidate regressors. In our case, as explained below, we have 160 candidate regressors to choose from, $K = 160$, which would results in a very large prior model size, $E[M_j] = 80$. Models of this size are uncommon as researchers and practitioners prefer smaller models. Therefore, instead of choosing one value for the prior model size, we specify a range of values for prior mean model sizes $\bar{k}$, with each variable having a prior inclusion probability of $\gamma = \bar{k}/K$, independent of the inclusion of other variables. We estimated our models for 9 different expected prior model sizes, $\bar{k} \in \{5, 7, 9, 11, 16, 22, 28, 40, 80\}$ resulting in the prior inclusion probabilities shown in Table 1.

We follow Sala-I-Martin et al. (2004) in assuming that most researchers strongly prefer models containing a large number of vari-

[3]For convenience we omit subscripts to PIP throughout this paper.

[4]See Mitchell and Beauchamp (1988).

$\bar{k}$	5	7	9	11	16	22	28	40	80
γ	0.031	0.044	0.056	0.069	0.1	0.14	0.18	0.25	0.5

Table 1: Prior model size and associated prior inclusion probabilities for the single variables.

ables so we will concentrate on models with prior model sizes between 5 and 16 variables. This is also in line with the fact that most empirical models on aggregated default rates (see Simons and Rolwes, 2009; Liu, 2009) use moderate numbers of explanatory variables. Our benchmark model will have the prior model size of $\bar{k} = 7$. While we calculate results for large models as well, we will not focus our attention on these cases when it comes to interpretation.

2.3 Estimation

In order to estimate our models we used Markov Chain Monte Carlo (MCMC) methods. In particular, the Gibbs sampler ran for 200,000 iterations, using a thinning of 10. The first 10,000 draws were discarded as burn-in period. This results in 19,000 draws from the posterior for each parameter of interest. All the computations are done using `JAGS` (Just another Gibbs sampler) and its `R` (see R Development Core Team, 2011) interface packages **rjags** (see Plummer, 2011). MCMC diagnostic is done with the package **coda** (see Plummer et al., 2010).

3 Data

We now apply the presented framework of BMA with ridge regression to analyse aggregate default probabilities in Austria. A common approach taken in the literature (see e.g. Simons and Rolwes, 2009; Foglia et al., 2009, among many others) is to use firm default frequencies as proxy for default probabilities. We follow this line by basing our analysis on quarterly corporate insolvency frequencies for the period between January 1987 and April 2011. These insolvency rates are aggregated over all Austrian corporate sectors and are

calculated by dividing the number of quarterly defaults by the total number of firms, which results in quarterly aggregated default rates, $\boldsymbol{pd}$. The number of firm defaults and the total number of firms were obtained from the Austrian creditor association Kreditschutzverband von 1870. As noted above, we transform default rates via the logit function, i.e.,

$$\boldsymbol{y} := \text{logit}(\boldsymbol{pd}).$$

The set of potential explanatory variables contain 32 different macroeconomic variables which are taken from the database of Oesterreichische Nationalbank (OeNB).[5] These macroeconomic variables are part of the Austrian Quarterly Forecast Model (AQM) and are used for forecasting by the OeNB twice a year. As this dataset reflects the variable set of a macroeconomic forecasting model, our results can be used to integrate the time-series of credit defaults into the macro-model, or implement a stress testing framework building on the respective macroeconomic forecasts.

Another advantage of using this dataset is that the list of candidate regressors covers multiple aspects of the economic environment. We consider financial regressors, like interest rates, the stock index and credit amount outstanding, private sector indicators, e.g., private consumption and disposable income, as well as general and external trade related variables, like GDP, exports and investment. Additionally, various price indicators, like the harmonized consumer price index or the oil price are included.

This large set is even further increased by adding lags up to 4 quarters of each candidate regressor, hence resulting in a design matrix $\boldsymbol{X}$ containing 160 explanatory variables each with 97 quarterly observations. The variable names, the applied transformation as well as two of their autocorrelation coefficients are illustrated in Table 2. The variables included were transformed as indicated in column 2 in Table 2 to ensure stationarity of the time-series. "YoY-Log-Difference"

[5] The only exception is the ATX, Austrian Traded Index, which was taken from Datastream.

equals a transformation of the original time-series, X_t, by $\log X_t - \log X_{t-4}$, "YoY-Difference" by $X_t - X_{t-4}$ and "YoY-Rel-Difference" by $(X_t/X_{t-4}) - 1$ where t is the time indicator in quarters.[6]

4 Results

In this section we present the results from the combined approach of BMA with ridge regression described in Section 2 applied to the Austrian dataset. To assess variable importance we calculate the posterior inclusion probabilities (PIP). These are a central quantity within BMA to measure a variable's importance (see Sala-I-Martin et al., 2004). In line with prior research (and intuition), we focus on variables with a higher PIP than their prior inclusion probability, i.e,. variables that are deemed more important after consideration of the data. Additionally, means and standard deviations of the coefficients — conditional on model inclusion — are displayed.

4.1 Macroeconomic predictors of firm failure rates: Baseline estimation

We are now ready to present the baseline estimation results with a prior model size[7] of 7. We find a posterior mean of 10.12, which is clearly above the prior model size and suggests that the posterior puts more importance on models with more explanatories.[8]

Table 3 presents the results of our analysis: The first column reports the PIP of the variables within the applied BMA framework. We

[6]Note that this transformation is followed by a standardisation (subtraction of mean and division by standard deviation) within the ridge regression.

[7]Note that as described in Section 2.2.2 a prior model size of 7 does not mean each model includes exactly 7 variables, but that each candidate regressor has a probability of inclusion, which yields on average a model size of 7.

[8]For the sake of completeness we provide here posteriors related to the shrinkage parameter (see Section 2.1). We find for the shrinkage parameter $\lambda = \sigma^2/\tau^2$ a posterior mean of 0.72012, whereby flat (uninformative) hyperpriors on τ^2 and σ^2 were assumed. The posterior means of the variances σ^2 and τ^2 are 0.11893 and 0.16516 respectively.

sorted the variables in decreasing order of PIP and print only those which have a PIP greater than the prior inclusion probability.

From Table 3 we infer that the 3-month interest rate lagged by four quarters (STI L4) has the highest PIP. Its posterior mean coefficient (P.MEAN) is positive and in line with standard economic theory that higher costs of funding imply higher PDs. Similarly, higher interest rates do not only increase the cost of funding but also prevent firms from receiving further funding due to bank lending standards, thus triggering firm failures. This finding is in line with previous literature. Vlieghe (2001), Liu and Wilson (2002) and Liu (2009) among many others report this strong and positive dependence between firm defaults and interest rates.

Interestingly, the second most important variable is the unlagged short term interest rate, STI, which has a negative posterior mean. The fact that the unlagged short term interest rate is negatively related to firm defaults is to the authors' knowledge a common puzzle in empirical works on aggregate credit risk (see e.g. Ali and Daly 2010 or Divino et al. 2008). However, there is an economic interpretation for this result. STI are usually highly correlated with central bank fund rates and these tend to be raised in economic boom phases to avoid overheating. Thus, higher short term interest rates are a timely measure for economic activity. Clearly, in economic good times PDs tend to decline.

On the third rank we find energy prices with a lag of one year, HEG L4. Energy prices constitute an essential determinant of factor prices and thus obviously pose a very relevant risk factor from the perspective of firms. Its posterior mean of 0.084 indicates a positive relationship between defaults and rising energy prices. This finding illustrates the power of BMA. While numerous papers have identified inflation as a determinant of aggregated default rates (see e.g. Foglia et al., 2009; Virolainen, 2004), we find a *component* of inflation, energy price rises, as one major factor. Owing to the application of BMA one is able to include components of indicators instead of the aggregates such as inflation or GDP yielding more precise conclusions.

The fourth most important variable according to its PIP is nominal import growth (MTN). To find imports among the top ranks is surprising as respective literature usually refrains to include it. However, the positive sign of the coefficient can be supported by several arguments. Firstly, imports by corporates are expenses. Ceteris paribus higher expenses increase the default probability. Secondly, more imports by private households could substitute domestic products which decrease the average revenue of domestic corporates. Thirdly, the time-series of imports might also catch exchange rate fluctuations to some extent, which in turn appear in papers as in Foglia et al. (2009) and Bhattacharjee et al. (2009).

The fifth (and tenth) highest PIPs can be observed for PRO L1 (PRO), the log differences of the average labour productivity. While the interpretation is less straight forward, it might be that an increase in labor productivity drives those firms out of the market which can not adopt such a productivity shock in their business strategy.

Furthermore, on the following ranks we find GON, gross operating surplus, and WIN, total compensation to employees. Both variables have the expected sign of the posterior mean. GON measures profits of firms which intuitionally lower firm defaults and is also reported in previous findings (see e.g. Liu, 2004; Liu and Wilson, 2002). WIN is the aggregate sum of wages paid out and according to our findings reduces the probability of a firm's default. It is important here to stress the difference to the variable WURYD, real compensation per employee, which appears on rank 13 with a lag of one year and a *positive* coefficient. While WURYD measures compensation per employee, WIN is the total sum across the economy. While seemingly related, there are important distinctions which also come apparent when regarding their opposite signs of posterior means. First, WURYD is measured *per employee* making it inversely related to the general employment level — or put differently, WIN is positively related to the general employment level, which constitutes another important variable at rank 19, LNN. Also, the finding of a positive coefficient on real compensation per employee confirms the results presented in Vlieghe (2001). Another difference is the time index with which both variables enter. While WIN enters without lag, thus reflecting more contemporary conditions, WURYD enters with a lag of 4.

On the ranks 8 to 9 we find government interest payments (GEI) and investment (ION). Both variables were selected in approximately 30% of the visited models. As the negative sign of ION indicates, investments reduce the number of defaults in the economy as also reported by e.g. Boss et al. (2007). Two channels may be responsible for this fact. First, investments reduce the number of firm defaults as they are a proxy for fresh equity induced into corporates. Second, it may also be that in times few firms default, managers decide to invest more, which results in a mutual dependance of both variables. However, the fact that investment enters with a lag of 2 (compare Table 3) speaks in favor of the first channel. The appearance of GEI is less anticipated. The positive sign of its posterior mean (together with a relatively small posterior standard deviation) tells us that in times of high interest payment from the side of the government defaults tend to increase. Potentially, this finding reflects the increased economic uncertainty when sovereign spreads rise. Koopman and Lucas (2005) report a positive dependence of default rates with aggregated corporate spreads. As such a variable is missing in our dataset GEI potentially acts as a proxy.

Beyond the "top 10", variables already discussed like other nominal investment, ION, and total employment, LNN, appear. In the majority of cases variables as well as posterior means are plausible from an economic perspective. However, private consumption, PCN, lagged by 4 quarters enters in most models with a positive sign. This puzzling finding may be explained by the fact that its unlagged cousin, private real consumption, PCR, enters with a negative sign at rank 16 (and further on rank 21 and 22). High private consumption one year ago might cause too optimistic turnover predictions on the side of firms, which begin to falter once stock levels do not sell. Such an interpretation is supported by the fact that contemporaneous (real) private consumption enters with the expected negative sign.

At the same time it is not only interesting to look at variables that were selected frequently, but also at variables that were *not* selected. Among those we find for example ATX, the Austrian Traded (stock market) Index. As the majority of firms are small enterprises, little (or no) dependence on stock market returns is plausible. Less anticipated is the fact that classical macroeconomic variables, especially GDP, or disposable income, play also a minor role. In our model setting the data do not support their inclusion, which confirms the findings of Simons and Rolwes (2009). However, this underlines the existence of model uncertainty and therefore the need for averaging over sets of possible models. Indeed, as noted above, by applying BMA we find components of general indicators, like investment of GDP and energy prices of inflation, as major risk drivers. BMA thus allows for a deeper insight into the matter of firm default determinants.

4.2 Model size robustness

So far, we presented the results of our baseline estimation with a prior model size of $\bar{k} = 7$. Although, we believe that models with 7 expected variables are reasonable, this choice is somewhat arbitrary and the effects of using different prior model sizes need to be explored. For the 30 most substantial variables in the baseline model, Tables 4 and 5 present the PIP and posterior means given inclusion of different prior model sizes. The prior inclusion probabilities are simply given by the choice of $\bar{k}$ divided by the number of possible variables, $K = 160$. For each prior model size, variables which appear within the 10 most substantial variables are printed in bold. Variables that are substantial in the baseline model but not when other priors are in use are printed in italics.

In total we find three variables which appear within the 10 most substantial variables for all considered model sizes. These are STI L4, WIN and GEI. From Table 5 we can infer that the signs of the variables are consistent across the different model sizes. Solely for two variables, PCR L4 and WURYD we find for the prior model size $\bar{k} = 80$ controversial signs compared to the other considered models[9]. A summary of Table 4 is also displayed in Figure 1 where we show

[9]These variables are ranked 52 respectively 35 in the baseline model and appear therefore not in Tables 4 and 5.

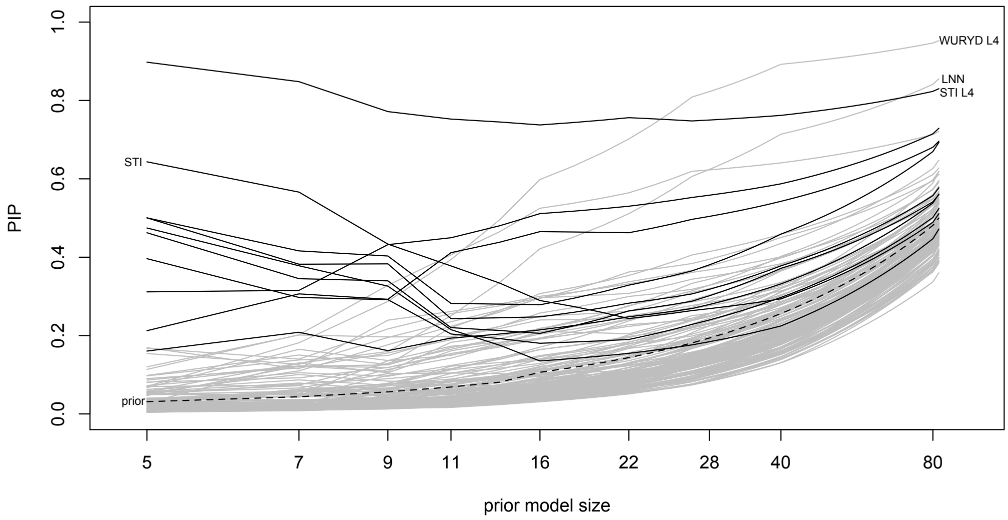

Figure 1: PIPs for top variables in the benchmark model (in black) and less important variables (in grey) across varying prior model sizes. The dotted line illustrates the prior inclusion probabilities γ. For the x-axis a logarithmic scale is used.

PIPs of the top variables in the benchmark model across the range of prior model sizes.

Variables tending to lose importance when increasing prior model size. Four variables tend to lose importance — by dropping out of the "top 10" — when increasing the prior model size. These are STI, MTN, PRO L1 and GON. Moreover, PRO L1 even becomes unsubstantial for prior model sizes above 40. This suggests that such variables could be acting as "catching-all" for various other effects (see Sala-I-Martin et al., 2004). That means in smaller models these variables capture several effects and mechanisms in a combined form, while in larger models, these effects are broken up as more regressors are added. As a matter of fact, this in turn implies that when interpreting coefficients one has to focus even more on the *partial* character of the coefficient, i.e., measuring the effect given the inclusion of other regressors.

A good example is STI, unlagged short term interest rates, which becomes less substantial as we increase the model size (see Figure 1). For prior model sizes $\bar{k} = 40$ respectively $\bar{k} = 80$, STI appears on the 34th respectively 44th rank, while for our benchmark model it appears on the 2nd rank. On the other hand, variables like PCN (nominal private consumption, lagged 4 quarters), PCR (real private consumption), ION (nominal investment) and LNN (total employment) become more important for larger prior model sizes. This nourishes the hypothesis that short term interest rates might be a "catch all substitute" for private consumption and investment. The fact that PCR, ION and LNN enter without lag, i.e., in their contemporaneous form, also supports the interpretation mentioned before – that STI is a proxy of economic activity in smaller models.

Variables becoming "top 10" when increasing prior model size. Within the most substantial variables we find some variables which do not appear within the "top 10" set of the baseline model, but seem to become "top 10" when changing the prior model size. Nevertheless, all these variables are substantial in the baseline model (that is, show a higher PIP than prior inclusion probability) and are mostly ranked between the 11th and 20th rank in the baseline model $\bar{k} = 7$. These are nominal private consumption (PCN L4), real compensation per employee (WURYD L4), both lagged by one year, total employment (LNN), total compensation to employees (WIN L3) lagged by 3 quarters, private consumption rate (PCR) and nominal investment (ION). Additionally, we find the variables real domestic demand (DDR), its one year lagged values[10] (DDR L4) and nominal total compensation to employees (WIN L4), lagged by one year, appearing as "top 10" variable for some considered prior model sizes.

[10]DDR L4 is ranked 39 for our baseline model and ranked 10 for the model $\bar{k} = 80$ with a PIP of 0.6193.

5 Conclusion and Discussion

In this paper we propose a fully Bayesian approach combining ridge regression and BMA to determine which macroeconomic variables are substantially related to aggregated probabilities of default. Compared to the literature, which mainly focuses on one single model, our approach addresses the problem of model uncertainty. Additionally, we propose ridge regression to deal with multicollinearity, an immanent problem in case lagged variables are included. In our benchmark model the most frequently selected candidate regressors indicate that firms' factor prices play a key role in determining defaults. Energy prices and interest rates lagged by one year are positively related to defaults. On the other hand, indicators of economic activity like investment and contemporaneous short term interest rates are associated with fewer firm defaults. As expected, firms' profits reduce the expected number of failures. Interestingly, classical macroeconomic variables, like GDP or disposable income are less frequently selected. This finding underlines the need for an approach capable of dealing with model uncertainty, a feature BMA perfectly provides.

Finally, we show that the results of our baseline model are fairly robust to the choice of the prior model size. More precisely, when increasing the prior model size, variables do not change the sign of their posterior mean (with only 2 exceptions in 54 substantial variables considered). Moreover, most of the "top 10" variables remain within the 20 most important variables for other estimated prior model sizes. However, the relative importance of some regressors does change. This finding suggests that some variables being of high relevance in smaller models act as proxy for multiple effects combined which can be successively split into its components when considering models of larger size.

Further research is needed to better understand the dynamics of firm failures, a highly relevant time-series for regulators and banks alike. On one hand, the application of statistical approaches robust to model uncertainty should be applied on a dataset of wider geographical coverage. In line with the findings of Simons and Rolwes

(2009) country specific circumstances need to be analysed. Also, our methodological framework allows for the considerations of (even) more candidate regressors. On the other hand, form a methodological perspective our approach could be revamped in a way that allows the examination of common sets of variables. This would allow to analyze substitutional and complementary effects between the explanatories. That is, asking not only which variables were selected, but also which variables were selected *together*.

NAME	MEANING	TRANSFORMATION	MEAN	SD
PD	Aggregated default rates	logit	-5.309	0.186
		No Transformation	0.005	0.001
			LAG1	LAG4
ATX	Austrian Traded Index	YoY-Log Difference	0.797	0.041
CPN	Nominal private credit, amount outstanding	YoY-Log-Difference	0.891	0.401
DDR	Real domestic demand	YoY-Log-Difference	0.923	0.454
GEI	Government interest payment	YoY-Log-Difference	0.881	0.033
GON	Nominal gross operating surplus	YoY-Log-Difference	0.862	0.001
HEG	HIC for energy	YoY-Difference	0.750	-0.263
HEX	HIC for non-energy	YoY-Difference	0.863	0.437
HIC	Harmonized index of consumption prices	YoY-Difference	0.830	0.164
IEN	Nominal equipment investment	YoY-Log-Difference	0.876	0.003
IER	Real equipment investment	YoY-Log-Difference	0.884	0.042
ION	Nominal other investment	YoY-Log-Difference	0.881	0.040
IOR	Real other investment	YoY-Log-Difference	0.881	0.040
ITR	Real total investment	YoY-Log-Difference	0.861	0.042
LNN	Total employment	YoY-Difference	0.868	0.200
LTI	Nominal long-term interest rate	No Transformation	0.970	0.834
MTN	Nominal imports	YoY-Log-Difference	0.887	-0.004
MTR	Real imports	YoY-Log-Difference	0.909	0.120
PCN	Nominal private consumption	YoY-Log-Difference	0.958	0.638
PCR	Real private consumption	YoY-Log-Difference	0.957	0.669
POIL	Oil price in EUR	YoY-Difference	0.702	-0.411
PRO	Average labor productivity	YoY-Log-Difference	0.869	0.093
PSN	Nominal private sector savings	YoY-Log-Difference	0.788	-0.082
PYN	Nominal private sector disposable income	YoY-Log-Difference	0.896	0.379
PYR	Private sector disposable income, real	YoY-Log-Difference	0.849	0.167
STI	Nominal short-term interest rate	No Transformation	0.969	0.780
URX	Unemployment rate	YoY-Rel-Difference	0.779	-0.205
WIN	Nominal total compensation to employees	YoY-Log-Difference	0.959	0.685
WURYD	Real compensation per employee	YoY-Log-Difference	0.866	0.302
XTN	Nominal exports	YoY-Log-Difference	0.871	-0.037
XTR	Real exports	YoY-Log-Difference	0.880	0.020
YEN	Nominal GDP	YoY-Log-Difference	0.914	0.291
YER	Real GDP	YoY-Log-Difference	0.892	0.188

Table 2: Set of candidate regressors. Additionally, up to four quarter lags are considered for each variable. The columns LAG1 and LAG4 display the autocorrelation coefficients for one and four quarters.

R	NAME	PIP	P.MEAN	P.SD	R	NAME	PIP	P.MEAN	P.SD
1	STI L4	0.84800	0.28701	0.10891	28	YEN	0.10055	-0.12546	0.07424
2	STI	0.56610	-0.24265	0.07130	29	WURYD L1	0.09220	0.06891	0.03197
3	HEG L4	0.41610	0.08353	0.02716	30	MTN L1	0.08680	0.12613	0.08042
4	MTN	0.38210	0.19201	0.06266	31	HEG L3	0.07895	0.06540	0.03486
5	PRO L1	0.37815	0.12603	0.03609	32	IOR L1	0.07675	-0.06694	0.03973
6	GON	0.34515	-0.14711	0.04303	33	IOR	0.07365	-0.06964	0.04922
7	WIN	0.31540	-0.14142	0.05773	34	HEG L2	0.07205	0.08202	0.04811
8	GEI	0.30640	0.06928	0.02431	35	WURYD	0.06620	0.07033	0.03893
9	ION L2	0.29725	-0.08130	0.02744	36	CPN L3	0.06325	-0.05976	0.02815
10	PRO	0.20840	0.11102	0.04294	37	STI L1	0.06205	-0.13917	0.12336
11	PCN L4	0.20200	0.16323	0.06898	38	HIC L4	0.06205	0.07822	0.04667
12	ION L1	0.20145	-0.08711	0.03077	39	DDR L4	0.05860	0.16775	0.14971
13	WURYD L4	0.18160	0.10794	0.04346	40	MTR L1	0.05790	0.11667	0.07839
14	DDR L1	0.16495	0.12784	0.05158	41	WIN L1	0.05780	-0.08299	0.07677
15	WIN L3	0.15040	-0.13192	0.06706	42	HEX L4	0.05720	0.07200	0.04427
16	PCR	0.14005	-0.09409	0.04336	43	GEI L1	0.05690	0.05129	0.03397
17	PRO L2	0.13775	0.08565	0.03721	44	MTR	0.05640	0.08118	0.06636
18	CPN L4	0.13020	-0.06264	0.02442	45	XTR	0.05590	0.07965	0.05714
19	LNN	0.12955	-0.07927	0.03024	46	PCR L3	0.05185	-0.10317	0.06654
20	ION	0.12670	-0.08061	0.03964	47	GON L1	0.05160	-0.06871	0.05737
21	PCR L1	0.12265	-0.09098	0.04514	48	DDR L2	0.05070	0.08072	0.05292
22	PCR L2	0.12205	-0.11796	0.05754	49	IOR L2	0.04900	-0.03837	0.05501
23	WIN L2	0.12170	-0.11908	0.05970	50	STI L3	0.04890	0.07022	0.08215
24	ITR	0.11715	0.10109	0.04509	51	XTN	0.04795	0.08639	0.05974
25	DDR	0.11510	0.14892	0.07090	52	PCR L4	0.04590	-0.12006	0.08160
26	WIN L4	0.11000	-0.16907	0.08902	53	PRO L4	0.04500	0.05605	0.03107
27	ION L3	0.10510	-0.08806	0.06439	54	ITR L1	0.04380	0.07015	0.04717

Table 3: Baseline model for all 54 variables with a PIP greater than the prior inclusion probability γ of $7/160 = 0.0437$. P.MEAN and P.SD denote the coefficient's mean and standard deviation given its model inclusion.

	$\bar{k}=5$	$\bar{k}=7$	$\bar{k}=9$	$\bar{k}=11$	$\bar{k}=16$	$\bar{k}=22$	$\bar{k}=28$	$\bar{k}=40$	$\bar{k}=80$
STI L4	**0.8977**	**0.8480**	**0.7716**	**0.7525**	**0.7376**	**0.7561**	**0.7478**	**0.7620**	**0.8301**
STI	**0.6434**	**0.5661**	**0.4334**	**0.3777**	0.2894	0.2434	0.2637	0.2938	0.5110
HEG L4	**0.5004**	**0.4161**	**0.4031**	**0.2821**	0.2786	0.3287	**0.3650**	**0.4589**	**0.6931**
MTN	**0.5003**	**0.3821**	**0.3830**	0.2432	0.2477	0.2827	0.3063	0.3778	0.5773
PRO L1	**0.4746**	**0.3781**	**0.3261**	0.2151	0.1351	0.1547	0.1762	*0.2245*	*0.4718*
GON	**0.4624**	**0.3452**	**0.3396**	0.2202	0.2056	0.2630	0.2877	0.3715	0.5605
WIN	**0.3115**	**0.3154**	**0.4320**	**0.4496**	**0.5112**	**0.5301**	**0.5524**	**0.5873**	**0.7291**
GEI	**0.2124**	**0.3064**	**0.2926**	**0.4116**	**0.4652**	**0.4625**	**0.4965**	**0.5425**	**0.6961**
ION L2	**0.3962**	**0.2973**	**0.2917**	0.2040	0.1804	0.1900	0.2284	0.2973	0.5243
PRO	0.1605	**0.2084**	0.1615	0.1935	0.2145	0.2480	0.2690	0.3317	0.5613
PCN L4	0.1137	0.2020	**0.3283**	**0.4061**	**0.5248**	**0.5644**	**0.6192**	**0.6407**	**0.7210**
ION L1	0.1202	0.2014	0.1966	0.1973	0.2496	0.2624	0.2861	0.3356	0.5424
WURYD L4	0.0702	0.1816	0.2873	**0.3939**	**0.5983**	**0.7018**	**0.8087**	**0.8923**	**0.9527**
DDR L1	0.0672	0.1650	0.1363	0.1727	0.1930	0.2235	0.2441	0.3200	0.5372
WIN L3	0.0983	0.1504	0.1263	0.2204	**0.3073**	**0.3534**	**0.4057**	**0.4584**	**0.6277**
PCR	0.1677	0.1401	0.2170	**0.2561**	**0.2958**	**0.3374**	**0.3663**	**0.4351**	**0.6478**
PRO L2	0.0973	0.1378	0.1492	0.1961	0.2399	0.2435	0.2691	0.3281	0.5444
CPN L4	0.0573	0.1302	0.1050	0.1610	0.2037	0.2424	0.2687	0.3141	*0.4890*
LNN	*0.0282*	0.1295	0.1593	**0.2505**	**0.4216**	**0.5112**	**0.6061**	**0.7140**	**0.8555**
ION	0.1541	0.1267	0.1953	0.2354	0.2950	**0.3623**	**0.3838**	**0.4503**	0.6117
PCR L1	0.0864	0.1226	0.1185	0.1542	0.1943	0.2410	0.2718	0.3648	0.5813
PCR L2	0.0470	0.1221	0.1086	0.1828	0.2170	0.2751	0.3332	0.3979	0.5917
WIN L2	0.0801	0.1217	0.1252	0.1694	0.2147	0.2766	0.3140	0.4021	0.6189
ITR	0.0901	0.1172	0.1556	0.1893	0.2575	0.2998	0.3181	0.3805	0.5955
DDR	**0.1691**	0.1151	0.2393	**0.2614**	**0.3001**	**0.3321**	0.3612	0.4008	0.5663
WIN L4	0.0521	0.1100	0.1766	0.2235	**0.2965**	0.3076	0.3438	0.3833	0.5508
ION L3	0.0725	0.1051	0.0961	0.1238	0.1642	0.1978	0.2215	0.2846	0.5123
YEN	0.0585	0.1006	0.0836	0.0825	0.1154	0.1704	0.2035	0.2783	0.5395
WURYD L1	0.0812	0.0922	0.1184	0.1509	0.1910	0.2281	0.2402	0.3040	0.5232
MTN L1	0.0625	0.0868	0.0732	0.1127	0.1382	0.1681	0.2152	0.2713	0.5088

Table 4: Posterior inclusion Probabilities with different model sizes sorted by the PIP derived in the baseline model, $\bar{k}=7$ for the 30 most substantial variables. Boldly printed variables appear within the 10 most important variables, while Italic printed variables do not appear substantial for the considered model size. Rows sorted by PIP estimated in the baseline model.

	$\bar{k}=5$	$\bar{k}=7$	$\bar{k}=9$	$\bar{k}=11$	$\bar{k}=16$	$\bar{k}=22$	$\bar{k}=28$	$\bar{k}=40$	$\bar{k}=80$
STI L4	0.3073	0.2870	0.2643	0.2324	0.1798	0.1523	0.1320	0.1053	0.0669
STI	-0.2608	-0.2426	-0.2311	-0.1931	-0.1404	-0.0973	-0.0750	-0.0453	-0.0201
HEG L4	0.0850	0.0835	0.0841	0.0770	0.0727	0.0679	0.0627	0.0576	0.0446
MTN	0.2062	0.1920	0.1859	0.1443	0.1042	0.0905	0.0758	0.0614	0.0335
PRO L1	0.1316	0.1260	0.1223	0.1066	0.0698	0.0565	0.0441	0.0303	0.0174
GON	-0.1526	-0.1471	-0.1409	-0.1201	-0.0957	-0.0860	-0.0755	-0.0611	-0.0329
WIN	-0.1490	-0.1414	-0.1512	-0.1438	-0.1274	-0.1150	-0.1004	-0.0828	-0.0531
GEI	0.0702	0.0693	0.0659	0.0672	0.0644	0.0592	0.0571	0.0516	0.0406
ION L2	-0.0825	-0.0813	-0.0801	-0.0770	-0.0661	-0.0572	-0.0517	-0.0420	-0.0260
PRO	0.1166	0.1110	0.1013	0.0966	0.0809	0.0722	0.0617	0.0498	0.0323
PCN L4	0.1563	0.1632	0.1653	0.1587	0.1446	0.1234	0.1134	0.0922	0.0532
ION L1	-0.0891	-0.0871	-0.0866	-0.0836	-0.0787	-0.0694	-0.0624	-0.0497	-0.0292
WURYD L4	0.1083	0.1079	0.1135	0.1105	0.1037	0.0964	0.0944	0.0885	0.0737
DDR L1	0.1200	0.1278	0.1120	0.1041	0.0861	0.0725	0.0630	0.0482	0.0263
WIN L3	-0.1248	-0.1319	-0.1115	-0.1302	-0.1207	-0.1057	-0.0955	-0.0732	-0.0417
PCR	-0.1032	-0.0941	-0.1055	-0.0938	-0.0851	-0.0771	-0.0694	-0.0589	-0.0408
PRO L2	0.0889	0.0856	0.0855	0.0829	0.0750	0.0665	0.0581	0.0478	0.0300
CPN L4	-0.0573	-0.0626	-0.0550	-0.0569	-0.0549	-0.0492	-0.0455	-0.0389	-0.0236
LNN	-0.0546	-0.0793	-0.0734	-0.0778	-0.0761	-0.0745	-0.0733	-0.0698	-0.0580
ION	-0.0919	-0.0806	-0.0858	-0.0831	-0.0785	-0.0746	-0.0685	-0.0582	-0.0360
PCR L1	-0.0915	-0.0910	-0.0861	-0.0838	-0.0740	-0.0670	-0.0616	-0.0529	-0.0337
PCR L2	-0.0824	-0.1180	-0.1070	-0.1070	-0.0919	-0.0791	-0.0736	-0.0589	-0.0356
WIN L2	-0.1158	-0.1191	-0.1112	-0.1171	-0.1065	-0.0928	-0.0790	-0.0649	-0.0405
ITR	0.1086	0.1011	0.0935	0.0943	0.0887	0.0817	0.0710	0.0567	0.0359
DDR	0.1776	0.1489	0.1580	0.1406	0.1167	0.1016	0.0867	0.0660	0.0334
WIN L4	-0.1522	-0.1691	-0.1771	-0.1723	-0.1390	-0.1086	-0.0942	-0.0654	-0.0289
ION L3	-0.0813	-0.0881	-0.0721	-0.0785	-0.0607	-0.0579	-0.0509	-0.0400	-0.0233
YEN	-0.1159	-0.1255	-0.0965	-0.0774	-0.0651	-0.0715	-0.0558	-0.0420	-0.0273
WURYD L1	0.0707	0.0689	0.0674	0.0676	0.0630	0.0578	0.0511	0.0417	0.0275
MTN L1	0.1362	0.1261	0.0879	0.1095	0.0822	0.0644	0.0579	0.0404	0.0212

Table 5: Posterior means conditional on inclusion with different prior model sizes. Rows sorted by PIP estimated in the baseline model.

References

A. Ali and K. Daly. Macroeconomic determinants of credit risk: Recent evidence from a cross country study. *International Review of Financial Analysis*, 19(3):165 – 171, 2010.

E. I. Altman. *Corporate Financial Distress*. John Wiley, New York, 1983.

E. I. Altman. The success of business failure prediction models. *Journal of Banking and Finance*, 8:171–198, 1984.

D. Avramov. Stock return predictability and model uncertainty. *Journal of Financial Economics*, 64:423–458, 2002.

A. Bhattacharjee, C. Higson, S. Holly, and P. Kattuman. Macroeconomic instability and business exit: Determinants of failures and acquisitions of UK firms. *Economica*, 76(301):108–131, 2009.

M. Boss, M. Fenz, J. Pann, C. Puhr, M. Schneider, and E. Ubl. Modeling credit risk through the austrian business cycle: An update of the oenb model. *Financial Stability Report*, (17):85–101, June 2007. URL `http://ideas.repec.org/a/onb/oenbfs/y2009i17b2.html`.

J. A. Divino, E. S. Lima, and J. Orrillo. Interest Rates and Default in Unsecured Loan Markets. Technical report, Catholic University of Brasilia, 2008.

C. Fernandez, E. Ley, and M. F. J. Steel. Model uncertainty in cross-country growth regressions. *Journal of Applied Econometrics*, 16(5):563–576, 2001.

A. Foglia, R. Fiori, and S. Iannotti. Beyond Macroeconomic Risk: The Role of Contagion in the Italian Corporate Default Correlation. Research Paper 12, CAREFIN Research, 2009.

J. Friedman, T. Hastie, and R. Tibshirani. *The Elements of Statistical Learning*. Springer, Berlin, 2 edition, 2009.

E. I. George and R. E. McCulloch. Variable selection via Gibbs sampling. *Journal of the American Statistical Association*, 88:1–10, 1993.

E. I. George and R. E. McCulloch. Approaches for Bayesian variable selection. *Statistica Sinica*, 7:339–373, 1997.

A. E. Hoerl and R. W. Kennard. Ridge regression: Biased estimation for nonorthogonal problems. *Technometrics*, 12(1):55–67, 1970a.

A. E. Hoerl and R. W. Kennard. Ridge regression: Applications to nonorthogonal problems. *Technometrics*, 12(1):69–82, 1970b.

J. A. Hoeting, D. Madigan, A. E. Raftery, and C. T. Volinsky. Bayesian Model Averaging: A Tutorial. *Statistical Science*, 14:382–417, 1999.

T. C. Hsiang. A Bayesian view on ridge regression. *Journal of the Royal Statistical Society. Series D (The Statistician)*, 24:267–268, 1975.

S. J. Koopman and A. Lucas. Business and default cycles for credit risk. *Journal of Applied Econometrics*, 20(2):311–323, 2005.

E. Ley and M. F. J. Steel. On the effect of prior assumptions in Bayesian model averaging with applications to growth regression. *Journal of Applied Econometrics*, 24, 2009.

J. Liu. Macroeconomic determinants of corporate failures: evidence from the UK. *Applied Economics*, 36(9):939–945, 2004.

J. Liu. Business failures and macroeconomic factors in the UK. *Bulletin of Economic Research*, 61(1):47–72, 2009.

J. Liu and N. Wilson. Corporate failure rates and the impact of the 1986 insolvency act: An econometric analysis. *Managerial Finance*, 28:1–11, 2002.

T. J. Mitchell and J. J. Beauchamp. Bayesian variable selection in linear regression. *Journal of the American Statistical Association*, 83(404):1023–1032, 1988.

M. Plummer. *rjags: Bayesian graphical models using MCMC*, 2011. URL `http://CRAN.R-project.org/package=rjags`. R package version 2.2.0-4.

M. Plummer, N. Best, K. Cowles, and K. Vines. *coda: Output analysis and diagnostics for MCMC*, 2010. URL `http://CRAN.R-project.org/package=coda`. R package version 0.14-2.

R Development Core Team. *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria, 2011. URL `http://www.R-project.org/`.

X. Sala-I-Martin, G. Doppelhofer, and R. I. Miller. Determinants of long-term growth: A Bayesian averaging of classical estimates (bace) approach. *American Economic Review:*, 94(4):813–835, 2004.

D. Simons and F. Rolwes. Macroeconomic default modeling and stress testing. *International Journal of Central Banking*, 5(3):177–204, 2009.

M. F. Steel. Bayesian model averaging and forecasting. *Bulletin of EU and US Inflation and Macroeconomic Analysis*, 200, forthcoming, 2011.

K. Virolainen. Macro stress testing with a macroeconomic credit risk model for Finland. Research Discussion Papers 18, Bank of Finland, July 2004. URL http://ideas.repec.org/p/hhs/bofrdp/2004_018.html.

G. W. Vlieghe. Indicators of fragility in the UK corporate sector. *Bank of England Working paper*, pages 1–32, 2001.

J. H. Wright. Bayesian Model Averaging and exchange rate forecasts. *Journal of Econometrics*, 146(2):329–341, 2008.