



THE UNIVERSITY *of* EDINBURGH

This thesis has been submitted in fulfilment of the requirements for a postgraduate degree (e.g. PhD, MPhil, DClinPsychol) at the University of Edinburgh. Please note the following terms and conditions of use:

This work is protected by copyright and other intellectual property rights, which are retained by the thesis author, unless otherwise stated.

A copy can be downloaded for personal non-commercial research or study, without prior permission or charge.

This thesis cannot be reproduced or quoted extensively from without first obtaining permission in writing from the author.

The content must not be changed in any way or sold commercially in any format or medium without the formal permission of the author.

When referring to this work, full bibliographic details including the author, title, awarding institution and date of the thesis must be given.

**3 Essays on Credit Risk Modeling
and the Macroeconomic
Environment**

Dimitrios Papanastasiou

Doctor of Philosophy
University of Edinburgh
2015

Declaration

I declare that this thesis was composed by myself and that the work contained therein is my own, except where explicitly stated otherwise in the text.

(Dimitrios Papanastasiou)

Acknowledgements

I would like to thank Professor Jonathan Crook for being an excellent supervisor and for providing valuable support throughout this long effort to bring this work to fruition.

I would also like to thank the Credit Research Centre for providing me with the Moody's Default & Recovery Database, without which this work would not have been possible.

Finally, I would like to thank the Engineering and Physical Sciences Research Council for financial support for payment of my tuitions.

To Eglè and my family for the moral support all these years

Abstract

In the aftermath of the recent financial crisis, the way credit risk is affected by and affects the macroeconomic environment has been the focus of academics, risk practitioners and central bankers alike. In this thesis I approach three distinct questions that aim to provide valuable insight into how corporate defaults, recoveries and credit ratings interact with the conditions in the wider economy.

The first question focuses on how well the macroeconomic environment forecasts corporate bond defaults. I approach the question from a macroeconomic perspective and I make full use of the multitude of lengthy macroeconomic time series available. Following the recent literature on data-rich environment modelling, I summarise a large panel of 103 macroeconomic time series into a small set of 6 dynamic factors; the factors capture business cycle, yield curve, credit premia and equity market conditions. Prior studies on dynamic factors use identification schemes based on principal components or recursive short-run restrictions. The main contribution to the body of existing literature is that I provide a novel and more robust identification scheme for the 6 macro-financial stochastic factors, based on a set of over-identifying restrictions. This allows for a more straightforward interpretation of the extracted factors and a more meaningful decomposition of the corporate default dynamics. Furthermore, I use a novel Bayesian estimation scheme based on a Markov chain Monte Carlo algorithm that has not been used before in a credit risk context. I argue that the proposed algorithm provides an efficient and flexible alternative to the simulation based estimation approaches used in the existing literature. The sampling scheme is used to estimate a state-of-the-art dynamic econometric specification that is able to separate macro-economic fluctuations from unobserved default clustering. Finally, I provide evidence that the macroeconomic factors can lead to significant improvements in default

probability forecasting performance. The forecasting performance gains become less pronounced the longer the default forecasting horizon.

The second question explores the sensitivity of corporate bond defaults and recoveries on monetary policy and macro-financial shocks. To address the question, I follow a more structural approach to extract theory-based economic shocks and quantify the magnitude of the impact on the two main credit risk drivers. This is the first study that approaches the decomposition of the movements in credit risk metrics from a structural perspective. I introduce a VAR model with a novel semi-structural identification scheme to isolate the various shocks at the macro level. The dynamic econometric specification for defaults and recoveries is similar to the one used to address the first question. The specification is flexible enough to allow for the separation of the macroeconomic movements from the credit risk specific unobserved correlation and, therefore, isolate the different shock transmission mechanisms. I report that the corporate default likelihood is strongly affected by balance sheet and real economy shocks for the cyclical industry sectors, while the effects of monetary policy shocks typically take up to one year to materialise. In contrast, recovery rates tend to be more sensitive to asset price shocks, while real economy shocks mainly affect secured debt recovery values.

The third question shifts the focus to credit ratings and addresses the Through-the-Cycle dynamics of the serial dependence in rating migrations. The existing literature treats the so-called rating momentum as constant through time. I show that the rating momentum is far from constant, it changes with the business cycle and its magnitude exhibits a non-linear dependence on time spent in a given rating grade. Furthermore, I provide robust evidence that the time-varying rating momentum substantially increases actual and Marked-to-Market losses in periods of stress. The impact on regulatory capital for financial institutions is less clear; nevertheless, capital requirements for high credit quality portfolios can be significantly underestimated during economic downturns.

Contents

Abstract	6
1 Introduction	17
1.1 Motivation and Aims	17
1.1.1 Background	17
1.1.2 Objective	20
1.2 Contributions	22
1.2.1 Large scale dynamic factor model to summarize the macro-financial environment	22
1.2.2 Forecasting corporate bond defaults	23
1.2.3 Identifying macro-financial shocks	24
1.2.4 Responses of corporate bond defaults and recoveries to macro- financial shocks	25
1.2.5 Duration dependence and business cycle effects in credit rating momentum effect	25
1.2.6 Impact of non-Markovian transition matrices on portfolio credit risk	26
1.3 Thesis Outline	27

2	Forecasting Corporate Defaults in a Data-Rich Environment	30
2.1	Introduction	30
2.2	Relevant Literature	32
2.3	Econometric Framework	33
2.3.1	The Corporate Default Model	33
2.3.2	Dynamic Factor Model	35
2.4	Estimation	42
2.5	Data	46
2.6	Results	49
2.6.1	Dynamic Factor Model	49
2.6.2	Corporate Default Forecasts	57
2.7	Conclusion	66
	Chapter 2 Appendix	69
2.8	Dynamic Factor Loadings	69
3	Macroeconomic Shocks and Credit Risk:	
	Monetary Policy and Real Economy Effects on Corporate Defaults	
	and Recoveries	72
3.1	Introduction	72
3.2	Relevant Literature	76
3.3	Econometric Framework	78
3.3.1	The Corporate Default Model	78
3.3.2	The Recovery Model for Defaulted Debt	82

3.3.3	The macroeconomic structural VAR model	85
3.3.4	Estimation	90
3.4	Data	94
3.5	Results	99
3.5.1	Macroeconomic VAR	99
3.5.2	Implications for Credit Risk	104
3.6	Conclusion	116
4	Momentum in Credit Rating Downgrades: Through-the-Cycle Evidence	118
4.1	Introduction	118
4.2	Relevant Literature	121
4.3	Econometric Framework	125
4.3.1	Specification	125
4.3.2	Estimation	130
4.4	Data	132
4.5	Empirical Analysis	139
4.5.1	Parameter Estimates	139
4.5.2	Impact on Credit Portfolio Losses	146
4.6	Conclusion	161
5	Conclusion	164
5.1	Summary	164
5.2	Contribution to knowledge	166

5.3	Implications for Risk Management and Policy Making	168
5.4	Limitations and Extensions	170
Technical Appendix		174
A	Markov chain Monte Carlo	174
A.1	Gibbs Sampling	174
A.2	Slice Sampling	175
B	Bayesian Analysis of Multivariate Regression Models	176
C	State Space Models	180
C.1	State Space Models and Bayesian Inference	180
C.2	The Kalman Filter and Smoother	182
References		187

List of Figures

Chapter 2	30
2.1 Historical Time Series - Macroeconomic Variables	48
2.2 Historical Time Series - Corporate Default Rates	50
2.3 Estimated Dynamic Factors	51
2.4 Estimated Frailty Factors	61
2.5 Default Rate Predictions Errors	63
Chapter 3	72
3.1 Historical Time Series - Macroeconomic Variables	95
3.2 Historical Time Series - Credit Variables	98
3.3 Macroeconomic Impulse Responses Functions	100
3.4 Corporate Default Impulse Responses Functions	107
3.5 Corporate Debt Recovery Impulse Responses Functions	113
Chapter 4	118
4.1 Quarterly number of rating migrations	135
4.2 Baseline intensity rates per industry sector	136

4.3	Rating momentum time varying effects	145
4.4	Predicted portfolio distributions	154

List of Tables

Chapter 2	30
2.1 Summary of Moody's Corporate Default Data	48
2.2 Estimated Dynamic Factor Transition Matrix	52
2.3 Dynamic Factor Forecast Error Variance Decomposition and R^2	54
2.4 Corporate Default Parameter Estimates	59
2.5 Default Rate Predictions RMSE and MAE	64
2.6 Estimated Dynamic Factor Loadings	71
Chapter 3	72
3.1 Macroeconomic VAR Forecast Error Variance Decomposition	103
3.2 Parameter Estimates for Default Rates	105
3.3 Default Forecast Error Variance Decomposition	110
3.4 Parameter Estimates for Recovery Rates	111
3.5 Recovery Forecast Error Variance Decomposition	115
Chapter 4	118
4.1 Migration event counts	133

4.2	Average Migration Rates	136
4.3	Average Migration Rates - Downward Momentum	138
4.4	Downgrade parameter estimates - Investment Grade	142
4.5	Credit portfolio loss simulation results - Investment Grade Portfolio . .	157
4.6	Credit portfolio loss simulation results - Speculative Grade Portfolio . .	158

Abbreviations

AIRB Advanced Internal Ratings Based.

CFNAI Chicago Fed National Activity Index.

EAD Exposure At Default.

FAVAR Factor Augmented Vector Autoregression.

FEVD Forecast Error Variance Decomposition.

FIRB Foundation Internal Ratings Based.

FRED Federal Reserve Economic Database.

GLMM Generalised Linear Mixed Model.

IRB Internal Ratings Based.

IRC Incremental Risk Charge.

LGD Loss Given Default.

MAE Mean Absolute Error.

MCMC Markov Chain Monte Carlo.

MtM Marked-to-Market.

OLS Ordinary Least Squares.

PD Probabability to Default.

PiT Point-in-Time.

RMSE Root Mean Squared Error.

RWA Risk Weighted Assets.

SVAR Structural Vector Autoregression.

TTC Through-The-Cycle.

VAR Vector Autoregression.

VaR Value-at-Risk.

Chapter 1

Introduction

1.1 Motivation and Aims

1.1.1 Background

In the aftermath of the recent financial crisis, the interaction of financial and real sides of the economy has come to the forefront of today's risk measurement practices, central banking policy debate and academic research in the broad field of financial economics. The notion of regulatory driven stress testing exercises has suddenly become a de-facto tool to assess the financial sector's strength and viability under adverse conditions; adverse conditions that a few years ago seemed implausible and yet they are now used as benchmarks for market and credit portfolio losses and capital requirements calculations.

The financial crisis highlighted vulnerabilities in the quantification of all risk types. The study of some of those risk types, such as counterparty and liquidity risk, is relatively new. On the contrary, research on other risk types such as market and credit risk has been active for many years; market risk certainly gained attention after the 1997-2002 period (Asian crisis, Russian default, Dot-Com bubble, September 11th attacks), while credit risk is an integral part of corporate finance for many decades. Nevertheless, the importance of market risk is somewhat masked by the ability to hedge and close-down positions in the short-term. At the same time, market risk is typically relevant for investment banks and asset managers, having medium to low

importance for commercial banks. It is really accumulated credit losses that lead to the deterioration of the overall banking sectors balance sheet, tightening of lending criteria and, consequently, amplification of economic downturns.

The literature in credit risk quantification is extensive. The study of Altman (1968) is perhaps the first to thoroughly explore the empirical determinants of bankruptcy, therefore linking credit risk to fundamental drivers. Merton (1974) is the first paper to structurally formulate the default likelihood and the price of corporate debt as functions of a firm's asset value and level of debt, while Jarrow and Turnbull (1995) provide the first reduced-form view on pricing derivative securities subject to credit risk. Despite the two approaches on credit risk being distinct, Duffie and Lando (2001) show that reduced form pricing can be viewed as a structural model with incomplete accounting information. A challenge for all credit models is the ability to perform well in both benign and stress periods. Das et al. (2007) and Duffie et al. (2009) introduce the notion of "frailty" in the credit risk literature to describe unobservable variables that are correlated across firms, therefore causing clustering in corporate failings; Duffie et al. (2007) use "frailty" terms to enhance the prediction of firm level defaults. More recently, Chen et al. (2009) adjust the structural credit risk models for market-varying prices of risk to help fit historical credit spread levels. Finally, the study of Chen (2010) is the first to provide a dynamic capital structure model that is able to capture the historical variation in credit spreads linking explicitly default losses and leverage to the business cycle.

The substantial increase in computational speed, combined with the availability of long time series of US corporate default data spanning over multiple credit cycles, enabled the estimation of high dimensional and possibly non-Gaussian systems of endogenous variables. Recent econometric studies on corporate credit risk incorporate complex correlation structures and dynamics. McNeil and Wendin (2007) apply a Bayesian Generalised Linear Mixed Model (GLMM) on quarterly corporate default counts, incorporating business cycle dynamics and latent, default specific, random effects. Koopman and Lucas (2008) focus more on the latent credit factors and present a non-Gaussian dynamic factor model to isolate the different dynamics from a panel of industry sector and rating grade corporate bond defaults. Koopman et al. (2012) extend the latent factor framework of Koopman and Lucas (2008) to include observed

macroeconomic variables and show that the observed covariates account for approximately one third of the systematic variation in defaults. Recognising the benefits of using many variables to capture all the aspects of the economic cycle, Koopman et al. (2011) use a large scale dynamic factor model based on principal components to reduce the dimensionality of the macro-financial stochastic movements; these dynamic factors are then used to forecast corporate bond defaults. A similar dynamic factor model based on principal components is employed by Boivin et al. (2013) to study the effect of credit shocks on the rest of the economic and financial environment. Finally, Creal et al. (2013) use a mixed-measurement and mixed-frequency dynamic factor framework to jointly model corporate defaults, corporate recoveries, rating migrations and key economic and financial variables.

In addition to the modelling of realised defaults, several studies explore the dynamic behaviour of credit ratings. Despite the criticism that the rating agencies have received throughout the recent financial crisis, credit ratings still play a central role in investment decisions on credit instruments, in setting covenants for credit contracts, and in calculating regulatory capital for financial intermediaries. Early studies on discrete time credit migration count modelling include Nickell et al. (2000) and Bangia et al. (2002). Koopman et al. (2008) use a continuous time semi-Markovian framework with a latent credit factor to model the transition matrix. A similar framework in discrete time is used by Stefanescu et al. (2009).

A particular topic of interest in credit migrations is the so-called momentum effect, or directional dependence. Directional dependence implies that a downgrade is more likely for a corporate that has been previously downgraded, and the reverse applies to upgrades. This momentum effect leads to non-Markovian stochastic behaviour for credit migrations and, therefore, has severe implications for the time-varying or time-invariant transition matrices typically used by investors and financial institutions alike for risk management purposes. Early studies on the non-Markovian nature of rating migrations include Kavvathas (2000) and Lando and Skødeberg (2002); both studies confirm the existence of positive correlation between successive migrations of the same type. In fact, Güttler and Raupach (2010) show that ignoring the rating momentum effect leads to a significant underestimation of Value-at-Risk (VaR) for credit portfolios. Furthermore, Güttler and Wahrenburg (2007) show that the momentum effect is shared

among rating agencies; a rating change from a single rating agency typically leads to a similar rating action from the other agencies.

1.1.2 Objective

Despite the extensive literature in credit risk modelling, there are still aspects that have not been thoroughly explored. This is especially true for the macroeconomic determinants and the cyclical dynamics of credit risk measures. The thesis focuses on three topics of particular importance: forecasting of the systematic movements in corporate defaults, the effect of structurally identified macro-financial shocks on credit risk measures and the dynamic properties of the momentum effect in credit rating migrations. The first two topics provide a macroeconomic view of credit risk, focusing on the times series aspect of the respective risk measures. In contrast, the third topic uses time-to-event continuous time methods to explore the micro-structure of rating migrations and defaults.

The first question that I ask in the thesis is how well can systematic movements in corporate bond defaults be forecasted by the current state of the economic cycle. The dependence of defaults on the economic and financial environment is not straightforward, as the corporate balance sheet can be affected by multiple macro-financial transmission channels. To avoid the problem of omitted variables from the econometric specification, I introduce a large scale dynamic factor model that is able to summarise the macro-financial dynamics into a small set of factors; the factors are then used to predict defaults. The studies of Duffie et al. (2007), Koopman et al. (2011), Duan et al. (2012) and Figlewski et al. (2012) also deal with the prediction of corporate failings. Nevertheless, only the work of Koopman et al. (2011) focuses on the systematic movements of corporate defaults and their correlation to the macroeconomic environment, with the remaining studies modelling the cross-section. Despite the similarities with the work of Koopman et al. (2011), I approach the question from a different perspective. While extraction of factors based on model a-theoretic principal component analysis efficiently summarises the variability from large panels of data, it does not necessarily give an intuitive interpretation for the factors. I argue that a set of identification restrictions based on an appropriate segmentation of the input variable panel leads to significant gains in corporate bond default forecasting, while providing an

intuitive decomposition high dimensional variable set into the different aspects of the macro-financial environment. In that sense, the approach is more similar to the work of Creal et al. (2013), but with a much higher dimensionality of the macro-financial input set and a particular focus on default forecasting.

The second question that I ask in the thesis is how corporate bond defaults and recoveries are affected by real economy and monetary policy shocks. The dependence of credit risk measures on the macroeconomic environment has attracted a lot of attention over the recent years. McNeil and Wendin (2007), Koopman et al. (2011), Koopman et al. (2012) and Creal et al. (2013) provide comprehensive dynamic factor structures that combine macroeconomic fluctuations with credit specific unobserved components. Concerning recoveries, the studies of Hu and Perraudin (2002), Altman et al. (2005), Acharya et al. (2007) and Jankowitsch et al. (2014) provide empirical evidence that the recovery amount from defaulted debt is linked to the default probability of the issuer and the aggregate default frequency within the issuer's industry sector. Despite the depth of research on the systematic link between credit measures and economic environment, no existing study attempts to explore the effects of structurally identified macro-financial shocks on corporate defaults and recoveries. This decomposition of the credit side of the economy to properly identified macro-financial shocks is of particular importance to regulators and policy makers alike. Stress testing exercises have become the norm in assessing the ability of the banking system to withstand extreme shocks. To formulate a coherent and plausible set of macro-financial indicators that quantifies the narrative of a specific scenario, it is much easier to work with properly identified fundamental shocks closely tied to economic theory. Furthermore, policy makers need to fully capture the different transmission mechanisms of the policy instrument. A change in money supply or the base rate affects corporates and consumers alike, but the impact is not straightforward. At a first stage, monetary policy shocks affect the assets and liabilities of corporate and household balance sheets through changes in demand and cost of debt. At a second stage, the deterioration or improvement of corporate and consumer balance sheets could have an impact on the conditions of further lending and that might in turn amplify the effect of a monetary policy shock.

The third and final question that the thesis addresses is what are the dynamic determinants of the momentum effect in credit ratings. The momentum effect in rat-

ing migrations is well documented. The studies of Kavvathas (2000) and Lando and Skødeberg (2002) use continuous time rating migration frameworks and report an overall positive directional dependence of each rating change to the type of the previous migration. Furthermore, Güttler and Raupach (2010) argue that ignoring the rating momentum effect leads to an underestimation of VaR by 107 basis points on average. Despite the existing work on the non-Markovian behaviour of rating transition matrices, no study has assessed the stability of the rating momentum effect through time. Therefore, the third question of the thesis focuses on how the momentum effect changes with the business cycle and the time since rating assignment. Once the stability of the rating momentum effect is statistically assessed, I quantify the impact of the various non-Markovian effects on credit portfolios.

1.2 Contributions

1.2.1 Large scale dynamic factor model to summarize the macro-financial environment

To quantify the forecasting power of the macroeconomic environment in predicting corporate bond defaults, the thesis uses a large scale dynamic factor model to reduce the dimensionality of the macro-financial dataset. While the use of dynamic factor modelling is very common in econometric studies, I move away from the existing literature by introducing a set of over-identifying restrictions applied on the matrix of factor loadings; that allows more robust isolation of the different dynamics and better interpretability of the extracted factors. To extract the factors, typically dynamic factor models use direct likelihood methods with suitable just-identifying restrictions on the loadings matrix or principal components, with possible Cholesky based rotations as in Boivin et al. (2013). I argue that since the set of macro-financial variables contains multiple series reflecting similar information (as for example the different interest rate variables) the zero restrictions on the loading matrix should be applied to the entire subset of economically similar inputs. By doing so, the factors more closely reflect the different aspects of the economy.

Based on the over-identification scheme described above, I extend the existing

literature by presenting a novel decomposition of the macroeconomic environment into 6 factors, reflecting business cycle fluctuations, cost of debt, yield curve risk, credit premia, asset returns and asset volatility. I provide detailed forecast error variance decomposition and correlation analysis to show that the proposed decomposition leads to intuitive factor dynamics.

1.2.2 Forecasting corporate bond defaults

The thesis provides new empirical evidence on how the macroeconomic environment helps to forecast corporate defaults. To assess the forecasting power of the macroeconomic environment, I combine the set of dynamic factors summarising the macro-financial conditions with a set of unobserved, default-specific, factors. The inclusion of macroeconomic and credit-specific factors in forecasting corporate defaults is not new, as numerous studies have explored the benefits of dynamic factor analysis in credit risk modelling. The difference with the existing literature is that I force the factors to have a clear interpretation based on a set of over-identifying restrictions. By doing so, the variability in default rates can be attributed to clear macro-financial dynamics as opposed to arbitrary rotations of factors based on principal components or recursive short-run identification restrictions that might be difficult to justify.

Furthermore, I base the estimation of the different non-Gaussian default models on a recently introduced Markov Chain Monte Carlo (MCMC) scheme. The MCMC algorithm combines a set of Gibbs sampling steps to transform the non-Gaussian observations into equivalent conditionally Gaussian data points; after transforming the specification into a conditionally Gaussian state space model, efficient Kalman filtering and smoothing is performed to sample the latent credit risk factors, combined with Bayesian linear regression sampling for the regression parameters. The sampling scheme follows the work of Frühwirth-Schnatter and Frühwirth (2007) and Frühwirth-Schnatter et al. (2009) and this is the first time it is used in a credit risk context. Unlike the existing literature that estimates dynamic credit event models either via simulation based maximum likelihood or generic MCMC sampling schemes, I argue that the data-augmentation algorithm used in this thesis offers an efficient and flexible alternative that provides both point and entire posterior estimates for parameters and dynamic factors.

1.2.3 Identifying macro-financial shocks

To quantify the effects of macro-financial shocks on corporate defaults and recoveries, I introduce a novel identification scheme to isolate the various shocks. The identification scheme combines a set of short- and long-run restrictions on the dynamic multipliers of a small-scale five variable Vector Autoregression (VAR) model. The VAR model is inspired by the credit channel/financial accelerator literature and it models output, inflation, leverage, asset value and interest rate variables. The suggested VAR analysis identifies five orthogonal shocks as aggregate supply, aggregate demand, corporate balance sheet, asset price and monetary policy. While based on existing views on macroeconomic dynamics, the suggested VAR approach differs from alternative models suggested in the literature in a number of ways.

First, I propose a new solution for isolating credit shocks in the economy, based on using the ratio of outstanding corporate debt to corporate profits as a proxy for the corporate sectors financial strength. This measure provides a highly cyclical indicator of the corporate leverage and helps to pinpoint the effect of balance sheet shocks. At the same time it justifies the short-run restrictions on the responses to asset price and monetary policy shocks and addresses the simultaneity concern between credit and macro-financial shocks. Studies like Boivin et al. (2013) that use bond spreads to proxy the credit conditions, suffer from the problem of simultaneity between bond yields and equity prices, even at a monthly frequency.

Second, I extend and strengthen the money neutrality argument of Bjørnland and Leitemo (2009) to a much longer time series of quarterly data. I find the resulting long-run restrictions to be particularly effective in separating asset price and monetary policy shocks.

Finally, I follow a more holistic view and I make sure that the responses of all the endogenous variables to each identified shock are intuitive. Typically, the VAR literature focuses solely on the small subset of structural shocks that are strongly identified (shocks with many restrictions). I show that even the weakly identified, real economy, shocks in my analysis produce intuitive impulse response functions.

1.2.4 Responses of corporate bond defaults and recoveries to macro-financial shocks

The thesis provides the first study that isolates the effects of independent macro-financial shocks on credit risk without relying on purely recursive identification schemes. This focus on the structural aspect of credit risk conditions renders the results of the thesis particularly important to policy makers and regulatory authorities. Monetary policy actions are transmitted via multiple financial channels to the economy and a failure to identify the links to the corporate and financial sectors can lead to a fresh balance sheet and net worth deterioration especially in the post-2008 crisis environment. Furthermore, financial stability and regulation of financial institutions move towards regular and comprehensive stress testing exercises, aiming to assess the solvency of the banking sector and identify the need for additional capitalisation. To form extreme but plausible shocks to drive the stress scenarios, based on quantitative grounds, a structural view of the economy is needed. A structural approach would enable one to express forward looking views on the economy in terms of independent, fundamental sources of macroeconomic activity.

1.2.5 Duration dependence and business cycle effects in credit rating momentum effect

The thesis contains the first effort to assess the stability of the directional serial dependence that is observed in credit ratings. In doing so, I test for business cycle and duration dependence of the so-called momentum effect. More specifically, I provide statistical significance tests for:

- Time-invariant rating momentum: This is a test for the well-documented serial dependence that is assumed to be constant. Unlike previous studies, I also adjust the rating momentum estimates for industry heterogeneity, by stratifying the baseline hazard rate.
- Rating momentum adjusted for duration time: This is a test for the interaction of serial dependence and time since rating assignment. For greater flexibility, I use fractional polynomials to model the duration dependence.

- Rating momentum adjusted for business cycle movements: This is a test for the interaction of serial dependence and macro-financial conditions. To proxy the business cycle fluctuations, I use the Chicago Fed National Activity Index (CFNAI) index.

In addition to the credit risk implications of the various non-Markovian effects that are tested, the econometric specification used also corrects the Cox proportional hazard model for continuous time-to-event data in two very important ways. First, adjusting the momentum effect for duration time corrects the specification for the non-proportionality of the rating momentum effect. Ignoring this non-proportional effect leads to the violation of the Cox model's assumptions and therefore to biased inference. Second, the parameter estimates are adjusted for possibly different baseline migration intensities across industry sectors. If the time profiles of the baseline hazard rates are fundamentally different across industries, failing to appropriately adjust the parameter estimates for the covariates leads to biased inference, see Kalbfleisch and Prentice (2002).

1.2.6 Impact of non-Markovian transition matrices on portfolio credit risk

Lending institutions and individual investors alike rely on published transition matrices by the major rating agencies to assess the credit quality of a debt instrument. The modelling of the transition matrices is typically based on the Markovian assumption. When additional effects such as serial dependence and duration time are taken into account, the Markov assumption is violated. Therefore, the final contribution of the thesis is to quantify the impact of the aforementioned non-Markovian transition matrix elements on credit portfolios.

To quantify the impact of the non-Markovian behaviour of the transition matrices, I examine a few key metrics, particularly focusing on the implications for financial institutions risk management. For a financial institutions banking book, I use the Basel II/III Risk Weighted Assets (RWA) prescribed formula to capture the impact rating momentum with duration and business cycle adjustments on regulatory capital. For provisions and economic capital I use jump-to-default expected and 99% tail losses.

Similarly, for the trading book I also cover expected and tail marked-to-market losses; for the trading book losses arise from a combination of credit spread movements and jump-to-default events. All loss estimates are based on two corporate bond portfolio structures; the first reflects primarily sub-investment grade exposures, while the second covers predominately investment grade risk exposures.

To quantify the impact of the non-Markovian behaviour of the transition matrices, I examine a few key metrics, particularly focusing on the implications for financial institutions risk management. The portfolio metrics cover actual and Marked-to-Market (MtM) losses, as well as regulatory capital measures. Adding to the existing literature on corporate credit ratings, I show that the presence of serial dependence on rating migrations leads to a substantial increase in actual and MtM losses, particularly in times of stress; this increase depends on vintage of the existing portfolio as rating momentum effect is not constant across a firm's survival time. Furthermore, capital requirements can be underestimated in times of stress, mainly for high credit quality portfolios.

1.3 Thesis Outline

The structure of the thesis mirrors the three questions that I am set to address. Therefore, the remainder of this thesis comprises of three chapters, structured to address each of the thesis' three main questions independently. The overall thesis follows an essay-type approach; each chapter is self-contained, with its own literature review, methodology, data and results.

More specifically, Chapter 2 uses a large scale dynamic factor model to quantify the forecasting power of the macro-financial environment in predicting corporate bond default rates. Section 2.2 provides a brief summary of the related literature on dynamic factor models and corporate default modelling. Section 2.3 explains the econometric model used to for the analysis: Section 2.3.1 covers the corporate default model while Section 2.3.2 is devoted to the macroeconomic dynamic factor model. Section 2.4 describes the Bayesian estimation procedure for both credit and macroeconomic models. Section 2.5 summarises the data used in the empirical analysis, while the Appendix provides the full set of macro-financial variables used in the analysis. Finally, Section 2.6 presents the results and Section 2.7 provides a summary of the chapter's main points

and some areas for future research.

Chapter 3 deals with the identification of macroeconomic shocks and the quantification of their effect on corporate bond defaults and recoveries. Sections 3.3.1 and 3.3.2 derive the econometric equations for corporate defaults and recoveries from well-established firm level structural default models. Section 3.3.3 covers the VAR model used for the macroeconomic part of the model and describes in detail the semi-structural identification scheme used to isolate the different macro-financial shocks. Section 3.3.4 describes the Bayesian estimation for defaults, recoveries and the macroeconomic VAR model, while Section 3.4 summarises the data used in the analysis. The results of the analysis are provided in Section 3.5 and comprise of the full set of impulse responses and variance decompositions for the macroeconomic (Section 3.5.1) and credit models (Section 3.5.2). Finally, Section 3.6 concludes the chapter and discusses possible extensions to the model used. Chapter 3 also contains an Appendix that contains the factor loadings for all the input macroeconomic variables.

Chapter 4 explores the properties of the momentum effect in credit ratings and shows how this is affected by duration and business cycle effects. Section 4.3.1 describes the continuous time model used to quantify and test the various effects in credit rating migrations, while Section 4.3.2 covers the maximum likelihood estimation. Section 4.4 summarises the data used in the analysis and depicts the effect of momentum on rating transition matrices. The empirical analysis is covered in Section 4.5. Section 4.5.1 provides the parameter estimates for the model and the output of the statistical tests that are used to determine the statistical significance of the various momentum effects. Section 4.5.2 quantifies the effect of the various momentum effects on credit portfolio losses, covering the Profit & Loss, MtM and Regulatory Capital aspects. Finally, Section 4.6 summarises the main findings of the chapter and highlights possible areas for future research.

Chapter 5 concludes. The chapter includes a summary of the findings presented in chapters 2-4, the implication of the results for credit risk management and policy making and recommends possible extensions and topics for future work.

The thesis also contains a technical appendix that contains a brief summary of Bayesian estimation methods and state space modelling. The topics covered in the

Appendix are applicable to both chapters 2 and 3.

Chapter 2

Forecasting Corporate Defaults in a Data-Rich Environment

2.1 Introduction

Following the recent economic downturn caused by the credit crisis, the systematic movement of default events has attracted additional attention, in both academic and practitioner circles. The systematic correlation of defaults during adverse economic conditions constitutes one of the major challenges in today's credit risk modelling research and various approaches have been proposed to address the default clustering due to, both, observed macroeconomic conditions and residual, credit cycle specific, stochastic movements.

Inspired by the strong link between macroeconomic environment and credit risk, I ask the question how well can systematic movements in corporate bond defaults be forecasted using the current state of the economic cycle. To capture the different aspects of the economic environment I use a large scale dynamic factor model, similar in nature to Koopman et al. (2011), and Boivin et al. (2013). I introduce a set of over-identifying restrictions to isolate the dynamic factors and represent business cycle fluctuations, yield curve movements (both level and slope), credit premia, and equity market conditions. The identification scheme is a significant contribution to the existing literature on dynamic factor modelling as I show that the extracted factors have a more

meaningful interpretation as compared to principal component analysis and short-run restrictions.

Using the extracted factors, I measure the forecasting power on corporate bond default rates across different industry sectors and forecasting horizons. Using a model with only Industrial Production as a base specification, I show that the dynamic factors can lead to substantial increases in forecasting performance for the majority of sectors; the performance gains can be as high as 90% in terms of Root Mean Squared Error (RMSE) and 80% in terms of Mean Absolute Error (MAE). The results broadly confirm the findings of Koopman et al. (2011) that use instead a principle component based approach. Nevertheless, unlike the work of Koopman et al. (2011), I assess the forecasting performance over multiple forecasting horizons. I find that the gains in predicting performance become less clear the lengthier the forecasting horizon.

Estimation is performed by means of a recently introduced MCMC sampling scheme based on data augmentation that expresses the non-Gaussian default events as conditionally Gaussian observations. Working with conditionally Gaussian observations allows the use of efficient Kalman filter based sampling for the various dynamic effects. The MCMC scheme follows the work of Frühwirth-Schnatter et al. (2009), and this is the first time it is used in a credit risk modelling context. I show that the data augmentation MCMC scheme is an efficient and flexible alternative to simulated maximum likelihood methods typically employed in the existing literature.

The remainder of this chapter is structured as follows. Section 2.3 provides the theoretical background and the econometric specification for both the macroeconomic dynamic factor model and the non-Gaussian stochastic default model. Section 2.4 deals with the estimation of both aspects of the combined model, describing MCMC estimation of the macroeconomic factor model, and the stochastic default model. Section 2.5 summarises the data used for the analysis, namely the macroeconomic variables sourced from the Federal Reserve Economic Database and the corporate bond default data sourced from Moodys Default & Recovery Database. Finally, section 2.6 presents the final estimation and forecasting results, while section 2.7 concludes.

2.2 Relevant Literature

The aggregated sector or economy wide default rates can be affected by both firm specific characteristics and systematic factors. By assuming the idiosyncratic risk for a homogeneous and relatively large portfolio of companies can be perfectly diversified, one can argue that default cycles are mainly affected by systematic observable (macroeconomic variables) or unobservable (latent) factors. The link between default cycles and systematic factors has been examined in many papers during the last decade, see Nickell et al. (2000), Bangia et al. (2002), Pesaran et al. (2006), Das et al. (2007), McNeil and Wendin (2007), Duffie et al. (2009), Figlewski et al. (2012). The results show that, although there is an interaction between observable real economy variables (such as output and interest rates) and credit cycles, the link is not straightforward and most of the variability in credit quality can be attributed to non observable factors. In fact, Koopman et al. (2012) report that the macroeconomic environment accounts for approximately 30% of the systematic variation in corporate defaults. Nevertheless, Koopman et al. (2011) find that aggregating the economic conditions into a small set of dynamic factors increases the prediction accuracy of corporate bond defaults.

Following the work of Koopman et al. (2011), and Creal et al. (2013), I use dynamic factor modelling to summarise the co-movements in the macroeconomic and financial variables. Early studies on dynamic factor models estimate systems of low dimensionality via maximum likelihood; see for example Engle and Watson (1981), Stock and Watson (1989), Sargent (1989), and Stock and Watson (1991). Straight maximisation of the likelihood function using the Kalman filter can be computationally infeasible for high dimensional systems. Bräuning and Koopman (2014), and Jungbacker and Koopman (2014) introduce more efficient estimation schemes to reduce the computational burden when the number of observable series is large. Alternatively, the work of Stock and Watson (2002a,b) on diffusion indices highlights the computational gains of using principal components to extract dynamic factors from large panels of macroeconomic variables, avoiding likelihood maximisation via the Kalman filter recursions. Despite the computational efficiency, using principal components for dynamic factor analysis does not necessarily lead to meaningful interpretations of the extracted factors. The semi-structural schemes described in Stock and Watson (2005) could be used to rotate the principal components but they are typically overly restrictive to allow for generic

identification of the factors. To overcome the limitations of the principal component analysis and avoid the computational burden typically associated with the maximum likelihood estimation, I use Bayesian MCMC techniques to estimate the macro-financial dynamic factor model. The MCMC scheme provides full posterior draws for the dynamic factors and links naturally to the estimation of the non-Gaussian specification for the corporate defaults.

The econometric model for the default events combines the dynamic factor analysis of the macroeconomic environment, with frailty serially correlated unobserved factors for the credit default cycle. The estimation is based on a recently introduced Gibbs sampling scheme, specifically tailored to non-Gaussian state space models. The MCMC algorithm transforms the non-Gaussian observation equation into an equivalent conditionally Gaussian model, for which Gibbs sampling is possible from the full conditionals of both parameters and hidden states. This so called auxiliary mixture sampling scheme, was first introduced by Shephard (1994) in the context of stochastic volatility models, and subsequently extended to models for count data by Frühwirth-Schnatter and Wagner (2006), and binary-binomial data by Frühwirth-Schnatter and Frühwirth (2007). These algorithms have been improved for efficiency by Frühwirth-Schnatter et al. (2009). Model selection based on marginal likelihood using auxiliary mixture models is explored in Frühwirth-Schnatter and Wagner (2008).

2.3 Econometric Framework

2.3.1 The Corporate Default Model

The econometric framework for a default process can be set either in a discrete or continuous time setting. Since the focus of this chapter is on economy-wide corporate defaults, with no firm-specific information used, I chose the discrete time setting. Continuous time to event models offer a greater level of granularity and predictive accuracy, as reported by Das et al. (2007), Koopman et al. (2008), and Koopman et al. (2009). Nevertheless, in the presence of frequent, monthly observations for the defaults and explanatory variables, the performance of discrete and continuous time models is not expected to diverge significantly (assuming that daily data provide a good proxy for

continuous time analysis).

I denote by D_{it}^h the default counts over h months for a cohort formed at month t with companies in sector $i \in \mathcal{S}$, where \mathcal{S} is the set of industry sectors used in the analysis, and $t = 1, \dots, T$. Conditioning on the available information set \mathcal{F}_{it}^h at time t , the default counts are assumed independent and binomially distributed:

$$D_{it}^h | \mathcal{F}_{it}^h \sim \text{Binom}(PD_{it}^h, N_{it}), \quad (2.1)$$

where N_{it} is the number of active companies in sector i at the beginning of month t , and PD_{it}^h is the sector specific probability of default over h months measured at month t . For the purpose of this study, the information set \mathcal{F}_{it}^h comprises of two sources of systematic correlation; a set of macroeconomic systemic factors \mathbf{F}_t^m , and a set of unobserved credit/default factors $\mathbf{F}_t^{d,h}$. The factors \mathbf{F}_t^m are derived as linear combinations of observed macroeconomic aggregates, while the factors $\mathbf{F}_t^{d,h}$ are unobserved and need to be estimated from the default data. The default specific factors $\mathbf{F}_t^{d,h}$ are designed to capture clustering in corporate defaults over and above what can be explained by observed covariates, and they are commonly referred to as frailty, see Koopman et al. (2011), and Koopman et al. (2012). Since the analysis is based on industry sector aggregate default information, I extract one frailty factor per sector and, therefore, the set of frailty factors can be written as $\mathbf{F}_t^{d,h} = \{f_{it}^{d,h} : i \in \mathcal{S}\}$.

Both observed and unobserved factors enter (2.1) via PD_{it}^h . Since this is a quantity bound to lie between 0 and 1, I use the logit function to make the link to the filtration process \mathcal{F}_{it}^h

$$\log\left(\frac{PD_{it}^h}{1 - PD_{it}^h}\right) = \alpha_{i,h} + \beta_{i,h}^m(L)\mathbf{F}_t^m + \beta_{i,h}^d f_{it}^{d,h}, \quad (2.2)$$

where $\alpha_{i,h}$ are the sector/forecast horizon specific intercepts, while $\beta_{i,h}^m(L)$ and $\beta_{i,h}^d$ are the sector/forecast horizon specific sensitivities to macroeconomic factors \mathbf{F}_t^m and industry frailty factor f_{it}^d respectively. The sensitivities $\beta_{i,h}^m(L)$ in (2.2) are allowed to be generic lag polynomials, therefore permitting defaults to be impacted by lagged values of the macroeconomic factors \mathbf{F}_t^m .

For the stochastic process of the frailty factors $\mathbf{F}_t^{d,h}$, I assume a 0 mean, p -th order

VAR of the form:

$$\mathbf{F}_t^{d,h} = \Phi_1^{d,h} \mathbf{F}_{t-1}^{d,h} + \dots + \Phi_p^{d,h} \mathbf{F}_{t-p}^{d,h} + \epsilon_t \Rightarrow \tilde{\mathbf{F}}_t^{d,h} = \Phi^{d,h} \tilde{\mathbf{F}}_{t-1}^{d,h} + \epsilon_t, \quad \epsilon_t \sim N(0, I_S), \quad (2.3)$$

where

$$\Phi^{d,h} = \begin{bmatrix} \Phi_1^{d,h} & \Phi_2^{d,h} & \dots & \Phi_{p-1}^{d,h} & \Phi_p^{d,h} \\ I_S & 0 & \dots & 0 & 0 \\ 0 & I_S & & 0 & 0 \\ 0 & 0 & \ddots & 0 & 0 \\ 0 & 0 & \dots & I_S & 0 \end{bmatrix}, \quad \tilde{\mathbf{F}}_t^{d,h} = \begin{bmatrix} \mathbf{F}_t^{d,h} \\ \vdots \\ \mathbf{F}_{t-p+1}^{d,h} \end{bmatrix}.$$

S is the number of industry sectors, and the variance-covariance matrix of the residuals is set to the identity matrix for identification purposes; if the variance-covariance is left unrestricted then it cannot be jointly identified with the sensitivities $\beta_{i,h}^d$ in (2.2). I assume that the VAR(p) is stable, or equivalently:

$$\det \left(I_S - \Phi_1^{d,h} z - \dots - \Phi_p^{d,h} z^p \right) \neq 0, \quad \text{for } |z| \leq 1,$$

which is satisfied if all the eigenvalues of the matrix $\Phi^{d,h}$ have modulus less than 1. The stable VAR(p) process implies a long-run variance covariance $P^{d,h}$ matrix for the frailty dynamics given by:

$$\text{vec}(P^{d,h}) = \left(I_{(Sp)^2} - \Phi^{d,h} \otimes \Phi^{d,h} \right)^{-1} \text{vec}(I_S), \quad (2.4)$$

where $\text{vec}(\cdot)$ is the vectorisation operator, and \otimes is the Kronecker product. The binomial model in (2.1)-(2.2), and the frailty dynamics described in (2.3)-(2.4) are an extension of the GLMM that McNeil and Wendin (2007) use in their analysis; while McNeil and Wendin (2007) restrict their analysis to independent AR(1) processes, I allow for more generic VAR dynamics for the unobservable factors.

2.3.2 Dynamic Factor Model

To define the effect of the economic environment on corporate defaults at any point in time, a wide variety of macroeconomic and financial variables can be used. Nevertheless, for any economic or financial concept there is no unique and indisputable

data series that can be used for all corporates in a given segmentation. Therefore, selecting individual variables to use in an econometric specification based on statistical significance and/or theoretical arguments, does not necessary capture the precise measure that agents use in their decision making process. Furthermore, variables not included in the econometric model may potentially bias forecasts and/or structural results (such as impulse responses in a systematic analysis), leading to the problem of omitted variables/selection bias.

Recent studies propose factor based models as a solution to the above mentioned problems. This type of high dimensional model is based on summarising the information contained in large panels of macroeconomic and financial variables into a small set of factors to be used for inference. The dimension reduction is typically based on principle components. For a high level overview of the dynamic factor literature, see Stock and Watson (2010). In the monetary policy literature, Bernanke et al. (2005) combine in a VAR model, the FED Funds Rate and a set of dynamic factors derived from a high dimensional panel of macro-financial variables. This so-called Factor Augmented Vector Autoregression (FAVAR) setup is able to better capture the complicated transmission channels of monetary policy in the US. More recently, Boivin et al. (2013) are interested in the effects of credit shocks in the economy. To identify the different shocks, they use a recursive scheme applied on a set of dynamic factors. The factors are derived from a principle component decomposition of a large macroeconomic panel of US data. Koopman et al. (2011) use a large scale dynamic factor model to forecast corporate bond default rates.

Instead of using principal component-based dynamic factor analysis and rely on asymptotic results, I follow Creal et al. (2013) and cast the factor model in its state space form and use likelihood methods for inference. This allows the use of efficient Kalman filter-based algorithms for factor extraction. To summarise the systematic effects across a big dataset of observed macroeconomic and financial variables, I assume that the economy adheres to a dynamic factor model, having the static form:

$$\mathbf{X}_t = \mathbf{\Lambda}\mathbf{F}_t^m + \mathbf{e}_t, \quad \mathbf{e}_t \sim N(0, \Sigma_e) \quad (2.5)$$

$$\mathbf{F}_t^m = \mathbf{\Phi}^m(L)\mathbf{F}_{t-1}^m + \boldsymbol{\eta}_t, \quad \boldsymbol{\eta}_t \sim N(0, I), \quad (2.6)$$

where \mathbf{X}_t is the vector of observable macroeconomic and financial variables at time t , and \mathbf{F}_t^m is the vector of dynamic factors. The linear dependence of the observed vector \mathbf{X}_t on the unobserved factors \mathbf{F}_t^m as described by (2.5) corresponds to the *observation equation* of the state space form. The stochastic evolution of the factor as described by (2.6) corresponds to the *state equation*. $\Phi^m(L)$ is a generic lag polynomial transition matrix that allows a VAR process of arbitrary order for the evolution of the dynamic factors. In a similar way to (2.3), the lag polynomial $\Phi^m(L)$ can be re-expressed as

$$\Phi^m = \begin{bmatrix} \Phi_1^m & \Phi_2^m & \dots & \Phi_{p-1}^m & \Phi_p^m \\ I_K & 0 & \dots & 0 & 0 \\ 0 & I_K & & 0 & 0 \\ 0 & 0 & \ddots & 0 & 0 \\ 0 & 0 & \dots & I_K & 0 \end{bmatrix},$$

where K is the number of dynamic factors and p is the order of the VAR process. The observation equation errors, \mathbf{e}_t , are assumed uncorrelated with the state equation errors, $\boldsymbol{\eta}_t$. The variance-covariance matrix for the error terms of the observation equation, Σ_e , is assumed diagonal, so that all the cross-correlations between the variables in \mathbf{X}_t are captured by the factors \mathbf{F}_t^m . The variance-covariance matrix for the state equation errors $\boldsymbol{\eta}_t$ is assumed to be the identity matrix for identification purposes; the matrix of factor loadings $\mathbf{\Lambda}$ and the variance of the dynamic factors cannot be jointly identified.

For simplicity, the variance-covariance matrix Σ_e in (2.5) is assumed to be diagonal. This assumption essentially forces the dynamic factors to capture the full cross-sectional correlation among the macroeconomic variables. In practice, the strong assumptions of uncorrelated residuals in both time and cross-section is too difficult to be met empirically. While Σ_e can be extended to be a full matrix and the residuals \mathbf{e}_t can be allowed to be autocorrelated by augmenting appropriately the state vector in (2.6), the sheer size of the set of input variables used in the analysis quickly renders these extensions computationally unfeasible. Since the purpose of the analysis is not to forecast each macroeconomic variable but to get estimates of the dynamic factors \mathbf{F}_t^m in (2.6), the empirical rejection of the assumptions for the behaviour of \mathbf{e}_t is less important due to the large cross-section and the length of the time series used. Furthermore, failing to properly account for the residual autocorrelation structure of \mathbf{e}_t can

only create problems when estimating factors in the presence of strongly trending variables. Nevertheless, introducing the identification assumptions below, I properly isolate the strongly trending variables in the sample (mainly interest rates) and therefore the dynamic factor estimates are not biased.

To forecast corporate bond defaults, I assume a-priori that the macro-economic conditions in a given economy can be grouped into business cycle effects, cost of debt effects, yield curve risk, credit conditions, asset value movements and asset volatility. Therefore the choice of variables to be included in the analysis is geared towards the above mentioned structure of the macro-financial environment. While the choice of variables is not exhaustive (variables that capture risk from the international environment could also be included), it covers a very wide spectrum of economic metrics. Based on the a-priori assumption on the structure of the economy and the corresponding choice of variables, I summarise the entire set of macro-financial variables into 6 dynamic factors. For a more formal selection process based on statistical theory, the method of Bai and Ng (2002) can be used to determine the number of factors in generic dynamic factor models. Nevertheless, it is less suited when semi-structural restrictions are used.

The high-level factor structure is designed to separate business cycles from financial shocks. Having removed the business cycle component from the financial variables, I then choose two separate sets of factors to approximate determinants of the assets and liabilities of a typical corporate balance sheet. The first set of factors captures the drivers of the cost of debt, while the second set of factors captures asset value movements and volatility. The 6 factors \mathbf{F}_t^m correspond to:

- *Business Cycle Factor* (F_t^{bs}): A factor that affects both real and financial sides of the economy. Essentially, to isolate the business cycle factor, I assume that financial shocks do not have a contemporaneous effect on the real economy. This is a typical assumption in both dynamic factor and VAR literature, see Bjørnland and Leitemo (2009) and Boivin et al. (2013). It essentially orders all financial shocks after the real economy variables. Following the terminology of Bernanke et al. (2005), the real economy variables are treated as slow-moving, while the financial variables are treated as fast-moving. Business cycle movements are very important for default prediction as they affect corporate profitability, which in

turn affects retained earnings and therefore the future corporate net worth. This is particularly true for the highly cyclical sectors, such as Capital and Consumer Industries.

- *Cost of Debt Factor* (F_t^{cd}): This factor loads exclusively on interest rate variables, including yield curve and credit spreads. The factor is designed to capture generic changes to the corporate funding cost, and proxies parallel shifts in the government and corporate yield curve. As changes in the cost of borrowing affect the investment decisions of corporates, an increase in the cost of debt is expected to have a positive impact on corporate defaults.
- *Yield Curve Risk Factor* (F_t^{yc}): A factor related to the term spread risk, loading on the yield curve and credit spread variables. A typical new corporate bond issuance involves maturities of more than 10 years and therefore the slope of the yield curve is very important. To separate the slope from the level of the yield curve, I assume that interest rate variables are not affected contemporaneously by yield curve slope and credit spread shocks. This assumption implies that yield curve slope and credit spread shocks are ordered after yield curve level shocks. The order of yield curve level and slope shocks can be reversed without materially changing the extracted factors. Yield curve level and slope changes typically exhibit low correlation, a structure that forms the base for the principle components and Nelson-Siegel decomposition of the yield curve, see Diebold and Li (2006) yield curve and the references therein. The directional effect of the yield curve slope on corporate defaults is not clear. On one hand, when the short end of the yield curve becomes more expensive this is typically regarded as a precursor to a recession, and therefore it should be positively correlated with future default rates. On the other hand, as corporate debt is typically issued on a medium to long term basis, it should be the relative changes of the longer end of the yield curve that are positively correlated with default rates.
- *Credit Risk Factor* (F_t^{cr}): This factor loads exclusively on credit spread variables and reflects shocks that are caused by deterioration in credit market conditions. Credit shocks have significant effects in economic activity, as shown by Gilchrist et al. (2009), Gilchrist and Zakrajšek (2012) and Boivin et al. (2013), and historically they are highly correlated with default events. By construction, I define credit

shocks as the incremental fluctuations in credit spreads after removing business cycle, interest rate level and yield curve slope movements. This decomposition is consistent with the work on the determinants of credit spreads by Collin-Dufresne et al. (2001). Even though the credit spreads are ordered after the yield curve spreads, the order can be reversed without affecting the resulting factors, as the correlation between yield curve slope and credit spreads is minimal. As credit spreads are positively correlated with corporate borrowing costs, an increase in F_t^{cr} should lead to an increase in default rates.

- *Asset Return Factor* (F_t^{ar}): This is an independent factor loading exclusively on equity return and equity volatility time series. As the equity market related variables are also allowed to depend on overall factor F_t^{bs} , the factor captures the fluctuations in stock prices over and above the business cycle. Structural credit models based on the work of Merton (1974) define default as the point the asset value of a corporate falls below its level of liabilities. This definition of default can be applied to public corporates to infer the unobserved asset values as the residual of the debt level and the market capitalisation, therefore making explicit the importance of equity price movements when studying credit risk. Other things equal, a drop in F_t^{ar} should lead to a drop in the asset value of corporates, pushing them closer to their default point, and therefore increasing their likelihood to default.
- *Asset Volatility Factor* (F_t^{av}): This is an additive adjustment to the factor F_t^{ar} to capture the volatility in equity markets. Therefore it only loads on equity volatility related time series. Again following the structural approach to credit risk, the asset value and the level of liabilities are not enough to characterise the probability to default. The remaining driving factor is the volatility of the asset value. In absence of observed asset values to base the calculation of the volatility on, I use equity returns and implied volatility series based on equity options. As the equity volatility series are allowed to depend on both F_t^{ar} and F_t^{av} factors, this structure implies that the asset volatility shocks are ordered after the asset return shocks. This choice is extremely robust to re-ordering and assuming that equity returns lead the equity volatility fluctuations leads to very similar factors. Other things equal, an asset volatility increase increases the likelihood of the asset

value process crossing the default point, and therefore the F_t^{av} factor should show positive correlation to the observed default rates.

The above factor definitions introduce a set of restrictions that structurally identify the model. Assuming that the input variables can be classified into Business Cycle (with corresponding loadings denoted by Λ_{bs}), Interest Rate (with corresponding loadings denoted by Λ_{ir}), Yield Curve Spreads (with corresponding loadings denoted by Λ_{ys}), Credit Spreads (with corresponding loadings denoted by Λ_{cs}), Equity Price Returns (with corresponding loadings denoted by Λ_{er}), and Equity Price Volatility (with corresponding loadings denoted by Λ_{ev}), then the factor loadings matrix can be expressed as:

$$\Lambda = \begin{bmatrix} \Lambda_{bs}^{F_t^{bs}} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} \\ \Lambda_{ir}^{F_t^{bs}} & \Lambda_{ir}^{F_t^{cd}} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} \\ \Lambda_{ys}^{F_t^{bs}} & \Lambda_{ir}^{F_t^{cd}} & \Lambda_{ir}^{F_t^{yc}} & \mathbf{0} & \mathbf{0} & \mathbf{0} \\ \Lambda_{cs}^{F_t^{bs}} & \Lambda_{cs}^{F_t^{cd}} & \Lambda_{cs}^{F_t^{yc}} & \Lambda_{cs}^{F_t^{cr}} & \mathbf{0} & \mathbf{0} \\ \Lambda_{er}^{F_t^{bs}} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \Lambda_{er}^{F_t^{ar}} & \mathbf{0} \\ \Lambda_{ev}^{F_t^{bs}} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \Lambda_{ev}^{F_t^{ar}} & \Lambda_{ev}^{F_t^{av}} \end{bmatrix}, \quad (2.7)$$

where $\Lambda_x^{F_t^*}$ denotes the vector of sensitivities of variable group x to the dynamic factor F_t^* .

The set of restrictions in (2.7) over-identifies the model. Under the generic factor model summarised by (2.5) and (2.6), if K is the number of factors, a total of K^2 non-redundant restrictions need to be imposed for the model to be identified. Forcing the state errors $\boldsymbol{\eta}_t$ to be uncorrelated with unit variance imposes $K(K+1)/2$ restrictions. For more details on dynamic factor model identification, see Bai and Ng (2013). Using a system with 6 dynamic factors, the variance restrictions impose 21 restrictions and for a just-identified system a total of 36 restrictions are needed. Clearly (2.7) imposes a much higher number of 0 restrictions than the 15 needed to just-identify the system, since each block of $\Lambda_x^{F_t^*}$ represents multiple series. Nevertheless, it is this over-identification scheme that distinguishes the dynamic factor analysis presented here from the rest of the literature. As factors represent clear macroeconomic and financial concepts, variables of similar nature should share the same identification restrictions; it is very difficult for example to argue that Industrial Production does not contemporaneously respond to shocks in the Cost of Debt, while the Unemployment Rate does.

2.4 Estimation

The dynamic factor model of section 2.3.2 and the corporate default model of section 2.3.1 can be jointly estimated by casting equations (2.1)-(2.3) and (2.5)-(2.6) into a combined partially-Gaussian state space, along the lines of Creal et al. (2013). Nevertheless, the sampling of the non-Gaussian part of the state space model and the overall dimensionality of the observation vector renders this sampling approach computationally infeasible. Instead, I choose to estimate the Gaussian dynamic factor model of section 2.3.2 independently and substitute the posterior means of the dynamic factors $\hat{\mathbf{F}}_t^m$ in (2.1)-(2.3) to estimate the corporate default model (2.1)-(2.3) of section 2.3.1. This is a similar approach to that used by Koopman et al. (2011).

The dynamic factor model described in section 2.3.2 has a Gaussian state space representation and it can therefore be estimated with a combination of Kalman filtering/smoothing for the dynamic factors and closed form posterior Gibbs sampling steps for the parameters. For the *observation equation* (2.5) I assume diffuse priors for the parameters of the form

$$p(\mathbf{\Lambda}, \Sigma_e) = p(\mathbf{\Lambda})p(\Sigma_e) = \prod_{i=1}^N p(\mathbf{\Lambda}_i) \prod_{i=1}^N p(\sigma_{ei}^2), \quad (2.8)$$

$$p(\mathbf{\Lambda}_i) = \text{constant}, \quad p(\sigma_{ei}^2) \propto 1/\sigma_{ei}^2$$

where i denotes the i -th observed variable in the model, $\mathbf{\Lambda}_i$ denotes the i -th row of $\mathbf{\Lambda}$ and σ_{ei}^2 denotes the variance of the residuals for the i -th equation. For the specific form of the priors I take advantage of the diagonal structure of Σ_e and the absence of cross equation restrictions in $\mathbf{\Lambda}$. For the transition matrix $\mathbf{\Phi}^m(L)$ in the *state equation* (2.6) I also use diffuse priors of the form $p(\mathbf{\Phi}^m(L)) = \text{constant}$.

Both dynamic factor model and the non-Gaussian corporate default specification can be estimated via maximum likelihood methods, see Koopman and Lucas (2008), Koopman et al. (2011) and Creal et al. (2013). Nevertheless, the high dimensionality of the observed time series and the full VAR dynamics for both macroeconomic dynamic factors in (2.6) and frailty factors in (2.3) make full likelihood maximisation computationally infeasible. To circumvent the computational requirements of maximum likelihood estimation, I choose to estimate the econometric specifications using

MCMC techniques. MCMC methods offer significant advantages in the estimation of high dimensionality Gaussian state space models as they break down the necessary sampling into 3 steps:

1. Sample the states \mathbf{F}_t^m conditional on $\mathbf{\Lambda}$, Σ_e and $\mathbf{\Phi}^m(L)$. For a Gaussian state space this involves simulated draws based on a forward pass using the Kalman filter and a backwards pass using the Kalman smoother, as per Carter and Kohn (1994). Details of the simulation smoother algorithm can be found in Durbin and Koopman (2002) and section C.2 of the Appendix. The Kalman filter is initialised from the stationary distribution for the dynamic factor

$$\mathbf{F}_0^m \sim N(\mathbf{0}, P^m), \quad \text{vec}(P^m) = (I_{(Kp)^2} - \mathbf{\Phi}^m \otimes \mathbf{\Phi}^m)^{-1} \text{vec}(I_K).$$

2. Sample the transition matrix $\mathbf{\Phi}^m(L)$ conditional on the states \mathbf{F}_t^m . Due to $\boldsymbol{\eta}_t \sim N(0, I_K)$, sampling of $\mathbf{\Phi}^m(L)$ can be performed on an equation-by-equation basis as follows:

$$p(\mathbf{\Phi}_k^m | \cdot) \sim N(\hat{\boldsymbol{\phi}}_k^m, \hat{\boldsymbol{\Phi}}_k^m), \quad \hat{\boldsymbol{\Phi}}_k^m = [(\tilde{\mathbf{F}}^m)' \tilde{\mathbf{F}}^m]^{-1}, \quad \hat{\boldsymbol{\phi}}_k^m = \hat{\boldsymbol{\Phi}}_k^m [(\tilde{\mathbf{F}}^m)' \mathbf{F}^m], \quad (2.9)$$

where k refers to the k -th factor, $\mathbf{F}^m = [\mathbf{F}_1^m, \dots, \mathbf{F}_T^m]'$, $\tilde{\mathbf{F}}^m = [\tilde{\mathbf{F}}_1^m, \dots, \tilde{\mathbf{F}}_T^m]'$ and $\tilde{\mathbf{F}}_t^m = [\mathbf{F}_t^m, \dots, \mathbf{F}_{t-p+1}^m]'$. (2.9) is the result of the diffuse priors and the unit variance assumption for the residuals of the stochastic process. For the derivation of the linear regression posterior moments under diffuse priors see Zellner (1971).

3. Sample factor loadings $\mathbf{\Lambda}$ and the variance-covariance matrix for the residuals Σ_e . Since Σ_e is diagonal and there are no cross-equation restrictions, sampling can be performed on an equation-by-equation basis. For each variable i , under the assumption of diffuse priors of the form (2.8), random draws from the posteriors for the loadings $\mathbf{\Lambda}_i$ and variances σ_{ei}^2 are obtained as follows:

$$p(\mathbf{\Lambda}_i | \cdot) \sim N(\hat{\boldsymbol{\lambda}}_i, \hat{\boldsymbol{\Lambda}}_i), \quad \hat{\boldsymbol{\Lambda}}_i = [(\mathbf{F}^m)' \mathbf{F}^m]^{-1}, \quad \hat{\boldsymbol{\lambda}}_i = \hat{\boldsymbol{\Lambda}}_i [(\mathbf{F}^m)' \mathbf{F}^m],$$

$$\sigma_{ei}^2 \sim \text{invGamma} \left(0.5 \cdot T, 0.5 \cdot \sum_{t=1}^T (X_{it} - \mathbf{\Lambda}_i \mathbf{F}_t^m)^2 \right), \quad (2.10)$$

where \mathbf{F}^m is defined in (2.9).

For the estimation of the non-Gaussian state space model in (2.1)-(2.3) I use a recently proposed Gibbs sampling scheme that has not been used before in a credit risk context. The sampling scheme is based on augmenting the binomial data with the latent variable D_{it}^{h*} , conditioning upon which, the observation equation (2.1) can be re-expressed as Gaussian. For the conditionally Gaussian system, closed form full conditional distributions exist for both states and parameters. Following Frühwirth-Schnatter and Frühwirth (2007) and Frühwirth-Schnatter et al. (2009), I augment the data with the latent variable D_{it}^{h*} as follows:

$$D_{it}^{h*} = Z_{it}^h + \varepsilon_{it}^{h*}, \quad \varepsilon_{it}^{h*} \sim \text{Logistic}(0, 1), \quad (2.11)$$

where $Z_{it}^h = \alpha_{i,h} + \beta_{i,h}^m(L) \mathbf{F}_t^m + \beta_{i,h}^d f_{it}^{d,h}$ and the error term's logistic distribution can be re-expressed as a negative log-Gamma, $\varepsilon_{it}^{h*} = -\log \xi_{it}^{h*}$, with $\xi_{it}^{h*} \sim \text{Gamma}(N_{it}, 1)$. Each ε_{it}^{h*} is approximated by a finite mixture of normal distributions. The mixture indicators r_{it}^{h*} are introduced as a second layer of data augmentation. The number of Normal densities, $R_{it}^*(N_{it})$, mixture weights, $w_{r_{it}^{h*}}(N_{it})$, means, $\mu_{r_{it}^{h*}}(N_{it})$, and variances $\sigma_{r_{it}^{h*}}^2(N_{it})$, are dependent on the total number of firms N_{it} in each sector i at quarter t . Conditional on r_{it}^{h*} , (2.11) reduces to the following linear Gaussian model:

$$D_{it}^{h*} = Z_{it}^h + \mu_{r_{it}^{h*}}(N_{it}) + \varepsilon_{it}^{h*}, \quad \varepsilon_{it}^{h*} | r_{it}^{h*} \sim N(0, \sigma_{r_{it}^{h*}}^2(N_{it})). \quad (2.12)$$

All mixture quantities are obtained by minimising the Kullback-Leibler divergence between the mixture approximation and the original negative log Gamma density. For more details, see Frühwirth-Schnatter et al. (2009).

Having defined the non-Gaussian observation equation in (2.1)-(2.2) as the equivalent conditionally Gaussian equation (2.12), and given starting values for D_{it}^{h*} and r_{it}^{h*} , the Gibbs sampling scheme consists of iterative draws from the full conditionals and can be summarised in the following steps:

1. Sample D_{it}^{h*} and r_{it}^{h*} conditional on $\alpha_{i,h}$, $\beta_{i,h}^m(L)$, $\beta_{i,h}^d$, $f_{it}^{d,h}$ and D_{it}^h , using the following steps for $t = 1, \dots, T$

(a) Sample D_{it}^{h*} conditional on Z_{it}^h and D_{it}^h as:

$$D_{it}^{h*} = -\log \left(\frac{U_{it}}{1 + e^{Z_{it}^h}} + \frac{V_{it}^h}{e^{Z_{it}^h}} \right)$$

where $U_{it} \sim \text{Gamma}(N_{it}, 1)$ and $V_{it}^d \sim \text{Gamma}(N_{it} - D_{it}^h, 1)$ are sampled independently.

(b) Sample the component indicators r_{it}^{h*} from the following discrete distribution:

$$Pr(r_{it}^{h*} = \mathbf{r}_i^* | D_{it}^{h*}, D_{it}^h, Z_{it}^h) \propto \frac{w_{\mathbf{r}^*}^*(N_{it})}{\sigma_{\mathbf{r}^*}^*(N_{it})} e^{-0.5 \left(\frac{D_{it}^{h*} - Z_{it}^h - \mu_{\mathbf{r}^*}^*(N_{it})}{\sigma_{\mathbf{r}^*}^*(N_{it})} \right)^2}$$

where $\mathbf{r}^* = \{1, \dots, R_{it}^*(N_{it})\}$.

- Sample $\alpha_{i,h}$, $\beta_{i,h}^m(L)$ and $\beta_{i,h}^d$ conditional on D_{it}^{h*} , r_{it}^{h*} , \mathbf{F}_t^m and $f_t^{d,h}$. For all the model parameters, I assume normal priors of the form $p(\alpha_{i,h}) \sim N(\alpha_{i,h}^0, \sigma_{\alpha_{i,h}^0}^2)$, $p(\beta_{i,h}^m(L)) \sim N(\beta_{i,h}^{m,0}(L), \sigma_{\beta_{i,h}^{m,0}(L)}^2)$ and $p(\beta_{i,h}^d) \sim N(\beta_{i,h}^{d,0}, \sigma_{\beta_{i,h}^{d,0}}^2)$. This choice of priors leads to conjugate Bayesian linear regression posteriors of the form:

$$p(\mathbf{b} | \cdot) \sim N(\mathbf{b}_1, \mathbf{B}_1), \text{ where} \quad (2.13)$$

$$\mathbf{B}_1 = (\mathbf{B}_0^{-1} + \mathbf{X}'\mathbf{X}/\sigma^2)^{-1}, \text{ and } \mathbf{b}_1 = \mathbf{B}_1(\mathbf{B}_0^{-1}\mathbf{b}_0 + \mathbf{X}'\mathbf{Y}/\sigma^2).$$

In (2.13), $\mathbf{X} = \text{vec}([\mathbf{1}_T, \mathbf{L}\mathbf{F}^m, \mathbf{f}_i^{d,h}])$ is the vectorised version of the covariates, $\mathbf{L}\mathbf{F}^m$ is the specific lag structure of \mathbf{F}^m , $\mathbf{Y} = \text{vec}(D_{it}^{h*})$, $\mathbf{b}_0 = [\alpha_{i,h}^0, \beta_{i,h}^{m,0}(L), \beta_{i,h}^{d,0}]$ stacks all the prior means, while $\mathbf{B}_0 = \text{diag}([\sigma_{\alpha_{i,h}^0}^2, \sigma_{\beta_{i,h}^{m,0}(L)}^2, \sigma_{\beta_{i,h}^{d,0}}^2])$ is a diagonal matrix, having the prior variances as diagonal elements.

- Sample the latent factors $\mathbf{f}_t^{d,h}$ conditional on all D_{it}^{h*} , $\alpha_{i,h}$, $\beta_{i,h}^m(L)$ and $\beta_{i,h}^d$. Working with the set of adjusted conditionally Gaussian observations $\tilde{D}_{it}^{h*} = D_{it}^{h*} - \alpha_{i,h} - \beta_{i,h}^m(L)\mathbf{F}_t^m$, sampling of $\mathbf{f}_t^{d,h}$ is performed by means of the Forward Filtering-Backward Sampling multi-move sampling algorithm, see Carter and Kohn (1994). The multi-move sampling scheme essentially samples from the distributions of the error terms in (2.3) and (2.12) and then makes a forward pass using the Kalman filter and a backward pass using the Kalman smoother. The detailed simulation smoother algorithm is presented in Durbin and Koopman (2002). The Kalman filter is initialised from the stationary distribution for $\mathbf{f}_t^{d,h}$,

$N(\mathbf{0}, P^{d,h})$, where $P^{d,h}$ is given by (2.4).

4. Sample the transition matrix $\Phi^{d,h}(L)$ conditional on the frailty factors $F_t^{d,h}$. Similarly to the dynamic factor econometric specification, I assume diffuse priors of the form $p(\Phi^{d,h}(L)) = \text{constant}$. Due to $\epsilon_t \sim N(0, I_S)$, sampling of $\Phi^{d,h}(L)$ is performed on an equation-by-equation basis as follows:

$$p(\Phi_i^{d,h} | \cdot) \sim N(\hat{\phi}_i^{d,h}, \hat{\Phi}_i^{d,h}), \quad (2.14)$$

$$\hat{\Phi}_i^{d,h} = [(\tilde{F}^{d,h})' \tilde{F}^{d,h}]^{-1}, \quad \hat{\phi}_i^{d,h} = \hat{\Phi}_i^{d,h} [(\tilde{F}^{d,h})' \mathbf{f}_i^{d,h}],$$

for sector $i \in \mathcal{S}$, where $\mathbf{f}_i^{d,h} = [f_{i1}^{d,h}, \dots, f_{iT}^{d,h}]'$, $\tilde{F}^{d,h} = [\tilde{F}_1^{d,h}, \dots, \tilde{F}_T^{d,h}]'$ and $\tilde{F}_t^{d,h} = [F_t^{d,h}, \dots, F_{t-p+1}^{d,h}]'$. The posterior moments in (2.14) are similar to the posterior moments for the dynamic factors in (2.9), and they are the combination of the diffuse priors and the identity matrix assumption for the covariance of ϵ_t .

2.5 Data

For the macroeconomic factors F_t^m , I construct a large panel of observed time series by sourcing 103 variables from the Federal Reserve Economic Database (FRED)¹. All variables are sourced at a monthly frequency from January 1982 to August 2014. January 1982 is chosen so that it allows for potential lags in the default rate forecasting equations. August 2014 coincides with the last monthly cohort for the default rate analysis (see below for more details). For all the variables reported on a daily basis, average values over corresponding month are used. The variables can be grouped into 12 macroeconomic concepts: Production, Income, Consumption, Employment, Prices, Inventories & Orders, Housing Starts, Bank Lending, Interest Rates, Yield Curve Spreads, Credit Spreads and Equity Market conditions. All non-stationary variables are transformed to covariance stationary by suitable transformations. The definitions for all variables along with the transformations applied are summarised in section 2.8 of the Appendix.

Figure 2.1 depicts 6 key macroeconomic variables that are used to pinpoint each of the 6 dynamic factors. Industrial Production (measured on a Year-on-Year change

¹<http://research.stlouisfed.org/fred2>

basis) provides a proxy for the business cycle and highly correlates with the overall Business Cycle Factor F_t^{bs} . The Federal Funds Rate is a proxy for the risk free rate and drives the Cost of Debt Factor F_t^{cd} . As a proxy for the yield curve slope I choose the difference between the 10Y Constant Maturity Treasury and the 3M Treasury Bill; this is one of the main determinants of the Yield Curve Risk Factor F_t^{yc} . The difference between the BAA corporate yield and the 10Y Constant Maturity Treasury helps to pinpoint the Credit Risk Factor F_t^{cr} . The Year-on-Year returns on the S&P 500 index are used as proxies for the Asset Return Factor F_t^{ar} . Finally the monthly volatility of the S&P 500 index (annualised) drives the Asset Volatility Factor F_t^{av} . It is worth noting that the equity price volatility is measured as a standard deviation of the daily observations within a given month.

For the industry sector specific corporate bond defaults D_{it}^h , I rely on Moody's Default & Recovery Database. For the sector classification I use Moody's 11 sector definition. Companies that have no recorded sector information are excluded from the analysis. From the 11 sector classification I also exclude Sovereign & Public Finance companies since their risk behaviour is very different from the rest of the corporates. Due to the scarcity of defaults, I group the Banking sector with the Finance, Insurance and Real Estate (FIRE) sector, and the Energy & Environment sector with the Utilities sector. The final 8 industry sectors used in the analysis are: Capital Industries², Consumer Industries³, Banks & Financials, Media & Publishing, Retail & Distribution⁴, Technology, Transportation and Utilities.

The default counts are based on monthly rolling cohorts. Within each cohort, only the defaults of active companies at the beginning of the cohort period are recorded. Using cohort based default data can lead to inaccuracies when companies that are active at the beginning of the period leave the sample due to rating withdrawal; this becomes an issue especially when having long forecasting horizons. The non credit related rating withdrawals can distort significantly the default rates when the number of withdrawals is large as compared to total number of companies. In order to mitigate this problem I assume that the rating withdrawals are on average uniformly distributed through time, and, following Moodys adjustment, I subtract half their number from the initial count

²broad sector that includes capital equipment, aerospace & defence, automotive, chemicals, containers & packaging, paper products, and pharmaceuticals

³including durable and non-durable consumer goods

⁴mainly retail and wholesale distribution

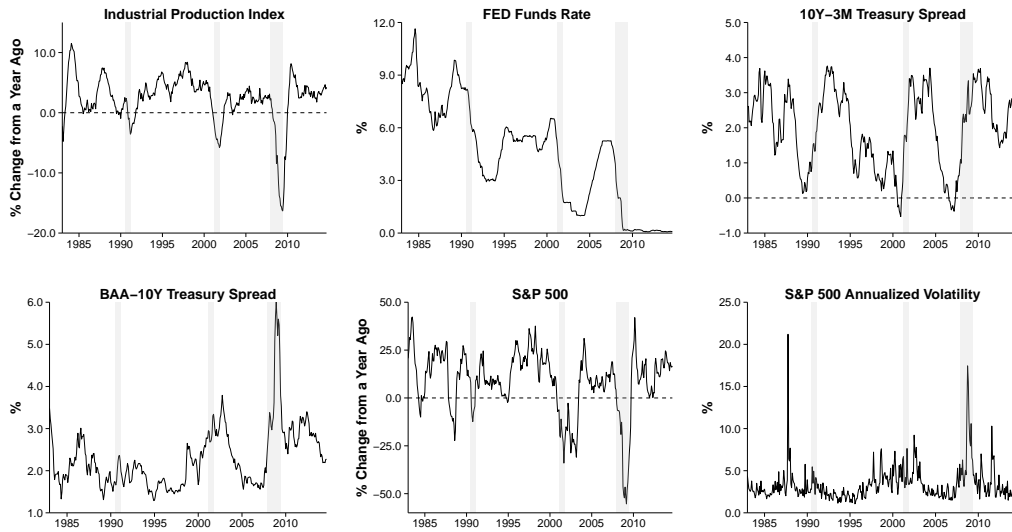


Figure 2.1: Historical Time Series - Macroeconomic Variables

Historical time series for key macroeconomic variables, covering the period January 1983-August 2014. All data sampled at monthly frequency. Industrial production and S&P 500 returns are expressed in Year-on-Year log differences. The yield curve slope is measured as the spread between the 10Y Treasury Constant Maturity and the 3M Treasury Bill. The credit spread is approximated as the difference between the BAA yield and the yield on the 10Y Treasury Constant Maturity. S&P 500 volatility is calculated on a monthly basis and then annualized. Federal funds rate, yield curve slope, credit spreads and S&P 500 volatility are all measured in percentage terms. Shaded areas correspond to NBER recession quarters.

	Capital Industries	Consumer Industries	Banks & Financials	Media & Publishing	Retail & Distribution	Technology	Transportation	Utilities
Avg. D^{1Y}	13	11	3	5	4	5	2	4
Avg. N	453	377	417	121	139	271	86	430
Avg. DR^{1Y}	2.7%	2.8%	0.9%	3.7%	2.9%	1.7%	2.8%	0.9%

Table 2.1: Summary of Moody's Corporate Default Data

Long-run average statistics for the corporate default dataset, for each of the 8 industry sectors used in the analysis. All reported figures are based on annual non-overlapping and forward-looking cohorts from 1983-2014 (2014 figures are calculated based on the period January 2014-September 2014). 'Avg. D^{1Y} ', refers to the average yearly default occurrence. 'Avg. N' refers to the average number of active companies at the beginning of each annual cohort. Finally, 'Avg. DR^{1Y} ' refers to the average annual default rate.

of companies.

When examining the default time series for corporate bonds, there is a clear structural break between the early 80s and late 70s, with default rates in the latter period significantly lower than the former. To exclude any possible outliers and to model a period as consistently as possible, I choose the starting point to be the January 1983. This starting point precedes the refinement of Moody's rating methodology to include an alphanumeric scale that took place in 1982. It is also consistent with other studies on corporate default modelling, such as Creal et al. (2013). The final period used in the estimation corresponds to September 2014. To evaluate the one month ahead forecasting performance, observations up to August 2014 are used.

Table 2.1 provides a summary of the corporate default data used in this study. Transportation, Media & Publishing and Retail & Distribution are the sectors with the lowest average firm count historically. Furthermore, Transportation is the sector with the fewest defaults, averaging 2 per year. Nevertheless, the sector's average annual default rate of 2.8% is relatively high as compared to the typically high credit quality sectors such as Banks & Financials and Utilities that exhibit average annual default rates of 0.9%. On the other hand, the sectors with the highest number of default counts are Capital and Consumer Industries, averaging 13 and 11 default per year respectively.

In addition to the high level summary provided in table 2.1, figure 2.2 depicts the time series of historical annual default rates for each of the 8 industry sectors used in the analysis. The annual default rates are provided on a monthly rolling frequency. Highly cyclical sectors such as Capital Industries, Consumer Industries and Retail & Distribution exhibit clear business cycle dependence, with increased default frequency over the 3 recession period in the sample (early '90s, early '00s and late '00s). Media & Publishing also shows strong cyclical dependence with higher default numbers during the recession periods. It is worth noting that the sector has been disproportionately hit by the recent credit crisis, with default rates over that period approximately 3 times the historical peak. Banks & Financials exhibit two major peaks in default rates: the first during the late '80s Savings & Loans crisis and the second during the late '00s credit crunch. The Technology sector was severely hit during the burst of the dot-com bubble in the early '00s. Finally, the Transportation and Utilities sectors exhibit a very mild dependence to the business cycle with default occurrence being highly irregular.

2.6 Results

2.6.1 Dynamic Factor Model

Using the MCMC algorithm described in section 2.4, the dynamic factor loadings for all macroeconomic variables are summarised in section 2.8 of Chapter 2 Appendix. As the variance across variables can be vastly different, I normalise all inputs to zero mean and unit standard deviation; not adjusting the inputs can lead to series with high variance dominating the estimated factors. Some of the loadings are very small in

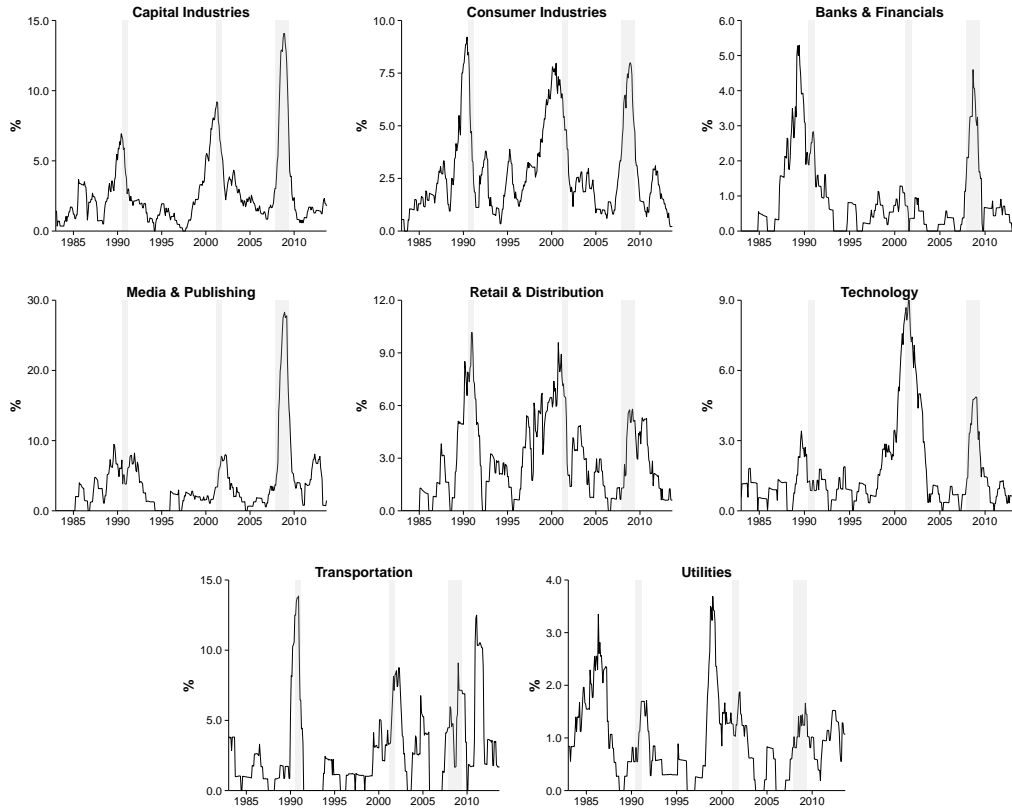


Figure 2.2: Historical Time Series - Corporate Default Rates

Historical time series of annual default rates for each of the 8 industry sectors used in the analysis, covering the period January 1983-September 2013. Annual default rates are calculated based on a monthly frequency and they are reported on a forward-looking basis (i.e. the default rate for September 2013 refers to the default count from September 2013-September 2014 as a proportion of the number of active companies at September 2013). Shaded areas correspond to NBER recession quarters.

absolute terms and not significantly different from 0. For some of those variables that exclusively load on a single dynamic factor, such as a number of unemployment duration variables, it is evident that the dynamic factor model employed in this analysis does not adequately capture the historical variation in those series. Furthermore, it becomes apparent that the error terms of those variables should be allowed to be autocorrelated, since the estimated dynamic factors fail to capture any of the historical autocorrelation in the time series. As I argue in section 2.3.2, the purpose of the analysis is not to forecast or accurately model every single input variable but rather to provide estimates of the dynamic factors \mathbf{F}_t^m . Therefore, I acknowledge that part of the extensive set of input variables will not be accurately captured, but I choose to maintain the model's tractability and use the 6 dynamic factor structure.

Without loss of generality, I use a simple VAR(1) structure for the factor stochastic evolution in (2.6). This choice of dynamics implies that the lag polynomial $\Phi(L)$

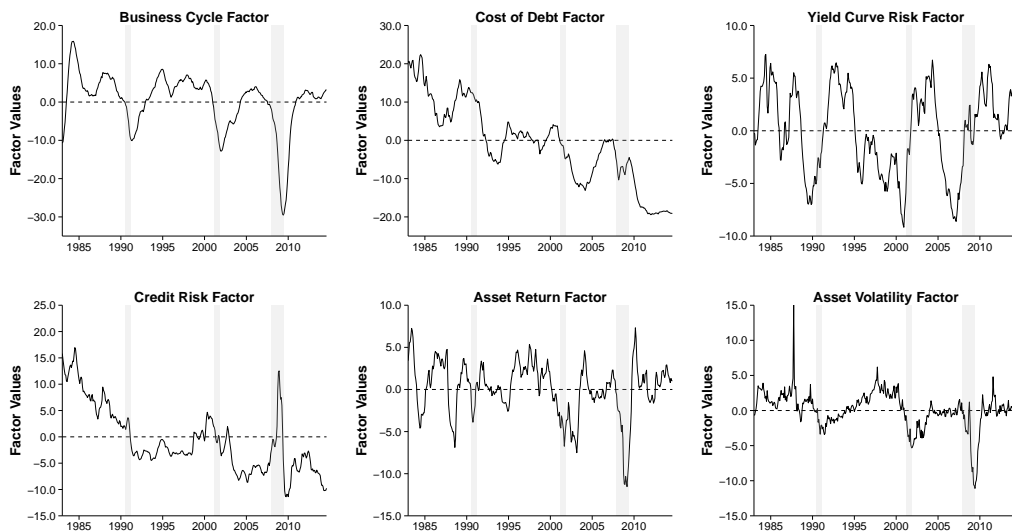


Figure 2.3: Estimated Dynamic Factors

Historical time series of the 6 estimated dynamic factors, covering the period January 1983-August 2014. All factors are estimated by means of the MCMC algorithm described in section 2.4. For each factor and month, the posterior mean is depicted, based on 50,000 MCMC draws. Shaded areas correspond to NBER recession quarters.

reduces to the full 6x6 matrix Φ . This choice is not overly restrictive, especially since the aim is not to forecast the dynamic factors. Table 2.2 provides point estimates and standard errors for the VAR(1) matrix Φ . The VAR dynamics in table 2.2 are bound to be marginally biased due to the simplifying assumptions of iid error terms in the observation equation of the state space representation. Estimates of both loadings and VAR coefficients are based on 60,000 draws of the MCMC sampling scheme, with the first 10,000 draws discarded to ensure inference is based on only posterior draws. For brevity, only the estimates based on the full sample January 1982-August 2014 are provided. Furthermore, figure 2.3 depicts the historical mean realisations of each dynamic factor.

The estimated loadings in section 2.8 and the historical realisation of F_t^{bs} in figure 2.3 confirm the interpretation of F_t^{bs} as a Business Cycle factor. The factor loads primarily on Industrial Production, Capacity Utilisation, Personal Income and changes in Employment variables. Equity Volatility variables also exhibit weights large in magnitude; nevertheless, the historical pattern of the factor seems unaffected by the equity market conditions. Prices, Inventories & Orders, Housing Starts and Bank Lending variables all have little impact on F_t^{bs} . This is not overly surprising, since Price variables are typically ordered after business cycle variables in VARs or dynamic factor models, see for example Boivin et al. (2013). Furthermore, Inventories & Orders and

	F_{t-1}^{bs}	F_{t-1}^{cd}	F_{t-1}^{yc}	F_{t-1}^{cr}	F_{t-1}^{ar}	F_{t-1}^{av}
F_t^{bs}	0.964 (0.020)	0.018 (0.012)	0.090 (0.016)	-0.036 (0.016)	0.150 (0.029)	0.024 (0.049)
F_t^{cd}	0.058 (0.016)	0.976 (0.010)	-0.020 (0.016)	0.010 (0.016)	-0.008 (0.029)	-0.088 (0.047)
F_t^{yc}	-0.016 (0.015)	-0.009 (0.010)	0.961 (0.016)	0.020 (0.016)	0.058 (0.029)	-0.050 (0.047)
F_t^{cr}	-0.049 (0.029)	0.041 (0.012)	0.044 (0.017)	0.907 (0.016)	-0.124 (0.036)	0.340 (0.054)
F_t^{ar}	0.026 (0.021)	0.004 (0.010)	-0.013 (0.016)	0.001 (0.016)	1.029 (0.031)	-0.211 (0.050)
F_t^{av}	0.113 (0.044)	-0.019 (0.010)	-0.007 (0.016)	0.042 (0.016)	0.279 (0.047)	0.437 (0.046)

Table 2.2: Estimated Dynamic Factor Transition Matrix

Parameter estimates for the 6 dynamic factor VAR(1) dynamics $\Phi^m(L)$ in the state equation (2.6), based on the period January 1983-August 2014. All estimates are based on the MCMC algorithm described in section 2.4 using 50,000 MCMC draws. Standard errors are also provided in brackets.

Housing Starts are typically treated as leading indicators of economic activity and they are not expected to show significant contemporaneous correlations with business cycle variables. In fact Bernanke et al. (2005) in their FAVAR classify Inventories & Orders and Housing Starts as 'fast-moving' together with stock price and interest rate variables. I refrain from grouping the former group of variables with the latter, since they represent fundamentally different information sets and the observed correlation is relatively low.

Interest rate variables mainly load on the Cost of Debt Factor F_t^{cd} . The loadings on F_t^{cd} are surprisingly similar, implying that the factor is approximately equally balanced across all interest rate variables. The positive contemporaneous dependence on the Business Cycle Factor F_t^{bs} , as implied by the uniformly positive loadings, is not overly surprising; as the economic conditions improve, prices are likely to rise (as indicated by the positive loadings of Price variables to F_t^{bs}) and as a response the FED will increase interest rates to suppress inflation. Exploring the VAR dynamics in table 2.2, this positive dependence of the interest rate variables to the business cycle fluctuations persists even at the 1 month lag.

The behaviour of the yield curve spread variables is less uniform as compared to the interest rate variables. The long-to-medium (10Y to 2Y) and long-to-short (10Y to 3M)

maturity spreads are negatively affected by changes in F_t^{bs} , while the difference between 2Y to 3M government yields and the TED spread show some positive correlation with F_t^{bs} , albeit statistically insignificant. The negative dependence of the former two spread variables also drives the negative correlation of the Yield Curve Risk Factor F_t^{yc} to the lagged Business Cycle Factor F_{t-1}^{bs} . These results are broadly consistent with the work of Diebold et al. (2006) that use a Nelson-Siegel dynamic model to decompose the yield curve and link it to macroeconomic variables. Diebold et al. (2006) report negative dependence of the slope factor to the capacity utilisation (in their paper, slope is defined as the difference between the short and the long end of the yield curve, which corresponds to minus the definition used in my study) both at lag 1 and lag 0. The dependence of the yield curve spread variables to the Cost of Debt Factor F_t^{cd} is mostly positive, with only the 10Y to 2Y spread having a negative loading. Furthermore, the overall dependence of F_t^{yc} to F_{t-1}^{cd} is negative, albeit highly insignificant. This appears to partially contradict the study of Diebold et al. (2006) that report a strongly positive dependence of the yield curve slope factor to the lagged value of the yield curve level factor and a strongly negative dependence of the yield curve slope to the lagged value of the FED Funds rate. Additionally, the contemporaneous correlations between slope and level factors is mildly positive, while the contemporaneous correlation between slope and FED Funds rate is strongly negative. As in the case of the analysis presented here the definition of Cost of Debt and Yield Curve Risk Factors is not the same with the work of Diebold et al. (2006), the comparison of the reported results is not straightforward.

Based on the results of section 2.8, credit spreads are negatively correlated to F_t^{bs} , F_t^{cd} and F_t^{yc} , while they respond positively to changes in F_t^{cr} . The signs of the credit spreads are intuitive. As the economic conditions improve (positive movements of the F_t^{bs} factor) and the corporate balance sheets strengthen, credit spreads are expected to be suppressed. Using the structural approach to valuing corporate debt, Longstaff and Schwartz (1995) notice that an increase in interest rates leads to an increase in the risk-neutral drift and they report a negative correlation between spreads and interest rates. This is consistent with the results of table 2.6 that indicate a negative relationship between credit spreads and F_t^{cd} . To the extent that increases in yield curve slope reflect increases in the expected future level of interest rates, positive movements in F_t^{yc} should negatively affect credit spreads. Furthermore, flattening or even inversion of the yield

Figures in %	F_t^{bs}	F_t^{cd}	F_t^{yc}	F_t^{cr}	F_t^{ar}	F_t^{av}	Σ_e
<u>FEVD</u>							
Industrial Production	13.8	1.3	7.5	2.5	37.6	11.9	25.4
Civilian Employment (Log Chg.)	16.4	1.5	8.9	3.0	44.7	14.2	11.3
FED Funds Rate	15.0	40.5	6.7	2.2	19.5	10.8	5.3
GS10-TB3MS	12.2	1.7	56.5	7.2	13.5	7.9	1.0
BAA Yield	11.3	35.4	4.9	1.6	12.1	7.9	26.8
BAA-GS10 Spread	11.1	25.5	10.6	14.2	26.0	12.0	0.6
S&P 500 Return	1.9	0.9	6.7	1.4	74.1	11.9	3.2
S&P 500 Volatility	8.4	1.4	4.9	2.1	61.0	21.8	0.3
<u>R²</u>							
Industrial Production	76.1	0.6	1.1	1.8	15.9	58.0	-
Civilian Employment (Log Chg.)	89.7	3.0	0.4	2.8	6.0	55.7	-
FED Funds Rate	11.5	92.8	11.8	60.9	1.0	7.6	-
GS10-TB3MS	9.5	1.8	92.0	0.4	2.4	8.7	-
BAA Yield	0.7	87.0	0.5	84.3	0.5	0.2	-
BAA-GS10 Spread	46.0	8.0	2.2	3.0	23.3	26.9	-
S&P 500 Return	20.8	0.1	0.7	3.9	96.4	46.1	-
S&P 500 Volatility	10.7	0.4	0.1	3.8	17.1	0.2	-

Table 2.3: Dynamic Factor Forecast Error Variance Decomposition and R^2

Proportion of the 60-step (5 years) forecast error variance of macroeconomic variables that is accounted for by the Business Cycle (F_t^{bs}), Cost of Debt (F_t^{cd}), Yield Curve Risk (F_t^{yc}), Credit Risk (F_t^{cr}), Asset Returns (F_t^{ar}) and Asset Volatility (F_t^{av}) factors. Variance decompositions are calculated using the posterior means of the Φ s in (2.6). Column Σ_e refers to the residual variance from the error terms in (2.5). R^2 are calculated using univariate regressions of each macroeconomic variable on each factor.

curve is typically associated with a deterioration in economic activity. Therefore credit spreads are expected to be negatively correlated to the slope of the yield curve (or factor F_t^{yc}); this is empirically supported by the work of Collin-Dufresne et al. (2001).

Both equity return and volatility series exhibit a strong dependency on the Asset Return Factor F_t^{ar} . While equity return variables have a positive dependence on F_t^{ar} , the loadings for the equity volatility variables are negative implying that volatility drops as Asset Returns improve. This is in line with intuition, since jumps in volatility are more likely during periods of stress. Equity volatility variables also show a negative dependence on the Business Cycle Factor F_t^{bs} , further supporting the view that volatility rises during downturns. Finally, equity return variables have marginally positive sensitivities to F_t^{bs} .

To complete the analysis of the dynamic factor model, table 2.3 depicts the Forecast Error Variance Decomposition (FEVD) for some key macroeconomic variables, alongside some univariate R^2 statistics for each of the 6 factors used in this study. The

FEVD is calculated as:

$$\omega_{ik,h}^m = \frac{\left[\sum_{z=1}^6 \left((\lambda_i^z)^2 \sum_{j=0}^{h-1} (\phi_{zk,j}^m)^2 \right) \right] + \sigma_{e_i}^2}{\sum_{n=1}^6 \left[\sum_{z=1}^6 \left((\lambda_i^z)^2 \sum_{j=0}^{h-1} (\phi_{zn,j}^m)^2 \right) \right] + \sigma_{e_i}^2}, \quad (2.15)$$

where i is the specific macroeconomic variable of interest, λ_i^z is the loading of variable i to factor z (for simplicity in the notation, I assume that the factors are ordered from 1 to 6), $\phi_{zk,j}^m$ is the (z, k) element of the matrix of dynamic multipliers at step j , $(\Phi^m)^j$, and $\sigma_{e_i}^2$ is the variance of residuals for variable i . For the calculation of the FEVD, the posterior mean values for the elements of matrix Φ^m are used. FEVDs are based on a 60-month horizon. The R^2 statistics are calculated by running univariate regression of each variable on each dynamic factor. In this particular case, the R^2 corresponds to the squared correlation between observed variable and estimated factor.

The FEVDs in table 2.3 indicate that Asset Return and Business Cycle shocks explain a significant proportion of the forecast variance for the majority of the variables. More specifically, the long-run forecast variance for real economy variables such as industrial production and changes in employment are explained by mainly Asset Return and Business Cycle shocks (approximately 40% for the former and close to 15% for the latter type of shocks) and, to a lesser extent, by Asset Volatility shocks. This relative weight of the Asset shocks might seem counterintuitive; nevertheless, Business Cycle shocks dominate the FEVD in the short-run. Government and corporate yields long-run forecasts are primarily driven by the the Cost of Debt Factor F_t^{cd} and to a lesser extent by the Business Cycle and Asset Return Factors F_t^{bs} and F_t^{ar} . The significant impact of Asset Return movements to yield curve level forecasts implies that FED and the bond market take into consideration equity market movements when determining the level of interest rates. More than 50% of the long-run forecast variance for the 10-year to 3-month government bond spread is attributed to the Yield Curve Risk Factor F_t^{yc} with Business Cycle movements and Asset Return fluctuations contributing by 12.2% and 13.5% respectively. This decomposition is not surprising, since the slope of the yield curve reflects expectations on future inflation and, as a market priced measure, it is likely to be affected by the investors' sentiment behind equity premia. Credit spread forecasts are more balanced, being mainly driven by Cost of Debt and Asset Return

shocks with the remaining variance decomposed uniformly across Business Cycle, Yield Curve Risk, Credit Risk and Asset Volatility Factor shocks. Finally, both return and volatility series of the S&P 500 index exhibit similar behaviour; the forecast error variance is predominately driven by Asset Return shocks and to a lesser extent Asset Volatility shocks.

The R^2 statistics reported in table 2.3 reflect the contemporaneous correlations between input variables and extracted factors. Despite the number of real economy variables used in the dynamic factor model, the Business Cycle Factor F_t^{bs} is able to capture 76% of the industrial production and 90% of the changes in employment time series. The Cost of Debt Factor F_t^{cd} is able to explain 93% and 87% of the historical variation in the two bond yields reported in table 2.3, namely FED funds rate and BAA yield respectively. Additionally, the BAA yield exhibits an 84% correlation with the Credit Risk Factor F_t^{cr} . As expected, 92% of the historical variation of the 10-year to 3-month government bond spread is explained by the Yield Curve Risk Factor F_t^{yc} alone. Somewhat surprisingly, the Credit Risk Factor F_t^{cr} does not explain a significant portion of the historical movements in the corporate bond spread, defined as the difference between the BAA yield and the 10-year government bond yield. This is mainly caused by the additive nature of the factor, that captures the variation over and above what is explained by F_t^{bs} , F_t^{cd} and F_t^{yc} . These 3 factors are jointly able to explain a significant proportion of the movements in corporate bond spreads pre-2000. If the R^2 is measured from 2000 onwards, then F_t^{cr} is able to explain 42% of the corporate bond spread variation. The Asset Return Factor F_t^{ar} explains 96% of the S&P 500 returns movements. Finally, the Asset Volatility Factor F_t^{av} does not seem to be correlated with the volatility in the S&P 500 index. If observed in isolation, this result appears somewhat counterintuitive. Nevertheless, the R^2 only picks up contemporaneous correlations between individual factors and variables without looking at the dynamic relationships between them. In this case, if the low R^2 is examined in conjunction with the FEVD, then the results of table 2.3 imply that the F_t^{av} factor mainly captures the long-run forecast dynamics of the equity volatility series rather than the short-run fluctuations. The short-run dynamics can be explained adequately by a combination of the Business Cycle and Asset Return Factors.

2.6.2 Corporate Default Forecasts

To assess the forecasting power of the extracted factors from section 2.6.1 for corporate bond defaults, I use 3 values for the forecasting horizon h in the corporate default econometric model (2.1)-(2.3): 3 months, 6 months and 1 year. The analysis can be generalised to cover default rates over shorter and longer time horizons. In practice 1-2 month default rates are very volatile and inherently difficult to explain and default rates of longer than 1 year do not strongly correlate with macroeconomic conditions at time t . Each sector and forecasting horizon is allowed to have a different mix of factors \mathbf{F}_t^m and lags, with the selection of the explanatory factors being determined solely on the basis of the statistical power for each sector/forecasting horizon segment. Finally, to ensure that the measurement of the forecasting power is robust, I sequentially estimate the model using windows of increasing size. Starting from the January 1983-August 2006 window, I re-estimate the corporate default specification in (2.1)-(2.3) using that same starting date, but sequentially increasing the sample ending date by 1 year; this results in 8 recursive samples within the period August 2006-August 2013.

In practice, the estimation of full VAR dynamics can lead to estimation problems. Experimenting with different lag structures indicates that a 2-month lag is the maximum lag that does not cause estimation problems. Therefore, I use 2 as the maximum lag order for forecasting purposes. To ensure the robustness of the parameter estimates, I use a full VAR(2) specification for the frailty factors when possible and I switch to single lag dynamics when the MCMC algorithm lead to numerical instabilities. For forecasting horizons of 3 months and 1 year, the MCMC algorithm provides numerically stable estimates for the VAR(2) specification from the sample ending dates from August 2009-August 2013. For the 6 months forecasting horizon, numerical stability is maintained only for the recursive samples ending from August 2010-August 2013. For the recursive samples ending before August 2009 for 3 months and 1 year forecasting horizons and before August 2010 for 6 months forecasting horizons, a simpler VAR(1) specification is used instead, both for estimation and out-of-sample forecasting.

The presence of the frailty factors in the specification (2.1)-(2.3) causes the sensitivities of corporate defaults to the macroeconomic factors \mathbf{F}_t^m to drop in absolute magnitude. While this is a desirable feature if inference is biased in the absence of

autoregressive dynamics in the econometric specification for defaults, it might not be desirable when assessing out of sample forecasting power for long horizons h . In the case of long forecasting horizons, the mean reverting property of the frailty factors might contribute disproportionately to the corporate default forecast and skew the outcome upwards or downwards. To assess the sensitivity of the forecast to the relative weight of the frailty factors, I estimate the model using two sets of priors. Each one is based on an equivalent econometric specification to (2.1)-(2.2), but without the presence of frailty factors $\mathbf{F}_t^{d,h}$ (from here on the "reference model"):

- Vague Priors: The priors for the sensitivities to the macroeconomic factors \mathbf{F}_t^m are centred around the estimates of the "reference model" and they have a prior variance of 1. While not completely diffuse, the priors are not overly restrictive and the MCMC algorithm of section 2.4 is able to freely estimate the frailty dynamics.
- Tight Priors: The priors for the sensitivities to the macroeconomic factors \mathbf{F}_t^m are centred around the estimates of the "reference model" and they have a prior variance equal to the variance of the "reference model" posterior draws. Since the choice of variables in the "reference model" is based on statistical significance, the standard errors of the "reference model" estimates are relatively small. Therefore, using the variance of the posterior draws from the "reference model" as the prior variance of the full specification (2.1)-(2.3) results in very tight priors that dominate the posterior and essentially force the frailty factors to reflect the residual fluctuation in corporate defaults after the macroeconomic dynamics are removed.

Under both choices for the priors of the sensitivities to the macroeconomic factors \mathbf{F}_t^m , the priors for the VAR dynamics in (2.3) are diffuse, as previously mentioned in section 2.4.

Based on the two set of priors mentioned previously, table 2.4 summarises the parameter estimates and their respective standard errors for the full sample period January 1983-August 2013. Estimates are provided for each forecasting horizon/industry sector/prior choice segment. All estimates are based on the MCMC algorithm described in section 2.4, using 50,000 draws for inference after discarding 10,000 draws

$\times 10^{-2}$		Tight Priors								Vague Priors							
F_t^m	Sec.	CA	CO	FI	ME	RE	TE	TR	UT	CA	CO	FI	ME	RE	TE	TR	UT
3M																	
F_t^{bs}		-5.1 (0.1)	-4.0 (0.1)	-6.4 (0.1)	-11.4 (0.2)	-4.3 (0.2)	-2.2 (0.1)	-6.6 (0.2)	-4.3 (0.1)	-5.3 (1.7)	-2.6 (1.8)	-1.9 (2.2)	-11.3 (2.5)	-0.7 (3.0)	-2.1 (1.9)	-2.9 (3.3)	-7.0 (2.0)
F_t^{cd}		- (-)	- (-)	6.1 (0.1)	- (-)	- (-)	- (-)	- (-)	- (-)	- (-)	- (-)	5.1 (1.1)	- (-)	- (-)	- (-)	- (-)	- (-)
F_t^{yc}		-7.3 (0.1)	-8.4 (0.1)	-5.6 (0.1)	- (-)	-5.5 (0.2)	-5.0 (0.2)	-3.7 (0.3)	- (-)	-6.4 (1.8)	-6.6 (2.0)	-2.9 (2.9)	- (-)	-3.5 (4.3)	-4.1 (2.8)	6.4 (4.2)	- (-)
F_t^{cr}		3.7 (0.1)	3.3 (0.1)	- (-)	- (-)	- (-)	- (-)	- (-)	4.9 (0.1)	1.6 (1.5)	3.7 (1.3)	- (-)	- (-)	- (-)	- (-)	- (-)	2.7 (2.1)
F_{t-11}^{cr}		- (-)	- (-)	- (-)	- (-)	- (-)	3.3 (0.1)	- (-)	- (-)	- (-)	- (-)	- (-)	- (-)	- (-)	3.5 (1.7)	- (-)	- (-)
F_t^{ar}		-6.9 (0.2)	-6.2 (0.2)	-10.9 (0.2)	-3.3 (0.3)	-3.9 (0.3)	-13.2 (0.2)	-5.5 (0.4)	- (-)	-6.7 (2.3)	-4.8 (2.1)	-17.3 (3.5)	0.1 (3.3)	-6.0 (4.1)	-13.0 (2.9)	-13.8 (5.8)	- (-)
F_t^{av}		- (-)	7.2 (0.4)	- (-)	- (-)	- (-)	- (-)	8.1 (0.8)	9.9 (0.3)	- (-)	5.3 (2.9)	- (-)	- (-)	- (-)	- (-)	6.7 (8.5)	11.5 (2.3)
F_{t-3}^{av}		- (-)	- (-)	12.8 (0.4)	15.8 (0.6)	8.3 (0.6)	- (-)	- (-)	- (-)	- (-)	- (-)	9.9 (3.6)	12.9 (3.9)	3.0 (6.1)	- (-)	- (-)	- (-)
6M																	
F_t^{bs}		-5.3 (0.1)	-3.9 (0.1)	-5.7 (0.1)	-9.2 (0.2)	-4.2 (0.2)	-2.8 (0.1)	-4.9 (0.2)	-3.7 (0.1)	-4.2 (1.5)	-0.5 (1.7)	1.4 (4.1)	-8.0 (3.1)	2.4 (2.8)	-3.9 (2.6)	-2.8 (2.8)	-6.0 (2.3)
F_t^{cd}		- (-)	- (-)	6.0 (0.1)	- (-)	0.9 (0.1)	- (-)	- (-)	- (-)	- (-)	- (-)	5.2 (1.5)	- (-)	1.7 (2.0)	- (-)	- (-)	- (-)
F_t^{yc}		-7.8 (0.1)	-8.3 (0.1)	-6.1 (0.1)	- (-)	-6.1 (0.2)	-7.2 (0.2)	-4.0 (0.3)	-2.0 (0.1)	-5.9 (2.0)	-3.0 (2.3)	0.1 (3.0)	- (-)	2.8 (2.9)	-3.8 (3.3)	-0.6 (4.8)	-2.8 (3.3)
F_t^{cr}		3.4 (0.1)	3.2 (0.1)	- (-)	2.9 (0.2)	- (-)	- (-)	- (-)	5.2 (0.1)	1.8 (1.7)	4.4 (1.8)	- (-)	0.4 (2.0)	- (-)	- (-)	- (-)	7.8 (1.9)
F_{t-11}^{cr}		- (-)	- (-)	- (-)	- (-)	- (-)	3.7 (0.1)	- (-)	- (-)	- (-)	- (-)	- (-)	- (-)	- (-)	1.2 (2.0)	- (-)	- (-)
F_t^{ar}		-9.5 (0.2)	-6.5 (0.2)	-15.0 (0.2)	-6.8 (0.4)	-6.4 (0.4)	-14.2 (0.2)	-3.7 (0.4)	- (-)	-3.1 (2.0)	0.3 (2.0)	-12.1 (4.1)	0.3 (3.3)	-3.8 (3.7)	-5.6 (3.4)	-4.2 (4.8)	- (-)
F_t^{av}		5.6 (0.4)	9.2 (0.4)	14.0 (0.4)	10.5 (0.7)	9.0 (0.7)	4.2 (0.5)	- (-)	7.9 (0.3)	-0.1 (2.4)	0.1 (2.5)	1.8 (3.4)	2.1 (3.4)	-4.8 (4.5)	-0.7 (3.6)	- (-)	6.4 (3.7)
1Y																	
F_t^{bs}		-3.8 (0.1)	-3.7 (0.1)	-4.2 (0.1)	-7.3 (0.2)	-4.8 (0.2)	-3.6 (0.1)	-5.4 (0.2)	-2.7 (0.1)	-0.6 (1.3)	2.9 (1.3)	3.2 (2.7)	-4.7 (2.4)	4.2 (1.8)	-1.9 (1.8)	-3.1 (1.9)	-7.0 (1.6)
F_t^{cd}		- (-)	- (-)	5.7 (0.1)	- (-)	0.9 (0.1)	- (-)	- (-)	- (-)	- (-)	- (-)	5.1 (0.7)	- (-)	1.1 (0.8)	- (-)	- (-)	- (-)
F_t^{yc}		-7.8 (0.1)	-8.4 (0.1)	-4.3 (0.1)	- (-)	-7.0 (0.2)	-9.0 (0.2)	-3.5 (0.3)	- (-)	-2.7 (1.2)	-0.8 (1.1)	-0.3 (1.7)	- (-)	0.8 (1.7)	-2.5 (1.7)	1.2 (2.5)	- (-)
F_{t-5}^{yc}		- (-)	- (-)	- (-)	- (-)	- (-)	- (-)	- (-)	-2.6 (0.1)	- (-)	- (-)	- (-)	- (-)	- (-)	- (-)	- (-)	0.7 (1.8)
F_{t-8}^{yc}		- (-)	- (-)	- (-)	-4.8 (0.2)	- (-)	- (-)	- (-)	- (-)	- (-)	- (-)	- (-)	-4.0 (1.4)	- (-)	- (-)	- (-)	- (-)
F_t^{cr}		2.8 (0.1)	3.1 (0.1)	- (-)	1.9 (0.2)	- (-)	2.0 (0.1)	- (-)	5.2 (0.1)	-0.1 (0.9)	0.9 (0.9)	- (-)	-1.4 (1.5)	- (-)	1.6 (1.3)	- (-)	-0.2 (1.6)
F_t^{ar}		-10.3 (0.2)	-5.5 (0.2)	-13.8 (0.2)	-8.7 (0.4)	-4.9 (0.3)	-10.1 (0.3)	-2.8 (0.4)	- (-)	-2.6 (1.4)	0.3 (1.3)	-3.8 (2.0)	0.6 (2.1)	2.0 (1.7)	0.2 (2.2)	-3.0 (2.6)	- (-)
F_t^{av}		5.3 (0.4)	11.2 (0.4)	12.9 (0.4)	8.7 (0.7)	10.5 (0.6)	5.5 (0.5)	- (-)	6.2 (0.3)	-1.6 (1.5)	1.7 (1.4)	-0.2 (2.1)	-0.9 (2.1)	-1.9 (2.6)	1.0 (2.5)	- (-)	4.3 (2.3)
F_{t-8}^{av}		- (-)	- (-)	- (-)	- (-)	- (-)	- (-)	5.2 (0.6)	- (-)	- (-)	- (-)	- (-)	- (-)	- (-)	- (-)	1.2 (3.6)	- (-)

Table 2.4: Corporate Default Parameter Estimates

Parameter estimates for the sensitivities of corporate defaults to the dynamic factors F_t^m in the econometric specification (2.1)-(2.3), based on the period January 1983-August 2013. Parameter estimates are provided for 3 forecasting horizons h : 3 months, 6 months and 1 year. The priors are loosely ("Vague Priors") or tightly ("Tight Priors") centered around the estimates of an equivalent model derived from (2.1)-(2.2) without frailty factors. All estimates are based on the MCMC algorithm described in section 2.4 using 50,000 MCMC draws.

as the burn-in sample. As a general remark, the choice of prior appears to increase in importance the longer the horizon h is. Comparing the parameter estimates across tight and vague priors indicates that there are only minor differences when forecasting defaults 3 months ahead, while there are major deviations and many sign reversions for the individual sensitivities when forecasting defaults 1 year ahead. This is not overly surprising. When forming rolling 1 year default rates on a monthly frequency, there is a large degree of overlap; at each time t the defaults count for the following 11 out of 12 months is the same with the default count from the period month $t - 1$. This creates a very high degree of (relatively long memory) autocorrelation. Additionally, corporate defaults over the period $t + 12$ months are less likely to exhibit high correlation with the macroeconomic conditions at time t , at least as compared to the corporate defaults over the period $t + 3$ months.

When examining the results of table 2.4 across sectors and the 3 forecasting horizons it is evident that there is a strong dependency of corporate defaults across all sectors to the Business Cycle Factor F_t^{bs} ; nevertheless, the sensitivity to F_t^{bs} appears to drop in magnitude with the length of the forecasting horizon h . It is worth noting that the Cost of Debt Factor F_t^{cd} plays a minor role, except for the defaults in the Financial and Retail sectors. For those sectors the effect is more likely to be linked to retail credit write-offs due to increased interest rates (drop in the affordability of household debt) and the drop in the demand for consumer goods. The remaining sectors depend on a combination of the Yield Curve F_t^{yc} and Credit Risk Factors F_t^{cr} to capture the debt financing conditions on the liability side of their balance sheet. While the effect of the F_t^{yc} factor appears stronger than the effect of F_t^{cr} , under vague priors most of the sensitivities to F_t^{yc} turn statistically insignificant, implying that the frailty factors are able to incorporate a large proportion of the yield curve risk. The Asset Return Factor F_t^{ar} has a strong effect across all sectors and that effect appears to strengthen with the length of the forecasting horizon h . Nevertheless, as the high standard errors in the specification with vague priors suggest, most of the variation in the F_t^{ar} factor in the longer horizons can be attributed to the frailty factors. Finally, the Asset Volatility Factor F_t^{av} appears to play a minor role in the short-run, with statistically significant estimates only for Financial, Media and Retail sectors, but it increases in significance as the forecasting horizon h increases.

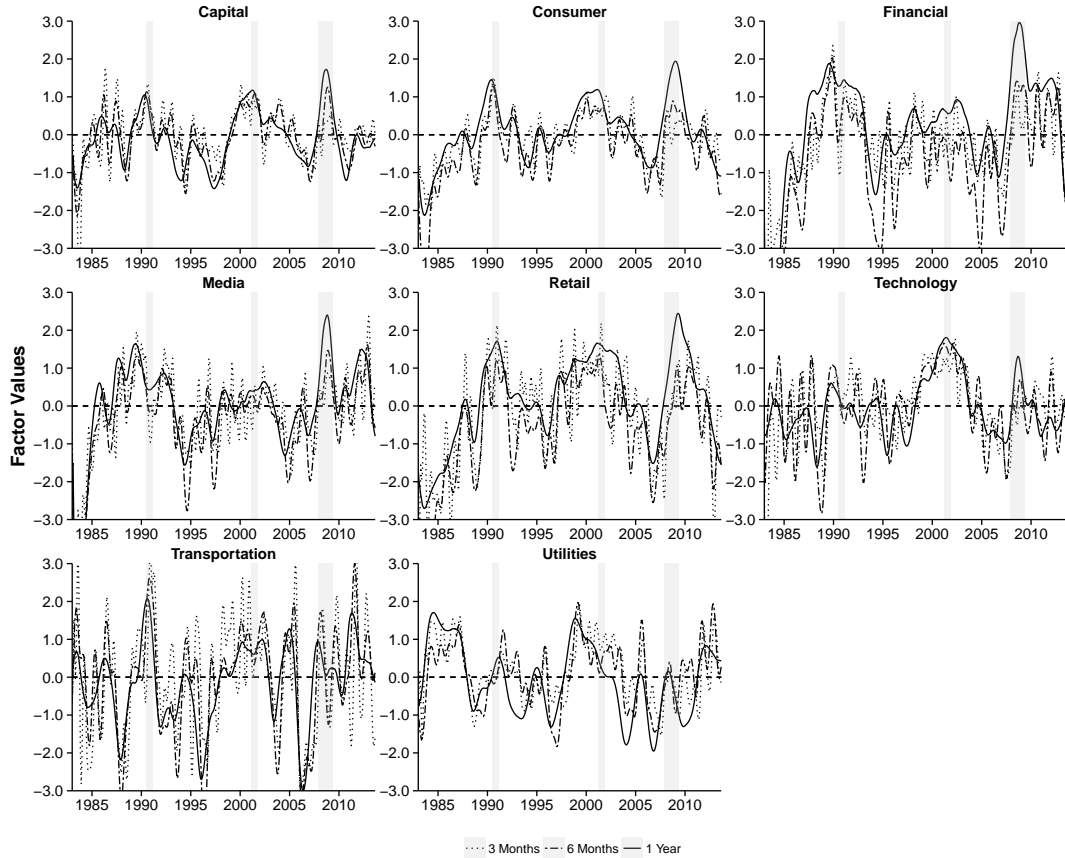


Figure 2.4: Estimated Frailty Factors

Historical time series of the estimated frailty factors for each of the 8 sectors, covering the period January 1983–August 2013. Frailty factors are provided for 3 forecasting horizons: 3 months, 6 months and 1 year. Factor estimates are extracted from the econometric specification (2.1)–(2.3) with priors loosely centered around the estimates from an equivalent model without frailty factors. All factors are estimated by means of the MCMC algorithm described in section 2.4. For each factor and month, the posterior mean is depicted, based on 50,000 MCMC draws. Shaded areas correspond to NBER recession quarters.

For the parameter estimates in table 2.4, figure 2.4 depicts the historical estimates for the frailty factors per industry sector and forecast horizon. The frailty estimates are based on the econometric specification with loose priors; that choice of priors provides less restrictions in the frailty dynamics and allows the frailty factors to be freely extracted from the observed data. For comparability across the different sectors, each estimated factor is multiplied by the respective loading $\beta_{i,h}^d$ in (2.2). The frailty factor estimates across the different forecasting horizons are similar, with the series becoming more volatile the shorter the horizon h . For the majority of sectors, the frailty factors exhibit a jump during the early '90s and late '00s recessions, especially in the 1-year forecasting horizon. This is not overly surprising since it is well documented that there is excess default correlation during economic downturns, see for example Koopman et al. (2012). Furthermore, the Technology sector frailty factor peaks around

the early '00s dot-com bubble and the subsequent recession, a finding consistent with the sector's performance during that period. Low default occurrence sectors such as Transportation and Utilities exhibit more idiosyncratic patterns. Transportation shows some minor cyclical movements but it is overall extremely volatile even at the 1-year forecasting horizon, reflecting the only 2 defaults per year on average for the sector, as reported in table 2.1. Finally, the Utilities frailty factor reflects the irregular default behaviour of the sector, with jumps in mid-'80s and during the Asian crisis/Russian default period.

To summarise the overall forecasting performance of the dynamic macroeconomic and frailty factor regression (2.1)-(2.3) against the "base model", figure 2.5 depicts the deviations of each specification from the historically realised default rate. Results are provided per industry sector for each of the 3 forecasting horizons using recursive 1-year estimations from August 2006-August 2013 and out-of-sample default rates over the period September 2006-September 2014. By examining the results of figure 2.5 it is clear that, overall, the full specification (2.1)-(2.3) provides closer forecasts to the realised default rates as compared to the "base model" with only Industrial Production as the explanatory variable. The forecasting performance difference between the full specification with dynamic and frailty factors and the "base model" is larger for Capital Industries, Technology and Transportation during the 2009 stress period, especially over the 3-month and 6-month forecasting horizons. For Consumer Industries, the full specification models with tight and vague priors provide less accurate forecasts over the period 2008-2010 as compared to the "base model", while the forecast performance improves substantially over the last 2 years in the out-of-sample window. Furthermore, for the Consumer Industries sector, the use of tight and vague priors gives materially different forecasts on the 1-year horizon, with the use of vague priors leading to a substantial over-estimation of the realised default rate. This is primarily caused by the relative low severity of the default activity for the sector during the recent economic downturn, in relative terms to the early '90s and early '00s recessions; that causes the frailty dynamics to overshoot the default rate forecast by a large margin. The same is also true for the Retail sector, for which the full specification with vague priors gives forecast different from the realised default rate 4 times higher than the forecasts generated by the same specification with tight priors. For Financials, the dynamic and frailty factors imply a much higher default rate as compared to the actual default

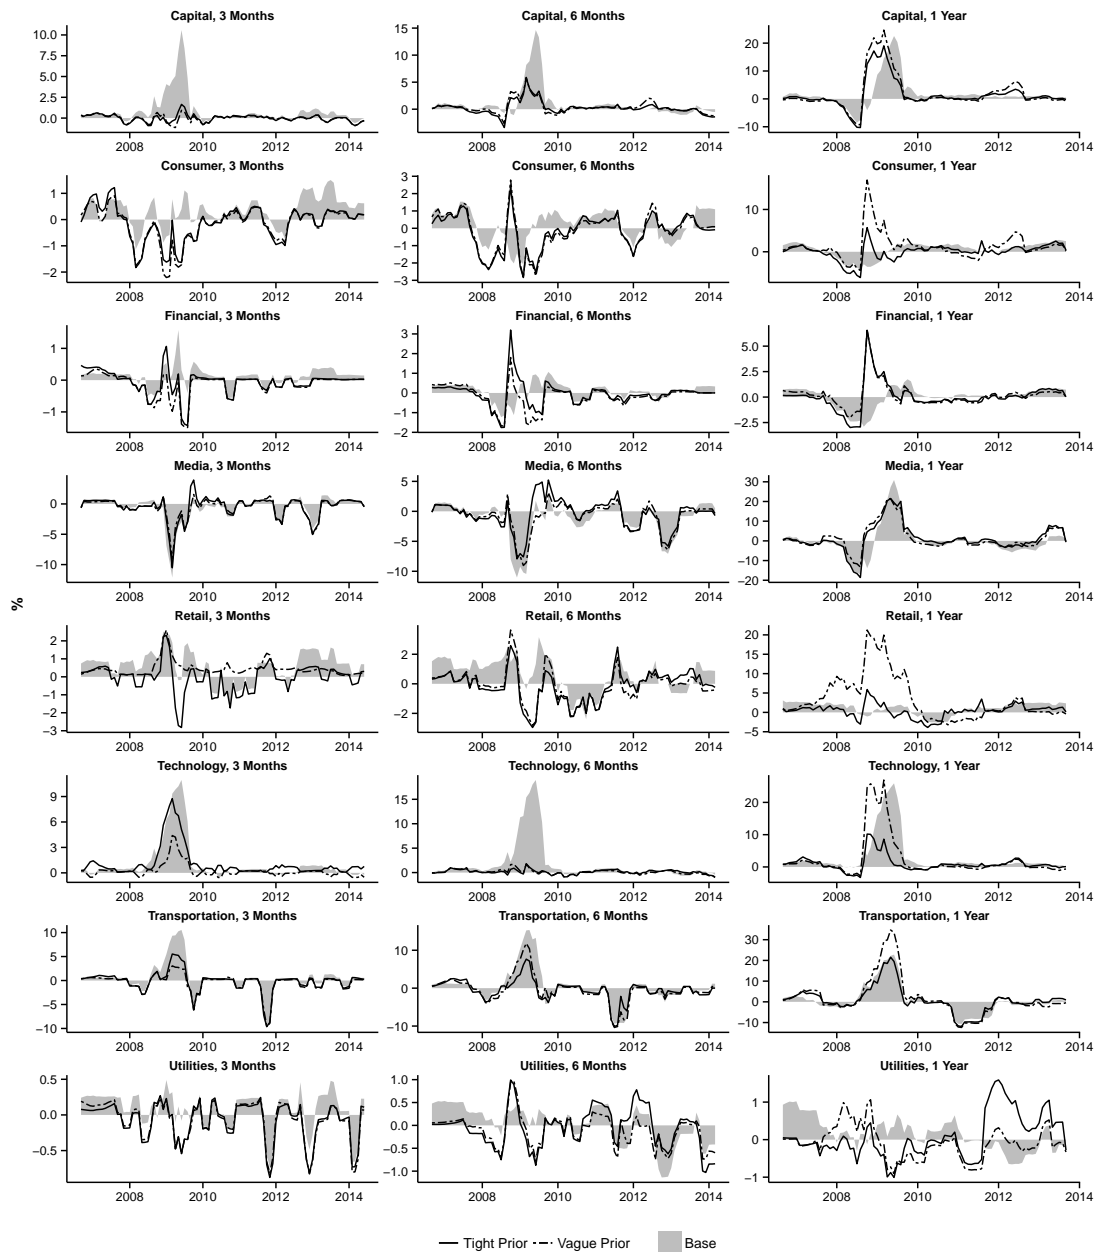


Figure 2.5: Default Rate Predictions Errors

Deviations from actual default rates using predictions from 3 models: "Vague Prior" refers to the econometric specification (2.1)-(2.3) with priors loosely centered around the estimates from an equivalent model without frailty factors, "Tight Prior" refers to the econometric specification (2.1)-(2.3) with priors tightly centered around the estimates from an equivalent model without frailty factors, while a binomial-logit model without frailty factors and only Industrial Production as an explanatory variable serves as the "Base" case. Predictions are provided for each sector across 3 forecasting horizons: 3 months, 6 months and 1 year. All predictions are based on recursive samples with starting date January 1983 and ending date sequentially augmented from September 2006-September 2013 at an annual frequency. All parameter estimates are based on the MCMC algorithm described in section 2.4 using 50,000 MCMC draws.

occurrence over 2009, while the forecasts for the remaining out-of-sample period are more accurate as compared to the "base model". Finally, the Utilities sector exhibits a relatively a-cyclical default pattern and all econometric specifications provide similar

Sector		3 Months			6 Months			1 Year		
		Base	Tight	Vague	Base	Tight	Vague	Base	Tight	Vague
Capital	RMSE	2.03	0.44	0.42	3.00	1.24	1.31	5.50	5.32	6.82
	MAE	0.96	0.33	0.32	1.32	0.76	0.81	2.91	2.78	3.53
Consumer	RMSE	0.64	0.67	0.71	0.89	1.03	1.07	1.85	1.87	3.43
	MAE	0.53	0.50	0.49	0.77	0.77	0.82	1.59	1.27	2.08
Financial	RMSE	0.31	0.37	0.38	0.45	0.61	0.59	0.98	1.36	1.19
	MAE	0.24	0.23	0.22	0.35	0.38	0.39	0.71	0.79	0.71
Media	RMSE	2.27	1.97	1.74	3.22	2.58	2.57	7.91	6.96	6.47
	MAE	1.29	1.20	1.04	2.00	1.84	1.63	4.75	4.16	4.10
Retail	RMSE	1.01	0.79	0.60	1.25	1.07	1.11	1.88	1.71	6.70
	MAE	0.86	0.57	0.44	1.04	0.80	0.80	1.71	1.33	4.26
Technology	RMSE	2.66	2.08	0.87	4.58	0.47	0.52	6.73	2.56	6.97
	MAE	1.19	1.09	0.45	1.95	0.35	0.37	3.08	1.50	2.99
Transportation	RMSE	3.15	2.26	2.00	4.52	2.76	3.55	6.97	6.33	9.55
	MAE	1.82	1.35	1.16	2.74	1.87	2.28	4.26	3.97	5.56
Utilities	RMSE	0.27	0.28	0.28	0.42	0.43	0.36	0.44	0.58	0.44
	MAE	0.22	0.20	0.21	0.34	0.33	0.26	0.35	0.43	0.34

Table 2.5: Default Rate Predictions RMSE and MAE

RMSE and MAE of predictions from 3 models: "Vague Prior" refers to the econometric specification (2.1)-(2.3) with priors loosely centered around the estimates from an equivalent model without frailty factors, "Tight Prior" refers to the econometric specification (2.1)-(2.3) with priors tightly centered around the estimates from an equivalent model without frailty factors, while a binomial-logit model without frailty factors and only Industrial Production as an explanatory variable serves as the "Base" case. Predictions are provided for each sector across 3 forecasting horizons: 3 months, 6 months and 1 year. All parameters are based on recursive samples with starting date January 1983 and ending date sequentially augmented from September 2006-September 2013 at an annual frequency. All parameter estimates are based on the MCMC algorithm described in section 2.4 using 50,000 MCMC draws. All RMSE and MAE are calculated based on rolling out-of-sample forecasts from September 2006-September 2013.

forecasting results.

Table 2.5 summarises the forecasting performance of the specifications with both vague and tight priors in terms of RMSE and the MAE. The RMSE is defined as:

$$RMSE_i^h = \sqrt{\frac{\sum_{t=1}^n (DR_{it}^h - PD_{it}^h)^2}{n}}, \quad (2.16)$$

where $DR_{it}^h = D_{it}^h/N_{it}$ is the forward-looking h -month default rate for sector i at month t . The MAE corresponds to:

$$MAE_i^h = \frac{\sum_{t=1}^n |DR_{it}^h - PD_{it}^h|}{n}. \quad (2.17)$$

Both RMSE and MAE give an indication of the deviation of the corporate default predictions from the observed default rate. While the MAE is a linear score that weights equally all the individual differences when calculating the average, the RMSE gives a relatively high weight to large errors. The RMSE is always larger or equal

to the MAE and the greater difference between the two, the greater the variance in the individual errors in the sample. Therefore, the RMSE is most useful when large errors are particularly undesirable. Both RMSE and MAE are calculated over the rolling out-of-sample forecasts from September 2006-September 2013. To put the results into perspective, I benchmark the model forecasts against binomial-logit regressions having Industrial Production as the single explanatory variable (from here on the "base model"). This specification corresponds to the econometric model (2.1)-(2.2) without including any frailty factors $F_t^{d,h}$ and substituting Industrial Production in the place of the macroeconomic factors F_t^m . The results from the "base model" are included in table 2.5.

The results from table 2.5 indicate a significant improvement in short horizon out-of-sample forecasting power when using dynamic factors F_t^m and frailty factors $F_t^{d,h}$, especially when interested at predicting default behaviour over 3-month horizons. For the 3-month forecasting horizon, the relative difference with respect to the "base model" specification is highest for Capital Industries, for which the RMSE improved by 78%/79% and MAE improved by 66%/67%(using tight and vague priors respectively). Other sectors with high RMSE and MAE difference for the 3-month horizon include Transportation (improvement in RMSE by 28%/37% for tight/vague priors respectively and reduction in MAE by 26%/36%) and Technology (reduction in RMSE by 22%/67% and in MAE by 8%/62%), with Media and Retail showing moderate improvement. On the other hand, the RMSE for Consumer Industries, Financials and Utilities slightly increases when using the full macroeconomic and frailty factors, but at the same time the MAE decreases. This indicates that the forecasts generated by the "base model" are less volatile and they do not exhibit extreme deviations, despite the fact that on average the "base model" does not forecast as well. The improvement in the forecasting performance for the 6-month horizon is similar to that for the 3-month horizon. The major difference is observed for the Technology sector for which the "base model" forecasts poorly; using the dynamic macroeconomic and frailty factor specification leads to a decrease in RMSE by 90%/89% for tight and vague priors respectively, while the MAE is improved by 82%/81%.

The benefits of using the dynamic macroeconomic and frailty factors is less clear for the 1-year forecasting horizon, as the use of the dynamic factors does not lead to

a substantial increase in performance as compared to the base case. Nevertheless, the use of dynamic and frailty factors does lead to an overall improvement of the RMSE (despite the "base model" performing marginally better for the Consumer Industries and Financials sectors) and better MAE for almost all individual sectors. An interesting observation is the extreme deviations of the full specification using vague priors for Consumer Industries, Retail and Transportation sectors. This indicates that allowing the frailty factors to dominate the regression specification (2.1)-(2.3) can provide better in-sample fit but it might also lead to problems out-of-sample for long forecasting horizons. As figure 2.5 shows, extreme forecasting errors when the frailty factors dominate the econometric equations are more likely during periods of stress; using the recent recession as an example, the mean reversion of the default rate time series seems to be much faster than what the historical dynamics imply, especially for the Consumer Industries, Retail, Technology and Transportation sectors.

The lack of clear performance gains at the 1-year horizon undermines the model's relevance in today's credit risk management (that is typically concerned about the 1-year+ horizons). This finding indicates that forecasting corporate defaults over long horizons can be quite difficult and the links to the macroeconomic environment are not that straightforward. Nevertheless, despite not substantially improving the forecasting performance, the model can still offer valuable insight into the main determinants of corporate defaults. Linking corporate defaults to a large set of macro-financial variables is much more informative than having links to a single business cycle variable. For example the sensitivity of a credit portfolio default rate on credit conditions and/or interest rate shocks can be of particular importance. Furthermore, the decomposition of the macroeconomic link for corporate defaults into major economic concepts allows financial institutions and investors alike to scenario test their credit portfolios using scenarios that cannot be captured by a single business cycle variable (such deflation, credit squeeze, monetary policy shock etc).

2.7 Conclusion

In this chapter I assess the forecasting power of the macroeconomic environment in predicting US corporate defaults. To address the issue of imperfect measurement of

macroeconomic concepts by single variables, I use a large panel of US macroeconomic time series. I summarise the macroeconomic movements into a small set of clearly distinguished factors, each describing a different side of the economic environment. The factors are identified and estimated via efficient MCMC techniques. The default predicting specification consists of the observed macroeconomic variation, summarised in the dynamic factors, and unobserved stochastic factors that correlate defaults through time, resulting in a non-Gaussian autoregressive random effects specification. I estimate the final non-Gaussian model, via a recently introduced and flexible MCMC sampling scheme based on data augmentation, that is first used in a credit risk context.

I show that forecasting with dynamic and unobserved frailty factors, can lead to an improvement of up to 90% in terms of RMSE and more than 80% in terms of MAE over base specifications that take into consideration only business cycle fluctuations. These substantial increases in forecasting power are observed for the Technology sector at the 6-month forecast horizon. Other sectors that the inclusion of dynamic and frailty factors leads to a dramatic increase in prediction accuracy include Capital Industries (improvement of 79% in terms of RMSE and 67% in terms of MAE when predicting defaults 3 months ahead) and Transportation (improvement of up to 67% in terms of RMSE and up to 62% in terms of MAE), while Media and Retail sectors also exhibit moderate improvements. Typically, the improvements in forecasting accuracy by using dynamic and frailty factors diminish when predicting default rates at the 1-year horizon. Finally, allowing the frailty factors to dominate the regression via the use of vague priors can lead to increased volatility in prediction accuracy, albeit providing better in-sample fit. Those results show that including macroeconomic conditions information can help in reliably predicting corporate default rates in the short term, while the use of frailty factors should be closely monitored when forecasting defaults over long time horizons.

The methodology presented in this chapter can be extended in a number of ways. A natural extension is to use more granular default time series, especially including a differentiation by credit quality. Nevertheless, using a credit grade split that is too granular can lead to low default occurrence and excessively volatile default time series that could render the forecasting power of the resulting model relatively poor. Furthermore, firm specific accounting/market data can be used; such a choice would increase dramatically the dimensionality of the econometric specification and the estimation of

the model can quickly turn unfeasible. Finally, more aspects of credit risk can be jointly modelled, along the lines of the work of Creal et al. (2013) that additionally include rating migrations and recoveries in a model of joint movements across corporate defaults and macroeconomic variables. Despite the attractiveness of joint modelling of all aspects of credit risk, forecasting rating migrations and recoveries out-of-sample is not straightforward. Rating agencies are known to follow a Through-the-Cycle approach to assigning ratings and therefore it is not entirely clear how much of the credit cycle Point-in-Time information they use in their rating process. Additionally, realised recovery rates are typically calculated over a lengthy emergence periods that often causes the recovery amount to be a function of different phases of the credit cycle.

Chapter 2 Appendix

2.8 Dynamic Factor Loadings

Description	Code	Tr. Code	F_t^{bs}	F_t^{cd}	F_t^{yc}	F_t^{cr}	F_t^{ar}	F_t^{av}
Production								
Industrial Production:								
-Index	INDPRO	2	0.115	0	0	0	0	0
-Manufacturing (NAICS)	IPMAN	2	0.114	0	0	0	0	0
-Durable Manufacturing (NAICS)	IPDMAN	2	0.113	0	0	0	0	0
-Consumer Goods	IPCONGD	2	0.098	0	0	0	0	0
-Durable Consumer Goods	IPDCONGD	2	0.086	0	0	0	0	0
-Nondurable Consumer Goods	IPNCONGD	2	0.078	0	0	0	0	0
-Business Equipment	IPBUSEQ	2	0.116	0	0	0	0	0
-Final Products (Market Group)	IPFINAL	2	0.119	0	0	0	0	0
-Materials	IPMAT	2	0.106	0	0	0	0	0
-Durable Materials	IPDMAT	2	0.107	0	0	0	0	0
-nondurable Materials	IPNMAT	2	0.077	0	0	0	0	0
ISM Manufacturing:								
-PMI Composite Index	NAPM	2	0.065	0	0	0	0	0
-Production Index	NAPMPI	2	0.049	0	0	0	0	0
Capacity Utilization:								
-Total Industry	TCU	2	0.104	0	0	0	0	0
-Manufacturing (NAICS)	MCUMFN	2	0.105	0	0	0	0	0
Income								
Real Personal Income:								
-Disposable	RPI	2	0.107	0	0	0	0	0
-Excluding current transfer receipts	DSPIC96	2	0.077	0	0	0	0	0
	W875RX1	2	0.115	0	0	0	0	0
Consumption								
Personal Consumption Expenditures*:								
-Durable Goods*	PCE	2	0.090	0	0	0	0	0
-Nondurable Goods*	PCEDG	2	0.068	0	0	0	0	0
-Services*	PCEND	2	0.074	0	0	0	0	0
	PCES	2	0.054	0	0	0	0	0
Unemployment								
Civilian Unemployment Rate								
Civilian Employment	UNRATE	1	-0.054	0	0	0	0	0
Employment Level - Nonagriculture	CE16OV	2	0.125	0	0	0	0	0
All Employees:	LNS12035019	2	0.124	0	0	0	0	0
-Total Private Industries	USPRIV	2	0.130	0	0	0	0	0
-Construction	USCONS	2	0.119	0	0	0	0	0
-Durable goods	DMANEMP	2	0.124	0	0	0	0	0
-Nondurable goods	NDMANEMP	2	0.110	0	0	0	0	0
-Financial Activities	USFIRE	2	0.090	0	0	0	0	0
-Goods-Producing Industries	USGOOD	2	0.128	0	0	0	0	0
-Government	USGOVT	2	0.029	0	0	0	0	0
-Manufacturing	MANEMP	2	0.124	0	0	0	0	0
-Mining and logging	USMINE	2	0.032	0	0	0	0	0
-Retail Trade	USTRADE	2	0.120	0	0	0	0	0
-Service-Providing Industries	SRVPRD	2	0.124	0	0	0	0	0

-Trade, Transportation & Utilities	USTPU	2	0.127	0	0	0	0	0
-Wholesale Trade	USWTRADE	2	0.122	0	0	0	0	0
Mean Duration of Unemployment	UEMPMEAN	1	-0.012	0	0	0	0	0
% Unemployed:								
-<5 Weeks	LNS13008397	1	0.049	0	0	0	0	0
-5 to 14 Weeks	LNS13025701	1	0.007	0	0	0	0	0
-15 to 26 Weeks	LNS13025702	1	-0.086	0	0	0	0	0
->15 Weeks	LNS13008517	1	-0.039	0	0	0	0	0
Prices								
Producer Price Index:								
-All Commodities	PPIACO	2	0.044	0	0	0	0	0
-Crude Materials: Further Processing	PPICRM	2	0.044	0	0	0	0	0
-Finished Consumer Foods	PPIFCF	2	0.023	0	0	0	0	0
-Finished Consumer Goods	PPIFCG	2	0.024	0	0	0	0	0
-Finished Goods Excluding Foods	PPIFLF	2	0.020	0	0	0	0	0
-Finished Goods	PPIFGS	2	0.022	0	0	0	0	0
-Finished Goods: Capital Equipment	PPICFE	2	-0.015	0	0	0	0	0
-Fuels & Related Products & Power	PPIENG	2	0.030	0	0	0	0	0
-Industrial Commodities	PPIIDC	2	0.043	0	0	0	0	0
-Intermediate Materials	PPIITM	2	0.045	0	0	0	0	0
Consumer Price Index:								
-All Items	CPIAUCSL	2	0.035	0	0	0	0	0
-All Items Less Food	CPIULFSL	2	0.034	0	0	0	0	0
-All Items Less Food & Energy	CPILFESL	2	0.013	0	0	0	0	0
-All Items Less Food	CUUR0000SA0L2	2	0.032	0	0	0	0	0
-All items less medical care	CUSR0000SA0L5	2	0.037	0	0	0	0	0
-Apparel	CPIAPPSL	2	0.016	0	0	0	0	0
-Commodities	CUSR0000SAC	2	0.041	0	0	0	0	0
-Durables	CUSR0000SAD	2	0.028	0	0	0	0	0
-Energy services	CUSR0000SEHF	2	0.002	0	0	0	0	0
-Medical Care	CPIMEDSL	2	-0.012	0	0	0	0	0
-Services	CUSR0000SAS	2	0.010	0	0	0	0	0
-Transportation	CPITRNSL	2	0.043	0	0	0	0	0
Inventories and Orders								
ISM Manufacturing:								
-Inventories Index	NAPMII	1	0.077	0	0	0	0	0
-New Orders Index	NAPMNOI	1	0.043	0	0	0	0	0
-Supplier Deliveries Index	NAPMSDI	1	0.043	0	0	0	0	0
Housing Starts								
Total New Priv. Housing Units	HOUST	1	0.059	0	0	0	0	0
Midwest Census Region	HOUSTMW	1	0.047	0	0	0	0	0
South Census Region	HOUSTS	1	0.051	0	0	0	0	0
West Census Region	HOUSTW	1	0.061	0	0	0	0	0
Bank Lending								
Loans and Leases in Bank Credit	LOANS	2	0.071	0	0	0	0	0
Commercial and Industrial Loans	BUSLOANS	2	0.061	0	0	0	0	0
Real Estate Loans	REALLN	2	0.037	0	0	0	0	0
Consumer Loans	CONSUMER	2	0.029	0	0	0	0	0
Bank Credit	LOANINV	2	0.065	0	0	0	0	0
Interbank Loans	IBLACBM027SBOG2		0.035	0	0	0	0	0
Interest Rates								
Treasury Bill Rates:								
-3-Month	TB3MS	1	0.033	0.078	0	0	0	0
-6-Month	TB6MS	1	0.033	0.078	0	0	0	0
Treasury Constant Maturity Rates:								
-1-Year	GS1	1	0.032	0.078	0	0	0	0
-2-Year	GS2	1	0.030	0.079	0	0	0	0
-3-Year	GS3	1	0.027	0.079	0	0	0	0
-5-Year	GS5	1	0.023	0.079	0	0	0	0
-10-Year	GS10	1	0.019	0.078	0	0	0	0
Effective Federal Funds	FEDFUNDS	1	0.030	0.078	0	0	0	0
Bank Prime Loan	MPRIME	1	0.028	0.078	0	0	0	0
3-Month LIBOR USD	USD3MTD156N	1	0.038	0.089	0	0	0	0
Moody's Corporate Bond Yield:								
-Aaa	AAA	1	0.011	0.078	0	0	0	0
-Baa	BAA	1	-0.003	0.078	0	0	0	0

Yield Curve Spreads								
GS10-GS2	Calculated	1	-0.047	-0.033	0.168	0	0	0
GS10-TB3MS	Calculated	1	-0.035	0.011	0.241	0	0	0
GS2-TB3MS	Calculated	1	0.003	0.063	0.186	0	0	0
USD3MTD156N-TB3MS	Calculated	1	0.003	0.049	-0.014	0	0	0
Credit Spreads								
AAA-GS10	Calculated	1	-0.056	-0.115	-0.055	0.116	0	0
BAA-GS10	Calculated	1	-0.084	-0.107	-0.042	0.178	0	0
Equity Market								
S&P 500	SP500	2	0.022	0	0	0	0.282	0
Dow Jones Industrial Average	DJIA	2	0.016	0	0	0	0.285	0
SP500 Volatility	Calculated	1	-0.193	0	0	0	-0.318	0.680
DJIA Volatility	Calculated	1	-0.190	0	0	0	-0.319	0.694
VIX S&P 500	VIXCLS	1	-0.183	0	0	0	-0.288	0.597
VIX Dow Jones Industrial	VXDCLS	1	-0.152	0	0	0	-0.260	0.528

Table 2.6: Estimated Dynamic Factor Loadings

Parameter estimates for macroeconomic input series and the dynamic factor loadings $\mathbf{\Lambda}$ in the econometric specification (2.5)-(2.6), based on the period January 1983-August 2014. All estimates are based on the MCMC algorithm described in section 2.4 using 50,000 MCMC draws. The input macroeconomic variables are either used untransformed ("Tr. Code"=1) or log Year-on-Year changes ("Tr. Code"=2).

Chapter 3

Macroeconomic Shocks and Credit Risk: Monetary Policy and Real Economy Effects on Corporate Defaults and Recoveries

3.1 Introduction

The 2008 financial crisis highlighted the central role that credit risk plays when assessing business cycle fluctuations, the strength of the financial sector, and the impact of policy actions. Periods of economic growth tend to lead to a relaxation of lending criteria and over-indebtedness, as both corporates and consumers seek additional return by leveraging. Over-leveraging, especially when combined with a slowdown in economic activity, usually leads to periods of increased default rates, deterioration of the banking sector's balance sheet and tightening of lending criteria. These business and leverage cycle fluctuations emphasise that credit risk events are far from being isolated and independent incidents, exhibiting strong systematic movements, clearly linked to macro-financial conditions. This strong dependence of credit risk on the economic envi-

ronment, has led regulators globally to adopt extreme but plausible economic scenarios to assess the solvency of the banking sector and identify the need for additional capital requirements.

Based on the strong linkages between macroeconomic and credit conditions, I make two major contributions in this chapter. First, I present a new semi-structural identification scheme to disentangle the macro-financial shocks, applied on a small-scale 5 variable VAR model. The VAR setup is inspired by the credit channel/financial accelerator literature and models output, inflation, leverage, asset value, and interest rate variables. The identification scheme comprises of a combination of short and long-run restrictions on the dynamic multipliers, and enables the identification of the 5 orthogonal shocks as aggregate supply, aggregate demand, corporate balance sheet, asset price and monetary policy. Second, I trace the impact of each structural shock on corporate bond defaults and recoveries, using a non-Gaussian specification that is able to capture both macroeconomic effects and credit market specific unobserved correlations. Finally, I undertake a detailed impulse response analysis using a flexible Bayesian MCMC estimation framework that provides simulated draws from the full conditional distribution of the impulse response functions.

The macroeconomic VAR model and the identification assumptions use elements of the financial accelerator and credit transmission channel literature. The advocates of the financial accelerator and the credit channel as shock transmission mechanisms stipulate that endogenous changes in external finance premium amplify the effect of real economy and monetary policy shocks. I follow Bernanke and Gertler (1989), Bernanke et al. (1996), and Kiyotaki and Moore (1997) and I assume that corporate financing is linked to the quality of the corporate balance sheet and the value of the available collateral. I, therefore, include in the VAR model measures of corporate leverage and equity returns to isolate balance sheet and asset price/collateral shocks. To separate the real economy effects into aggregate demand and supply, I follow Blanchard and Quah (1989) and I restrict the aggregate demand shock to only have a transitory effect on output. Finally, I adopt the idea of money neutrality in the long-run and I build upon the analysis of Bjørnland and Leitimo (2009) to assume that monetary policy shocks do not have long-run effects on real output and real equity prices.

The choice of variables in the VAR model and the identification scheme used,

produce intuitive responses for all macroeconomic variables to each of the 5 structural shocks. Positive aggregate supply shocks that increase quarterly real output by 0.4%, lead to a reduction of 0.3% in price inflation and a lower level of interest rates by 0.25pp. On the other hand, aggregate demand shocks that increase output by 0.6%, lead to inflationary pressures that ultimately push interest rates up, to peak at a level 0.75pp higher than the long-run average. When examining the impact of balance sheet, asset price, and monetary policy shocks, I provide robust evidence for the presence of a working capital/cost channel transmission mechanism, along the lines of Barth and Ramey (2002), Chowdhury et al. (2006), and Ravenna and Walsh (2006). This transmission mechanism implies that corporates are able to partially pass on increases in their cost of debt to consumers. The presence of a cost channel causes some inflationary pressures when corporates struggle to attract funding due to balance sheet shocks, or face increasing yields in their debt as a result of monetary policy shocks. Following adverse balance sheet shocks, output drops by approximately 0.1%, while price inflation rises by 0.05% and interest rates increase by 0.15pp, largely due to the effect of the cost channel transmission mechanism. Contractionary policy shocks linked to an increase in the FED Funds rate by 0.6pp, take approximately 1 year to adversely affect output (0.15% quarterly decline), and, despite the initial drop in inflation by 0.15%, the cost channel creates some minor inflationary pressures in 3-4 quarters. Finally, the cost channel drives the responses of corporate leverage, inflation and interest rates to asset price shocks; following an increase of 5% in real equity prices, interest rates increase sharply and persistently by 0.4pp, leading to a 0.1% higher corporate leverage level and 0.05% positive rate of inflation.

Consistent with Koopman et al. (2012), I find that the macroeconomic environment plays a significant role in explaining default rate dynamics. The FEVD suggests that the macroeconomic shocks account for approximately 30% of the 1 quarter forecast variance and approximately 45% of the long-run forecast variance. The remaining proportion of the forecast error variance is attributed to unobserved credit specific factors. Following Das et al. (2007), McNeil and Wendin (2007), Koopman and Lucas (2008), Duffie et al. (2009), and Koopman et al. (2011), I interpret the latent credit factors as proxies of frailty in the corporate environment. Quantifying the impact of the 5 structural shocks on credit measures of default likelihood, I report a sharp increase of the cyclical sector default rates by 5-10 basis points following a balance sheet shock, and a drop of similar

magnitude following aggregate supply/demand shocks. While the level of default rates remains lower than the long-run average following an aggregate supply type shock, an aggregate demand shock leads to a higher interest rate level, and that causes default rates to increase in the mid to long-run. I find the effect of asset price shocks to be similar to the effect of aggregate demand shocks, for the default rates in most sectors; main differences are observed for sectors strongly dependent on the equity market, such as Technology firms. Finally, monetary policy tightening shocks lead to a 5-10 basis points increase in default rates, approximately 1 year after the initial shock.

The impact of the 5 structural shocks on the fractional recovery of face value is much less pronounced. The forecast error variance analysis indicates that macroeconomic shocks account for only 24% of the short-run and 34% of the long-run forecast variance, with recovery rates mainly affected by asset price shocks in the short-run (irrespective of seniority, asset price shocks lead to short-run increases in recoveries rates of 3%, with reversion to the long -run average taking place 1 year after the initial shock). Secured debt recoveries strongly depend on aggregate supply/demand shocks, particularly in the long-run; the forecast error variance decomposition over a 5 year period indicates that, depending on the sector, the aggregate supply shocks accounts for 4%-18% of the variance, while the aggregate demand shocks account for 10%-38%. Finally, balance sheet and monetary policy shocks (both leading to increased cost of debt and discounting of future cash flows for the calculation of recoveries) both lead to an initial drop in recovery rates of approximately 1% on average, while the effect quickly dies out, typically less than a year after the initial shock.

The remainder of this chapter is organised as follows. Section 3.2 gives a high level summary of the relevant literature and highlights the gap that this chapter addresses. Section 3.3 presents the econometric models used in this chapter, and covers the 3 sub-components of my analysis: the corporate default econometric equations in sub-section 3.3.1, the recovery rate model in sub-section 3.3.2, and the overarching macroeconomic VAR specification, including the identification of the 5 structural shocks, in sub-section 3.3.3. Details on the estimation of the various econometric equations are provided in sub-section 3.3.4. Section 3.4 summarises the macroeconomic, corporate default, and corporate debt recovery data used in the analysis. Section 3.5 presents the findings: sub-section 3.5.1 focuses on the impulse response functions and variance decomposition

of the macroeconomic VAR model, while sub-section 3.5.2 summarises the parameter estimates and the responses of defaults and recoveries to each of the 5 structural shocks. Finally, section 3.6 concludes.

3.2 Relevant Literature

Linking credit cycles to macroeconomic fluctuations has attracted a lot of attention over the recent years, see Koopman et al. (2009), and also Pesaran et al. (2006) for a global perspective. While correlating corporate defaults with the macroeconomic environment has received the necessary attention, very few papers deal with the fundamental economic shocks that impact corporate default occurrence. The works of Boivin et al. (2013), and Creal et al. (2013) attempt to give an answer to this question, the former using a large scale macro-financial dynamic factor model and the latter using a smaller scale dynamic factor model but with richer dynamics. Nevertheless, both papers follow a rather simplistic approach to identification, that relies on sequential impact of each economic shock to the observed macroeconomic quantities. In this chapter I present a new semi-structural identification scheme that uses a combination of economic theory founded short- and long-run restrictions on the residuals of a small-scale VAR, and therefore relaxes the reliance on Cholesky decomposition to disentangle the macro-financial shocks.

Empirical evidence suggests that, despite the explanatory power of observed macroeconomic variables in predicting the default behaviour of corporates, there is a significant amount of additional variation that cannot be directly linked to the macroeconomic environment. For a decomposition of the variation into its different sources, see Koopman et al. (2012). A number of papers try to explain the additional systematic correlation not captured by observed macroeconomic variables via the use of latent factors that reflect frailty in the corporate environment. Examples from the literature include, Das et al. (2007), McNeil and Wendin (2007), Koopman et al. (2008), Koopman and Lucas (2008), Duffie et al. (2009), and Koopman et al. (2011). Another strand of research interprets (part of) the additional correlation as steaming from the contagion/domino effect a default of single firm can have on the rest of the corporate environment, see Giesecke (2004), Jorion and Zhang (2007), and Lando and Nielsen (2010). In practice

it is very difficult to robustly decompose default from contagion effects. Especially for contagion, detailed balance sheet and trade data are necessary to capture business links and disentangle the excess correlation from frailty.

I follow the frailty literature and I provide a decomposition into observed and unobserved fluctuations in corporate bond defaults. For the observed fluctuations in corporate defaults, I use the structurally identified macro-financial shocks and I provide full impulse response analysis and variance decompositions. At the same time, I improve on the work of Creal et al. (2013), by providing model consistent confidence bounds for the impulse response functions that allow the isolation of statistically significant effects.

In addition to the clear macroeconomic effects on corporate failures, there are numerous studies suggesting that recovery rates of defaulted debt are also driven by the state of the economy. In some of the early works, Frye (2000a) and Frye (2000b) use a simple structural model to link recoveries to defaults and ultimately the state of the economy, while Jarrow (2001) explicitly links equity prices and recovery rates, using a reduced-form approach. For the time period used, recovery rates are shown to be up to 25% lower during periods of stress. The strongly negative relationship between recoveries and defaults (which implies a positive relationship between recoveries and the macroeconomic conditions) is also confirmed by Hu and Perraudin (2002), Altman et al. (2005), and Altman and Kalotay (2014). Two recent studies, Acharya et al. (2007) and Jankowitsch et al. (2014) provide further evidence of the strong correlation between defaults and recoveries, that is statistically significant even after accounting for firm and debt issuance characteristics. Furthermore, Jankowitsch et al. (2014) show that, in addition to firm/debt characteristics and distress information at market and industry level, the FED Funds rate provides a statistically significant determinant of corporate bond recovery rates.

Using the identified structural shocks, I extend the above mentioned literature by providing the first link of corporate recoveries to macroeconomic fundamentals. In addition to the explicit link to macroeconomic conditions, I also include unobserved systematic factors that help to better fit the historical time series of recovery rates. Finally, I provide a detailed impulse response and variance decomposition for the two sourced of systemic variation in recovery rates.

3.3 Econometric Framework

3.3.1 The Corporate Default Model

Following the popular firm value approach to credit risk, as introduced by Merton (1974) and later extended by Black and Cox (1976), a default occurs when the value V_{it} of a firm i at time t becomes negative. Assuming that a firm's value is the difference between assets A_{it} and liabilities B_{it} , default can be re-expressed as an event taking place when $A_{it} < B_{it}$ or, equivalently, $\log(A_{it}) < \log(B_{it})$. The focus of this chapter is not the firm specific default processes, but rather the systematic movements of credit events. Therefore, I impose suitable assumptions and express firm specific probabilities to default, PD_{it} , as a function of the systematic, cyclical drivers, \mathbf{F}_t .

For the stochastic evolution of A_{it} , I assume that its logarithmic return, A_{it}^* , can be decomposed into a systematic component $\mathbf{F}_t^{A^*}$, that affects the asset values of all firms at a given time, and an idiosyncratic component ϵ_{it}^* , that is firm specific:

$$A_{it}^* = \log(A_{it}) - \log(A_{it-1}) = \boldsymbol{\omega}_i \mathbf{F}_t^{A^*} + \varpi_i \epsilon_{it}^*, \quad \epsilon_{it}^* \sim N(0, 1), \quad (3.1)$$

where $\boldsymbol{\omega}_i$ is the sensitivity to the systematic factors, and ϖ_i is the sensitivity to the idiosyncratic factor. The systematic factor $\mathbf{F}_t^{A^*}$ in (3.1) is a linear combination of the economy-wide fluctuations, \mathbf{F}_t , and unobserved credit cycle specific movements:

$$\mathbf{F}_t^{A^*} = \boldsymbol{\beta}^{A^*}(L) \mathbf{F}_t + \mathbf{b}^{A^*}(L) \mathbf{f}_t^{A^*} \quad (3.2)$$

where $\mathbf{f}_t^{A^*}$ is the set of latent credit cycle factors, and $\boldsymbol{\beta}^{A^*}(L)$ and $\mathbf{b}^{A^*}(L)$ are vector lag polynomials for the sensitivities of $\mathbf{F}_t^{A^*}$ to the economy-wide systematic factors \mathbf{F}_t and frailty factors $\mathbf{f}_t^{A^*}$ respectively. Both \mathbf{F}_t and $\mathbf{f}_t^{A^*}$ are assumed to be Gaussian processes. By construction, $\mathbf{f}_t^{A^*}$ captures additional correlation in asset values that macroeconomic and financial aggregates cannot explain. This additional correlation in corporate bond default rates is strongly supported by empirical evidence and can be attributed to omitted variables that drive the co-movements in corporate balance sheets. This source of correlation is usually referred to as "frailty", see Das et al. (2007), McNeil and Wendin (2007), Duffie et al. (2009), and Koopman et al. (2011). Additional correlation can also result from the contagion/domino effects that single defaults have

on the rest of the corporate environment. Studies exploring the contagion effects as a source of default correlation include Giesecke (2004), Jorion and Zhang (2007), and Lando and Nielsen (2010). In practice it is very difficult to robustly decompose frailty from contagion effects. Especially for contagion, detailed balance sheet and trade data are necessary to capture business links and disentangle the excess correlation from frailty. Since this is not the purpose of the current chapter and such data are not available, I choose to interpret all additional correlation as frailty.

I do not make explicit assumptions about the stochastic properties of the liabilities B_{it} , but I rather focus on their movement relative to the asset value. Without loss of generality, I restrict the dynamics of the leverage ratio $\text{Lev}_{it} = \log(\frac{B_{it}}{A_{it-1}})$ to be a deterministic function of, possibly lagged, economic factors:

$$\text{Lev}_{it} = \log\left(\frac{B_{it}}{A_{it-1}}\right) = \mu_i^{\text{Lev}} + \boldsymbol{\beta}^{\text{Lev}}(L)\mathbf{F}_t \quad (3.3)$$

where μ_i^{Lev} is the long-run average of the leverage ratio, \mathbf{F}_t is the set of systematic factors, and $\boldsymbol{\beta}^{\text{Lev}}(L)$ is a vector lag polynomial for the sensitivities to systematic factors \mathbf{F}_t . $\boldsymbol{\beta}^{\text{Lev}}(L)$ allows additional flexibility in modelling the impact that the economic environment has on a firm's liabilities. In the absence of lags in $\boldsymbol{\beta}^{\text{Lev}}(L)$, the economic environment has only a contemporaneous effect on Lev_{it} , and this assumption might be overly restrictive. As mentioned earlier, \mathbf{F}_t is a set of cyclical, stationary systematic factors. By making leverage a function of stationary economic drivers, (3.3) implicitly assumes that corporate leverage mean reverts to the target leverage ratio of a given firm. This assumption is supported by both theoretical and empirical evidence in the literature, see Collin-Dufresne and Goldstein (2001), Flannery and Rangan (2006), Geanakoplos (2010), and Flannery et al. (2012). In fact, the deterministic function (3.3) can be seen as a discretised version of the stochastic process for corporate debt used by Collin-Dufresne and Goldstein (2001). In their chapter, Collin-Dufresne and Goldstein (2001) model log debt k_t as $dk_t = \lambda(y_t - v - k_t)dt$, where y_t is the log asset value level. In my specification, I assume a value of 1 for λ (which implies instant adjustment to the constant leverage ratio v) and I add residual cyclical movements to the leverage ratio, in the form of the systematic factors \mathbf{F}_t . Equation (3.3) can be extended to include a frailty term (and therefore become stochastic). For the benefit of the mathematical derivation of the PD equation below and given that the empirical analysis is based on

a reduced form model, I refrain from including a frailty term. Nevertheless, the nature of the reduced form results that follow would not change under the presence of a frailty term for the leverage process.

I define PD_{it} as the probability of a firm i defaulting at time t , providing that a default did not occur before t . I further condition PD_{it} on the filtration $\mathcal{F}_t = \{\mathbf{F}_1, \dots, \mathbf{F}_t, \mathbf{f}_1^{A^*}, \dots, \mathbf{f}_t^{A^*}\}$. Using the asset value dynamics from (3.1) and (3.2), and the leverage dynamics from (3.3), PD_{it} is given by:

$$\begin{aligned}
PD_{it} &= P(\log(A_{it}) < \log(B_{it}) | \mathcal{F}_t, A_{ij} \geq B_{ij}, \forall j = 1, \dots, t-1) = \\
&= P(A_{it}^* < \log(B_{it}) - \log(A_{it-1}) | \cdot) = \\
&= P(A_{it}^* < \text{Lev}_{it} | \cdot) = \\
&= P\left(\epsilon_{it}^* < \frac{\mu_i^{\text{Lev}} + \boldsymbol{\beta}^{\text{Lev}}(L)\mathbf{F}_t - \boldsymbol{\omega}_i(\boldsymbol{\beta}^{A^*}(L)\mathbf{F}_t + \mathbf{b}^{A^*}(L)\mathbf{f}_t^{A^*})}{\varpi} \mid \cdot\right) = \\
&= P(\epsilon_{it}^* < \mu_i + \boldsymbol{\beta}^d(L)\mathbf{F}_t + \boldsymbol{\rho}_i^d \mathbf{f}_t^{A^*} | \cdot) \\
&= \Phi\left(\mu_i + \boldsymbol{\beta}^d(L)\mathbf{F}_t + \boldsymbol{\rho}_i^d \mathbf{f}_t^{A^*}\right), \tag{3.4}
\end{aligned}$$

where $\mu_i = \frac{\mu_i^{\text{Lev}}}{\varpi}$ is the reduced form mean of $\Phi^{-1}(PD_{it})$, $\boldsymbol{\beta}^d(L) = \frac{\boldsymbol{\beta}^{\text{Lev}}(L) - \boldsymbol{\omega}_i \boldsymbol{\beta}^{A^*}(L)}{\varpi}$ is the reduced form sensitivity of the default probability to the dynamic economic factors, $\boldsymbol{\rho}_i^d = -\frac{\boldsymbol{\omega}_i \mathbf{b}^{A^*}(L)}{\varpi}$ is the reduced form sensitivity to the frailty factor, and $\Phi(\cdot)$, $\Phi^{-1}(\cdot)$ are the cumulative and inverse cumulative functions of the standard normal distribution. Since the objective is to forecast the default process, I work with the reduced form (3.4) without separating $\boldsymbol{\beta}^d(L)$ to the individual effects coming from the asset and leverage processes in (3.2) and (3.3).

The credit risk framework presented so far can be re-expressed as a multi-factor version of popular Merton based portfolio credit models, such as Vasicek (2002) and CreditMetrics (2007). Working with standard normal versions of the stochastic factors in (3.1) and (3.2), $\tilde{\mathbf{F}}_t^{A^*}$, $\tilde{\mathbf{F}}_t$, and $\tilde{\mathbf{f}}_t^{A^*}$, the asset value process can be re-written as:

$$\begin{aligned}
A_{it}^* &= \log(A_{it}) - \log(A_{it-1}) = \omega_i \left(\boldsymbol{\beta}_i^{A^*} \mathbf{F}_t + \mathbf{b}_i^{A^*} \mathbf{f}_t^{A^*} \right) + \varpi_i \epsilon_{it}^* \\
&= \sigma_i \left[\sqrt{r_i} \left(\tilde{\boldsymbol{\beta}}_i^{A^*} \tilde{\mathbf{F}}_t + \tilde{\mathbf{b}}_i^{A^*} \tilde{\mathbf{f}}_t^{A^*} \right) + \sqrt{1-r_i} \epsilon_{it}^* \right], \quad \tilde{\boldsymbol{\beta}}_i^{A^*} \left(\tilde{\boldsymbol{\beta}}_i^{A^*} \right)' + \tilde{\mathbf{b}}_i^{A^*} \left(\tilde{\mathbf{b}}_i^{A^*} \right)' = 1,
\end{aligned}$$

where σ_i is the volatility of the asset value process A_{it}^* and the restriction $\tilde{\boldsymbol{\beta}}_i^{A^*} \left(\tilde{\boldsymbol{\beta}}_i^{A^*} \right)' +$

$\tilde{\mathbf{b}}_i^{A^*} (\tilde{\mathbf{b}}_i^{A^*})' = 1$ is imposed for $\tilde{\beta}_i^{A^*} \tilde{\mathbf{F}}_t + \tilde{\mathbf{b}}_i^{A^*} \tilde{\mathbf{f}}_t^{A^*}$ to have unit variance. The pairwise asset correlation in this multi-factor model equals to

$$\text{Asset Correlation} = \sqrt{r_i} \sqrt{r_{i^*}} \left(\tilde{\beta}_i^{A^*} (\tilde{\beta}_{i^*}^{A^*})' + \tilde{\mathbf{b}}_i^{A^*} (\tilde{\mathbf{b}}_{i^*}^{A^*})' \right),$$

for two distinct companies i and i^* . For a homogeneous portfolio, consisting of companies with equal $r_i = r$, $\tilde{\beta}_i^{A^*} = \tilde{\beta}^{A^*}$, and $\tilde{\mathbf{b}}_i^{A^*} = \tilde{\mathbf{b}}^{A^*} \forall i$, the above expression for the asset correlation reduces to

$$\text{Asset Correlation} = r.$$

For the purpose of this chapter, the focus is on the systematic movements of the default probabilities and, therefore, I do not use company specific information. The chosen aggregation level corresponds to the intersection of industry sectors $j \in \mathcal{S}$ and rating grades $k \in \mathcal{G}$. To estimate the reduced form parameters in (3.4), I use quarterly observed default counts, denoted as D_{jkt} , for each sector j , rating grade k , and time t . Conditioning on the available information set \mathcal{F}_t at time t , I assume the default counts to be independent and binomially distributed:

$$D_{jkt} | \mathcal{F}_t \sim \text{Binom}(PD_{jkt}, N_{jkt}), \quad (3.5)$$

where N_{jkt} is the number of active companies at grade j and sector k at the beginning of quarter t , and PD_{jkt} is the sector/grade specific probability of default for quarter t . PD_{jkt} are linked to the systematic factors \mathbf{F}_t and the industry specific frailty factors f_{jt}^d via the probit function:

$$\Phi^{-1}(PD_{jkt}) = \alpha_{jk} + \beta_j^d(L) \mathbf{F}_t + \rho_{jk} f_{jt}^d, \quad (3.6)$$

where α_{jk} is the industry sector/rating grade specific intercept, while $\beta_j^d(L)$ are the industry-wide sensitivities to the systematic factors \mathbf{F}_t and ρ_{jk}^d are the industry/rating level sensitivities to the industry specific frailty factors f_{jt}^d . The granularity of the specification could be further increased, possibly allowing for rating specific sensitivities to the systematic factors \mathbf{F}_t or rating specific frailty factors; nevertheless, due to the infrequent nature of defaults, it is very difficult to find statistical significance when estimating sensitivities at the sector/rating level. It is worth noting that, unlike the analysis of

Chapter 2 that uses a logit specification to model corporate defaults, (3.6) describes a probit specification. The probit specification for this chapter is chosen as the natural result of the well-known Merton model that helps to provide a more structural link of corporate defaults to macroeconomic and frailty factors. On the other hand, the choice of the logit specification in Chapter 2 is dictated by the MCMC algorithm used to estimate the large scale model in a very efficient manner.

The composition and the stochastic evolution of the systematic factors \mathbf{F}_t is described in great detail in section 3.3.3. The last piece in the default process specification is the stochastic process for the frailty factors f_{jt}^d . To avoid over-parameterising the model, I assume that all f_{jt}^d factors are mutually independent and each one follows an AR(1) process. These simplifications lead to:

$$f_{jt}^d = \vartheta_j^d \cdot f_{jt-1}^d + e_{jt}^d, \quad e_{jt}^d \sim N(0, 1), \quad \forall j \in \mathcal{S}, \quad (3.7)$$

where ϑ_j^d is the AR(1) coefficient and $-1 < \vartheta_j^d < 1$, $\forall j \in \mathcal{S}$ to ensure stationarity. The variances of the errors e_{jt}^d are set to unity since they cannot be jointly identified with the ρ_{jk}^d sensitivities in (3.6). Each of the factors f_{jt}^d is initialized from its unconditional distribution:

$$f_{j0}^d \sim N\left(0, \frac{1}{1 - (\vartheta_j^d)^2}\right). \quad (3.8)$$

3.3.2 The Recovery Model for Defaulted Debt

In the event a corporate defaults on its debt, the recovery rate is a function of the residual value of the firm's assets. For simplicity, I assume the (partial) recovery of the debt's face value takes place a quarter after default (the exposition can be generalised to accommodate a longer recovery period). Using the notation of section 3.3.1 for the value of a firm's assets, A_{it} , and liabilities, B_{it} , the recovery rate RR_{it} can be expressed as $RR_{it} = \frac{A_{it+1}}{B_{it}} | A_{it+1} < B_{it}$. Taking logs and using the assumed stochastic processes for asset value and leverage in (3.1), (3.2), and (3.3), leads to the econometric specification for the recovery rate:

$$\begin{aligned}
\log(RR_{it}) &= \log(A_{it+1}) - \log(B_{it}) | \mathcal{F}_t, \quad \forall i : A_{it} < B_{it} = \\
&= \log(A_{it}) + \boldsymbol{\omega}_i \mathbf{F}_{t+1}^{A^*} + \varpi_i \epsilon_{it+1}^* - \log(A_{it-1}) - \mu_i^{\text{Lev}} - \boldsymbol{\beta}^{\text{Lev}}(L) \mathbf{F}_t = \\
&= A_{it}^* + \boldsymbol{\omega}_i \mathbf{F}_{t+1}^{A^*} + \varpi_i \epsilon_{it+1}^* - \mu_i^{\text{Lev}} - \boldsymbol{\beta}^{\text{Lev}}(L) \mathbf{F}_t = \\
&= \boldsymbol{\omega}_i \mathbf{F}_t^{A^*} + \varpi_i \epsilon_{it}^* + \boldsymbol{\omega}_i \mathbf{F}_{t+1}^{A^*} + \varpi_i \epsilon_{it+1}^* - \mu_i^{\text{Lev}} - \boldsymbol{\beta}^{\text{Lev}}(L) \mathbf{F}_t = \\
&= \boldsymbol{\omega}_i (\boldsymbol{\beta}^{A^*}(L) \mathbf{F}_t + \boldsymbol{\beta}^{A^*}(L) \mathbf{F}_{t+1} + \mathbf{b}^{rr}(L) \mathbf{f}_t^{rr}) - \mu_i^{\text{Lev}} - \boldsymbol{\beta}^{\text{Lev}}(L) \mathbf{F}_t + \varpi_i \epsilon_{it}^{**} = \\
&= \mu_i^{rr} + \boldsymbol{\beta}^{rr}(L) \mathbf{F}_t^{rr} + \boldsymbol{\rho}_i^{rr} \mathbf{f}_t^{rr} + \epsilon_{it}^{rr}, \quad \epsilon_{it}^{rr} \sim N(0, \varpi_i^2) \tag{3.9}
\end{aligned}$$

where $\mu_i^{rr} = -\mu_i^{\text{Lev}}$ is the mean of the log recovery rate, $\mathbf{F}_t^{rr} = [\mathbf{F}_t, \mathbf{F}_{t+1}]'$ are the reduced form systematic factors for the recovery rate, $\boldsymbol{\beta}^{rr}(L) = [\boldsymbol{\omega}_i \boldsymbol{\beta}^{A^*}(L) - \boldsymbol{\beta}^{\text{Lev}}(L), \boldsymbol{\omega}_i \boldsymbol{\beta}^{A^*}(L)]$ are the reduced form sensitivities to the systematic factors \mathbf{F}_t^{rr} , $\mathbf{f}_t^{rr} = \mathbf{f}_t^{A^*} + \mathbf{f}_{t+1}^{A^*}$ are the reduced form latent factors driving the recoveries, $\boldsymbol{\rho}_i^{rr} = \boldsymbol{\omega}_i \mathbf{b}^{rr}(L)$ are the reduced form sensitivities to the frailty factors, and $\epsilon_{it}^{rr} = \varpi_i \epsilon_{it}^{**} = \varpi_i (\epsilon_{it}^* + \epsilon_{it+1}^*)$ are the reduced form residuals.

It is very apparent from (3.9) that if a firm's assets form the collateral in case of default on its debt, then there is a natural link between the probability to default and the recovery amount in case of default. Following the structural approach to recovery rate modelling, there are multiple studies that provide an explicit link between collateral and default. For some of the early structural based studies see Frye (2000b), and Pykhtin (2003), while for a good overview of both structural and reduced form approaches to modelling recovery rates, see Altman (2008). While it is appealing to jointly model defaults and recoveries, the purpose of this chapter is not to capture the tail loss in a portfolio of credit exposures (in which case the co-dependence of default and recovery processes would be crucial), but rather to capture the systematic effects of the macroeconomic environment on defaults and recoveries. Therefore, I simplify the econometric framework by treating recovery and default as independent processes, and I work with the purely reduced form version of (3.9).

For the analysis of recovery rates I do not include any bond or company specific information, as the focus of this chapter is the time series element of recoveries rather than the cross-section. I choose to model the systematic effects on recovery rates at the industry/seniority level. While the seniority is an obvious determinant of a bonds

recovery rate (so that bondholders of secured debt should expect to recover more than the holders of unsecured or subordinated debt), industry is a more ambiguous factor. Acharya et al. (2007) provide empirical evidence that support the usage of industry classification as an important determinant of debt recoveries; the evidence suggests that the industry level effect can be attributed to both the number of defaults in a given sector (economic downturn/distress effect) and the liquidity level of the assets that firms within a sector typically pledge as collateral. Finally, the inclusion of industry and seniority level information seems to be robust to the addition of bond and firm specific characteristics as well as macroeconomic variables, as evidenced by Jankowitsch et al. (2014).

For a given sector $j \in \mathcal{S}$ and seniority level $s \in \mathcal{C}$, the reduced form recovery rate equation for the i th default takes the form:

$$\log(RR_{it}) = \alpha_{js}^{rr} + \beta_s^{rr}(L)\mathbf{F}_t + \rho_{js}^{rr}f_{st}^{rr} + \epsilon_{it}^{rr}, \quad \epsilon_{it}^{rr} \sim N(0, (\sigma_s^{rr})^2) \quad (3.10)$$

where α_{js} are the industry and seniority level specific averages of the log recovery rate, $\beta_s(L)$ are the seniority level specific sensitivities to the systematic factors \mathbf{F}_t , and ρ_{js}^{rr} is the sector/seniority specific sensitivity to the frailty factors f_{st}^{rr} . Finally, for each default i , ϵ_{jst}^{rr} correspond to the model's residuals, that are normally distributed with a seniority specific constant variance $(\sigma_s^{rr})^2$. Equation (3.10) implies that seniority is chosen as the primary aggregation level for the various effects, with industry differentiation only captured via the intercepts α_{js}^{rr} and the sensitivities to the latent factors f_{st}^{rr} . This simplified specification is driven by the lack of adequate recovery data to robustly model the full cross-section of industry sectors and seniority level. Finally, similarly to the default process specification, I allow the frailty factors f_{st}^{rr} to follow independent AR(1) processes of the form:

$$f_{st}^{rr} = \vartheta_s^{rr} \cdot f_{st-1}^{rr} + e_{st}^{rr}, \quad e_{st}^{rr} \sim N(0, 1), \quad \forall s \in \mathcal{C}, \quad (3.11)$$

where ϑ_s^{rr} dictates the speed of the mean reversion for the AR(1) processes, the variance of the residuals e_s^{rr} is set to unity for identification purposes, and each of the seniority

specific frailty factors f_{st}^{rr} is initialised from the unconditional distribution:

$$f_{s0}^{rr} \sim N\left(0, \frac{1}{1 - (\vartheta_s^{rr})^2}\right). \quad (3.12)$$

3.3.3 The macroeconomic structural VAR model

For the 3 sources of systematic fluctuations in the broad macroeconomic and credit specific conditions, \mathbf{F}_t , \mathbf{f}_t^d and \mathbf{f}_t^{rr} , I impose an autoregressive Gaussian process of finite order:

$$\Gamma(L) \begin{bmatrix} \mathbf{F}_t \\ \mathbf{f}_t^d \\ \mathbf{f}_t^{rr} \end{bmatrix} = \gamma + u_t, \quad u_t \sim N(0, \Sigma_u), \quad (3.13)$$

where $\mathbf{f}_t^d = \text{vec}(f_{jt}^d)_{\{j:j \in \mathcal{G}\}}$ is the stacked vector of sector specific unobserved default factors, $\mathbf{f}_t^{rr} = \text{vec}(f_{st}^{rr})_{\{s:s \in \mathcal{C}\}}$ is the stacked vector of seniority specific latent recovery factors, γ is the vector of intercepts and u_t is the vector of residuals. Unlike the unobserved frailty factors \mathbf{f}_t^d and \mathbf{f}_t^{rr} , the systematic factors \mathbf{F}_t correspond to a fully observable set of macro-financial variables following a finite order VAR process. More specifically, I assume the 3 constituents of the filtration process \mathcal{F}_t in (3.6) and (3.10) to be orthogonal. This assumption helps to keep the model tractable and not very cumbersome to estimate. This simplifying assumption implies a block diagonal structure for both $\Gamma(L)$ and Σ_u in (3.13):

$$\Gamma(L) = \begin{bmatrix} A(L) & 0 & 0 \\ 0 & \Theta^d(L) & 0 \\ 0 & 0 & \Theta^{rr}(L) \end{bmatrix}, \quad \Sigma_u = \begin{bmatrix} \Sigma_e^m & 0 & 0 \\ 0 & \Sigma_e^d & 0 \\ 0 & 0 & \Sigma_e^{rr} \end{bmatrix}, \quad (3.14)$$

where $A(L)$, $\Theta^d(L)$, and $\Theta^{rr}(L)$ are the lag polynomial matrices for \mathbf{F}_t , \mathbf{f}_t^d and \mathbf{f}_t^{rr} respectively, while Σ_e^m , Σ_e^d , and Σ_e^{rr} correspond to the residual covariance matrices for the macroeconomic and default/recovery unobservable factors. The block diagonality assumption allows modelling separately the dynamics of the 3 processes, avoiding the cumbersome estimation of a high dimensional non-Gaussian dynamic model. Furthermore, it provides a natural interpretation of the unobserved factors \mathbf{f}_t^d and \mathbf{f}_t^{rr} as credit specific frailty, or, in other words, additional default and recovery rate clustering above and beyond what is captured by the macroeconomic variability \mathbf{F}_t . Such interpretation would be difficult if the two effects are not properly identified.

To model the macroeconomic environment, as captured by the evolution of the observed macroeconomic variables \mathbf{F}_t , I use a small scale, closed economy, VAR of 5 variables. A key element of the system of endogenous variables is that I follow Bernanke and Gertler (1989), Bernanke and Gertler (1995), and Bernanke et al. (1996) and assume that the balance sheet channel plays a major role in the monetary policy transmission process; at the same time, it also determines the availability of funding for the corporate sector. The 5 variables that I choose to model correspond to real GDP growth (y_t), price inflation (π_t), corporate leverage (l_t), real equity price inflation (eq_t), and interest rate (ir_t). The variables capture the main macroeconomic determinants of credit risk, such as business cycle fluctuations, strength of the corporate balance sheet, distance to default, value of posted collateral and cost of debt. More specifically, real GDP growth and price inflation are essential in identifying aggregate demand and supply shocks, while corporate leverage and real equity prices help to identify balance sheet and asset price shocks respectively. Government interest rates play a major role in isolating the monetary policy shock; nevertheless, as section 3.5.1 suggests, monetary policy shocks are also strongly driven by business cycle fluctuations.

For the $k = 5$ variable vector of observables $\mathbf{F}_t = [y_t, \pi_t, l_t, eq_t, ir_t]'$, the reduced form p -th order VAR takes the form:

$$A(L)\mathbf{F}_t = \nu + \varepsilon_t^m, \quad (3.15)$$

where $A(L)$ is the lag polynomial $A(L) = I_k - \sum_{i=1}^p A_i L^i$, ν is the vector of intercepts, and ε_t^m is the k -dimensional vector of reduced form residuals. I assume the residual vector is Gaussian white noise, implying that $E(\varepsilon_t^m) = 0$, $E(\varepsilon_t^m (\varepsilon_t^m)') = \Sigma_\varepsilon^m$, and $E(\varepsilon_t^m (\varepsilon_s^m)') = 0$ for $s \neq t$. I assume that the covariance matrix Σ_ε^m is not singular, and that the VAR is stable, and therefore the roots of $|A(L)|$ lie outside the unit circle. Using Wold's Decomposition Theorem, the VAR model (3.15) can be written in the structural moving average form:

$$\mathbf{F}_t = \mu^m + \sum_{i=0}^{\infty} \Phi_i \eta_{t-i}, \quad (3.16)$$

where $\mu^m = A(L)^{-1}\nu$ is the unconditional mean, Φ_i is the matrix of the dynamic multipliers, $\Psi_\infty = \sum_{i=0}^{\infty} \Phi_i$ is the long-run impact matrix and $\eta = [\eta_s, \eta_d, \eta_{bs}, \eta_{ap}, \eta_m]'$ is

the vector of white noise orthogonal structural residuals, reflecting aggregate supply, aggregate demand, corporate balance sheet, asset prices, and monetary policy respectively.

I identify the structural innovations η_t in (3.16) from the reduced form residuals ε_t^m in (3.15), by assuming that the reduced form residuals are linear functions of the structural innovations. For a k^2 -dimensional matrix B , the relationship between reduced form and structural innovations can be written as $\varepsilon_t^m = B\eta_t$, which implies $\Sigma_\varepsilon^m = B\Sigma_\eta B'$. Normalising the variance of the structural innovations to unity, $\Sigma_\eta = E(\eta_t\eta_t') = I_k$, leads to the first set of identifying restrictions:

$$\Sigma_\varepsilon^m = BB'. \quad (3.17)$$

Since Σ_ε^m is symmetric, (3.17) provides only $k(k+1)/2 = 15$ restrictions. To exactly identify all $k^2 = 25$ elements of B , $k(k-1)/2 = 10$ additional restrictions are needed. In principal, more than necessary restrictions can be imposed, but I choose a just-identified VAR model to avoid restrictions on the lag coefficients $A(L)$.

For the additional 10 restrictions, I make the following assumptions on the short and long-run behaviour of the impulse responses:

- An aggregate demand shock has only a transitory effect on output. This assumption is along the lines of Blanchard and Quah (1989) type of long-run restriction, and is used to distinguish an aggregate demand from an aggregate supply shock.
- Balance sheet shocks have a lagged effect on real activity. An increase of the external finance premium is likely to cause a reduction in investment, production, employment and prices. Nevertheless, theoretical arguments support the view that the financing constraints resulting from a deterioration in credit conditions will not immediately erode the firm's net worth and, therefore, the effects on economic activity will not be observed instantaneously. This assumption is consistent with the existing literature on Structural Vector Autoregression (SVAR) or FAVAR containing both real economy and credit condition variables, see Gilchrist et al. (2009), and Gilchrist and Zakrajšek (2012).
- Monetary policy and asset price changes do not have an instantaneous effect on

real economy. This is a typical assumption in macroeconomic VARs, whereby monetary policy and financial shocks are ordered after business cycle shocks in recursive identification schemes.

- Monetary policy and asset price shocks do not affect corporate leverage at the same quarter. As described in section 3.4, I use a forward looking book value measure of corporate leverage, defined as the ratio of corporate debt (measured at historical cost) to corporate profits. As neither determinant of corporate leverage is MtM, the absence of an instantaneous impact of asset price shocks on corporate leverage naturally follows from my choice of variables. For the short-run restriction of the monetary policy effect on corporate leverage, I assume that only a negligible proportion of companies roll-over their debt at each quarter, so that the effect of changing interest rates (and, consequently, bond yields and loan rates) takes time to affect corporate leverage.
- While balance sheet, asset price, and monetary policy shocks do not affect real activity at the same quarter, a change in monetary policy is allowed to have a contemporaneous effect on price inflation. Balance sheet and asset price shocks affect price inflation with a lag. Under the assumption of partially sticky prices, all 3 financial shocks should have a lagged effect on price inflation. Nevertheless, due to total number of restrictions in my identification scheme (see below for more details), I need to leave the effect of monetary policy shocks on inflation unrestricted in order to have a just-identified semi-structural VAR.
- Monetary policy shocks have no long-run effects on real output. This assumption implies neutrality of money in the long-run, see Bernanke and Mihov (1998).
- To distinguish asset price from monetary policy shocks, I follow Bjørnland and Leitemo (2009) and I assume that monetary policy has only transitory effects on real asset prices. This restriction is consistent with the money neutrality assumption. It is imposed due to the inherent difficulty in justifying short-run restrictions in either the interest rate or the real asset price equations, especially at a quarterly frequency.

The above restrictions also satisfy the necessary conditions for global and local identification, that require the stacked matrix of short and long-run restrictions to have

exactly one column with j zero restrictions, for each $j = 0, \dots, k - 1$. The restriction matrix M can be written as (. denotes unrestricted element, while 0 denotes element restricted to be zero):

$$M = \begin{matrix} & \eta_s & \eta_d & \eta_{bs} & \eta_{ap} & \eta_m \\ \begin{matrix} y_t \\ \pi_t \\ l_t \\ eq_t \\ ir_t \end{matrix} & \left(\begin{array}{ccccc} \cdot & \cdot & 0 & 0 & 0 \\ \cdot & \cdot & 0 & 0 & \cdot \\ \cdot & \cdot & \cdot & 0 & 0 \\ \cdot & \cdot & \cdot & \cdot & \cdot \\ \cdot & \cdot & \cdot & \cdot & \cdot \\ - & - & - & - & - \end{array} \right) & , & (3.18) \\ \begin{matrix} y_t \\ \pi_t \\ l_t \\ eq_t \\ ir_t \end{matrix} & \left(\begin{array}{ccccc} \cdot & 0 & \cdot & \cdot & 0 \\ \cdot & \cdot & \cdot & \cdot & \cdot \\ \cdot & \cdot & \cdot & \cdot & \cdot \\ \cdot & \cdot & \cdot & \cdot & 0 \\ \cdot & \cdot & \cdot & \cdot & \cdot \end{array} \right) \end{matrix}$$

where the upper part corresponds to the short-run restrictions and the lower part to the long-run restrictions. For a detailed exposition of the technical requirements for global and local identification of SVARs, see Rubio-Ramírez et al. (2010).

The short-run restrictions from (3.18) are applied directly to the matrix B , by setting the respective elements to 0, while the long-run restrictions are applied to the long-run impact matrix $\Psi_\infty = A(L)^{-1}B$. Defining the long-run restriction matrix $J = \mathbf{1}_{\{M_{LR=.\}}\}$, the second set of identifying restrictions can be summarised as follows:

$$\text{vec}(J)'[I_5 \otimes A(L)^{-1}]\text{vec}(B) = 0, \quad B = \begin{bmatrix} b_{11} & b_{12} & 0 & 0 & 0 \\ b_{21} & b_{22} & 0 & 0 & b_{25} \\ b_{31} & b_{32} & b_{33} & 0 & 0 \\ b_{41} & b_{42} & b_{43} & b_{44} & b_{45} \\ b_{51} & b_{52} & b_{53} & b_{54} & b_{55} \end{bmatrix}. \quad (3.19)$$

I solve the combined set of equations from (3.18) and (3.19) by Newton's method¹.

The extraction of the frailty terms \mathbf{f}_t^d , \mathbf{f}_t^{rr} and the estimation of the macroeconomic VAR model, the default equations (3.5)-(3.7), and the recovery rate equations (3.10)-

¹see rootSolve package in R

(3.11) are presented in detail in section 3.3.4.

3.3.4 Estimation

Under the block diagonal assumption for the dynamics of \mathbf{F}_t , \mathbf{f}_t^d and \mathbf{f}_t^{rr} in (3.14), the model consists of 3 components that can be independently estimated: the full VAR dynamics for the Gaussian process \mathbf{F}_t , the parameters and AR(1) frailty factors for the non-Gaussian set of equations (3.5) and (3.6), and the parameters and AR(1) frailty factors for the Gaussian equations for the recovery rate (3.10)-(3.11).

The VAR process for the systematic factor \mathbf{F}_t in (3.15) can be estimated by applying Ordinary Least Squares (OLS) equation-by-equation. In order to obtain the full distribution of the impulse response functions, I choose to estimate the VAR model using Gibbs sampling, assuming fully diffuse priors for the parameters. This choice of priors provides point estimates for the VAR parameters identical to OLS, while obtaining full posterior samples. The VAR equation (3.15) can be re-expressed using the concise matrix notation:

$$\mathbf{F}_t = (I_M \otimes X)\tilde{A} + \epsilon^m, \quad \epsilon^m \sim N(0, \Sigma_\epsilon^m \otimes I_T), \quad (3.20)$$

where

$$X = \begin{bmatrix} 1 & \mathbf{F}_{p-1} & \dots & \mathbf{F}_1 \\ 1 & \mathbf{F}_p & \dots & \mathbf{F}_2 \\ \vdots & \vdots & \ddots & \vdots \\ 1 & \mathbf{F}_{T-1} & \dots & \mathbf{F}_{T-p} \end{bmatrix}, \quad \tilde{A} = \text{vec}[(\nu, A_1, \dots, A_p)'].$$

Using the concise matrix notation (3.20), the posterior distribution for \tilde{A} is multivariate normal, while the posterior for Σ_ϵ^m is an inverse Wishart distribution. The moments for the posterior distributions are given by:

$$\begin{aligned} \tilde{A} | \Sigma_\epsilon^m, \mathbf{F} &\sim N(\text{vec}(\hat{A}), \Sigma_\epsilon^m \otimes (X'X)^{-1}), \quad \hat{A} = (X'X)^{-1}X'\mathbf{F}, \\ (\Sigma_\epsilon^m)^{-1} | \mathbf{F} &\sim W([\mathbf{F} - X\hat{A}]'[\mathbf{F} - X\hat{A}])^{-1}, T - K - M - 1, \end{aligned} \quad (3.21)$$

where $\mathbf{F} = [\mathbf{F}_1, \dots, \mathbf{F}_T]$ if the stacked matrix of observed macroeconomic variables. For more details, see section B of the Appendix at the end of the thesis.

One strand of the credit risk literature estimates via Maximum Likelihood non-

Gaussian state space models with AR(1) latent factors, similar in form with the specification summarised by (3.5),(3.6) and (3.7). Examples include Koopman and Lucas (2008), Koopman et al. (2009), and Koopman et al. (2011), that use simulated Maximum Likelihood estimation techniques based on importance sampling. In models with a large number of parameters to be estimated, Maximum Likelihood techniques are time consuming, and the accuracy and robustness of the final estimates rely largely on starting values and possible flat regions in the likelihood function. This is particularly true when simulation is used to approximate the likelihood function. Furthermore, Maximum Likelihood gives point estimates and obtaining confidence intervals for highly non-linear impulse response functions is not possible. I choose, therefore, to estimate the econometric model by MCMC techniques that give full posterior distributions of all the desired quantities. McNeil and Wendin (2007) use MCMC techniques to estimate a similar setup to (3.5)-(3.7), although the aim of their work is different from this chapter.

Denoting $\boldsymbol{\psi}^d = (\alpha_{j\kappa}^d, \beta_j^d, \rho_{j\kappa}^d, \vartheta_j^d)_{\{j \in \mathcal{S}, \kappa \in \mathcal{G}\}}$ the parameter vector, $\mathbf{f}^d = [\mathbf{f}_1^d, \dots, \mathbf{f}_T^d]$ the stacked vector of the unobserved factors, the posterior can be written as:

$$p(\mathbf{f}^d, \boldsymbol{\psi}^d | \mathbf{D}) \propto p(\mathbf{D} | \mathbf{f}^d, \mathbf{F}, \boldsymbol{\psi}^d) p(\mathbf{f}^d | \boldsymbol{\psi}^d) p(\boldsymbol{\psi}^d),$$

which can be re-expressed, using the Markovian structure of the latent factors \mathbf{f}^d , as:

$$p(\mathbf{f}^d, \boldsymbol{\psi}^d | \mathbf{D}) \propto p(\mathbf{f}_0^d | \boldsymbol{\psi}^d) p(\boldsymbol{\psi}^d) \prod_{t=1}^T p(\mathbf{D}_t | \mathbf{f}_t^d, \mathbf{F}_t, \boldsymbol{\psi}^d) p(\mathbf{f}_t^d | \mathbf{f}_{t-1}^d, \boldsymbol{\psi}^d), \quad (3.22)$$

where $p(\mathbf{f}_0^d | \boldsymbol{\psi}^d)$ is the prior distribution for the initial state of the latent factors, $\mathbf{f}_0^d = \text{vec}(f_{j0}^d)_{\{j \in \mathcal{S}\}}$, given by (3.8), and $p(\mathbf{D}_t | \mathbf{f}_t^d, \mathbf{F}_t, \boldsymbol{\psi}^d)$ is the Binomial density given by (3.5). The posterior distribution (3.22) is not of closed form and analytical solutions are not available. The posterior distributions of the parameters are obtained by means of MCMC numerical techniques. Gibbs sampling is used to sequentially get draws for the parameter set $\boldsymbol{\psi}^d$ and the factors \mathbf{f}^d from the full conditional distributions. Full

conditionals for the individual components of $\boldsymbol{\psi}^d$ can be written as:

$$\begin{aligned}
p(\alpha_{jk}^d | \cdot) &\propto \prod_{t=1}^T p(D_{jkt} | \alpha_{jk}^d, \cdot) p(\alpha_{jk}^d), \\
p(\beta_j^{d,m} | \cdot) &\propto \prod_{t=1}^T \prod_k p(D_{jkt} | \beta_j^{d,m}, F_t^m, \beta_j^{d,-m}, \mathbf{F}_t^{-m}, \cdot) p(\beta_j^{d,m}), \\
p(\rho_{jk}^d | \cdot) &\propto p(D_{jkt} | \rho_{jk}^d, \cdot) p(\rho_{jk}^d), \\
p(f_{jt}^d | \cdot) &\propto \begin{cases} \prod_k p(D_{jkt} | f_{jt}^d, \cdot) p(f_{jt+1}^d | f_{jt}^d, \vartheta_j^d) p(f_{jt}^d | f_{jt-1}^d, \vartheta_j^d), & \text{if } t \neq 0, T \\ \prod_k p(D_{jkt} | f_{jT}^d, \cdot) p(f_{jT}^d | f_{jT-1}^d, \vartheta_j^d), & \text{if } t = T \\ p(f_{j1}^d | f_{j0}^d, \vartheta_j^d) p(f_{j0}^d), & \text{if } t = 0, \end{cases} \quad (3.23) \\
p(\vartheta_j^d | \cdot) &\propto \prod_{t=1}^T p(f_{jt}^d | f_{jt-1}^d) p(\vartheta_j^d),
\end{aligned}$$

for sectors $j \in \mathcal{S}$, rating grades $k \in \mathcal{G}$, and observed macroeconomic variables $m = 1, \dots, 5$.

For the parameters $\alpha_{jk}^d, \beta_j^{d,m}$ and $\rho_{jk}^d, \forall j \in \mathcal{S}, \kappa \in \mathcal{S}$ and $i = 1, \dots, 5$, I choose normal priors of the form $p(\alpha_{jk}^d) \sim N(\alpha_{jk,0}^d, \sigma_{\alpha_{jk,0}^d}^2)$, $p(\beta_j^{d,m}) \sim N(\beta_{j,0}^{d,m}, \sigma_{\beta_{j,0}^{d,m}}^2)$, and $p(\rho_{jk}^d) \sim N(\rho_{jk,0}^d, \sigma_{\rho_{jk,0}^d}^2)$ respectively. To sample from the posterior distributions, the slice sampling algorithm of Neal (2003) is used. The slice sampling algorithm is also used to sample the unobserved factors \mathbf{f}^d . To ensure stationarity, the parameters $\vartheta_j^d, \forall j \in \mathcal{S}$ need to be constrained to values less than 1 in absolute value. Therefore, I choose truncated normal priors of the form $p(\vartheta_j^d) \sim N(\vartheta_{j,0}^d, \sigma_{\vartheta_{j,0}^d}^2) I(-1, 1)$, and I sample from the posterior via the slice sampling algorithm. For more information on Gibbs and slice sampling, see sections A.1 and A.2 of the Appendix at the end of the thesis.

The priors and posteriors for the recovery rate equations (3.10)-(3.11) are very similar to those of the default process, with the difference of the likelihood function $p(\log(RR_{it}) | \cdot)$ being now normal instead of binomial, providing closed form posteriors for some of the parameters. Denoting $\widetilde{RR}_{it} = \log(RR_{it})$, the full conditionals of the

parameters used for Gibbs sampling take the form:

$$\begin{aligned}
p(\alpha_{js}^{rr}|\cdot) &\propto \prod_{t=1}^T \prod_{i:j,s} p(\widetilde{RR}_{it}|\alpha_{js}^{rr}, \cdot) p(\alpha_{js}^{rr}), \\
p(\beta_s^{rr,m}|\cdot) &\propto \prod_{t=1}^T \prod_{i:s} p(\widetilde{RR}_{it}|\beta_s^{rr,m}, F_t^m, \beta_s^{rr,-m}, \mathbf{F}_t^{-m}, \cdot) p(\beta_s^{rr,m}), \\
p(\rho_{js}^{rr}|\cdot) &\propto \prod_{t=1}^T \prod_{i:j,s} p(\widetilde{RR}_{it}|\rho_{js}^{rr}, f_{st}^{rr}, \cdot) p(\rho_{js}^{rr}), \\
p(f_{st}^{rr}|\cdot) &\propto \begin{cases} \prod_{i:s} p(\widetilde{RR}_{it}|f_{st}^{rr}, \cdot) p(f_{st+1}^{rr}|f_{st}^{rr}, \vartheta_s^{rr}) p(f_{st}^{rr}|f_{st-1}^{rr}, \vartheta_s^{rr}), & \text{if } t \neq 0, T \\ \prod_{i:s} p(\widetilde{RR}_{it}|f_{sT}^{rr}, \cdot) p(f_{sT}^{rr}|f_{sT-1}^{rr}, \vartheta_s^{rr}), & \text{if } t = T \\ p(f_{s1}^{rr}|f_{s0}^{rr}, \vartheta_s^{rr}) p(f_{s0}^{rr}), & \text{if } t = 0 \end{cases} \\
p(\vartheta_s^{rr}|\cdot) &\propto \prod_{t=1}^T p(f_{st}^{rr}|f_{st-1}^{rr}) p(\vartheta_s^{rr}), \\
p((\sigma_s^{rr})^2|\cdot) &\propto \prod_{t=1}^T \prod_{i:s} p(\widetilde{RR}_{it}|\sigma_s^{rr}, \cdot) p((\sigma_s^{rr})^2),
\end{aligned} \tag{3.24}$$

for each default i , industry sector $j \in \mathcal{S}$, seniority level $s \in \mathcal{C}$, and observed macroeconomic variables $m = 1, \dots, 5$. I denote by $i : j, s$ the subset of defaults i that belong to industry sector j and seniority level s , and by $i : s$ the subset of defaults that belong to seniority level s .

Since $p(\widetilde{RR}_{it}|\cdot)$ is a Gaussian density, choosing independent normal-inverse gamma priors for intercepts, sensitivities, and variances leads to closed form posteriors for the aforementioned parameters. More specifically, I assume priors of the form $p(\alpha_{js}^{rr}) \sim N(\alpha_{js,0}^{rr}, \sigma_{\alpha_{js,0}^{rr}}^2)$, $p(\beta_j^{rr,m}) \sim N(\beta_{j,0}^{rr,m}, \sigma_{\beta_{j,0}^{rr,m}}^2)$, and $p(\rho_{js}^{rr}) \sim N(\rho_{js,0}^{rr}, \sigma_{\rho_{js,0}^{rr}}^2)$. Conditional on the observed macroeconomic variables \mathbf{F}_t , unobserved frailty factors \mathbf{f}_{st}^{rr} and the variances $(\sigma_s^{rr})^2$, this choice of priors leads to the conjugate Bayesian linear regression posteriors:

$$p(\mathbf{b}|\cdot) \sim N(\mathbf{b}_1, \mathbf{B}_1), \text{ where}$$

$$\mathbf{B}_1 = (\mathbf{B}_0^{-1} + \mathbf{X}'\mathbf{X}/\sigma^2)^{-1}, \text{ and } \mathbf{b}_1 = \mathbf{B}_1(\mathbf{B}_0^{-1}\mathbf{b}_0 + \mathbf{X}'\mathbf{Y}/\sigma^2).$$

The posteriors for α_{js}^{rr} , $\beta_j^{rr,m}$, and ρ_{js}^{rr} are obtained by substituting in the expressions above the appropriate covariates and residual variances (namely, a vector of 1s, the macroeconomic factors \mathbf{F}_t , the frailty terms \mathbf{f}_{st}^{rr} , and the variances $(\sigma_s^{rr})^2$ respectively).

Just like the default process posteriors in (3.23), I sample the latent factors \mathbf{f}_{st}^{rr} and AR(1) coefficients ϑ_s^{rr} using slice sampling. For the latter, I use truncated normal priors of the form $p(\vartheta_s^{rr}) \sim N(\vartheta_{s,0}, \sigma_{\vartheta_{s,0}}^2)I(-1, 1)$. Finally, choosing conjugate Inverse Gamma priors for the variances $(\sigma_s^{rr})^2$ of the form $p((\sigma_s^{rr})^2) \sim IG(c_{(\sigma_s^{rr})^2,0}, C_{(\sigma_s^{rr})^2,0})$, gives closed form posteriors $p((\sigma_s^{rr})^2) \sim IG(c_{(\sigma_s^{rr})^2}, C_{(\sigma_s^{rr})^2})$, with $c_{(\sigma_s^{rr})^2} = c_{(\sigma_s^{rr})^2,0} + 0.5 \sum_i \mathbf{1}_{i:s}$ (where $\sum_i \mathbf{1}_{i:s}$ is the total number of seniority level s defaults) and $C_{(\sigma_s^{rr})^2} = C_{(\sigma_s^{rr})^2,0} + 0.5 \cdot \sum_t \sum_{i:s} (\log(\widetilde{RR}_{it}) - \alpha_{j_s}^{rr} - \beta^{rr} \mathbf{F}_t - \rho_{j_s}^{rr} f_{st}^{rr})^2$. Again more details for the Bayesian estimation of a linear regression model can be found in section B of the Appendix at the end of the thesis.

3.4 Data

For the macroeconomic VAR analysis I use quarterly data spanning from Q1 1959 until Q2 2013. All data are sourced from the FRED. Output growth is defined as $y_t = \Delta \log(Y_t)$, where the real GDP series Y_t is measured in billions of chained 2009 dollars using the GDPC96 series. Inflation, $\pi_t = \Delta \log(P_t)$, is defined as quarterly changes in (logarithmic) consumer prices. For P_t I use the consumer price index series CPIAUCSL. For the policy rate ir_t I use the FED Funds rate (series FEDFUNDS). As a proxy for the indebtedness/leverage (l_t) I use the ratio of outstanding debt of business corporate (non farm) sector (series BCNSDODNS) to corporate profits after tax with inventory valuation and capital consumption adjustments (series CPATAX). Corporate profits reflect the internal funds available to a company, and the proportion not distributed to shareholders contributes towards future net worth. Therefore dividing debt by profits provides a more forward view of top down corporate leverage, as compared to the typical definition of leverage that involves debt and equity/net worth. I have also considered using the market value of the firm in the definition of leverage. Unfortunately, such a choice invalidates the identification scheme. The market value of a firm is strongly correlated with the aggregate equity price index I use in the VAR, and they should both share the same short and long-run restrictions. A different option is to use the book value of net worth; unfortunately, using the book net worth value leads to a highly non-stationary time series, with unclear cyclical dynamics. Finally, I define the real equity market returns $eq_t = \Delta \log(SP_t/P_t)$ as the quarterly change in logarithmic

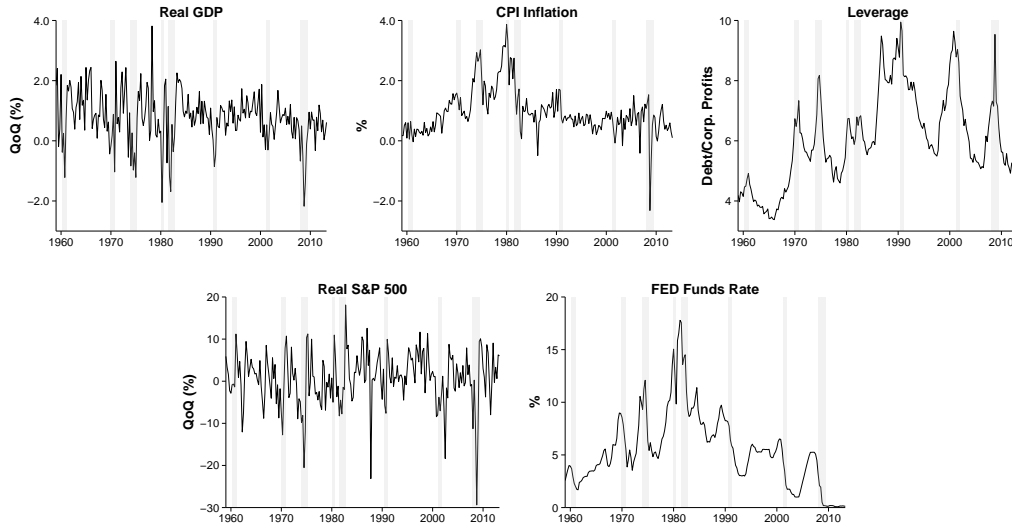


Figure 3.1: Historical Time Series - Macroeconomic Variables

Historical macroeconomic time series for the period Q1 1959-Q2 2013. All data sampled at quarterly frequency. Output growth, inflation and equity prices are expressed in 1st order quarterly log differences of real GDP, CPI and S&P 500 respectively. S&P 500 returns are deflated using the CPI. The FED Funds rate is measured in percentage terms. Leverage corresponds to the ratio of the total credit market debt owed by the non-financial corporates to corporate profits after tax. Shaded areas correspond to NBER recession quarters.

values of the S&P 500 index SP_t , deflated by the consumer price index P_t . Historical time series plots of the 5 macro-financial variables are depicted in figure 3.1.

For corporate bond default data I rely on Moody's Default & Recovery database. When examining the default time series for corporate bonds, there is a clear structural break between the early 80s and late 70s; default rates in the latter period are significantly lower than the former. Furthermore, Moody's introduced their alphanumeric rating scale above Caa in mid-1982. To exclude any possible outliers and to model a period as consistently as possible, I choose the first quarter of 1983 as the starting point. The series ends in Q2 2013, which coincides with the last quarter of the macroeconomic data. The different starting point of macroeconomic and credit data does not pose a problem since the credit analysis of sections 3.3.1 and 3.3.2 is conditionally independent of the macroeconomic model of section 3.3.3. Therefore, while the macroeconomic VAR is estimated over the period 1959-2013, the credit models only use the macroeconomic time series from 1983 till 2013. To define the number of firms N_{jit} in (3.5) I use the cohort definition: number of active companies at the beginning of each quarter. To adjust for firms that have their rating withdrawn within a quarter, I subtract half the number of withdraws during the quarter. Therefore, I make the implicit assumption that withdraws are distributed uniformly within quarters.

To capture the credit quality I use 4 rating groups. The highest credit quality group corresponds to the investment grade rated firms (firms rated Aaa-Baa). Those firms rarely default directly and, therefore, introducing additional granularity on that end of the rating scale is subject to large estimation errors. For the sub-investment grade rated firms, I use Moody's letter grading system to classify firms into Ba, B and Caa-C. Further split by Moody's alphanumeric system is also possible; I refrain from using an overly granular split to avoid picking up noise in the data as the default count for each of the alphanumeric grades becomes low. I choose to group the firms rated Caa and below, since firms of such low credit quality default in irregular patterns and it is difficult to empirically maintain the discriminatory order of the rating grades. Based on this 4 rating group segmentation, I record 14 Investment Grade firm defaults in the sample, 78 Ba firm defaults, 465 B firm defaults, and 895 Caa-C firm defaults.

For the industry sector dynamics of corporate defaults I use Moody's 11 segment classification as the base. From the analysis I exclude Sovereigns and Public Finance firms and companies with missing or unclassified industry information. Finally, to avoid low default counts, I merge some sectors due to similar default patterns: the Media & Publishing with the Consumer Industries segment, the Banking with the Finance & Real Estate segment, and the Energy & Environment with the Utilities segment. The final 6 industries are: Capital Industries, Consumer Industries (including Media & Publishing), Financials (consisting of Banking, and Finance & Real Estate), Technology, Transportation, and Utilities (including Energy & Environment). This industry segmentation results in 388 company defaults for Capital Industries, 607 defaults for Consumer Industries, 107 defaults for Financials, 156 defaults for Technology, 70 defaults for Transportation, and 124 defaults for Utilities.

For the recovery rates I rely on market-implied measures provided in Moody's Default & Recovery database. The recorded recovery rates are constructed as the percentage change in traded prices of defaulted debt, 30 days post the default date. For corporate bonds, recoveries are typically measured using trading prices 30 days post default or by appropriately discounting the recovered cash flow until resolution. While the question of which of the two measures is better often gets asked, the answer is not straightforward. If the purpose of the recovery rate estimates is to allocated capital for a portfolio of liquid instruments that must be sold shortly after default,

then the 30 days post default price is the better recovery rate estimate. If, on the other hand, the purpose of estimating recovery rates is to allocate capital for a portfolio of illiquid assets that must be held until maturity, then the ultimate recoveries should be preferred. Furthermore, despite not being a realised recovery rate, using post default price data has distinct advantages. First, Moody's Ultimate Recovery database that contains actual recoveries has a much shorter history of data (historical data go back to 1987), unlike 30-day post default price data that go back to the early 1900's. Second, defaulted bonds can have long and highly diverse emergence periods making it very difficult to assign macroeconomic variables to observed recoveries. A thorough analysis of the macroeconomic determinants of ultimate recoveries will need to include the effect of the economic environment on both the length of the emergence period and the level of recoveries at emergence; this is beyond the scope of this chapter. However, Metz et al. (2012) report that recovery rates based on 30-day post default trading prices are very good predictors of ultimate recovery rates, both at the average and instrument level. Despite this strong correlation between the two measures of recoveries, Metz et al. (2012) note that trading prices tend to be close to 3% lower than ultimate recoveries when the latter are discounted at the coupon rate. Despite all the useful properties, it is worth pointing out that using trading prices for typically creates a higher correlation between recoveries and the equity market. While recovery rates are a function of the residual asset value of company after default and, therefore, there is a natural link to the equity market, using trading prices for recovery rates might inflate this dependence. This becomes more relevant when interpreting the result of section 3.5.2.

I segment recovery rates across two dimensions: industry sector and seniority level. Consistently with the default rate analysis, I use the following 6 industry sectors for the industry level segmentation: Capital Industries, Consumer Industries, Financials, Technology, Transportation, and Utilities. There are an adequate number of unique default events with recorded recoveries across all industry sectors, with 411 events for Capital Industries, 647 events for Consumer Industries, 112 events for Financials, 154 events for Technology, 67 events for Transportation, and 136 events for Utilities. To ensure that an adequate number of observations are available for inference, I opt for 3 seniority levels: Secured, Unsecured, and Subordinated. Secured debt includes First, Second, and Third Lien issuances, while the Subordinated segmentation level includes both Junior and Senior Subordinated debt. This 3 level seniority segmentation leads

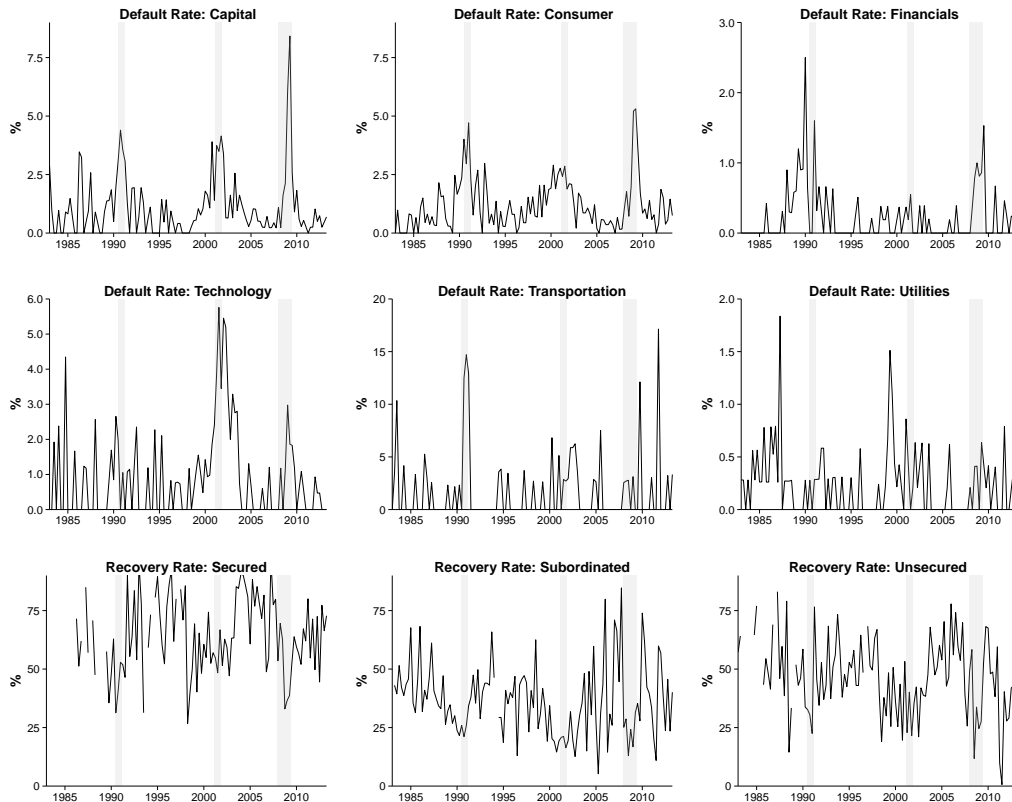


Figure 3.2: Historical Time Series - Credit Variables

Historical default and recovery time series for the period Q1 1983-Q2 2013. All data sampled at quarterly frequency. Default rates reported per industry group and calculated using the cohort approach. They correspond to number of defaults in a given quarter divided by the number of active companies at the beginning of the quarter. Recovery rates are reported per seniority level, using 30-day post default trading price information. Shaded areas correspond to NBER recession quarters.

to 579 default events with recorded recoveries for the Secured segment, 618 events for the Unsecured segment, and 741 events for the Subordinated segment.

Figure 3.2 depicts the time series of historical data for both defaults and recoveries. Default information is provided for the aforementioned industry sector segmentation, while recovery rates are averaged across the 3 seniority levels. From figure 3.2 it is very apparent that highly cyclical sectors like Capital and Consumer Industries exhibit a substantial increase in default intensity during recessionary periods. For other sectors default behaviour seems to be affected by particular historical events, like the dot-com bubble in the early '00s for the Technology sector, and the Savings & Loans crisis in the late '80s for the Financial sector. As compared to the default occurrence, recoveries seem to be much less cyclical. Despite the reduced cyclicity, it is evident that during the early and late '00s recessions, recovery rates are lower than the long-run trend. Finally, as expected, Secured debt has an average recovery level of 64%, Unsecured

debt has an average recovery rate of 47%, and subordinated debt exhibits the lowest recoveries, with an average recovery rate of 36%.

3.5 Results

3.5.1 Macroeconomic VAR

Figure 3.3 depicts the impulse response functions for each of the 5 macroeconomic variables to each of the 5 structural shocks. The impulse responses are performed in terms of the structural moving average form (3.16) for the 5 variable VAR, assuming 1 standard deviation independent shocks at the beginning of the forecast horizon. Results are based on 100,000 MCMC draws from the full conditional distribution (see section 3.3.4 for details) and alongside the median responses, 1 standard deviation error bands are also provided.

First observing the impact of aggregate supply and demand shocks, figure 3.3 indicates that both shocks lead to a comparable instantaneous increase in output. While aggregate demand shocks disappear after less than a year, it takes almost two years for output to return to the equilibrium growth path after an aggregate supply shock. An aggregate demand shock shifts the demand curve to the right leading to an increased level of output above the natural rate, creating inflationary pressures. Consistent with expectations, the response of prices is instantaneous, as inflation peaks in the first quarter of the shock and then slowly moves back to equilibrium. Furthermore, in line with neo-classical models, an aggregate supply shock has an immediate negative effect on prices over the first quarter (indicating no price stickiness), while output also rises instantly. This sharp contrast in price behaviour under the 2 structural shocks, also leads to very different time profiles for the FED Funds rate and by extension to the corporate leverage. To tackle the inflationary pressures resulting from an aggregate demand shock, FED raises interest rates by 0.75% approximately one year after the initial shock, and this in turn makes debt issuance more costly for corporates and results in higher than normal leverage ratios. Finally, due to the inflationary pressures in the case of an aggregate demand shock, the real asset price level drops rapidly after the instantaneous quarterly increase of 1%, to reach a minimum of -1% quarterly growth

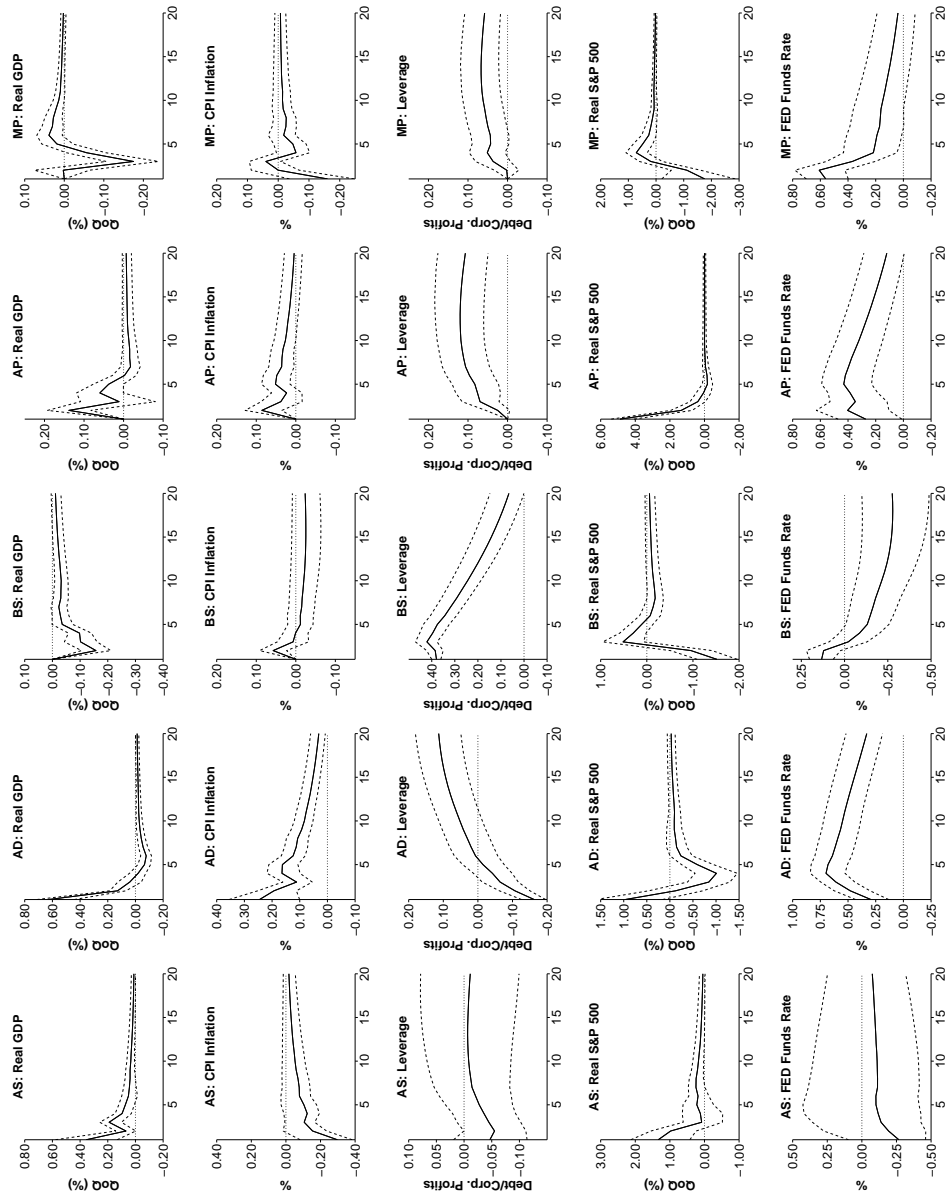


Figure 3.3: Macroeconomic Impulse Responses Functions

Impulse-Response functions of the 5 macroeconomic series to each of the 5 structural shocks. The magnitude of the shocks is normalized to 1σ . AS denotes an Aggregate Supply shock, AD an Aggregate Demand shock, BS a Balance Sheet shock, AP an Asset Price shock, and MP a Monetary Policy shock. Horizontal scales depict time in quarters. Median impulse responses (solid lines) are reported with 1σ error bands (dotted lines), computed from 100,000 MCMC samples.

rate approximately a year after the initial shock. On the other hand, the low level of price inflation and interest rates in the case of an aggregate supply shock helps to maintain a persistent real stock price growth rate throughout the 20 quarters forecast horizon.

The balance sheet shock leads to a decrease in real output, consistent with an expected gradual decrease in investment (and ultimately production) while the leveraged corporate sector struggles to finance its capital growth. This decrease in investment also translates to a lower demand for assets, leading to a drop in asset values. This decrease in equity prices is also the result of the stock market participants pricing in the future prospects and viability of the corporate sector. Perhaps counterintuitively, the balance sheet shock leads to a temporary increase in price inflation and policy rate. Nevertheless, following the working capital or cost channel strand of monetary policy literature, in addition to financing long term projects, (short term) debt is also used to finance a firm's working capital. As the cost of debt becomes part of the firm's marginal cost, an increase in funding costs resulting from a credit squeeze could lead to an aggregate supply type of shock, whereby prices for goods increase as producers pass on their increased costs to the consumers. In order for that to happen, the cost channel needs to be stronger than the demand channel, so that prices rise while output drops (and consequently, interest rates also rise temporarily to respond to increased price inflation). It is worth pointing out that this cost channel effect of a balance sheet shock appears to be present even if I re-estimate the model using different sub-samples. For more details on the cost channel, see Barth and Ramey (2002), Chowdhury et al. (2006), and Ravenna and Walsh (2006). Despite the initial rise in inflation and interest rates, the drop in investment and output eventually brings down prices and policy rate, as the FED tries to facilitate the de-leveraging of corporates.

The effects of an asset price shock are easier to explain. An increase in equity prices leads to an increase in wealth that eventually pushes up consumption and stimulates investment via a Tobin Q effect. These two effects positively affect aggregate demand and this is evident by the effect on real GDP, which stays positive for more than a year following the shock. Price inflation also rises following an asset price shock, indicating nominal rigidities and a slow reaction in the price formation. This relatively persistent rise in inflation (price inflation reaction stays mostly positive throughout the 20 quarter

horizon) leads to an increase in the FED Funds rate, consistent with an inflation-targeting central bank. With this higher policy rate, debt issuance for corporates becomes more costly, leading to a higher leverage level. The level of persistence for inflation, policy rate, and corporate leverage suggests again the presence of a cost channel in the transmission of debt cost changes. The increased level of interest rates leads to an increased leverage level, that ultimately pushes prices up, as corporates pass on their costs to the consumers; this increased price level keeps interest rates at a relatively high level, and this feeds again to the corporate leverage.

A tightening monetary policy shock is associated with an increase in the FED Funds rate by approximately 0.6%. As a result of the higher level of interest rates, output drops, to reach a minimum approximately a year after the monetary policy shock. Following that point, output bounces back. This sharp increase after the first year, might not be identical to VARs using recursive identification schemes, see for example Bernanke et al. (2005) for a FAVAR with Cholesky decomposition based shocks. Nevertheless, this is a direct result of my identification scheme, that assumes no long-run effect of monetary policy shocks on real output. Despite the argument of the presence of a cost channel transmission mechanism, inflation drops following the tightening of monetary policy. This indicates that the deflationary effect of the monetary policy shock is higher than the cost corporates are able to pass on to consumers, and that results in a drop in consumer price inflation. Furthermore, it avoids the so called "inflation puzzle" of Sims (1992), and Eichenbaum (1992), whereby inflation rises following an contractionary monetary policy shock. The increased level of interest rates causes a rise in the cost of debt, leading to a higher corporate leverage level that could also lead to a partial sell-off of financial assets. Combined with the increase in the discount rate for future dividends and the drop in output, contractionary monetary policy shocks cause equity returns to drop sharply.

Table 3.1 contains the typical forecast error variance decomposition for the macroeconomic VAR. The forecast error variance decomposition captures the proportion of the h -step forecast error variance of variable j that is accounted for by each of the structural shocks η_{kt} , $k = 1, \dots, 5$, and is based on:

$$\omega_{jk,h} = \sum_{i=0}^{h-1} \phi_{jk,i}^2 / \sum_{i=0}^{h-1} \sum_{m=1}^5 \phi_{jm,i}^2,$$

Figures in %	Quarter 1					Quarter 20				
Variable\Shock	AS	AD	BS	AP	MP	AS	AD	BS	AP	MP
Output	24.3	75.7	0.0	0.0	0.0	27.0	56.4	6.9	3.9	5.8
Inflation	52.9	33.5	0.0	0.0	13.6	39.6	48.5	1.8	3.8	6.3
Leverage	1.0	17.0	82.4	0.0	0.0	0.9	9.0	74.6	12.0	3.5
S&P 500 Returns	6.4	2.3	6.8	75.2	9.3	9.2	7.8	8.5	61.6	13.0
FED Funds Rate	18.2	18.9	2.7	9.9	50.3	7.6	58.0	7.7	16.1	10.6

Table 3.1: Macroeconomic VAR FEVD

Proportion of the h -step forecast error variance of each of the 5 macroeconomic variables that is accounted for by the Aggregate Demand (AD), Aggregate Supply (AS), Balance Sheet (BS), Asset Price (AP), and Monetary Policy (MP) structural shocks. Variance decompositions are calculated using the posterior means of the Φ_i s in the moving average VAR representation (3.16), and are provided for the 1st and 20th forecast quarter.

where Φ_i is the matrix of the dynamic multipliers, and $\phi_{jm,i}$ is the (j, m) element of Φ_i . For the Φ_i , the posterior means of the MCMC draws are used for the calculations. Variance decompositions are provided for 1 and 20 quarters forecast horizon.

As expected, aggregate demand and supply shocks explain most of the forecast error variance of output and inflation, both in the short and the long-run. Monetary policy proves effective in reducing price inflation in the short-run, as evidenced by the 13.6% portion or the forecast error variance at quarter 1. In the long-run this proportion drops to 6%, as the response of inflation to an aggregate demand dominates. The contemporaneous big drop resulting from an expansionary aggregate demand shock contributes significantly in explaining the short-run forecast errors of corporate leverage, in addition to the balance sheet shock (17% for the former effect and 82% for the later). Interestingly, the asset price shock, despite having 0 initial impact on corporate leverage, explains 12% of the forecast error variance in the long-run. The asset price shock also dominates the real equity price movements. Due to the inflation adjustment, the long-run forecast error variance of real equity returns is also attributed to the other 4 structural shocks, with proportions close to 10%, leaving a 60% share for the asset price shock. Finally, the monetary policy shock contributes the most to the short-run forecast error variance of the FED Funds rate, with a proportion of 50%, while aggregate demand and supply shocks each contribute 18%. Interestingly, the contribution of the monetary policy shock drops to 10% in the long-run, while the aggregate demand shock dominates with 58%. The asset price shock also helps explaining 16% of the long-run forecast error variance for the FED Funds rate.

To assess the robustness of my results, I apply the identification scheme in various sub-periods in the sample. The results are remarkably stable to the sample used for the analysis, and remain similar to those I report even if the last recession is excluded

from the estimation. Nevertheless, it is worth noting that excluding the '70s period leads to problems when trying to separate aggregate demand and supply shocks; this is not overly surprising, since the 1973 oil crisis and 1979 energy crisis are 2 of the most characteristic periods of supply type of shocks. Furthermore, when sample only covers the period from 1980 onwards, the cost channel of monetary policy transmission is very strong and causes a sharp rise in price inflation following a tightening in monetary policy.

3.5.2 Implications for Credit Risk

Corporate Defaults

For the corporate default specification, I work with highly uninformative priors for all the parameters, centred at 0. More specifically, for the observation equation (3.6), I use the normal priors $p(\alpha_{jk}^d) \sim N(0, 10^4)$, $p(\beta_j^{d,m}) \sim N(0, 10^4)$, and $p(\rho_{jk}^d) \sim N(0, 10^4)$, while for the latent factor equation (3.7) I reduce the variance of the prior for the AR(1) coefficient to $p(\vartheta_j^d) \sim N(0, 10)I(-1, 1)$. Inference is based on 100,000 MCMC draws, using the MCMC algorithm described in section 3.3.4, following a 10,000 burn-in sample that ensures only posterior draws are kept for inference. To narrow the confidence intervals for the impulse response analysis, the lag polynomial of macroeconomic sensitivities $\beta_j^d(L)$ in (3.6) takes the simplified form $\beta_j^d(L) = [\beta_j^{d,1}L^{\ell_1^d}, \dots, \beta_j^{d,5}L^{\ell_5^d}]$, where ℓ_m^d is the lag operator order for variable m . The specific set of observed macroeconomic variables \mathbf{F}_t and their respective lag are based on a statistical model choice, using the 5% significance level as the cut-off point for the inclusion of a variable in the final specification.

Table 3.2 contains the resulting parameter estimates and their corresponding standard errors. Consistent with the existing literature, I report highly persistent frailty factors across all sectors (the only exception is the Transportation sector, for which most of the autocorrelation in the observed time series is removed by the observed macroeconomic variables). The degree of persistency, which ranges from 0.56 to 0.93 (excluding Transportation), is similar to that reported by Koopman and Lucas (2008), and Koopman et al. (2012). Corporate leverage is the dominant factor for the majority of sectors, and its impact on defaults is instantaneous. Especially for the Utilities sec-

	<u>Capital</u>	<u>Consumer</u>	<u>Financials</u>	<u>Technology</u>	<u>Transportation</u>	<u>Utilities</u>					
α_{IG}^d	-	α_{IG}^d	-	α_{IG}^d	-6.315 (1.483)	α_{IG}^d	-	α_{IG}^d	-	α_{IG}^d	-3.700 (0.145)
α_{Ba}^d	-3.214 (0.112)	α_{Ba}^d	-3.285 (0.364)	α_{Ba}^d	-3.818 (0.533)	α_{Ba}^d	-4.038 (0.868)	α_{Ba}^d	-3.135 (0.511)	α_{Ba}^d	-3.148 (0.314)
α_B^d	-2.596 (0.103)	α_B^d	-2.634 (0.360)	α_B^d	-2.602 (0.154)	α_B^d	-2.780 (0.083)	α_B^d	-2.653 (0.225)	α_B^d	-2.808 (0.358)
α_{Caa}^d	-1.584 (0.102)	α_{Caa}^d	-1.671 (0.149)	α_{Caa}^d	-1.305 (0.098)	α_{Caa}^d	-1.730 (0.157)	α_{Caa}^d	-1.845 (0.206)	α_{Caa}^d	-1.653 (0.109)
ρ_{IG}^d	-	ρ_{IG}^d	-	ρ_{IG}^d	1.609 (0.737)	ρ_{IG}^d	-	ρ_{IG}^d	-	ρ_{IG}^d	-0.012 (0.092)
ρ_{Ba}^d	0.016 (0.083)	ρ_{Ba}^d	0.133 (0.056)	ρ_{Ba}^d	0.722 (0.319)	ρ_{Ba}^d	0.676 (0.440)	ρ_{Ba}^d	0.308 (0.431)	ρ_{Ba}^d	0.290 (0.142)
ρ_B^d	0.198 (0.050)	ρ_B^d	0.169 (0.043)	ρ_B^d	0.325 (0.122)	ρ_B^d	0.036 (0.086)	ρ_B^d	0.554 (0.197)	ρ_B^d	0.405 (0.123)
ρ_{Caa}^d	0.196 (0.049)	ρ_{Caa}^d	0.067 (0.022)	ρ_{Caa}^d	0.140 (0.098)	ρ_{Caa}^d	0.347 (0.092)	ρ_{Caa}^d	0.710 (0.187)	ρ_{Caa}^d	0.049 (0.094)
$\beta_{y_{t-1}}^d$	-0.124 (0.055)	$\beta_{y_{t-1}}^d$	-0.140 (0.038)	$\beta_{y_{t-2}}^d$	-0.169 (0.087)			$\beta_{y_{t-1}}^d$	-0.323 (0.156)		
$\beta_{l_t}^d$	0.159 (0.041)	$\beta_{l_t}^d$	0.079 (0.030)			$\beta_{l_t}^d$	0.154 (0.046)			$\beta_{l_t}^d$	0.117 (0.046)
$\beta_{eq_{t-2}}^d$	-0.009 (0.005)	$\beta_{eq_{t-2}}^d$	-0.009 (0.003)	$\beta_{eq_{t-1}}^d$	-0.018 (0.007)	$\beta_{eq_{t-1}}^d$	-0.022 (0.007)				
$\beta_{ir_{t-2}}^d$	0.065 (0.024)			$\beta_{ir_{t-1}}^d$	0.167 (0.026)	$\beta_{ir_t}^d$	0.059 (0.027)	$\beta_{ir_t}^d$	0.095 (0.048)		
ϑ^d	0.746 (0.122)	ϑ^d	0.934 (0.048)	ϑ^d	0.555 (0.193)	ϑ^d	0.671 (0.149)	ϑ^d	0.320 (0.205)	ϑ^d	0.834 (0.104)

Table 3.2: Parameter Estimates for Default Rates

Parameter estimates and corresponding standard errors for the corporate default non-Gaussian specification (3.6)-(3.7). All estimates are based on 100,000 draws from the full posterior distribution, using the MCMC algorithm described in section 3.3.4.

tor, corporate leverage is the only macroeconomic variable that I find to be statistically significant; defaults in this sector are few, highly non-cyclical, and mainly driven by changes in cost and availability of credit, as Utilities companies typically operate with large amounts of debt. GDP growth plays a very important role for the cyclical sectors, such as Capital Industries, Consumer Industries, and Transportation. Unlike leverage, the effect of a drop in aggregate output takes 1 quarter to increase defaults in the aforementioned sectors. GDP growth also affects defaults in the Financial sector, but with a 2 quarter lag. This additional delay is possibly caused by the fact that Financial firms indirectly feel the effects of a drop in output through increase in mortgage and corporate lending charge-offs. Equity prices affect the net worth of corporates, which in turn affects the availability of credit or funding in general. The effect on Capital and Consumer industries comes with a 2 quarter lag, while for Financial and Technology firms it comes with a single quarter lag. Financial firms typically have direct trading

activities and a shock on the equity market is more likely to affect them sooner than the rest of the corporates. The Technology sector includes a high number of companies that do not yet produce profits or cash flows and, therefore, short term viability is more closely linked to the performance of the equity market. Finally, the FED Funds rate has typically a low impact for most sectors, except Financials. This somehow counter-intuitive finding, is the result of concentrated defaults during the late '80s Savings & Loans crisis that followed the gradual increase in interest rates throughout the decade.

Figure 3.4 contains the impulse response functions for the corporate probabilities to default PD_{jkt} defined in (3.6). The impulse response functions are provided for each of the 6 industry sectors, assuming a B rating for the corporates. Selecting a single rating grade is necessary for brevity reasons, since PD_{jkt} is a non-linear function of \mathbf{F}_t and the responses to structural macroeconomic shocks depend on the long-run average level of PD_{jkt} . Assuming a given structure for the systematic factors \mathbf{F}_t , the response functions for structural shock k are calculated as:

$$\text{IRF}_{jk,h}^d = N \left(\alpha_{j,B}^d + \sum_{m=1}^5 \beta_j^{d,m} \cdot \mathbf{1}_{h \geq \ell_m} \cdot \phi_{mk, \max(h-\ell_m, 0)} \right) - N \left(\alpha_{j,B}^d \right), \quad (3.25)$$

for sector j and a forecast horizon h . $\phi_{mk,h}$ is the (m, k) element of Φ_h . Equation (3.25) is evaluated at each of the 100,000 MCMC draws for $\alpha_{j,Ba}^d$ and $\beta_j^{d,m}$ to get the full distribution of the impulse response functions. Figure 3.4 depicts the median responses; 1 standard deviation error bands are also provided.

It is worth noting that unlike figure 3.3, the impulse response functions of figure 3.4 are highly non-linear and depending on the average PD level for a given sector shocks might appear overly persistent (for example the aggregate demand and asset price shocks for capital industries). In this sense the impulse response functions could also be provided in terms of the $N^{-1}(PD_{jkt})$ transformation that would ensure linearity and shocks that converge to the long-run average (probit) level. While I acknowledge this shortcoming, I argue that measuring the impact on the probit scale does not give an intuitive interpretation of the magnitude. In contrast, measuring the impulse response functions on the actual PD scale can give risk managers a clear indication of the impact of each shock to a quantity that directly affects their portfolio loss distribution.

A positive aggregate supply shock leads to a maximum decrease in corporate default




Figure 3.4: Corporate Default Impulse Responses Functions

Impulse-Response functions of the 6 sector specific probabilities to default (PD) to each of the 5 structural shocks. The magnitude of the shocks is normalized to 1σ . AS denotes an Aggregate Supply shock, AD an Aggregate Demand shock, BS a Balance Sheet shock, AP an Asset Price shock, and MP a Monetary Policy shock. All Impulse-Response functions are provided for B rated corporates. Vertical scales depict the difference from the long-run average level of default rates for B rated corporates in percentage terms. Horizontal scales depict time in quarters. Median impulse responses (solid lines) are reported with 1 standard deviation error bands (dotted lines), computed from 100,000 MCMC samples.

rates by 5-10 basis points during the first year, with persistently lower than average default rates thereafter, due to the combination of the shock's long-run positive impact on output, and the resulting low leverage and interest rate level. Nevertheless, the impact after the first year is not statistically different from 0 for all sectors. As opposed to an aggregate supply shock, a positive aggregate demand shock leads to an increase in inflation that pushes up interest rates. This results in fundamentally different time profiles for the 2 shocks when quantifying the effect on default rates. A positive aggregate demand shock leads to an eventual rise in default rates in the long-run due to the increased cost of debt. Following a short-run decrease in defaults of approximately 10 basis points for the more cyclical sectors, such as Capital, Consumer, and Transportation, and close to 0 for the rest of the sectors, defaults reach a level of 5 (for Technology firms) and 20 (for Financial firms) basis points above the long-run average and then very slowly revert to the through-the-cycle level. This cost of debt side effect is close to 0 for Utilities companies (as the effect of business cycles is very small for the sector), and is very dampened for Consumer Industries (as the elevated profits from the increased economic activity compensate for the increase in funding costs).

Responses to asset price shocks have very similar time profiles to the responses to aggregate demand shocks; following a positive asset price shock, there is a sharp short-run decrease in default rates over the first year, followed by a persistent rise in default occurrence over the remaining quarters in the observation period. The magnitude of the responses is less pronounced than those following an aggregate demand (peak drop of default rates by 5 basis points for Capital, Consumer, Financials, and Technology, followed by a peak rise of 5-10 basis points), while for Transportation the response of default rates is positive throughout the 20 quarter horizon.

Finally, as expected, following positive balance sheet (shocks that lead to increased corporate leverage) and monetary policy shocks (shocks that lead to an increase in the FED Funds rate), there is a substantial increase in default rates across all the sectors. Default rates increase for 2-4 quarters (2 quarters for Financials and Technology, and 4 quarters for all other sectors), after which they slowly mean revert, for the case of monetary policy shocks, or decline below the long-run average, for the case of balance sheet shocks. The peak response to a balance sheet shock ranges from 4 basis points for Utilities companies to 17 basis points for Capital Industries, with the rest of the

sectors experiencing an increase in default rates of 5-10 basis points. Monetary policy shocks lead to milder increases in default occurrence, with sector default rates reaching a peak of 5-10 basis points, except Financials for which default rates increase by 20 basis points.

According to the assumptions in equations (3.5)-(3.6), corporate defaults are neither normally distributed nor linear functions of observed and unobserved systematic factors. Therefore the typical variance decomposition cannot be applied. To get an approximate variance decomposition of the systematic default rate dynamics, I use instead of the observation equation (3.5), the definition for the transformed probabilities of default (3.6), that is both Gaussian and linear function of the systematic factors \mathbf{F}_t and f_{jt}^d . Using the independent assumption between factors from (3.14), and the AR(1) structure of the frailty factors from (3.7), I define the proportion of the h -step forecast error variance of the probability to default for corporates in sector j and rating grade g that is accounted for by each of the structural shocks and frailty factors as:

$$\omega_{gjk,h}^d = \frac{\mathbf{1}_{k \neq FR} \cdot \left[\sum_{z=1}^5 \left((\beta_j^{d,z})^2 \sum_{i=0}^{h-1} (\phi_{zk,i}^d)^2 \right) \right] + \mathbf{1}_{k=FR} \cdot \left[\rho_{gj}^2 \sum_{i=0}^{h-1} (\vartheta_j^d)^{2h} \right]}{\sum_{m=1}^5 \left[\sum_{z=1}^5 \left((\beta_j^{d,z})^2 \sum_{i=0}^{h-1} (\phi_{zm,i}^d)^2 \right) \right] + \rho_{gj}^2 \sum_{i=0}^{h-1} (\vartheta_j^d)^{2h}}, \quad (3.26)$$

where $\phi_{zm,h}^d = \mathbf{1}_{h \geq \ell_m^d} \cdot \phi_{zk, \max(h-\ell_m^d, 0)}$ is a binary indicator selecting the appropriate element of Φ_h according to the sector specific lag structure provided in table 3.2, and $\mathbf{1}_{k \neq FR} / \mathbf{1}_{k=FR}$ are binary indicators that take the value 1 in case of structural/frailty shocks respectively, and the value 0 otherwise.

Table 3.3 provides the proxy FEVD results for defaults, based on the definition (3.26). Overall, there is clear evidence that the frailty factors play a central role in explaining corporate bond defaults. This finding is not surprising, as there are substantial empirical findings supporting the view that the macroeconomic environment captures only a small portion of corporate default dynamics; Koopman et al. (2012) report that macroeconomic factors account for approximately one third of the systematic variation in corporate defaults. Table 3.3 suggests that, in the short-run, macroeconomic shocks account for approximately 30% of the forecast error variance across grade and sector corporate defaults (portion measured as a weighted average across rating grades/industry sectors, using the total firm count for each segment). When exploring

Figures in %		Quarter 1						Quarter 20					
Variable	Shock	AS	AD	BS	AP	MP	FR	AS	AD	BS	AP	MP	FR
Capital	<i>Ba</i>	1.3	13.8	79.2	0.0	0.0	5.8	6.0	17.3	59.4	12.1	4.3	1.0
Capital	<i>B</i>	0.1	1.5	8.4	0.0	0.0	89.9	2.5	7.3	24.9	5.1	1.8	58.5
Capital	<i>Caa</i>	0.1	1.5	8.5	0.0	0.0	89.8	2.5	7.3	25.1	5.1	1.8	58.2
Consumer	<i>Ba</i>	0.1	0.8	4.8	0.0	0.0	94.3	2.3	5.3	6.5	2.4	0.9	82.5
Consumer	<i>B</i>	0.0	0.5	3.0	0.0	0.0	96.4	1.5	3.5	4.3	1.6	0.6	88.4
Consumer	<i>Caa</i>	0.3	2.8	16.1	0.0	0.0	80.8	5.9	13.8	16.9	6.3	2.4	54.7
Financials	<i>IG</i>	0.0	0.0	0.0	0.0	0.0	100.0	0.2	3.8	0.7	1.3	0.9	93.1
Financials	<i>Ba</i>	0.0	0.0	0.0	0.0	0.0	100.0	0.8	14.9	2.7	5.2	3.4	73.0
Financials	<i>B</i>	0.0	0.0	0.0	0.0	0.0	100.0	2.0	35.6	6.4	12.5	8.2	35.4
Financials	<i>Caa</i>	0.0	0.0	0.0	0.0	0.0	100.0	2.8	50.0	9.0	17.5	11.5	9.3
Technology	<i>Ba</i>	0.0	0.2	0.8	0.0	0.2	98.8	0.2	2.4	4.2	2.3	0.8	90.0
Technology	<i>B</i>	2.5	12.2	49.9	2.6	14.6	18.1	1.8	23.5	41.3	22.9	8.0	2.4
Technology	<i>Caa</i>	0.1	0.7	2.8	0.1	0.8	95.4	0.5	7.2	12.5	7.0	2.4	70.4
Transportation	<i>Ba</i>	2.5	5.4	1.3	3.9	21.5	65.3	1.2	51.4	9.6	16.2	11.6	10.0
Transportation	<i>B</i>	1.0	2.2	0.5	1.6	8.7	85.9	0.9	42.0	7.9	13.3	9.5	26.4
Transportation	<i>Caa</i>	0.7	1.4	0.3	1.0	5.6	90.9	0.8	36.0	6.7	11.4	8.1	37.0
Utilities	<i>IG</i>	1.3	13.7	78.9	0.0	0.0	6.1	0.5	7.1	76.5	10.5	3.4	2.0
Utilities	<i>Ba</i>	0.0	0.4	2.3	0.0	0.0	97.3	0.0	0.6	6.5	0.9	0.3	91.7
Utilities	<i>B</i>	0.0	0.2	1.2	0.0	0.0	98.6	0.0	0.3	3.5	0.5	0.2	95.6
Utilities	<i>Caa</i>	0.7	7.3	42.0	0.0	0.0	50.1	0.4	5.5	59.6	8.2	2.7	23.7

Table 3.3: Default FEVD

Proportion of the h -step forecast error variance of each time series of default rates that is accounted for by the Aggregate Demand (AD), Aggregate Supply (AS), Balance Sheet (BS), Asset Price (AP), and Monetary Policy (MP) structural shocks, as well as shocks to the frailty systematic factors (FR). Variance decompositions are calculated using the proxy equation (3.26), and are provided for the 1st and 20th forecast quarter.

the results of table 3.3 at the grade/sector level, it is apparent that it is mainly the balance sheet shocks that help to explain a significant portion of the default forecast error variance (79% for Ba rated Capital Industry corporates, 50% for B rated Technology firms, and 79%/42% for Investment Grade/Caa rated Utilities companies respectively). Monetary policy shocks only capture 20% of the forecast error variance for Ba rated Transportation corporate defaults and 15% of the forecast error variance for B rated Technology corporate defaults.

When moving to long-run effects of the different shocks, the findings suggest that the importance of macroeconomic shocks increases substantially; on average across rating grades/industry sectors, the macroeconomic environment accounts for approximately 45% of the forecast error variance (again the weighted average is based on using the aggregate firm count as weights). Typically, aggregate supply shocks play a minor role in explaining the forecast error variance. Aggregate demand shocks dominate the forecasts for the Transportation sector (with proportions of total variance ranging from 36%-51%), and the sub-investment grade Financial firms (with proportions of total variance ranging from 15%-50%). Balance sheet shocks are very important in explaining long-run fluctuations in Capital Industries (proportion of total variance between 25% and 59%), high risk Technology (B rated Technology firms variance proportion of 41%,

	<u>Secured</u>	<u>Subordinated</u>	<u>Unsecured</u>		<u>Secured</u>	<u>Subordinated</u>	<u>Unsecured</u>
α_{cap}^{rr}	-0.568 (0.060)	-1.500 (0.086)	-1.058 (0.069)	ρ_{cap}^{rr}	0.227 (0.058)	0.027 (0.190)	0.389 (0.067)
α_{cons}^{rr}	-0.567 (0.056)	-1.545 (0.110)	-1.218 (0.111)	ρ_{cons}^{rr}	0.260 (0.044)	0.707 (0.100)	0.785 (0.084)
α_{fin}^{rr}	-0.524 (0.191)	-1.992 (0.224)	-0.716 (0.172)	ρ_{fin}^{rr}	0.081 (0.185)	1.101 (0.237)	1.192 (0.173)
α_{tech}^{rr}	-0.614 (0.092)	-1.709 (0.128)	-1.557 (0.076)	ρ_{tech}^{rr}	-0.059 (0.096)	0.090 (0.264)	0.099 (0.180)
$\alpha_{transp.}^{rr}$	-0.426 (0.112)	-1.304 (0.220)	-1.358 (0.091)	$\rho_{transp.}^{rr}$	0.191 (0.132)	-0.076 (0.360)	-0.153 (0.117)
α_{util}^{rr}	-0.436 (0.080)	-1.401 (0.153)	-0.635 (0.088)	ρ_{util}^{rr}	0.201 (0.122)	0.377 (0.268)	0.399 (0.083)
$\beta_{y_{t-1}}^{rr}$	0.149 (0.052)	- -	- -	ϑ^{rr}	0.273 (0.226)	-0.095 (0.358)	0.093 (0.141)
$\beta_{y_{t-2}}^{rr}$	- -	0.173 (0.076)	0.165 (0.057)	σ^{rr}	0.653 (0.015)	1.212 (0.028)	0.761 (0.012)
$\beta_{eq_t}^{rr}$	0.011 (0.005)	0.039 (0.008)	0.028 (0.005)				

Table 3.4: Parameter Estimates for Recovery Rates

Parameter estimates and corresponding standard errors for the recovery rate Gaussian specification (3.10)-(3.11). All estimates are based on 100,000 draws from the full posterior distribution, using the MCMC algorithm described in section 3.3.4. Due to the majority of sensitivities and AR(1) coefficients of the frailty factors not being statistically significant, the statistically significant estimates for those quantities are bolded.

Caa rated Technology firms variance proportion of 13%), and Investment Grade/Caa rated Utilities firms (variance proportion of 77% and 60% respectively). Finally, asset price and monetary policy shocks are mainly important in explaining long-run fluctuations in Transportation sector defaults (11%-16% variance proportion for asset price shocks, and 8%-12% variance proportion for monetary policy shocks), and to a lesser extent high risk Financial firms (B-Caa rated Financial firms variance proportion 13% and 18% respectively for asset price shocks), and B rated Technology firms (23% variance proportion for asset price shock).

Recovery Rates

Moving to the recovery rate findings, table 3.4 provides the parameter estimates for the set of equations (3.10)-(3.11). Similarly to the corporate default specification, the results are based on the highly uninformative Normal priors for intercepts and sensitivities to systematic factors $p(\alpha_{j_s}^{rr}) \sim N(0, 10^4)$, $p(\beta_j^{rr}) \sim N(0, 10^4)$, and $p(\rho_{j_s}^{rr}) \sim N(0, 10^4)$, and moderately informative truncated Normal priors for the AR(1) coefficients $p(\vartheta_s^{rr}) \sim N(0, 10)I(1, 1)$. Finally, for the error variances, I assume the relatively uninformative priors $p((\sigma_s^{rr})^2) \sim IG(10^{-2}, 10^{-2})$. Inference is based on 100,000 MCMC draws, fol-

lowing a burn-in sample of 10,000 draws. Similarly to the default specification, the lag polynomial for the macroeconomic sensitivities $\beta_s^{rr}(L)$ in (3.10) takes the simplified form $\beta_j^{rr}(L) = [\beta_j^{rr,1} L^{\ell_1^{rr}}, \dots, \beta_j^{rr,5} L^{\ell_5^{rr}}]$, where ℓ_m^{rr} is the lag operator order for variable m . The seniority specific structure of the F_t matrix is chosen using the 5% significance level as the cut-off point for the inclusion of a variable in the final specification.

A first remark about the results in table 3.4 is that the parameter estimates for the intercepts across the 3 seniority levels are monotonically decreasing when moving from high to low collateralisation debt, across all sectors. The relative order of recovery rates across the 3 seniority levels is a desirable feature, as it highlights that the long-run average recoveries are intuitive within each sector. The second remark is that, unlike defaults, the importance of the macroeconomic environment in explaining recovery rates appears significantly weaker, resulting in a much smaller set of variables included in the final specification. This is not entirely unexpected, since recovery rates do not typically refer to a well defined period, but strongly depend on the type of debt and the length of the resolution process (or, in this case, the market's expectation of the time to recover). Since the time to recover is not uniform, recovery rates might correspond to multiple states of the business cycle. Despite this misalignment between recoveries and business cycle, I find strong statistical evidence for including real GDP and equity price growth rates in the econometric specification (3.10); the former enters with 1 quarter lag for secured debt recoveries, and 2 quarter lag for recoveries on unsecured and subordinated debt. While for secured debt it is mainly the GDP growth that drives the results, the recovery rates for unsecured and subordinated debt appear to depend more on the equity market movements.

In addition to the macroeconomic environment, the dependence of recovery rates on the frailty factors is fundamentally different across secured and unsecured/subordinated debt. The frailty factor for secured debt loads almost uniformly on the different sectors (the only major exception is the Technology sector, which has a marginally negative loading). On the contrary, for unsecured/subordinated debt (loadings are remarkably similar across these 2 seniority levels), it is the Financial and Consumer sectors that dominate the construction of the frailty factors, with Capital Industries and Utilities being the only other sectors that have statistically significant loadings for unsecured recovery rates. Finally, the macroeconomic variables are able to remove the autocorre-

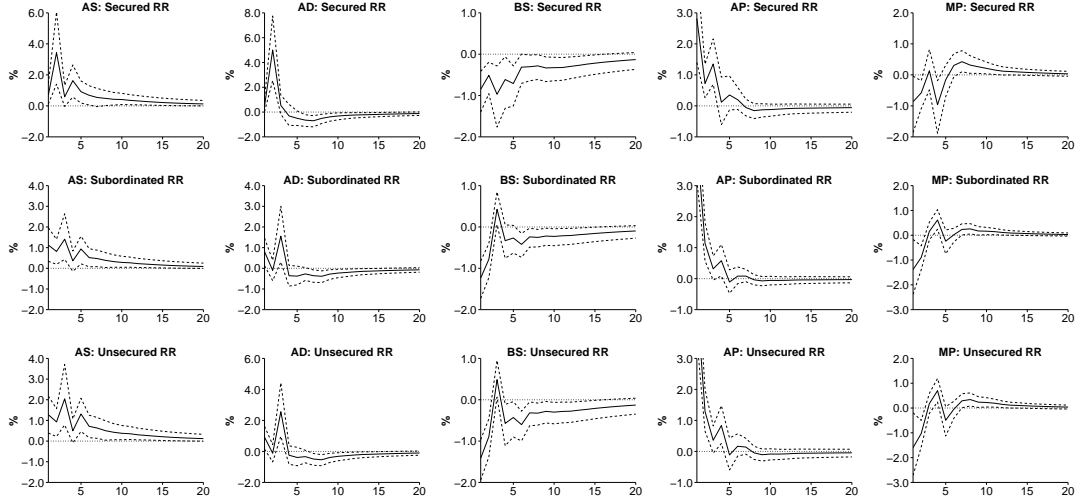


Figure 3.5: Corporate Debt Recovery Impulse Responses Functions

Impulse-Response functions of the log recovery rates across the 3 seniority levels (RR) to each of the 5 structural shocks. The magnitude of the shocks is normalized to 1σ . AS denotes an Aggregate Supply shock, AD an Aggregate Demand shock, BS a Balance Sheet shock, AP an Asset Price shock, and MP a Monetary Policy shock. Vertical scales depict the difference from the long-run average level of log recovery rates in percentage terms. Horizontal scales depict time in quarters. Median impulse responses (solid lines) are reported with 1 standard deviation error bands (dotted lines), computed from 100,000 MCMC samples.

lation from the time series of recovery rates as evidenced by the reported estimates for the AR(1) coefficients ϑ_s^{rr} (for unsecured and subordinated debt the estimates are very close to 0, while for secured debt the estimate of 27% appears to be not statistically significant).

Figure 3.5 provides the impulse response functions for the log of corporate recovery rates $\log(RR_{it})$, for each of the 3 seniority levels. The impulse response functions are calculated as:

$$\text{IRF}_{sk,h}^{rr} = \sum_{m=1}^5 \beta_s^{rr,m} \mathbf{1}_{h \geq \ell_m^{rr}} \phi_{mk, \max(h - \ell_m^{rr}, 0)}, \quad (3.27)$$

where s is the seniority level, h is the forecast horizon, $\phi_{mk,h}$ is the (m, k) element of the matrix Φ_h , and $\mathbf{1}_{h \geq \ell_m^{rr}}$ is a binary indicator that takes the value 1 when the forecast horizon is higher or equal to the lag for variable m in the specification for seniority level s , and the value 0 otherwise. Equation (3.27) is evaluated at each MCMC draw for $\beta_s^{rr,m}$, and in figure 3.5 I depict median responses and 1 standard-error bands. The shape of the responses is quite similar across the 3 seniority levels. As expected based on the estimates of table 3.4, the asset price shock has the biggest impact, with recovery rates rising by approximately 3%.

Both aggregate supply and demand shocks lead to a substantial increase in recovery

rates over the first year; for secured debt, the instantaneous response is a level of recovery rates approximately 2% higher than the long-run average, while for unsecured and subordinated debt, recovery rates rise approximately 1% higher than the long-run average. Recovery rates following an aggregate supply shock stay persistently positive, while the effect of the aggregate demand shock becomes statistically not different from 0 from the second year onwards.

A positive balance sheet shock and the resulting drop in corporate net worth, as firms struggle to find funding for new investment, cause an immediate drop in recovery rates to a level approximately 1% lower than the historical average. The time profiles of the responses to balance sheet shocks are significantly different across secured and non secured bonds; while for the former type of bonds the drop in recovery rates is persistent through time, for the latter type of bonds the recovery rates temporarily rise 1 year after the initial shock. This difference is based on the higher sensitivity of non secured debt to equity market movements; as the cash flows for unsecured and subordinated debt are more uncertain, traded prices for defaulted bonds are more likely to be driven by fluctuations in aggregate equity prices. The moderate increase in real equity prices 1 year after a balance sheet shock, due to a combination of falling inflation and interest rates, appears to drag upwards the recovery rates. Nevertheless, it is worth pointing out that this relationship might be biased by the use of traded prices for defaulted debt as a proxy for recovery rates; results based on ultimate recoveries could potentially shed more light.

Finally, the effect of a monetary policy shock is relatively mild, having a statistically significant impact for the first 3 quarters. The initial decrease in recovery rates following the higher level of interest rates ranges from 0.9%-1.5% and reflects the different discounting the market applies to future cash flows across the 3 seniority levels.

For the forecast error variance of the recovery rates for sector j and seniority level s I use the following definition:

$$\omega_{jsh}^{rr} = \frac{\mathbf{1}_{k \neq FR} \cdot \left[\sum_{z=1}^5 \left((\beta_s^{rr,z})^2 \sum_{i=0}^{h-1} (\phi_{zk,i}^{rr})^2 \right) \right] + \mathbf{1}_{k=FR} \cdot \left[\rho_{js}^2 \sum_{i=0}^{h-1} (\vartheta_s^{rr})^{2h} \right]}{\sum_{m=1}^5 \left[\sum_{z=1}^5 \left((\beta_s^{rr,z})^2 \sum_{i=0}^{h-1} (\phi_{zm,i}^{rr})^2 \right) \right] + \rho_{js}^2 \sum_{i=0}^{h-1} (\vartheta_s^{rr})^{2h}}, \quad (3.28)$$

where $\phi_{zm,h}^{rr} = \mathbf{1}_{h \geq \ell_m^{rr}} \cdot \phi_{zk, \max(h - \ell_m^{rr}, 0)}$. For comparability with the default rates I

Figures in %	Quarter 1						Quarter 20					
Variable\Shock	AS	AD	BS	AP	MP	FR	AS	AD	BS	AP	MP	FR
Secured <i>Cap.</i>	0.3	0.1	0.5	4.5	0.5	94.1	5.5	11.7	2.2	4.4	1.7	74.6
Secured <i>Cons.</i>	0.2	0.1	0.4	3.5	0.3	95.5	4.4	9.5	1.8	3.6	1.4	79.4
Secured <i>Fin.</i>	1.8	0.8	2.7	25.1	2.5	67.2	15.7	33.4	6.2	12.5	4.9	27.3
Secured <i>Tech.</i>	2.6	1.2	3.9	36.7	3.7	52.0	18.0	38.3	7.2	14.4	5.6	16.6
Secured <i>Trans.</i>	0.4	0.2	0.7	6.2	0.6	91.8	7.0	15.0	2.8	5.6	2.2	67.3
Secured <i>Util.</i>	0.4	0.2	0.6	5.7	0.6	92.6	6.5	14.0	2.6	5.2	2.0	69.6
Subordinated <i>Cap.</i>	5.3	2.5	8.0	75.1	7.6	1.7	11.5	20.4	9.6	47.8	9.7	1.0
Subordinated <i>Cons.</i>	0.4	0.2	0.6	6.0	0.6	92.1	1.5	2.7	1.3	6.3	1.3	87.0
Subordinated <i>Fin.</i>	0.2	0.1	0.3	2.6	0.3	96.6	0.7	1.2	0.6	2.8	0.6	94.2
Subordinated <i>Tech.</i>	4.5	2.1	6.8	64.2	6.5	16.0	10.4	18.6	8.7	43.5	8.8	9.9
Subordinated <i>Trans.</i>	4.7	2.2	7.1	67.4	6.8	11.8	10.8	19.1	9.0	44.9	9.1	7.1
Subordinated <i>Util.</i>	1.2	0.6	1.9	17.7	1.8	76.8	4.0	7.1	3.3	16.6	3.4	65.6
Unsecured <i>Cap.</i>	0.7	0.3	1.1	9.9	1.0	87.0	3.3	6.3	2.2	9.8	2.2	76.3
Unsecured <i>Cons.</i>	0.2	0.1	0.3	2.7	0.3	96.5	1.0	1.9	0.7	2.9	0.6	92.9
Unsecured <i>Fin.</i>	0.1	0.0	0.1	1.2	0.1	98.4	0.4	0.8	0.3	1.3	0.3	96.8
Unsecured <i>Tech.</i>	3.7	1.7	5.7	53.4	5.4	30.1	11.5	21.9	7.8	34.1	7.5	17.2
Unsecured <i>Trans.</i>	2.6	1.2	4.0	37.6	3.8	50.8	9.3	17.7	6.3	27.5	6.1	33.2
Unsecured <i>Util.</i>	0.7	0.3	1.0	9.5	1.0	87.6	3.2	6.0	2.1	9.4	2.1	77.2

Table 3.5: Recovery FEVD

Proportion of the h -step forecast error variance of each time series of recovery rates that is accounted for by the Aggregate Demand (AD), Aggregate Supply (AS), Balance Sheet (BS), Asset Price (AP), and Monetary Policy (MP) structural shocks, as well as shocks to the Frailty systematic factors (FR). Variance decompositions are calculated at the average level of recovery rates using the equation (3.28), and are provided for the 1st and 20th forecast quarter.

exclude $(\sigma_s^{rr})^2$ from (3.28), interpreting variance decomposition with respect to only observed and unobserved systematic factors (that implies a variance decomposition on the average level of recovery rates). Based on the definition (3.28), table 3.5 summarises the FEVD at the industry sector and seniority level. Macroeconomic shocks account for approximately 24% of the forecast error variance for recovery rates on average in the short-run (weighted average across all seniority/sector segments, using the number of defaults as weights), and approximately 34% in the long-run. These findings clearly show that the recovery rates are much less impacted by macroeconomic conditions as compared to corporate defaults, and that the systematic frailty factors are essential when trying to capture short and long-run dynamics of losses following corporate defaults. Exploring the variance decomposition at the sector/seniority level suggests that recovery rates are mainly impacted by asset price and aggregate demand shocks; the former type of shocks is more evident in unsecured and subordinated debt recoveries, while the latter type of shocks mainly drives secured debt recoveries. The effect of asset price shocks is particularly strong in the short-run, though the results might be biased due to the use of traded prices of default rates as a proxy for ultimate recoveries.

As described in section 3.4, one of the drawbacks of using 30 days post default prices as proxies for the recovery rates is that it might inflate the dependence on the equity market. This is definitely a factor that should be taken into consideration when

comparing the strong dependence of corporate defaults to the real economy and balance sheet shocks and the increased dependence of recovery rates on the asset price shock. Nevertheless, the depressed dependence of recoveries on real economy and balance sheet shocks should not be solely attributed to the use of trading prices. The use of ultimate recoveries instead does not necessarily lead to higher dependence on the business cycle. This is mainly because recovery rates measured at resolution time would typically depend on the economic conditions more than a year from the time of default (corporate debt workout periods post bankruptcy can be extremely lengthy). Therefore, finding any statistical significance in correlations with the economic activity at time of default might be as difficult, if not even more difficult, when moving away from trading price recovery rates.

3.6 Conclusion

In this chapter I introduce a new identification scheme to isolate real economy and macro-financial shocks. The scheme combines a set of short and long-run restrictions on the residuals of a 5 variable VAR, and leads to intuitive response functions to aggregate supply, aggregate demand, corporate balance sheet, asset price, and monetary policy shocks. The identification of the corporate balance sheet shock leads to credit channel/financial accelerator type of effects in the endogenous macro-financial system. In addition, I have reported strong evidence that corporates are able to partially pass on changes in their funding costs to consumers, therefore creating a cost channel mechanism in the economy. This cost channel mechanism causes balance sheet, asset price, and monetary policy shocks to have a material impact on inflation.

The identification scheme has been then used to quantify the impact of the structural shocks on corporate default likelihood and recovery rate post default. On average, default rates are more sensitive to balance sheet and aggregate demand shocks, especially for the cyclical sectors. Contractionary monetary policy shocks typically increase defaults by 5%-20%, with the effect taking 3-4 quarters to materialise. The effect of monetary policy shocks is more pronounced for financial firms that have their income stream tied to mortgage lending. Recovery rates are more sensitive to asset price shocks; this sensitivity might be partially biased by the use of traded price of defaulted debt as

a proxy for ultimate recoveries. On average, macroeconomic shocks account for approximately 30% and 45% of the short-run and long-run forecast error variance of corporate defaults respectively, while for the recovery rates, the macroeconomic shocks explain approximately 24% and 34% of short and long-run forecast error variance respectively.

The work presented in this chapter can be extended in a number of ways. For the identification of the macroeconomic shocks, using large scale FAVAR or dynamic factor models, similar to Bernanke et al. (2005) and Boivin et al. (2013), can help to address the issue of omitted variables from the VAR. The transmission mechanism for many shocks can be complicated and the information set available to agents can be very large as compared to the, typically low, dimensionality of VARs. Therefore, principal component based factor models have been suggested in the literature as appealing and computationally feasible alternatives to traditional VAR analysis. Despite the appealing features of high dimensionality and low computational burden, further research is required to incorporate long-run impact restrictions of structural shocks on individual input variables when extracting factors via principal components.

Concerning the credit risk analysis, more complex non-Gaussian dynamic factor structures can be used to jointly model defaults, recoveries, and potentially rating migrations, along the lines of Creal et al. (2013). This type of joint modelling of defaults, recoveries and migrations captures more accurately the complicated correlation structures across the different components of credit losses in a portfolio setting. Furthermore, when combined with macroeconomic variables, the mixed-measurement dynamic factor models of Creal et al. (2013) can simultaneously capture the interactions of credit specific shocks with the macroeconomic environment. Nevertheless, when moving away from recursive identification schemes for structural shocks, the size of the resulting state space model and the complexity of the required calculations, quickly turn the estimation computationally infeasible.

Chapter 4

Momentum in Credit Rating Downgrades: Through-the-Cycle Evidence

4.1 Introduction

External credit agencies and their ratings continue to play a very important role in today's corporate finance and risk management, despite the criticism they have been subject to in the aftermath of the recent credit crisis. By providing an independent assessment of relative credit risk for corporate debt, credit agencies are key determinants of credit instrument yields, credit portfolio profit & loss distributions, and regulatory capital for banks and financial institutions. The Markov assumption offers computational advantages and, therefore, is typically used in practice to model the stochastic evolution of the entire set of discrete ratings. According to the Markov model, the information needed to predict a rating migration does not extend beyond the current rating of a firm. Nevertheless, this Markovian assumption is highly questionable empirically, an observation that was highlighted more than 20 years ago by Altman and Kao (1992).

In this chapter I focus on a specific type of non-Markovian behaviour in corporate credit ratings, the momentum effect or serial dependence. While estimating the mo-

momentum effect has attracted some interest over the past decade this is the first study that explores the stability of this rating momentum effect over time. More specifically, I am interested in testing whether the business cycle and the time since rating assignment enhance or weaken the directional serial dependence that is observed in credit ratings. Following the vast body of literature, I base the analysis on continuous time, as it allows to better capture the probabilities of rare transitions and to fully reflect the available information provided in the data. Using migrations from Moody's Default & Recovery Database, I statistically test for 3 types of effects:

- Rating momentum: This is the Through-the-Cycle serial dependence effect, measuring the average increase in likelihood of observing a down(up)grade for firms that reached their current rating grade from a higher(lower) grade. While, this is the same effect as in Lando and Skødeberg (2002), in this study I adjust the estimates for industry heterogeneity. Consistently with prior studies, I report a very strong downward drift (that can lead to an average increase of 4 times the baseline intensity) while the upward drift is significantly milder (with the average increase to be approximately twice the baseline intensity).
- Rating momentum adjusted for duration time: This effect adjusts the directional dependence for the amount of time a firm has spent in a given rating grade. I show that the duration dependence is highly non-monotonic and I model it with flexible fractional polynomials. The duration adjustment is very strong for both downgrades and upgrades, albeit exhibiting different time profiles across investment and speculative grade firms. The downward drift for investment grade firms is approximately exponentially decaying with time, starting from a point 6 times higher than the baseline downgrade intensity. On the contrary, the time profile of the upward drift for investment grade firms is monotonically increasing at a diminishing rate of change (negative second derivatives). For sub-investment grade firms, both upward and downward drift have similar time profiles, with a rapid rise of the momentum effect to peak within 2 months following a rating change. This peak level is much higher for the downward drift (7.5 times the baseline intensity) as compared to the upward drift (4 times the baseline intensity).
- Rating momentum adjusted for business cycle: This effect captures cyclical movements in rating momentum. Without loss of generality, I use the CFNAI index

as a proxy for aggregate business cycle fluctuations. The analysis shows that, for investment grade firms, the time profile of both upward and downward rating momentum remains largely unaffected across the business cycle. On the contrary, for speculative grade firms, business cycle dependence of the momentum effect is statistically significant for both upgrades and downgrades. As the economic conditions improve, the downward momentum effect peaks earlier and decays faster than during periods of stress. Furthermore, the peak of the downward drift becomes higher as the economy deteriorates and can exceed 8 times the baseline intensity level, during periods of extreme stress. Finally, for speculative grade upgrades, the upward drift has a clear positive dependence to the state of the business cycle, with the time profile shifting proportionally to approach a level of 6 times the baseline under extremely benign economic conditions.

All 3 effects are tested in addition to the base business cycle effect. Due to significantly smaller sample coverage, I refrain from using watch-lists and outlooks. Despite their clear additional information, they are available for a much shorter historical time period and cover a significantly smaller set of firms as compared to the actual ratings.

To quantify the impact of this chapter's findings, I examine a few key metrics, particularly focusing on the implications for financial institutions risk management. For a financial institution's banking book, I use the Basel II/III RWA prescribed formula to capture the impact on regulatory capital, while the results for provisions and economic capital are based on jump-to-default expected and 99% tail losses. Finally, for the trading book I also cover expected and tail marked-to-market losses, resulting from a combination of credit spread movements and jump-to-default events. I examine the behaviour of each of the metrics across 2 corporate bond portfolio structures, that reflect high (primarily sub-investment grade bonds) and low (predominately investment grade bonds) risk exposures respectively.

Using a business cycle adjusted transition matrix as the base case, I show that the regulatory capital does not strongly depend on the presence of momentum in credit ratings; large differences are mainly observed for high credit quality portfolios during periods of extreme stress. Unlike capital requirements, expected actual losses are significantly affected by the presence of rating momentum. During periods of stress, momentum adjusted 1-year actual losses are close to 11% higher than the momentum

insensitive case; over the 2-year period, the relative impact of rating momentum is close to the 20% mark. A significant proportion of this difference can be attributed to the business cycle and duration time dependence of the momentum effect; using the typical proportional rating momentum effect as in Lando and Skødeberg (2002) can lead to 1.5-4 times lower losses as compared to the non-proportional specification of this chapter. Furthermore, rating momentum has a higher absolute impact on marked-to-market losses during the severe stress period 2008-2009 as compared to jump-to-default losses (3.5 times the absolute loss amount for 1-year jump-to-default losses and 1.8 times the 2-year loss figure). Finally, the absolute momentum impact on 99% tail losses is of similar magnitude to that for the average losses; it is only during periods of stress that the absolute impact for tail losses can be 15%-20% higher than the impact on average losses.

The econometric framework and the details of the estimation process are described in section 4.3. Section 4.4 details the data I use for the analysis and summarises the key features of historical credit rating transitions. I provide the parameter estimates of the econometric specification in section 4.5.1. The credit portfolio assumptions and the impact of rating momentum on portfolio credit losses are presented in section 4.5.2. Finally, section 4.6 concludes.

4.2 Relevant Literature

Deviations from the Markovian behaviour typically take the form of directional serial dependence in rating actions and duration dependence from the point a rating is assigned. Serial dependence is normally associated with a downward drift in subsequent rating migrations. This rating drift is a consequence of the dynamic evolution of a firm's financial strength. If a firm's financial condition is deteriorating, then eventually it will be downgraded. If it does not take the necessary actions to increase profitability and reduce debt then it will cause its rating to be changed downwards again. In parallel, the rating change itself might have an adverse effect on this domino effect. Avramov et al. (2007) show that there is a strong momentum in profitability for low rated public firms with short selling strategies generating statistically significant positive payoffs. This empirical evidence provides an indication that as firms are downgraded to lower rating

grades, market pressure causes their asset values to decrease, which in return might cause further deterioration of credit quality. This feedback effect is further empirically supported by Campbell and Taksler (2003) and Campbell et al. (2008) who show that low rated firms exhibit abnormally low stock returns and high equity volatility that can cause as much cross-sectional variation in corporate bond yields as can rating changes. Duration dependence is the result of two opposing effects. First, maintaining a stable financial condition becomes less likely the longer a firm stays in a given rating grade; therefore downgrades (for cases of deteriorating financial strength) and upgrades (for cases of improving financial strength) become more possible. These types of duration effects are closely related to firm ageing and life-cycle effects, see Agarwal and Gort (2002) and the references therein. A second type of duration effects, this time negative, is related to the reluctance of rating agencies for multi-notch rating changes. In this case, firms that will eventually be downgraded or upgraded by multiple notches, spend a short amount of time in the intermediate grades. Which of the two types of effects prevails at the end and what is the actual time profile of the duration effect remains unclear, despite some evidence from Lando and Skødeberg (2002) for an overall negative effect.

Early studies on the non-Markovian rating migrations literature, including Kavathas (2000), and Lando and Skødeberg (2002), use continuous time rating migration and default frameworks to test for the presence of duration and rating momentum effects. The 2 studies report an overall negative dependence to duration time and a strongly positive correlation between successive migrations of the same type. Christensen et al. (2004) use a hidden Markov approach to incorporate the rating momentum as "excited states" when estimating rating transition and default probabilities conditional on previous downgrades. Using a similar framework, Güttler and Raupach (2010) show that ignoring the rating momentum effect leads to an underestimation of VaR by 107 basis points on average. Generalising the serial dependence literature, Güttler and Wahrenburg (2007) show that momentum effects are not only observed in single agency ratings. Using joint data from Moody's and S&P, they provide evidence that the rating momentum is shared among the rating agencies; the direction and the magnitude of a rating change from one agency causes a similar action from the other rating agencies.

Introducing non-Markovian elements in transition matrices makes the generation

of long run predictions a computationally intensive task. Another approach suggested in the literature is to expand the data to also include watch-lists and rating outlooks in addition to rating changes. Hamilton and Cantor (2004) show that adjusting credit ratings under review by two notches and adjusting ratings in the outlook list by one notch helps to increase the accuracy of rating predictions. Evidence provided by Hull et al. (2004) from the credit default swap market support this finding. They show that while additions to the watch-list are informative for explaining credit default swap movements, the eventual rating downgrades are not, indicating that the market prices the rating migrations the moment a company is watch-listed. Bannier and Hirsch (2010), using data from Moody's, argue that watch-lists are used to enhance the information flow for better rated firms, while for worse rated firms, they are used as a controlling mechanism to incentivise towards a lower risk taking behaviour.

Since portfolio level rating migrations should reflect movements in the aggregate default rates, transition matrices are expected to exhibit cyclical variations. Past historical experience shows that the credit cycle follows approximately the business cycle and, therefore, business cycle variables could help explain stochastic behaviour of rating migrations. Early studies that identified the increase in downgrade frequency during economic downturns include Nickell et al. (2000), and Bangia et al. (2002). In a more recent study, Stefanescu et al. (2009), using S&P data, show that including observed business cycle and equity price index variables increases the out of sample performance of rating migration and default rate forecasting models. Fei et al. (2012) use a hidden Markov chain as a proxy for the state of the business cycle and model the rating migration matrix as a Mixture of Markov chains. At the same time they assess the implications for financial institutions and they argue that properly incorporating the cyclical dynamics of rating migrations helps to decrease the level of procyclicality in capital requirements. Finally, Figlewski et al. (2012) explore the effect of general macroeconomic conditions on default and rating migrations using a continuous time Cox proportional hazards specification by controlling for rating momentum and ageing effects. For their sample, covering all corporate issuers over 1981-2002, they report unstable macroeconomic sensitivities that are strongly affected by the variable selection process.

While the concept of serial dependence in credit ratings has been explored in great

length by the aforementioned studies, there is currently no study that looks into how this momentum effect changes with time. Therefore, this chapter directly extends the work of Lando and Skødeberg (2002) and Figlewski et al. (2012) to account for a time variation in the momentum effect, both in survival (the time it takes a rated company to migrate from one rating to another) and calendar time. For the calendar time variation I employ the CFNAI economic index and I explore the interaction of the momentum effect with the business cycle, extending in that way the work of Nickell et al. (2000), Bangia et al. (2002) and Stefanescu et al. (2009) that do not consider the serial dependence in the analysis.

The second contribution of this chapter is to I extend the Cox proportional hazard specification of Lando and Skødeberg (2002) in 2 very important ways in order to make statistical inference more robust. First, including the interaction of rating momentum with the time a firm has spent in its current rating grade corrects the semi-parametric specification for the non-proportionality of the rating momentum effect. Ignoring this non-proportional effect leads to the violation of the model's assumptions and therefore to biased inference. Second, I allow for industry sector heterogeneity among firms. The unadjusted Cox model assumes the same term structure of rating migrations for all firms. As I show, the time profiles across different industries can deviate significantly and, therefore, adjusting for sector heterogeneity enhances the robustness of the reported results.

The final contribution of this chapter is to quantify the impact of the rating momentum effects on credit portfolios. This is an area that academic studies typically ignore. The two most relevant studies that quantify the different effects for credit risk management are Fei et al. (2012) and Figlewski et al. (2012). Nevertheless, the first study does not deal with serial dependence, while the second study treats only static momentum effects. This is the first study that quantifies momentum and business cycle effects (and their interaction) in a model consistent manner. Clearly, relaxing the Markovian assumption underlying the stochastic movements of rating grades and introducing duration and business cycle dependence has implications for credit portfolio risk measurement and management. For lending institutions, losses for loan exposures need to be forecasted for accurate provisioning, pricing, risk allocation and regulatory compliance. Traded credit positions need to be marked-to-market to reflect spread

movements and calculate appropriate hedges. Credit ratings are used to both differentiate the default likelihood of different obligors and determine the market prices of credit products. Therefore, the findings in this chapter have implications for both financial institutions and individual investors alike.

4.3 Econometric Framework

4.3.1 Specification

The econometric analysis is based on historically observed rating changes between the discrete set of rating grades $\mathbb{K} = \{1, 2, \dots, K\}$ for a set of $i = 1, \dots, n$ firms. The alphanumeric rating grades (Aaa, Aa1, ..., C) are mapped to the numeric set of grades \mathbb{K} in order of decreased creditworthiness, so that Aaa corresponds to grade 1, Aa1 corresponds to grade 2 etc. To make full use of the available data, I use continuous time, as rating changes are recorded on a daily basis.

I assume that for each firm i there exists a right continuous counting process $\{N_{jki}(t) : t \geq 0\}$, starting from value 0 when firm i is assigned the rating j at $t = 0$. The counting process $N_{jki}(t)$ increases by 1 each time firm i migrates from grade j to grade k and its value set is finite, $N_{jki}(t) < \infty$. The specific time that firm i experiences a rating change to k is defined as T_{ki} (T_{ki} is a vector when firm i has multiple rating changes to grade k). Finally, I denote by $i \in \mathcal{R}_j(t)$ that firm i has rating j at time t and therefore is "at risk" for all migrations from grade j . Using these definitions, I express both firm-specific, $N_{jki}(t)$, and aggregate rating migration type, $N_{jk}(t)$, counting processes as follows:

$$N_{jk}(t) = \sum_{i=1}^n N_{jki}(t), \tag{4.1}$$

$$N_{jki}(t) = I\{T_{ki} \leq t : i \in \mathcal{R}_j(t)\}.$$

In (4.1), I slightly abuse the notation for i , to take a separate value each time firm $i \in \mathcal{R}_j(t)$, therefore forcing each of the individual point processes $N_{jki}(t)$ to take values in $\{0, 1\}$. All point processes $N_{jki}(t)$ are adapted to the filtration \mathcal{F}_t , which summarises all the available information at time t . In a survival analysis setting, the filtration \mathcal{F}_t

typically consists of the known covariate values at time t and the number of events prior to t . Conditional on the filtration \mathcal{F}_t , rating events, and consequently the counting processes $N_{jki}(t)$, are independent. Since $N_{jki}(t)$ captures by construction at most 1 event, past events do not enter the filtration and I express $\mathcal{F}_t = \mathbf{Z}_{jki}(t)$, for a covariate vector $\mathbf{Z}_{jki}(t)$.

Each point process $N_{jki}(t)$ has a corresponding intensity $\lambda_{jki}(t)$, capturing the instantaneous migration probability from rating grade j to rating grade k . In terms of the point process $N_{jki}(t)$, the intensity $\lambda_{jki}(t)$ is defined as

$$\lambda_{jki}(t) = \lim_{\Delta t \rightarrow 0} \frac{P(N_{jki}(t + \Delta t) - N_{jki}(t) = 1 | \mathcal{F}_t)}{\Delta t}.$$

I assume a multiplicative form to link the covariate vector $\mathbf{Z}_{jki}(t)$ to the migration intensity and therefore $\lambda_{jki}(t)$ can be written as follows:

$$\lambda_{jki}(t) = \lambda_{jk}^0(t) e^{\beta_{jk} \mathbf{Z}_{jki}(t)}, \quad (4.2)$$

where $\lambda_{jk}^0(t)$ is the baseline intensity (intensity when explanatory variables are 0). The intensity form in (4.2) has the property that changes in covariates result in intensities proportional to the baseline level (commonly referred to as proportional hazard model in the literature). The resulting density function for the time spent in a given rating grade takes the form

$$f_{jki}(t | \mathbf{Z}_{jki}(t)) = \lambda_{jk}^0(t) \exp(\beta_{jk} \mathbf{Z}_{jki}(t)) \exp\left(-\exp(\beta_{jk} \mathbf{Z}_{jki}(t)) \int_0^t \lambda_{jk}^0(u) du\right),$$

where $\Lambda_{jk}^0(t) = \int_0^t \lambda_{jk}^0(u) du$ is the baseline cumulative intensity up to time t for the migration from rating j to rating k . The corresponding survival function can be expressed as:

$$S_{jki}(t | \mathbf{Z}_{jki}(t)) = S_{jk}^0(t) e^{\beta_{jk} \mathbf{Z}_{jki}(t)},$$

where $S_{jk}^0(t) = \exp(-\Lambda_{jk}^0(t))$ is the baseline survival function for the migration from rating j to rating k . For a more detailed exposition of the counting process approach to survival analysis, see Andersen et al. (1993), Aalen et al. (2008), and Fleming and Harrington (2011).

Since there is no theory behind the shape of the baseline intensity for rating migra-

tions, I leave it unspecified. Leaving the baseline unspecified circumvents any estimation bias caused by choosing a parametric form that does not fully reflect the empirical data. Despite not giving a parametric form to the baseline, due to the proportionality assumption, intensity ratios can be used to draw useful conclusions. From the estimates of β_{jk} in (4.2) I derive the relative migration intensity across different values of an explanatory factor. For example, if the explanatory factor of interest in (4.2) is the rating momentum indicator Z_{jki}^m (reflecting whether a given firm has reached its current rating via a downgrade), then the ratio of the intensities for 2 discrete values of the indicator leads to the relative intensity

$$\frac{\lambda_{jk}^0(t)e^{\beta_{jk}^m\{Z_{jki}^m=1\}}}{\lambda_{jk}^0(t)e^{\beta_{jk}^m\{Z_{jki}^m=0\}}} = e^{\beta_{jk}^m(1-0)} = e^{\beta_{jk}^m}, \quad j > k \quad (4.3)$$

Therefore, if the coefficient β_{jk}^m is 0.7 then firms that have reached the current rating grade via a downgrade are twice more likely to be downgraded again as compared to firms that have not been previously downgraded (since $\exp(0.7) \approx 2$).

The rating migration pair specific baseline intensities $\lambda_{jk}^0(t)$, $j, k \in \mathbb{K}$ are assumed to be different across industry sectors $g = 1, \dots, G$, despite being left unspecified. Differentiating the baseline intensities by industry sector adjusts the parameter estimates for sector heterogeneity without increasing the parameter set. Alternatively, sector heterogeneity can be introduced via industry specific fixed or random effects (intercepts or slopes). Nevertheless, both fixed and random effects approaches assume there are enough observations for each sector and migration pair to make inference. The benefit of introducing industry heterogeneity via baseline intensity stratification is that elements of β_{jk} are still estimated jointly across sectors; therefore, getting parameter estimates in cases of missing observations for a subset of sectors in a given migration pair is still feasible.

Once the parameters β_{jk} have been estimated, I obtain smoothed, non parametric, estimates of the common baseline intensities $\lambda_{jk}^0(t)$ by using the standard kernel smoothing methodology

$$\lambda_{jk}^0(t) = \frac{1}{b} \sum_{\{i: T_{ki} \leq t\}} \mathcal{K}\left(\frac{t - T_{kl}}{b}\right) \frac{1}{|\{h \in \mathcal{R}_j(T_{ki})\}|}, \quad (4.4)$$

where $|\{h \in \mathcal{R}_j(T_{ki})\}|$ denotes the size of the risk set at each transition time, the kernel \mathcal{K} is a bounded function on $[-1, 1]$, which is 0 outside $[-1, 1]$ and integrates to 1 over this interval, and b is the kernel bandwidth that determines how much rating migration events away from t would influence the intensity at t . For \mathcal{K} I choose the Epanechnikov kernel

$$\mathcal{K} = \begin{cases} \frac{3}{4} (1 - \frac{1}{5}x^2) / \sqrt{5} & \text{if } |x| < \sqrt{5} \\ 0 & \text{otherwise.} \end{cases}$$

For the bandwidth b I choose the value that minimises the mean integrated squared error if the data were Gaussian and a Gaussian kernel had been used. For further information on kernel smoothing and bandwidth selection, see Wand and Jones (1995).

The main focus of the chapter is to explore whether the rating momentum effect depends on business cycle fluctuations and time since rating assignment. The addition of those two effects leads to a covariate vector $\mathbf{Z}_{jki}(t)$ in (4.2) that varies across both calendar and survival time, therefore introducing non-proportional elements in the base specification (4.2) that assumes a time invariant \mathbf{Z}_{jki} . More specifically, the covariate vector $\mathbf{Z}_{jki}(t)$ comprises of 4 types of explanatory factors:

- Explanatory factors that are constant: After a certain migration event takes place, the rating momentum effect is constant across both survival and calendar time. The rating momentum indicator is denoted by Z_{jki}^m

$$Z_{jki}^m = \begin{cases} I\{j > k\} & \text{if previous migration is downgrade,} \\ I\{j < k\} & \text{if previous migration is upgrade,} \\ 0 & \text{otherwise.} \end{cases} \quad (4.5)$$

Z_{jki}^m takes the value 1 if the previous rating action was a downgrade(upgrade) and $j > k$ ($j < k$), and the value 0 otherwise. Once a migration occurs, this variable is fixed.

- Explanatory factors that change with calendar time τ : To capture systematic movements in rating migrations, I use the CFNAI aggregate diffusion index as a proxy for the business cycle fluctuations. The CFNAI covariate takes the form

$$Z_{cfnai,i}(\tau) = CFNAI(\tau), \quad (4.6)$$

where τ is the calendar time corresponding to each i . This corresponds to the base specification, since it does not adjust the systematic fluctuations for the rating momentum effects.

In addition to the systematic effect $Z_{cfnai,i}(\tau)$, I include in the model its interaction with rating momentum. This interaction term takes the form

$$Z_{cfnai,jki}^m(\tau) = Z_{cfnai,i}(\tau)Z_{jki}^m, \quad (4.7)$$

where the rating momentum Z_{jki}^m is defined in (4.5).

- Explanatory factors that change with survival time t : To capture variations of the rating momentum effect with the time spent in a given rating grade, I include in $\mathbf{Z}_{jki}(t)$ the interaction of Z_{jki}^m from (4.5) with survival time. There is no prior information to dictate the functional form of time and, therefore, I use flexible 2nd order fractional polynomials. The fractional polynomial form uses two terms, each one containing survival time raised to a power

$$\mathbf{Z}_{jki}^m(t) = [Z_{jki}^m t^{p_1} \quad Z_{jki}^m t^{p_2}], \quad (4.8)$$

where t is the survival time corresponding to each i and p_1, p_2 are the degrees of the two polynomial terms. The set of powers is restricted to $\{-2, -1, -0.5, 0, 0.5, 1, 2, 3\}$. For each migration pair, I try all possible permutations, and I select the set of powers that leads to the highest likelihood value. For more information on fractional polynomials see Royston and Altman (1994).

- Explanatory factors that change with survival time t and calendar time τ : To adjust the time-varying rating momentum effect for business cycle fluctuations, I include in $\mathbf{Z}_{jki}(t)$ the interaction of $Z_{jki}^m(t)$ in (4.8) with the CFNAI index, in the form of variable

$$\mathbf{Z}_{cfnai,jki}^m(t) = \mathbf{Z}_{jki}^m(t)Z_{cfnai,i}(\tau) = [Z_{jki}^m t^{p_1} Z_{cfnai,i}(\tau) \quad Z_{jki}^m t^{p_2} Z_{cfnai,i}(\tau)], \quad (4.9)$$

where t is the survival time for each i and τ is the calendar time corresponding to each i . Without loss of generality, I assume the fractional polynomial powers in (4.8) and (4.9) are the same.

Allowing the rating momentum to depend on survival time has a dual role. First, it sheds light into how the rating momentum effect varies across firms with different durations in their current rating grade. Second, it renders the maximum likelihood inference of section 4.3.2 valid, if the true effect of rating momentum is non-proportional. The hazard ratio example in (4.3) and the semi-parametric analysis of section 4.3.2 are based on the proportionality assumption. If an effect is non-proportional but the survival time is omitted from the corresponding covariate, then the shape of the baseline intensity in (4.3) cannot be ignored and the partial likelihood estimates in section 4.3.2 are invalid. Revisiting the hazard ratio example in (4.3), the inclusion of time varying effects in (4.2) renders the proportional effect only valid for a given t .

4.3.2 Estimation

As Cox (1972) and Cox (1975) show that the semi-parametric survival model can be estimated by maximising the so-called partial likelihood function. The partial likelihood for the observed migration from rating grade j to rating grade k is defined as

$$L_{jk}(\boldsymbol{\beta}_{jk}) = \prod_{i=1}^n \left[\frac{e^{\boldsymbol{\beta}_{jk} \mathbf{Z}_{jki}(T_i)}}{\sum_{h \in \mathcal{R}(T_i)} e^{\boldsymbol{\beta}_{jk} \mathbf{Z}_{jkh}(T_i)}} \right]^{\delta_i} \quad (4.10)$$

where, using the notation of (4.1), $i = 1, \dots, n$ iterates over the distinct firm observations in grade j , T_i is the time firm i leaves grade j , $\mathcal{R}(T_i)$ is the set of active firms at time T_i and δ_i is the censoring indicator such that $\delta_i = 1$ if T_i is a migration to grade k and $\delta_i = 0$ if T_i otherwise.

Equation (4.10) corresponds to the exact partial likelihood only when there are no tied rating actions at any given time point T_i (and therefore the rating migrations can be ordered in terms of survival time). Under a continuous time framework there can be no events occurring at precisely the same time. Therefore, firms have the same survival times only because the measurement is not accurate and granular enough, resulting in information loss. When there are tied rating actions the, so called, exact partial likelihood calculation involves the evaluation of the likelihood under all possible orderings of the tied event times, see Kalbfleisch and Prentice (1973) and Kalbfleisch and Prentice (2002). Given the number of rating actions available in the sample, this approach is

not computationally feasible. Breslow (1974) shows that, assuming random ordering and using (4.10) for every tied event (each time including all the other tied events at the risk set), leads to a reasonable approximation when the ratio of events to risk set size at each tie is small. Since under the Breslow approximation the denominator in (4.10) is slightly overweight towards the observations exhibiting a rating change (the set of tied events is included multiple times in the risk set), estimates of β_{jk} may be biased towards zero. The approximation proposed by Efron (1977) addresses this issue by introducing ordering into the partial likelihood and reducing the denominator weight in (4.10). Hertz-Picciotto and Rockhill (1997) provided evidence that the Efron approximation is more accurate than the Breslow approximation, a view shared by some statisticians, especially for small sample sizes. Nevertheless, there is no general consensus as to whether the increase in accuracy under the Efron approximation is statistically significant to compensate for the additional complexity, especially in problems with a large number of events. For the purpose of this study, the existence of ties is not a material issue since time is measured in days and only a very small proportion of migration events are tied (much less than 1%). Therefore, I use the Breslow approximation for the empirical analysis in section 4.5, in which case the partial likelihood in (4.10) can be used unadjusted.

To find the maximum likelihood estimates of the parameter vector β_{jk} in (4.10), I work with the log partial likelihood function:

$$\log L_{jk}(\beta_{jk}) = \sum_{i=1}^n \delta_i \left[\beta_{jk} \mathbf{Z}_{jki}(T_i) - \log \left(\sum_{h \in \mathcal{R}(T_i)} e^{\beta_{jk} \mathbf{Z}_{jkh}(T_i)} \right) \right]. \quad (4.11)$$

The maximum of the log partial likelihood function (4.11) is maximised by setting the first partial derivative with respect to the parameter vector β_{jk} (score vector) to zero. Assuming there are $p = 1, \dots, P$ covariates, the p 'th element of the score vector is given by:

$$U_{jk}^p(\beta_{jk}) = \frac{\partial \log L_{jk}(\beta^p)}{\partial \beta_{jk}^p} = \sum_{i=1}^n \delta_i (Z_{jki}^p(T_i) - \overline{Z_{jki}^p(T_i)}),$$

$$\overline{Z_{jki}^p(T_i)} = \sum_{h \in \mathcal{R}(T_i)} w_{ih} Z_{jki}^p(T_i), \quad w_{ih} = \frac{e^{\beta_{jk} \mathbf{Z}_{jkh}(T_i)}}{\sum_{h^* \in \mathcal{R}(T_i)} e^{\beta_{jk^*} \mathbf{Z}_{jkh^*}(T_i)}}$$

The negative matrix of second derivatives (Fisher information matrix) $\mathbf{V}_{jk}(\beta_{jk})$ has

elements:

$$\begin{aligned}
V_{jk}^{pp*}(\boldsymbol{\beta}_{jk}) &= -\frac{\partial^2 \log L_{jk}(\boldsymbol{\beta}^p)}{\partial \beta_{jk}^p \partial \beta_{jk}^{p*}} \\
&= \sum_{i=1}^n \delta_i \left[\sum_{h \in \mathcal{R}(T_i)} w_{ih} (Z_{jkh}^p(T_i) - \overline{Z_{jkh}^p(T_i)}) (Z_{jkh}^{p*}(T_i) - \overline{Z_{jkh}^{p*}(T_i)}) \right]
\end{aligned}$$

$\mathbf{V}_{jk}(\boldsymbol{\beta}_{jk})$ is used to calculate the parameter standard errors, since asymptotically the maximum likelihood estimate $\hat{\boldsymbol{\beta}}_{jk}$ is multivariate normally distributed

$$\hat{\boldsymbol{\beta}}_{jk} \sim N(\tilde{\boldsymbol{\beta}}_{jk}, \mathbf{V}_{jk}(\boldsymbol{\beta}_{jk})^{-1}),$$

where $\tilde{\boldsymbol{\beta}}_{jk}$ is the unknown true parameter vector.

The intensity parameterisation in (4.2), the partial likelihood in (4.10) and the partial log likelihood in (4.11) assume that all observations have the same baseline intensity and any heterogeneity is only reflected in the different covariate values. In a sample of rating actions across multiple industry sectors, the common baseline intensity and the proportionality assumptions are difficult to justify. Therefore, I introduce some additional heterogeneity by allowing the baseline intensities to differ across sectors. Stratifying the analysis by the $g = 1, \dots, G$ sectors involves collapsing the the partial log likelihood function (4.11) into G independent sums:

$$\log L_{jk}(\boldsymbol{\beta}_{jk}) = \sum_{g=1}^G \sum_{i \in g} \delta_i \left[\boldsymbol{\beta}_{jk} \mathbf{Z}_{jki}(T_i) - \log \left(\sum_{h \in \mathcal{R}(T_i)} e^{\boldsymbol{\beta}_{jk} \mathbf{Z}_{jkh}(T_i)} \right) \right], \quad (4.12)$$

where for each of the G sums, only observations from the same sector are included.

4.4 Data

For the empirical analysis, I use corporate bond rating data sourced from Moody's Default & Recovery Database (DRD), covering the period from January 1983 to December 2012. Moody's rating scale currently consists of 21 ordered alphanumeric ratings, spanning from Aaa for the low credit risk firms to C for the extremely high credit risk firms. Moody's introduced the alphanumeric scale for ratings above Caa in April 1982 and therefore the rating migration behaviour is bound to be fundamentally different pre

	Aaa	Aa1	Aa2	Aa3	A1	A2	A3	Baa1	Baa2	Baa3	Ba1	Ba2	Ba3	B1	B2	B3	Caa-C
Aaa	.	54	41	9	8	5	2	0	0	0	0	0	0	0	0	0	0
Aa1	24	.	88	53	13	2	0	3	0	0	0	0	1	0	0	0	0
Aa2	14	57	.	192	62	17	3	0	2	0	1	0	0	0	1	0	0
Aa3	4	25	84	.	269	111	19	4	5	1	0	0	0	2	0	0	0
A1	2	4	14	160	.	334	107	19	5	4	6	7	2	3	0	0	0
A2	0	1	8	50	209	.	459	150	31	11	3	4	3	3	0	1	1
A3	1	3	3	3	38	293	.	384	173	52	6	2	6	3	1	0	0
Baa1	2	4	6	2	3	51	250	.	426	119	28	9	8	4	2	0	0
Baa2	2	3	3	2	8	28	88	303	.	465	91	30	13	5	7	1	1
Baa3	1	1	2	1	3	7	16	84	376	.	312	143	53	34	11	3	7
Ba1	1	0	0	1	1	3	4	19	64	290	.	226	160	47	26	4	9
Ba2	0	0	1	0	0	0	0	2	15	72	264	.	274	99	95	27	9
Ba3	0	1	0	0	0	5	4	2	11	17	85	278	.	294	279	108	16
B1	0	1	1	0	2	4	1	3	3	4	16	90	339	.	388	285	74
B2	0	0	1	0	0	0	0	4	4	3	5	17	82	348	.	627	295
B3	0	0	1	0	0	0	1	1	3	6	6	5	15	86	323	.	828
Caa-C	0	0	0	0	0	0	0	1	4	2	2	2	13	20	69	310	.

Table 4.1: Migration event counts

The table contains the number of historically observed migrations from each alphanumeric rating grade (rows) to a different alphanumeric rating grade (columns). For presentation convenience, withdrawn rating migrations and defaults are excluded, since they do not affect the analysis. To test for rating momentum, I limit the analysis to 2 notches away from the current grade and I use all migration pairs with more than 60 historical events (shaded cells in the table). The migration from Aaa to Aa1 is excluded from the analysis since this type of downgrades are not affected by rating momentum.

and post 1982. I choose to start the analysis from January 1983 to allow for a gradual adoption of the new rating methodology, implicitly assuming that Moody's does not immediately change firm ratings to reflect the more granular rating scale. Despite setting the starting point for the analysis at January 1983, the rating history prior to 1983 is also taken into consideration and therefore the rating momentum effect is quantified without delay at the beginning of the sample. Moody's introduced the alphanumeric Caa ratings in July 1997. The limited rating activity prior to 1997 prevents any statistical separation of the rating momentum effects before and after 1997 and therefore I group the Caa-C into a single grade and treat the entire sample as a whole.

Including all the rating data available in Moody's DRD provides an increased sample size, but also introduces unnecessary heterogeneity in the dataset when trying to make inference across different regions. To construct a sample as homogenous as possible, I limit the analysis to US rated firms, since the US data history is sufficiently rich to draw robust conclusions. The US sample period contains a total of 8689 downgrades and 5181 upgrades across 6589 distinct firms. To limit the high dimensionality of the migration matrix and to exclude migration pairs with very few events through time, I analyse migrations of up to 2 notches and more than 60 events historically. Downgrades from rating grade Aaa are excluded since there can be no rating momentum for such rating moves. Table 4.1 summarises the number of events for the entire transition matrix, highlighting the cells that have been used for the empirical analysis of section 4.5.

To proxy the business cycle I use the CFNAI as published by the Federal Reserve Bank of Chicago^{1,2}. Using an aggregate index instead of individual time series allows to capture many aspects of the economic environment without losing degrees of freedom by including many, highly correlated, covariates. The CFNAI is essentially the first principal component of a panel of 85 economic indicators, sampled at a monthly frequency. The indicators cover a wide spectrum of the real economic activity, by including multiple measures of production, income, employment, consumption, housing, sales and inventories. The index is standardised to have mean 0 and unit standard deviation. The CFNAI index has been used in previous studies on default and rating migration modelling, see McNeil and Wendin (2006), McNeil and Wendin (2007), and Figlewski et al. (2012). I average and use the CFNAI index at a quarterly frequency, despite the index also being available monthly. There are numerous studies showing that rating agencies rate "Through-the-Cycle" to achieve the desired rating stability, see Amato and Furfine (2004), Altman and Rijken (2004), Löffler (2004), and Löffler (2013). Furthermore, as Cantor (2001) argues, rating agencies suppress rating changes when they are likely to be reversed within a relatively short period of time. Therefore, rating changes do not adjust immediately to changes in the economic environment and using monthly data would require rating grade specific lag optimisation. Using quarterly data smoothes out the lag effects while maintaining a sufficiently high granularity level.

Figure 4.1 provides times series graphs for the quarterly number of rating migrations. For illustration purposes, I summarise the time series by aggregate downgrades and upgrades and I provide a split by investment and speculative starting grades. Quarterly frequency is only used for illustration purposes, as the model is estimated using daily data. It is very apparent that downgrades follow very closely the NBER recessions, even though the investment grade downgrades seem to be relatively unaffected by the early '90s recession and lag the early '00s recession. In contrast, upgrades are less affected by the business cycle, with only speculative grade upgrades strongly affected over the recent economic downturn.

¹<http://chicagofed.org/webpages/publications/cfnai/index.cfm>

²The concept behind the CFNAI index is similar to the dynamic factor analysis of Chapter 2. While in principle the aggregate factor from Chapter 2 can be used instead (the correlation between the two factors is close to 80%), I use the CFNAI to make the results of this chapter independent and allow for an easier comparison with prior studies

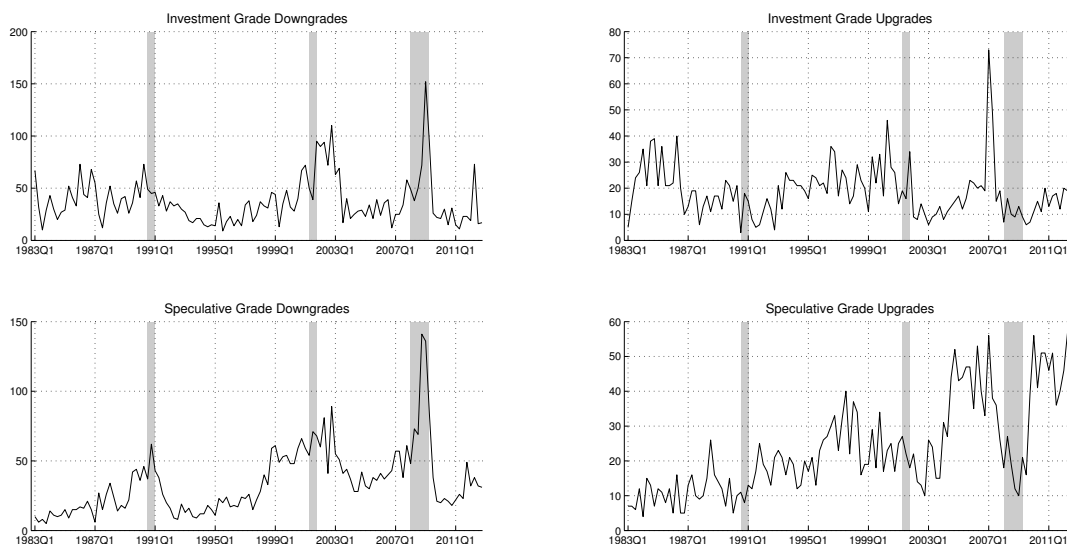


Figure 4.1: Quarterly number of rating migrations

Number of upgrades and downgrades per quarter for the period 1983:Q1-2012:Q3. Investment Grade rating actions refer to migrations from rating grades Aaa-Baa3, while Speculative Grade rating actions refer to migrations from rating grades Ba1-C. Shaded areas correspond to NBER recession quarters.

Moody’s DRD database contains a variety of sector classifications. I choose the Moody’s 11 sector broad industry classification to avoid having many segments with very few or no rating migration events. The 11 sectors correspond to Banking, Capital Industries, Consumer Industries, Energy & Environment, Finance, Insurance & Real Estate Finance (FIRE), Media & Publishing, Retail & Distribution, Sovereign & Public Finance, Technology, Transportation, and Utilities. I exclude from the analysis the Sovereign & Public Finance entities, since their behaviour is likely not to follow the corporate entities dynamics. Finally, I also exclude companies with missing sector information.

Figure 4.2 plots the daily baseline intensities for aggregate downgrades and upgrades for each of the 10 sectors used in the analysis (reported in bp). For illustration purposes the cut-off point for the graphs has been set to 5 years. Despite a clear upward trend for the first year since rating assignment for both downgrades and upgrades, it is apparent that the term structure has significant differences across sectors. Energy & Environment, Transportation and Utilities do not exhibit many similarities with other sectors. Intuitively, Banking and FIRE are strongly correlated, while Capital and Consumer related sectors (including Media & Publishing, Retail & Distribution) also show some similarities in their term structures. Technology firm downgrades seem to peak at 3 years post rating assignment, an observation somewhat counterintuitive, since small

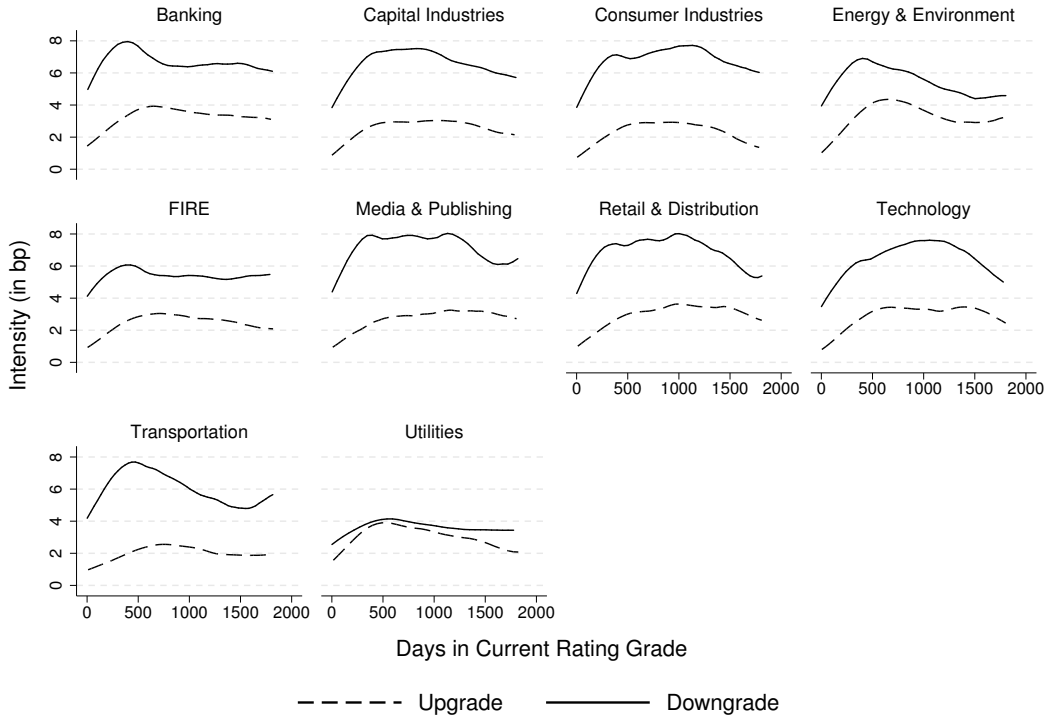


Figure 4.2: Baseline intensity rates per industry sector

Daily baseline intensity rates λ_t^0 (in bp) for each of the 11 industry sectors. Intensities are provided for aggregate downgrades and upgrades, for up to 5 years from the point a firm is assigned a new rating.

	Aaa	Aa1	Aa2	Aa3	A1	A2	A3	Baa1	Baa2	Baa3	Ba1	Ba2	Ba3	B1	B2	B3	Caa-C
Aaa	88.2%	5.3%	4.4%	0.8%	0.9%	0.4%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%
Aa1	2.1%	81.3%	7.1%	6.2%	1.8%	1.1%	0.1%	0.3%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%
Aa2	0.9%	3.3%	79.3%	9.9%	3.9%	1.5%	0.6%	0.2%	0.3%	0.0%	0.1%	0.0%	0.0%	0.0%	0.0%	0.0%	0.1%
Aa3	0.1%	0.9%	3.0%	80.7%	8.7%	4.1%	1.4%	0.5%	0.3%	0.1%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%
A1	0.1%	0.1%	0.5%	4.3%	82.1%	7.9%	3.1%	0.9%	0.3%	0.2%	0.2%	0.2%	0.1%	0.1%	0.0%	0.0%	0.0%
A2	0.0%	0.0%	0.1%	0.8%	4.0%	82.2%	7.6%	3.1%	1.0%	0.5%	0.1%	0.1%	0.1%	0.1%	0.0%	0.0%	0.1%
A3	0.0%	0.1%	0.1%	0.1%	1.0%	6.6%	79.6%	6.8%	3.1%	1.4%	0.4%	0.1%	0.3%	0.2%	0.1%	0.1%	0.0%
Baa1	0.0%	0.1%	0.1%	0.1%	0.1%	1.5%	6.0%	80.0%	6.9%	2.8%	0.9%	0.5%	0.3%	0.5%	0.1%	0.0%	0.1%
Baa2	0.0%	0.0%	0.0%	0.0%	0.2%	0.8%	2.4%	6.2%	80.1%	6.4%	1.5%	0.6%	0.6%	0.4%	0.2%	0.2%	0.2%
Baa3	0.0%	0.0%	0.0%	0.0%	0.1%	0.2%	0.5%	2.6%	8.2%	78.1%	4.7%	2.6%	1.2%	0.9%	0.4%	0.2%	0.3%
Ba1	0.0%	0.0%	0.0%	0.0%	0.1%	0.2%	0.3%	0.8%	2.4%	10.2%	72.4%	5.5%	3.8%	1.5%	1.4%	0.9%	0.6%
Ba2	0.0%	0.0%	0.1%	0.0%	0.0%	0.1%	0.0%	0.4%	0.8%	3.1%	8.5%	71.9%	7.1%	2.9%	3.0%	1.5%	0.7%
Ba3	0.0%	0.0%	0.0%	0.0%	0.0%	0.1%	0.1%	0.1%	0.4%	0.9%	2.5%	6.3%	74.9%	6.0%	5.0%	2.4%	1.3%
B1	0.0%	0.0%	0.0%	0.0%	0.0%	0.1%	0.1%	0.1%	0.1%	0.3%	0.4%	2.3%	6.7%	77.6%	5.4%	4.6%	2.2%
B2	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.1%	0.1%	0.1%	0.2%	0.9%	3.0%	5.9%	75.7%	8.9%	5.1%
B3	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.1%	0.1%	0.1%	0.3%	0.3%	0.8%	3.1%	6.1%	77.5%	11.4%
Caa-C	0.0%	0.0%	0.0%	0.0%	0.0%	0.1%	0.0%	0.2%	0.7%	0.2%	0.4%	0.1%	1.4%	2.0%	2.5%	6.8%	85.6%

Table 4.2: Average Migration Rates

The table depicts the average 1 year historical migration rates, using the entire sample 1983-2012. The average is calculated using rolling quarterly windows. Migrations to default and rating withdrawals are excluded and each row is re-normalized to sum to unity.

to mid-sized tech firms' financial strength is generally volatile and the margin for business model failure is high. The long-run term structures depicted in figure 4.2 provide evidence that adjusting the estimations for different baseline shapes and levels for the intensities, enhances robustness against any sector heterogeneity bias.

Lando and Skødeberg (2002) report an overall negative duration effect in their analysis of rating migrations. Based on the baseline term structure profiles in figure

4.2, the overall negative effect reported appears questionable. Even at the 5 years horizon for most of the sectors there seems to be a positive relationship between time spent in a given grade and rating migration, implying that the longer a firm stays in a given rating grade the higher is the probability it will experience a rating migration. Even if the baselines are extended to more than 5 years and the average effect across the term structure is negative, it is apparent that additional flexibility is needed if the duration effect is to be captured properly.

Corporate bond ratings are designed to ensure some long term stability, trying to smooth out temporary changes in credit quality. Rating agencies achieve this stability in their ratings by adopting a more Through-the-Cycle view when assessing the creditworthiness of a given company. This long term stability is depicted in Table 4.2, where I calculate the average 1 year migration matrix across the entire sample used in the study. The average is calculated using rolling quarterly windows, instead of the annual non-overlapping cohorts that credit agencies usually report in their default studies. Furthermore, since the focus of this study is the behaviour of rating agencies, I exclude defaults and rating withdrawals and I re-normalise each row of the transition matrix to sum to 1 (defaults are not controlled by the rating agencies and rating withdrawals are not credit risk related events). By examining table 4.2 it is apparent that close to 80% of the companies retain their existing rating a year after observation, and this percentage is clearly higher for investment grade firms as compared to speculative grade firms. Finally, the migration matrix provides evidence of overall deterioration in credit quality through time, since the downgrade rate is significantly higher than the upgrade rate. This observation indicates that unless there is a constant addition of highly rated corporates in the sample, the overall portfolio distribution shifts towards sub-investment grade as time passes.

Various studies show that the effect of rating momentum is very strong, especially for downgrades. To provide a high level impact of the momentum effect, I depict in table 4.3 the same Through-the-Cycle migration matrix as in table 4.2 restricted to firms that have been previously downgraded within 2 years prior to the observation point (Aaa rated firms are excluded, since firms in that grade cannot have been previously downgraded). The downward drift is very clear, since the downgrade rates are on average 80% higher than those reported in Table 4.2, with the number of firms retaining

	Aaa	Aa1	Aa2	Aa3	A1	A2	A3	Baa1	Baa2	Baa3	Ba1	Ba2	Ba3	B1	B2	B3	Caa-C
Aaa	0.0%	65.7%	17.8%	13.8%	1.9%	0.8%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%
Aa1	0.0%	1.1%	77.8%	10.9%	5.4%	2.8%	1.3%	0.7%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%
Aa2	0.0%	0.3%	0.5%	81.3%	4.4%	8.8%	4.0%	0.3%	0.2%	0.0%	0.2%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%
Aa3	0.0%	0.0%	0.3%	2.1%	77.4%	11.6%	5.0%	2.2%	0.5%	0.0%	0.7%	0.0%	0.1%	0.0%	0.0%	0.0%	0.0%
A1	0.0%	0.0%	0.2%	0.6%	2.5%	75.7%	11.9%	6.0%	1.7%	0.9%	0.2%	0.2%	0.0%	0.0%	0.0%	0.1%	0.0%
A2	0.0%	0.0%	0.0%	0.2%	0.3%	4.1%	75.9%	7.4%	6.3%	3.0%	0.7%	0.1%	1.0%	0.4%	0.3%	0.0%	0.3%
A3	0.0%	0.1%	0.0%	0.2%	0.1%	1.4%	3.0%	74.2%	10.4%	5.8%	1.7%	1.3%	0.6%	0.9%	0.2%	0.1%	0.1%
Baa1	0.0%	0.0%	0.0%	0.1%	0.0%	0.5%	2.2%	2.0%	75.1%	11.0%	3.5%	2.0%	1.4%	1.2%	0.3%	0.1%	0.4%
Baa2	0.0%	0.0%	0.2%	0.1%	0.0%	0.1%	0.6%	2.3%	5.4%	67.0%	8.5%	7.5%	2.8%	2.3%	1.0%	0.8%	1.5%
Baa3	0.0%	0.0%	0.0%	0.2%	0.0%	0.0%	0.0%	0.3%	1.3%	8.0%	59.0%	11.4%	9.9%	3.1%	3.6%	1.4%	1.8%
Ba1	0.0%	0.0%	0.2%	0.0%	0.0%	0.0%	0.0%	1.2%	0.4%	2.7%	6.4%	59.2%	10.1%	7.3%	6.8%	3.0%	2.6%
Ba2	0.0%	0.0%	0.0%	0.0%	0.0%	0.5%	0.1%	0.1%	0.5%	1.2%	3.8%	4.9%	61.9%	11.9%	8.5%	3.7%	2.9%
Ba3	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.1%	0.1%	0.7%	0.6%	2.3%	5.4%	68.7%	7.4%	8.3%	6.3%
B1	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.1%	0.0%	0.1%	0.0%	0.2%	1.2%	4.1%	6.6%	65.7%	12.7%	9.3%
B2	0.0%	0.0%	0.0%	0.0%	0.0%	0.1%	0.0%	0.1%	0.0%	0.1%	0.2%	1.2%	4.1%	6.6%	65.7%	12.7%	9.3%
B3	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.1%	0.2%	0.2%	0.2%	0.3%	0.0%	0.7%	4.9%	5.2%	70.9%	17.3%
Caa-C	0.0%	0.0%	0.0%	0.0%	0.0%	0.3%	0.0%	0.0%	1.4%	0.2%	0.1%	0.0%	0.3%	1.4%	1.6%	6.6%	88.0%

Table 4.3: Average Migration Rates - Downward Momentum

The table depicts the average 1 year historical migration rates for firms that have been previously downgraded, up to 2 years prior to the observation point. The average is calculated using rolling quarterly windows and I include all firms across the entire sample from Q1 1983 to Q4 2012. Migrations to default and rating withdrawals are excluded and each row is re-normalized to sum to unity. The migrations from Aaa to Aa1 are excluded from the analysis since they are not affected by rating momentum.

their existing rating within the 1 year horizon being approximately 10% lower.

There is no commonly agreed threshold for the classification of subsequent migrations as rating drift. If momentum is assumed to be caused by the unwillingness of rating agencies for multi-notch rating changes, then a short period threshold should be applied, as multi-notch migrations are not caused by any subsequent deterioration/improvement of the corporate balance sheet. On the other hand, if momentum is attributed to the countercyclical effect of rating migrations or to the trend of balance sheet quality, then a longer period threshold should be adopted in order for new financial information to be incorporated in the rating process. In this chapter I choose 2 years as the cut-off point. This is higher than the 210-day threshold used by Güttler and Raupach (2010), but lower than the unrestricted rating momentum used in some early studies, such as Lando and Skødeberg (2002). I choose the 2-year threshold for multiple reasons: 210 days are not enough to allow for a robust modelling of the momentum time profile, as the sample size of momentum observations is small. Furthermore, a 7 month period might be too short to allow for rating agencies to incorporate changes in corporate balance sheets and, therefore, the firm ageing effect might not be properly captured.

It is worth noting that the analysis captures the behaviour of rating agencies, which is intrinsically a combination of quantitative and qualitative assessments for a given bond issuer. Due to the relative importance of credit ratings, Duan and Van Laere (2012) present a framework to construct independent risk assessment based on

purely transparent quantitative methods. The framework is based on the work of Duan et al. (2012) and provides default probability estimates for all public firms around the world. While such estimates can be very useful for risk management, it is out of scope of the analysis presented in this chapter.

4.5 Empirical Analysis

4.5.1 Parameter Estimates

The hazard rate specification described in (4.2) and (4.5)-(4.9) provides the most generic case of rating pair specific coefficients for all covariates. The number of migration events is too small to allow for such generic specification to be robustly estimated at a rating pair level. Therefore, without loss of generality, I assume 4 different groups for the coefficients in (4.2): 1 set of coefficients for the investment grade downgrades, 1 set of coefficients for the speculative grade downgrades, and 2 additional sets of coefficients to describe investment and sub-investment grade upgrades. Using the notation of section 4.3.1, the migration intensity of the estimated model takes the form:

$$\lambda_{jki}(t) = \lambda_{jkg}^0(t) \cdot \underbrace{\exp\left(\beta_c^m Z_{jki}^m + \beta_c^{\text{cfnai}} Z_{\text{cfnai},i}(\tau) + \beta_c^{\text{cfnai},m} Z_{\text{cfnai},jki}^m(\tau)\right)}_{\text{Survival Time-Invariant Part}} \cdot \underbrace{\exp\left(\beta_c^{m(t)} \mathbf{Z}_{jki}^m(t) + \beta_c^{\text{cfnai},m(t)} \mathbf{Z}_{\text{cfnai},jki}^m(t)\right)}_{\text{Survival Time-Varying Part}} \quad (4.13)$$

where $c = \{I_{\{j \leq 10 \wedge j > k\}}, I_{\{j > 10 \wedge j > k\}}, I_{\{j \leq 10 \wedge j < k\}}, I_{\{j > 10 \wedge j < k\}}\}$ represents each of the 4 sets of coefficients, and $\lambda_{jkg}^0(t)$ is the baseline intensity per industry sector g , which is left unspecified. The survival time-invariant covariates are described in (4.5)-(4.7), and correspond to the rating momentum effect (Z_{jki}^m), the business cycle effect ($Z_{\text{cfnai},i}(\tau)$) and their interaction ($Z_{\text{cfnai},jki}^m(\tau)$). The survival time-varying part of the specification corresponds to the fractional polynomials for rating momentum ($\mathbf{Z}_{jki}^m(t)$) and rating momentum adjusted for the business cycle ($\mathbf{Z}_{\text{cfnai},jki}^m(t)$), as described in more detail in (4.8)-(4.9).

To test the statistical significance of the various effects in (4.13) I use the likelihood

ratio test. For hypotheses of type:

$$H_0 : \boldsymbol{\beta}_c = \boldsymbol{\beta}_c^r$$

$$H_1 : \boldsymbol{\beta}_c = \boldsymbol{\beta}_c^u,$$

where $\boldsymbol{\beta}_c^r$ is a restricted subset of $\boldsymbol{\beta}_c^u$, the likelihood ratio test is calculated as:

$$LR_{\boldsymbol{\beta}_c^r/\boldsymbol{\beta}_c^u} = -2\log\left(\frac{L_{jk}(\boldsymbol{\beta}_c^r)}{L_{jk}(\boldsymbol{\beta}_c^u)}\right) = 2\log L_{jk}(\boldsymbol{\beta}_c^u) - 2\log L_{jk}(\boldsymbol{\beta}_c^r). \quad (4.14)$$

Asymptotically, the likelihood ratio statistic in (4.14) is χ^2 distributed, with degrees of freedom equal to the number of restrictions imposed on $\boldsymbol{\beta}_c^u$ to get $\boldsymbol{\beta}_c^r$. Having obtained parameter estimates for the 3 covariates used in the study, I test 5 main hypotheses:

- The proportional rating momentum effect is significant: Test to assess whether including the rating momentum indicator to a model with only aggregate systematic effects, results in a model with higher explanatory power. The likelihood ratio test statistic in (4.14) takes the form

$$LR_{\hat{\beta}_c^{\text{cfnai}}/\hat{\beta}_c^m} = 2\log L_{jk}(\hat{\beta}_c^{\text{cfnai}}, \hat{\beta}_c^m) - 2\log L_{jk}(\hat{\beta}_c^{\text{cfnai}}),$$

having a $\chi^2(1)$ distribution. For notational convenience, I only report the additive part of the specification as the unrestricted model, i.e. $LR_{\hat{\beta}_c^{\text{cfnai}}/\hat{\beta}_c^m} = LR_{\hat{\beta}_c^{\text{cfnai}}/(\hat{\beta}_c^{\text{cfnai}}, \hat{\beta}_c^m)}$. This test of proportional rating momentum effects is similar to Lando and Skødeberg (2002). Unlike Lando and Skødeberg (2002) that assume the base case to be a long run average intensity model, I also consider systematic movements in the base case.

- Rating momentum depends on time spent in current rating grade: Test to assess whether adding the interaction of the rating momentum indicator with the time in current rating grade results in a statistically significant improvement in model fit. The likelihood ratio test statistic in (4.14) takes the form

$$LR_{\hat{\beta}_c^m/\hat{\beta}_c^m(t)} = 2\log L_{jk}(\hat{\beta}_c^{\text{cfnai}}, \hat{\beta}_c^m(t)) - 2\log L_{jk}(\hat{\beta}_c^{\text{cfnai}}, \hat{\beta}_c^m),$$

having a $\chi^2(2)$ distribution, due to the second order fractional polynomials used

for the time varying effects. For notational convenience I assume $\hat{\beta}_c^m(t) = \{\hat{\beta}_c^m, \hat{\beta}_c^m(t)\}$.

- Rating momentum depends on business cycle: Test to assess whether adding the interaction of the proportional rating momentum effect with the CFNAI diffusion index results in a statistically significant improvement in model fit. For a proportional dependence of the rating momentum on the business cycle, the likelihood ratio test statistic in (4.14) takes the form

$$LR_{\hat{\beta}_c^m(t)/\hat{\beta}_c^{\text{cfnai},m}} = 2\log L_{jk}(\hat{\beta}_c^{\text{cfnai}}, \hat{\beta}_c^m(t), \hat{\beta}_c^{\text{cfnai},m}) - 2\log L_{jk}(\hat{\beta}_c^{\text{cfnai}}, \hat{\beta}_c^m(t)).$$

This likelihood ratio statistic has a $\chi^2(1)$ distribution. For notation convenience, I assume that $\hat{\beta}_c^m(t) = \{\hat{\beta}_c^m, \hat{\beta}_c^m(t)\}$.

- Rating momentum depends non-proportionally on business cycle: Test to assess whether adding the interaction of the time-varying rating momentum effect with the CFNAI diffusion index results in a statistically significant improvement in model fit. In this case, the non-proportional rating momentum effect changes non-proportionally with the CFNAI index. The likelihood ratio test statistic in (4.14) takes the form

$$LR_{\hat{\beta}_c^{\text{cfnai},m}/\hat{\beta}_c^{\text{cfnai},m(t)}} = 2\log L_{jk}(\hat{\beta}_c^{\text{cfnai}}, \hat{\beta}_c^m(t), \hat{\beta}_c^{\text{cfnai},m(t)}) - 2\log L_{jk}(\hat{\beta}_c^{\text{cfnai}}, \hat{\beta}_c^m(t), \hat{\beta}_c^{\text{cfnai},m}),$$

where $\hat{\beta}_c^m(t) = \{\hat{\beta}_c^m, \hat{\beta}_c^m(t)\}$ and $\hat{\beta}_c^{\text{cfnai},m(t)} = \{\hat{\beta}_c^{\text{cfnai},m}, \hat{\beta}_c^{\text{cfnai},m(t)}\}$ for notational convenience. This likelihood ratio statistic has a $\chi^2(2)$ distribution.

- Business cycle and duration dependence effects are jointly significant: Test to assess whether adding the time varying effects and the interaction with the CFNAI diffusion index jointly results in a statistically significant improvement in model fit. The likelihood ratio test statistic in (4.14) takes the form

$$LR_{\hat{\beta}_c^m/\hat{\beta}_c^{\text{cfnai},m(t)}} = 2\log L_{jk}(\hat{\beta}_c^{\text{cfnai}}, \hat{\beta}_c^m(t), \hat{\beta}_c^{\text{cfnai},m(t)}) - 2\log L_{jk}(\hat{\beta}_c^{\text{cfnai}}, \hat{\beta}_c^m),$$

having a $\chi^2(5)$ distribution. Just like in all other cases, I assume for notational convenience that $\hat{\beta}_c^m(t) = \{\hat{\beta}_c^m, \hat{\beta}_c^m(t)\}$ and $\hat{\beta}_c^{\text{cfnai},m(t)} = \{\hat{\beta}_c^{\text{cfnai},m}, \hat{\beta}_c^{\text{cfnai},m(t)}\}$.

		IG Downgrades	SG Downgrades	IG Upgrades	SG Upgrades
Parameter Estimates	$p1, p2$	[0.5,.]	[-0.5,0.5]	[-1,.]	[-0.5,0.5]
	$\hat{\beta}_c^m$	1.8966 (0.1610)	3.4806 (0.2782)	0.9780 (0.0951)	2.4518 (0.4040)
	$\hat{\beta}_c^{cfnai}$	-0.3372 (0.0232)	-0.3513 (0.0217)	0.3345 (0.0457)	0.1960 (0.0343)
	$\hat{\beta}_c^{cfnai,m}$	-0.0732 (0.0755)	-0.4065 (0.1540)	-0.0946 (0.1242)	0.1102 (0.5464)
	$\hat{\beta}_c^{m(t)}$	[-0.0402,.] (0.0089,.)	[-4.4155, -0.1223] (1.1249, 0.0123)	[-39.7542,.] (16.1552,.)	[-4.7200, -0.0542] (2.0749,0.0159)
	$\hat{\beta}_c^{cfnai,m(t)}$	[0.0048,.] (0.0047,.)	[1.6561, 0.0163] (0.6257, 0.0074)	[17.4266,.] (26.2302,.)	[0.1266, 0.0017] (2.9799, 0.0209)
LR statistic p-values	$LR_{\hat{\beta}_c^{cfnai}/\hat{\beta}_c^m}$	<0.0001	<0.0001	<0.0001	<0.0001
	$LR_{\hat{\beta}_c^m/\hat{\beta}_c^{m(t)}}$	<0.0001	<0.0001	0.0009	0.0008
	$LR_{\hat{\beta}_c^{m(t)}/\hat{\beta}_c^{cfnai,m}}$	0.8812	0.4576	0.6727	0.0441
	$LR_{\hat{\beta}_c^{cfnai,m}/\hat{\beta}_c^{cfnai,m(t)}}$	0.3045	0.0191	0.5007	0.9959
	$LR_{\hat{\beta}_c^m/\hat{\beta}_c^{cfnai,m(t)}}$	<0.0001	<0.0001	0.0085	0.0026

Table 4.4: Downgrade parameter estimates - Investment Grade

Parameter estimates with standard errors (in brackets) for proportional rating momentum ($\hat{\beta}_c^m$), baseline CFNAI effect ($\hat{\beta}_c^{cfnai}$), proportional interaction of rating momentum and the CFNAI index ($\hat{\beta}_c^{cfnai,m}$), rating momentum with duration dependence ($\hat{\beta}_c^{m(t)}$), and rating momentum with business cycle dependence ($\hat{\beta}_c^{cfnai,m(t)}$). The table also provides p-values for 5 Likelihood Ratio (LR) tests: presence of proportional rating momentum ($LR_{\hat{\beta}_c^{cfnai}/\hat{\beta}_c^m}$), incremental benefit of adding duration dependence ($LR_{\hat{\beta}_c^m/\hat{\beta}_c^{m(t)}}$), additional benefit of including proportional business cycle dependence ($LR_{\hat{\beta}_c^{m(t)}/\hat{\beta}_c^{cfnai,m}}$), additional benefit of allowing for non-proportional business cycle dependence ($LR_{\hat{\beta}_c^{cfnai,m}/\hat{\beta}_c^{cfnai,m(t)}}$), and finally, the joint significance of adding time-varying duration and business cycle effects to the benchmark model of proportional rating momentum ($LR_{\hat{\beta}_c^m/\hat{\beta}_c^{cfnai,m(t)}}$).

In table 4.4 I provide the parameter estimates and the corresponding standard errors for the 5 effects in my specification. For the statistical hypothesis testing, table 4.4 also includes the likelihood ratio statistics from (4.14) and the corresponding p-values to assess the significance of the tested effect. P-value is the probability of obtaining the observed likelihood increase when including the respective effect, under the assumption that the true effect is zero (null hypothesis). If the probability of obtaining the observed difference in likelihood under the null hypothesis is close to zero, there is very strong statistical evidence to reject the null hypothesis, and therefore, the tested effect is statistically significant. To determine the strength of the statistical evidence presented in table 4.4, the typical significance levels of 1% and 5% can be used as the upper bounds for the p-value to indicate strong evidence against the null hypothesis. P-values lower than 5% and especially 1% indicate a statistically significant effect, while p-values higher than 5% indicate the effect tested is not statistically different from zero.

Examining the parameter estimates for downgrades in table 4.4, the baseline effect

of the business cycle, as captured by the CFNAI index, is negative and highly significant for both investment (-0.3372) and sub-investment grade firms (-0.3513). Overall, the business cycle effect appears to be stronger for sub-investment grade firms, although the difference is not statistically significant. Confirming earlier studies on rating migrations, the baseline rating momentum effect is positive and highly significant. Furthermore, the effect is approximately 1.5 times higher for sub-investment grade firms as compared to investment grade (3.4806 and 1.8966 respectively). This finding might be linked to the overall higher level of migration activity for the speculative grade part of the rating scale, but might also indicate that the acceleration in balance sheet deterioration is much more pronounced as firms approach their default point. The time-varying rating momentum estimates are not very informative as they depend on the specific fractional polynomial form. More informative are the p-values. The extremely low p-values for the $LR_{\hat{\beta}_c^m/\hat{\beta}_c^m(t)}$ statistic show that the rating momentum is strongly dependent on time since rating is assigned, for both investment and sub-investment grades. On the contrary, business cycle dependency of the momentum effect is less clear. The $\hat{\beta}_c^{\text{cfnai},m}$ effect is on average negative (-0.0732 for investment grade and -0.4065 for sub-investment grade), implying that as the economic conditions worsen, the rating momentum effect becomes more pronounced. Nevertheless, $\hat{\beta}_c^{\text{cfnai},m}$ is only statistically supported for sub-investment grade firms. Furthermore, for sub-investment firms, this dependence of rating momentum on business cycle is highly non proportional, implying that the time profile of the rating drift changes shape across different phases of the business cycle.

The parameter estimates for upgrades in table 4.4 indicate that the business cycle effects are not as pronounced as for downgrades. The proportional baseline sensitivity to the CFNAI index implies a decrease of 40% (based on the 0.3345 estimate) and 22% (based on the 0.1960 estimate) in investment and sub-investment upgrade intensity respectively, if the index decreases in value by 1. The respective sensitivities for downgrades imply an increase in intensity of 40% and 42% for investment and sub-investment rated firms respectively. Furthermore, the baseline upward momentum effect is also milder than the downward momentum, with an implied increase in upgrade intensity approximately 60% lower than the increase in downgrade intensity. Despite being milder than the downward drift, this time-invariant upward momentum effect is strong, in contrast to the findings of Lando and Skødeberg (2002), that report a mostly

insignificant momentum effect for upgrades. Nevertheless, the findings of Lando and Skødeberg (2002) refer to a much shorter historical period and do not include business cycle effects or differentiation of the baseline intensity across industry sectors. Similar to the downward momentum, the low p-values for the $LR_{\hat{\beta}_c^m / \hat{\beta}_c^m(t)}$ statistic show that the rating momentum depends on time since rating is assigned, for both investment and sub-investment grades. On the contrary, the upward momentum's dependence on the business cycle is only marginally significant at the 5% level for sub-investment grade firms, and only on a proportional basis; the time profile of the upward momentum for sub-investment grade firms seems to shift proportionally across the different phases of the business cycle.

To visualise the reported findings in table 4.4, figure 4.3 provides a graphical summary of the rating momentum effect as a function of time in current grade, using the full specification with the 2nd order polynomial terms for the time dependence. To show the possible dependence on the state of the economy, the rating momentum effect is summarised for 3 values of the CFNAI index: a value of 0 to capture the Through-the-Cycle effect, a value of 2 to capture the effect during periods of good economic conditions, and a value of -2 to capture the effect during period of stress. The graphs on figure 4.3 ignore the baseline CFNAI effect $\hat{\beta}_c^{\text{cfnai}}$, since this is not specific to firms that have been previously migrated.

Exploring the graphs on figure 4.3 clearly highlights the non-linearity of the rating momentum effect, both across survival time and the state of the economy. The Through-the-Cycle dynamics show that for investment grade firms previously downgraded, there is an immediate increase in re-downgrade intensity of approximate 6 times the baseline and this effect almost exponentially decays with time the firms stay in their current grade. For sub-investment grade firms, the nature of time dependence changes drastically. After the firms are downgraded, there is a sharp rise in re-downgrade intensity to reach the maximum in approximately 2 months after the current rating is assigned. After the peak is reached, there is a gradual decrease over the 2 year period. This peak corresponds to approximately 7.5 times the baseline. For investment grade firms, the intensity for the upward momentum monotonically increases with the time spent in current grade; nevertheless, after approximately 3 quarters since the current rating is assigned, the upward drift can be considered as constant at a level twice the

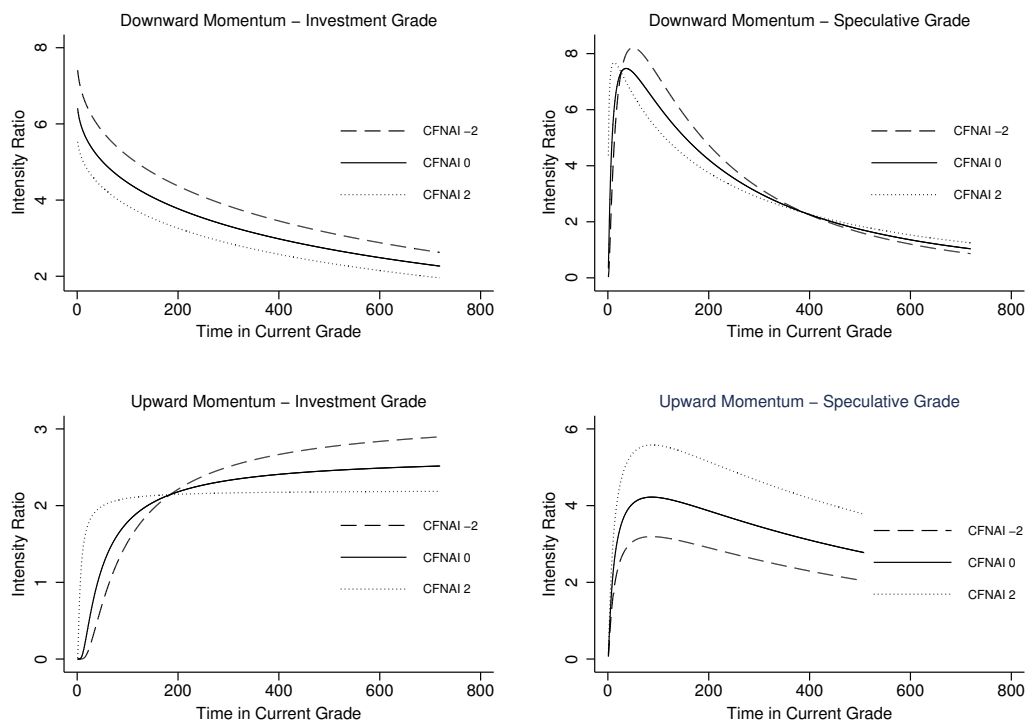


Figure 4.3: Rating momentum time varying effects

Time varying rating momentum effects for aggregate downgrades and upgrades. The effect is depicted as the intensity ratio over the case of 0 rating momentum. Intensity ratios are depicted for 3 values of CFNAI: 1 corresponding to severe conditions (CFNAI value of -2), 1 corresponding to average conditions (CFNAI value of 0), and 1 corresponding to mild conditions (CFNAI value of 2). Baseline business cycle effects are ignored and only the interaction terms of the CFNAI index with rating momentum are used. Duration time is capped at 2 years since current grades has been assigned. Investment Grade rating actions refer to migrations from rating grades Aaa-Baa3, while Speculative Grade rating actions refer to migrations from rating grades Ba1-C.

baseline upgrade intensity. Finally, for sub-investment grade firms, the shape of the time dependence of the upward momentum effect is similar to the downwards momentum. The upgrade intensity monotonically increases until a maximum of 4 times the baseline intensity is reached after 2 months in the current grade. After this peak is reached, the upward rating momentum effect decays with time, but a much slower rate as compared to the downward drift.

In addition to the clear duration effects, business cycle fluctuations also have an impact on the rating momentum's time profile. For investment grade downgrades, there is an inverse relationship between the state of the business cycle and the strength of the downward momentum effect; as economic conditions worsen, firms that have been previously downgraded are more likely to be downgraded again. This shift in downgrade intensity appears proportional. Nevertheless, a careful examination of the results in table 4.4 shows that the number of downgrades is not high enough to provide

statistical evidence that this inverse relationship between the CFNAI index and the strength of the momentum effect is significant. For speculative grade firms, the business cycle clearly changes the shape of the downward momentum survival time distribution. As the economic conditions improve, the time profile of the downward momentum effect for sub-investment grade firms resembles that of investment grade firms: starting from a level 7.5 times higher than the baseline, the effect exponentially declines with time spent in current grade. As the economic conditions worsen, the peak of the downward momentum is reached with increasing delay and the decay of the effect after having reached this peak is slower as compared to benign economic conditions. Shifting the focus to the upward momentum, the upgrade probability for investment grade firms that have been previously upgraded rises much faster in periods of benign economic conditions as compared to periods of stress. Nevertheless, as the economic conditions improve, the time profile of the upward momentum tends to flatten at the point the effect becomes twice as strong as the baseline. Examining the results presented in table 4.4 shows that this change in shape is not statistically significant. On the other hand, for speculative grade firms, there is a very strong positive relationship between the state of the business cycle and the strength of the upward momentum and this relationship is clearly proportional, as evidenced by both the p-value for the $LR_{\hat{\beta}_c^{cfnai,m} / \hat{\beta}_c^{cfnai,m(t)}}$ statistic in table 4.4 and the shift in the time profile depicted in figure 4.3.

4.5.2 Impact on Credit Portfolio Losses

Bank regulators require financial institutions to provide capital to cover potential worst-case portfolio losses. The regulatory requirements take the form of a minimum amount of Tier I and Tier II capital, calculated by aggregating all types of risk (namely credit, market and operational) according to the Basel II or, the newly introduced, Basel III standards (with Basel III further adding a leverage target ratio and capital conservation and countercyclical buffers). For a banking institution credit is the dominant risk type and it is the component most severely affected by the momentum in credit ratings. In this chapter I only focus on corporate exposures, and therefore, for the rest of the section, credit risk refers to the corporate banking book only.

For the calculation of the credit risk capital (banking book) two approaches are provided under Basel II/III: the standardised and the Internal Ratings Based (IRB).

The standardised approach is based on regulatory pre-determined risk weights for each asset class. Under the IRB approach, banks use their internally calculated ratings to estimate the Probability to Default (PD) for each counterparty. The PDs are supplemented by internally estimated Loss Given Default (LGD) and Exposure At Default (EAD) measures under the Advanced Internal Ratings Based (AIRB), or regulatory prescribed levels under the Foundation Internal Ratings Based (FIRB). Internal ratings might or might not coincide with external agency ratings. It is well known that rating agencies assign ratings based on the long-run creditworthiness of a counterparty, putting more weight on rating stability over short term accuracy³. The alignment between the two rating sources largely depends on whether a bank uses a Through-The-Cycle (TTC) assessment of risk for its regulatory capital calculations, or it prefers a more Point-in-Time (PiT) approach. If banks under the IRB approach aim (or they are constrained by the regulators) to have stable capital through time, then they would need to use TTC ratings with long run PDs. In such case, I assume that the agency ratings are good proxies for the internal ratings and therefore the rating momentum affects banking capital in a similar fashion under standardised and IRB approaches. Of course, a more thorough analysis of the impact rating momentum has on banking capital would have to use actual internal ratings from financial institutions, but this is well out of scope for this chapter. For completeness, the reader is referred to Jacobson et al. (2006) for a comparison of internal rating models. Finally, it is worth mentioning that, in the aftermath of the recent credit crisis, the Dodd-Frank Act in the US⁴ explicitly removes references to credit ratings under the standardised approach to credit risk capital calculation. Furthermore, even though the AIRB to calculating capital is unchanged, the Collins Amendment Floor⁵ requires the large banking institutions that are subject to the advanced capital calculations to floor their capital ratios to that of the standardised approach. In light of these recently introduced changes, my results concerning the impact of rating momentum on regulatory capital might be less relevant for the US.

To assess the magnitude of the impact rating momentum has on minimum regula-

³See www.moodys.com, <http://www.standardandpoors.com>, and <http://www.fitchratings.com> for the rating approaches of the 3 major rating agencies

⁴Dodd-Frank Wall Street Reform and Consumer Protection Act of 2010, section 939A

⁵Section 171 of the Dodd-Frank Act, 12 U.S.C. 5371

tory requirements for credit capital, I use the Basel II/III formula:

$$\begin{aligned}
RWA &= 12.5 * UL * EAD * \text{Maturity Adjustment} \\
UL &= LGD \times \left[\Phi \left(\sqrt{\frac{1}{1-R}} \times \Phi^{-1}(PD) + \sqrt{\frac{R}{1-R}} \times \Phi^{-1}(0.999) \right) - PD \right] \\
\text{Mat. Adj.} &= \frac{1 + (M - 2.5) + (0.11852 - 0.05478 * \log(PD))^2}{1 - 1.5 + (0.11852 - 0.05478 * \log(PD))^2} \\
R &= 0.12 \times \frac{1 - e^{-50 \times PD}}{1 - e^{-50}} + 0.24 \times \left(1 - \frac{1 - e^{-50 \times PD}}{1 - e^{-50}} \right),
\end{aligned} \tag{4.15}$$

where UL is the unexpected loss for each unit of exposure, R is the asset correlation, Φ and Φ^{-1} are the cumulative and the inverse cumulative normal functions, and M is the maturity of the exposure. Under Basel II, the minimum capital that a financial institution holds should be no less than 8% of its RWA. The regulatory formula in (4.15) is based on the limiting distribution for an infinitely granular and homogeneous portfolio, essentially defining capital as the difference between the worst case loss at the 99.9% confidence level (potential losses beyond this level are considered too expensive to hold capital against) and the expected loss (which reflects the cost of doing business and it is covered by appropriate provisioning and pricing with no additional capital held against it). For the derivation of the limiting distribution for the infinitely granular and homogeneous portfolio see Vasicek (2002). The asset correlation within the one factor structure assumed by the underlying model is provided explicitly as a decreasing function of the PD, taking values from 12% to 24%. The maturity adjustment component provides a conservative overlay to allow for potential future downgrades of the exposures and therefore it is an increasing function of the maturity and a decreasing function of the PD.

The regulatory capital provides a balance sheet view of credit risk. Credit ratings and PDs also drive the expected loss metrics, that ultimately feed into the Profit & Loss (P&L) calculation. Since the expected loss is a projection of the actual losses a bank expects to incur at a given time horizon, the PDs used for the calculations need to reflect as close as possible the default rate of the portfolio at hand. Therefore, long run PDs are no longer appropriate and PiT equivalents are needed. Even though the derivation of PiT PDs is not the focus of this chapter, I provide a simplified set of semi-parametric regressions that link continuous time PDs to the CFNAI index, consistent with the

rating momentum econometric framework of section 4.3. Due to the low number of defaults, regressions at individual rating grade level show low statistical power and therefore I group ratings at the Letter level. Meaningful results are obtained only for the groups B and Caa-C, with the remaining grades receiving only the baseline intensity contributions. After stratifying by rating grade, the common parameter estimate for the sensitivity of defaults to the CFNAI is $-0.297(0.042)$.

The importance of corporate bond ratings in capital requirements and loss estimates of financial intermediaries is well documented, see Varotto (2012) and the references therein. Varotto (2012) provides a historical perspective of capital requirements using Moody's corporate default, recovery and migration statistics. By benchmarking against realised credit losses Varotto (2012) finds that Basel II regulatory capital is only sufficient to absorb Great Depression-style losses for high credit quality portfolios and 1 year horizon; low credit quality portfolios and longer holding periods could generate losses that cannot be covered by the available capital. The additional buffers introduced under Basel III help to address this concern. Nevertheless, even under Basel III a recapitalisation would be needed as the capital buffers would be almost completely drained.

In addition to the credit risk component, market risk capital (trading book) is also affected by a non-Markovian behaviour in credit ratings via the Incremental Risk Charge (IRC). The IRC regulatory requirements were introduced in response to the recent credit market turmoil, where a number of financial institutions experienced large losses on their trading books, attributed not to default of the underlying credit instruments (a risk that can be quantified by the banking book approach described above), but rather to rating migrations and widening of credit spreads (MtM losses). To cover the potential effect that the rating momentum can have on MtM losses I also quantify the MtM impact on zero coupon corporate bonds, using corporate yields for Letter grades from Moody's. Bond prices take the simple form:

$$P = FV/e^{rt}, \tag{4.16}$$

where FV is the bond's face value, r is the bond's yield and t is the time to maturity.

Due to the non-Markovian nature of the rating momentum effects, transition ma-

trices do not have an analytically tractable form. Therefore, to quantify the impact rating momentum has on portfolio losses, I resort to simulation. To simulate rating migration and default times, I use the inverse CDF approach:

- I let the event arrival time t to be determined by solving the equation

$$S_{ji}(t|\mathbf{Z}(t)) = U, \quad (4.17)$$

where U is a simulated, uniformly distributed random variable. $S_{ji}(t|\mathbf{Z}(t)) = \exp(-\Lambda_{ji}(t|\mathbf{Z}(t)))$ is the survival probability for firm i , when it is currently in grade j . The cumulative intensity function of firm i exiting grade j , $\Lambda_{ji}(t|\mathbf{Z}(t))$, is calculated by summing the intensities across all possible grades (including default) except j :

$$\Lambda_{ji}(t|\mathbf{Z}(t)) = \int_0^t \sum_{k \neq j} \lambda_{jki}(u|\mathbf{Z}(t)) du, \quad (4.18)$$

where $\lambda_{jki}(t|\mathbf{Z}(t))$ corresponds to the intensity of the migration pair $j \rightarrow k$ (where k includes default as a state).

- Given the simulated survival time t , the specific credit event is drawn from the univariate Multinomial distribution with probabilities:

$$\pi_{jki}(t|\mathbf{Z}(t)) = \frac{\lambda_{jki}(t|\mathbf{Z}(t))}{\sum_{k \neq j} \lambda_{jki}(t|\mathbf{Z}(t))} \quad (4.19)$$

Equations (4.17)-(4.19) in the 2 step approach described above, simulate independent events and arrival times. Independence in default occurrence is a very strong assumption on the portfolio level, since observed and unobserved correlation in defaults is well documented. Default correlation can be attributed to multiple sources, see Duffie et al. (2009) for an explanation based on unobserved frailty, and Jorion and Zhang (2009) for an analysis from a credit contagion viewpoint. Instead of explicitly modelling the stochastic correlation process, I use instead a copula function to introduce dependence in default arrival times. Copulas can capture a wide range of dependence structures and have been used extensively in credit portfolio modelling. For more information on copulas and their application to credit risk management see Embrechts et al. (2003), and Rosenberg and Schuermann (2006) and the references therein. To be consistent with the regulatory formula for calculating credit capital and without

a loss of generality, I adopt the Gaussian copula with a common pairwise correlation parameter ρ_c

$$C_{\rho_c}^{Gauss}(u) = \Phi_{\rho_c}(\Phi^{-1}(u_1), \dots, \Phi^{-1}(u_d)),$$

where d is the number of variables, Φ^{-1} is the inverse cumulative distribution function of a standard normal and Φ_{ρ_c} is the joint cumulative distribution function of a multivariate normal distribution with mean vector zero and correlation matrix with a common pairwise correlation parameter ρ_c . Simulating correlated default times is performed by:

- Simulate $i = 1, \dots, n$ correlated standard normal variables $\{X_1, \dots, X_n\}$ via the Cholesky decomposition of the covariance matrix Σ_c that is constructed by using the common pairwise correlation parameter ρ_c for all the off-diagonal elements.
- Calculate $U_i = \Phi(X_i)$, where $\Phi(\cdot)$ is the standard normal cumulative function.
- Substitute U_i in (4.17) and continue as in the independent case.

In addition to the default arrival correlation, I also include a dependence on the times of rating migration. It is intuitive to think that rating agencies do not review firms in isolation and downgrade or upgrade more than one entity at a time. For simplicity I assume a Gaussian copula with the same pairwise correlation as in the default case.

To quantify the impact on a portfolio of credit exposures, I use 2 hypothetical portfolio structures. The first hypothetical portfolio mainly consists of investment grade bonds at a ratio of 80%-20% against sub-investment grade bonds. For simplicity I assume that the exposures are uniformly spread across the rating grades within the investment and speculative grade spectrum. The second portfolio represents a less risk averse lender and consists of primarily speculative grade bonds at a proportion of 20%-80% against investment grade exposures. Just like the high credit quality portfolio, the exposures are assumed to be uniformly distributed across the rating grades. For simplicity I assume that the portfolios are homogeneous each consisting of 1000 bonds, with all bonds par valued at \$100, with no coupon payments and a maturity of 10 years. For each of the 2 hypothetical portfolios, I assume that 30% of the bonds within each of the grades have been previously downgraded, while 20% of the bonds within each of the grades have been previously upgraded. The historical ratio of downward to upward momentum is broadly consistent with the 30%/20% split, while leaving 50% of

the exposures unaffected by previous rating movements allows the end results not to be dominated by the effects of rating drift. To assess the sensitivity of the results to the CFNAI index, I identify 2 periods that loosely correspond to the peak and the trough of the business cycle; the first period corresponds to 01/01/2004-01/01/2005, a year of benign economic conditions, with the real economy growing and a low number of defaults/downgrades, while the second period corresponds to 01/01/2008-01/01/2009, a period of severe stress, with a contracting economy, and an increased number of defaults/downgrades especially after the Lehman Brothers collapse. Finally, to properly assess the impact from adding time varying rating momentum effects, I use two separate assumptions for the time in current grade each bond has spent at the beginning of the period of interest. The first assumption is that all bonds are re-rated at the beginning of the period of interest (i.e. 0 survival time at the start of benign and stress periods), and the second assumption is that all bonds have spent a year in the current grade at the beginning of the period of interest. Each of these 2 different assumptions are applied to each of the 2 hypothetical portfolios for each of the 2 periods that I analyse, resulting in 8 individual cases in total.

I first examine the shift in the portfolio rating distribution by including the, possibly time varying, rating momentum. Figure 4.4 depicts the impact of the rating momentum effect on the portfolio composition, for each of the 2 hypothetical portfolios, across the stress and benign economic periods. For presentation purposes, I only include the case where the portfolio is re-rated at the beginning of the respective period, and therefore, the starting survival time for all exposures is 0. In addition to the distribution at the start of each of the 2 periods, I provide the portfolio distributions 1 year from the starting snapshot, based on 50,000 simulations. The business cycle adjusted transition matrix serves as the base case for predicted 1 year portfolio distribution. This business cycle adjusted transition matrix is obtained by using the baseline duration profiles for each of the rating grades and including the CFNAI index as the only covariate. To show the effect of the rating momentum I include it both proportionally, with β_c^m as the sensitivity in (4.13), and non-proportionally, using the full specification in (4.13). Examining the rating distributions in Figure 4.4 indicates that the presence of momentum has a significant impact on the final portfolio composition, especially in periods of stress. During those periods of adverse economic conditions, rating momentum causes the portfolio to shift towards the worse end of the rating scale.

As an extreme example, during the 2008 stress period, the proportion of Caa-C rated bonds in the speculative grade hypothetical portfolio is approximately 15% higher than the base case when the full specification for the rating drift is used. During benign economic periods, the effect of rating momentum is less pronounced as the overall number of migrations is low, irrespective of the specification used. Nevertheless, even using the 2004 CFNAI index values, there is a clear shift of Aa exposures towards Baa, for the investment grade range, and from Ba to Caa-C, for the speculative grade range. This overall negative effect in creditworthiness also leads to default rates significantly higher when either proportional or non-proportional rating momentum is included in the specification. Further analysis of the differentiation between proportional and non-proportional rating momentum effects shows significant deviations at the boundaries of the rating scale and the Ba grades, where the downward momentum is much higher than the upward momentum (indicating the speculative grade firms are more likely to get downgraded rather than move to the investment grade range). This observation and the significantly higher default rate caused by the non-proportional momentum during stress periods, indicates that choosing the wrong specification for the rating drift can lead to a severe underestimation of credit risk.

I then analyse the impact on portfolio loss and capital requirements. The analysis is split between regulatory long run capital, actual PiT loss, and MtM losses. For capital I use the regulatory formulae (4.15) with average PDs per grade calculated over the entire sample and I report the RWAs. Simulated defaults do not affect the calculations as I assume that defaulted positions are replaced by exposures of the same credit quality. An alternative assumption could be that the portfolio shrinks if defaulted positions are not replaced. I refrain from choosing the latter assumption since large fluctuations in regulatory capital would be the result of defaults rather than migrations. For actual PiT losses, PiT PDs are used based on the parameters of table 4.4. Actual losses refer to default only losses. For MtM losses, the same PiT PDs as in the case of default only losses are used. Losses are based on the change in portfolio market value due to bond yields and default. The non-defaulted bonds, are priced using (4.16). In the case of default, the bond position loses the LGD fraction of its face value. Using the bond prices from equation (4.16), the MtM losses from period 0 to period 1, $L_{0 \rightarrow 1}^{MtM}$, can be

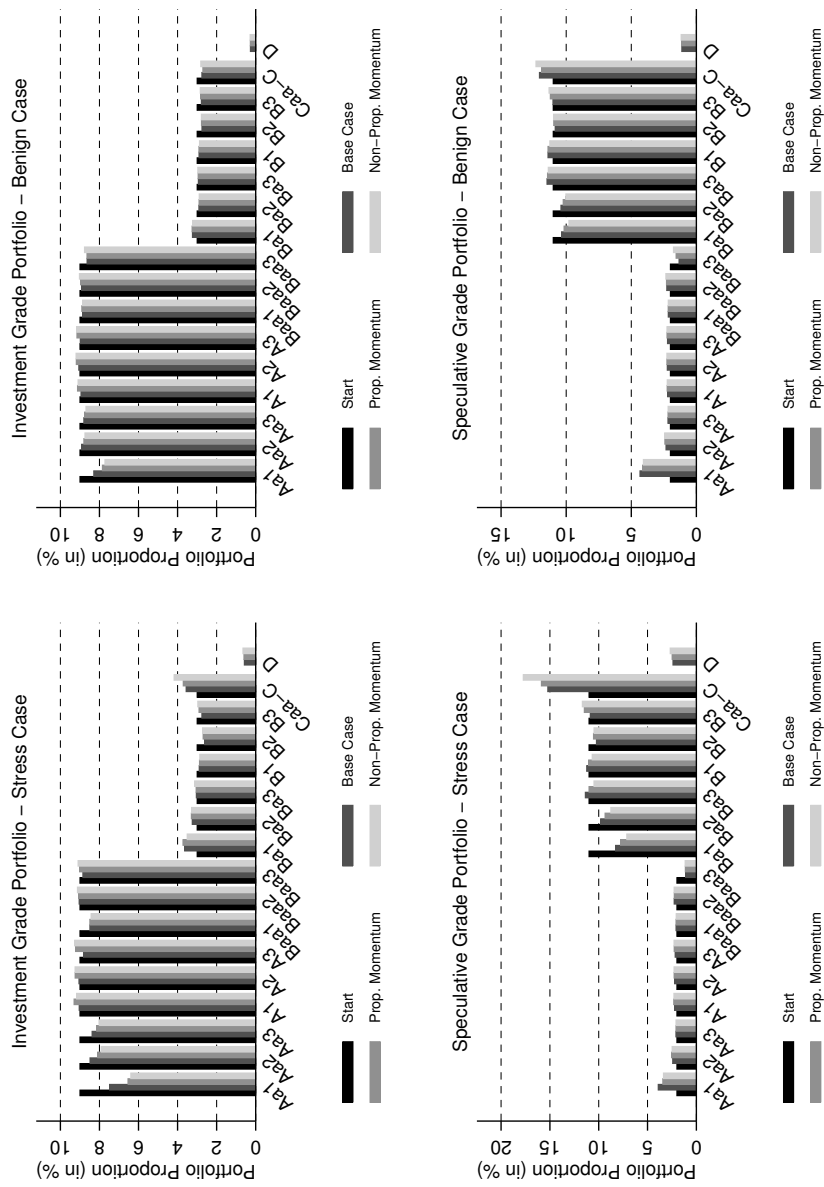


Figure 4.4: Predicted portfolio distributions

Predicted 1 year distributions for investment grade and speculative grade hypothetical concentrated portfolios, under stress and benign economic conditions. All portfolios contain 1000 bonds. The investment grade concentrated portfolio has 80%-20% investment/sub-investment grade bonds, spread uniformly across grades, while the speculative grade concentrated portfolio has a 20%-80% split across investment/sub-investment grade bonds, with uniform spread across grades. The Aaa rating grade is excluded since it cannot have a downward drift. I define as "benign" economic conditions the period 01/01/2004-01/01/2005, and as "stress" economic conditions the period 01/01/2008-01/01/2009. "Start" refers to the portfolio at the start of each period. "Base Case" is the portfolio distribution at the end of each period, when only the CFNAI index is used as a covariate. For the "Prop. Momentum" case I add the proportional rating momentum effect to the baseline business cycle fluctuations of the "Base Case". Finally, the "Non-Prop. Momentum" case refers to the full model, which allows the rating momentum to vary with time spent in current grade and business cycle movements.

expressed as:

$$L_{0 \rightarrow 1}^{MtM} = [FV(1 - LGD)]_{D=1} + [FV/e^{r_1(t-\tau_1+\tau_0)}]_{D=0} - FV/e^{r_0 t}, \quad (4.20)$$

where $\tau_1 - \tau_0$ refers to the time passed between period 1 and period 0, t is the maturity of the bond at period 0, and r_0, r_1 are the bond yield at period 0 and 1 respectively. For all 3 types of losses, baseline migration and default intensities are calculated using the kernel smoother (4.4), described in section 4.3.1.

The naive assumption that all migrations only receive baseline intensity contributions adjusted for the point in the economic cycle serves as the base case for all rating momentum results. Therefore, the base case corresponds to duration effects proportionally shifted according to the value of the CFNAI index. To further differentiate the source of deviation from the base case, the results are split by proportional rating momentum effects (with β_c^m as the sensitivity in (4.13)), and rating momentum with duration and business cycle dependence (the full specification in (4.13)). The process of deriving PiT PDs is the same under both base and rating momentum cases, since the PiT PDs are not related to rating drift. To better capture the effect of duration, I analyse the portfolio's behaviour when all the bonds are re-rated at the beginning of the observation period and when all bonds have spend 1 year in their current rating grade. Finally, to quantify the impact of default and migration clustering, I provide results based on a Gaussian copula with 50% pairwise correlation, in addition to the case of independent credit event occurrence.

Tables 4.5 and 4.6 provide the portfolio impact results for the 2 hypothetical portfolios chosen for the analysis, based on 50,000 simulations. As mentioned above, the first portfolio primarily consists of investment grade bonds (at a ratio of 80%-20% against speculative grade bonds), while the second portfolio reflects a less risk averse investor/financial institution and consists of 80% sub-investment grade bonds. Both tables report the RWA for the respective portfolios, average and 99% actual losses, and average and 99% MtM losses. All figures, except the RWAs, reflect losses and negative results indicate a gain. The MtM losses assume that the investor is long in the portfolio of bonds and that the yields rise within the observation window (which corresponds to a price drop of the bonds in the portfolio). Conversely, losses can also occur if the investor is short in the portfolio of bonds and the yields drop within the

observation window. Results are presented for 1 and 2-year observation windows, starting from January 2004 and January 2008. The former period reflects benign economic conditions, while the latter covers the peak of the recent economic downturn, including Lehman Brothers collapse in September 2008. The periods are chosen so that the peak of the economic activity and the peak of the economic downturn approximately coincide with the the end of the 1 year window for the respective periods. Therefore, the 2-year observation window corresponds to a deterioration in economic activity and an increase in corporate bond yields for the period starting in January 2004, and an improvement in the economic activity combined with lower bond yields for the period starting in January 2008. Results for the momentum effects are reported as percentage deviations from the base specification.

For the portfolio loss impact analysis additional parameters are needed. The Gaussian copula for dependent default and migration times is parameterised using 50% as the pairwise correlation value. LGD is set across all periods to its stress value of 66%. The stress value is based on 30 days post default market prices and corresponds to the observed level over the 2008-2009 period. Basel II/III regulatory capital is recommended to be calculated using "downturn" LGDs and for comparability with the additional loss measures, I apply the same stress LGD value to all results. Alternatively a stochastic LGD could be used instead, with the Beta distribution frequently chosen in practice. I refrain from adding complexity to the LGD, since the focus of this chapter is the PD component. I source rating specific yields from Moody's. I use the median yield of regular coupon bonds, maturing within 6-8 years and having an outstanding value of more than \$50 million. The median is calculated based on 1000-1200 bonds each month. Finally, CFNAI data are sourced from the FRED⁶.

The results for the base case in both tables 4.5 and 4.6, confirm the choice of 01/01/2008-01/01/2010 as the stress period; both average and quantile losses for the 2008-2010 period are approximately twice as high as the losses for the more benign economic period 01/01/2004-01/01/2006. Examining the MtM losses, the relative impact of the 2008-2009 period is dramatic: for the investment grade portfolio, MtM losses jump from approximately \$500 in 2004 to approximately \$14,000 in 2008 (an MtM loss implying that 14% of the face value is wiped off during the peak of the stress), while

⁶<http://research.stlouisfed.org/fred2/>

	1 Year Horizon					2 Year Horizon						
	RWA	Loss-Av	Loss-99%	MtM-Av	MtM-99%	RWA	Loss-Av	Loss-99%	MtM-Av	MtM-99%		
January 2004												
0% Asset Correl.	Base											
	0 Starting Dur.	88,003	181	463	473	763	88,411	315	661	3,886	4,235	
	1Y Starting Dur.	88,258	316	661	587	966	88,703	460	925	3,980	4,409	
	Prop. Drift											
	0 Starting Dur.	0.58	1.59	0.00	2.08	2.44	1.56	4.89	10.00	1.10	2.14	
	1Y Starting Dur.	1.62	3.76	9.09	6.52	7.26	2.76	8.20	0.00	2.72	3.55	
0% Asset Correl.	Non-Prop. Drift											
	0 Starting Dur.	0.85	4.83	0.00	8.68	7.67	1.48	8.76	10.00	1.79	2.93	
	1Y Starting Dur.	1.44	5.60	9.09	6.93	6.34	1.91	8.04	0.00	1.52	2.56	
	50% Asset Correl.	Base										
		0 Starting Dur.	88,003	181	2,181	473	4,138	88,411	315	3,239	3,886	8,468
		1Y Starting Dur.	88,258	316	3,305	587	5,573	88,703	460	4,165	3,980	9,540
Prop. Drift												
0 Starting Dur.		0.58	1.59	0.00	2.08	1.90	1.56	4.89	2.04	1.10	1.65	
1Y Starting Dur.		1.62	3.76	1.89	6.52	1.06	2.76	8.20	1.59	2.72	1.88	
50% Asset Correl.	Non-Prop. Drift											
	0 Starting Dur.	0.85	4.83	0.00	8.68	2.15	1.48	8.76	2.04	1.79	3.11	
	1Y Starting Dur.	1.44	5.60	1.89	6.93	0.69	1.91	8.04	1.59	1.52	1.80	
	January 2008											
	0% Asset Correl.	Base										
		0 Starting Dur.	89,693	386	793	13,962	14,170	91,597	744	1,256	-1,824	1,609
1Y Starting Dur.		90,945	492	925	14,261	14,520	92,878	866	1,454	-1,569	2,016	
Prop. Drift												
0 Starting Dur.		1.63	3.44	0.00	1.77	2.00	3.48	7.87	5.26	-27.17	16.49	
1Y Starting Dur.		3.30	9.28	7.14	3.94	4.91	5.13	16.19	9.09	-55.97	28.30	
0% Asset Correl.	Non-Prop. Drift											
	0 Starting Dur.	2.09	12.18	8.33	3.52	3.92	3.23	20.24	15.79	-39.51	26.56	
	1Y Starting Dur.	2.47	11.61	7.14	3.48	3.91	3.39	15.65	9.09	-43.10	21.67	
	50% Asset Correl.	Base										
		0 Starting Dur.	89,693	386	793	13,962	16,358	91,597	744	1,256	-1,824	1,609
		1Y Starting Dur.	90,945	492	925	14,261	16,742	92,878	866	1,454	-1,569	2,016
Prop. Drift												
0 Starting Dur.		1.63	3.44	0.00	1.77	4.37	3.48	7.87	0.69	-27.17	57.07	
1Y Starting Dur.		3.30	9.28	0.84	3.94	6.09	5.13	16.19	3.66	-55.97	64.67	
50% Asset Correl.	Non-Prop. Drift											
	0 Starting Dur.	2.09	12.18	2.08	3.52	9.46	3.23	20.24	4.17	-39.51	108.08	
	1Y Starting Dur.	2.47	11.61	0.84	3.48	8.70	3.39	15.65	4.88	-43.10	84.31	

Table 4.5: Credit portfolio loss simulation results - Investment Grade Portfolio

Simulation results for regulatory capital, actual loss and MtM loss for the hypothetical portfolio consisting of 80%-20% investment and sub-investment grade bonds respectively. Portfolio results are provided for the periods January 2004-January 2006, and January 2008-January 2010. The first period corresponds to benign economic conditions, while the second period is representative of a stress economic environment. To properly capture the impact of both baseline and rating momentum duration, I differentiate between all bonds starting from 0 duration and all bonds having spent 1 year in their current grade. For the "Base" case I only use the CFNAI index as a covariate. For the "Prop. Drift" case I add the proportional rating momentum effect to the baseline business cycle fluctuations of the "Base" case. Finally, the "Non-Prop. Drift" case refers to the full model, which allows the rating momentum to vary with time spent in current grade and business cycle movements. All results correspond to loss estimates and, therefore, negative figures indicate a gain. All rating momentum results are reported as differences from Base case, expressed in %.

for the sub-investment grade portfolio an MtM gain of \$500 in 2004 compares to a \$21,000 loss in 2008 (an MtM loss level corresponding to 21% of the face value wiped off). During the second year of each observation window the MtM differences are much less pronounced, as corporate bond yields return to levels similar to those observed at the beginning of each window. For the investment grade portfolio, a loss of \$4,000 in 2005 is compared to a \$1,800 gain in 2009, while a loss of \$3,600 reduces to \$2,000 for the sub-investment grade portfolio. From a regulatory capital perspective the results are less clear; the investment grade portfolio shows an increase in capital requirements in 2008-2009 as compared to 2004-2005, but the reverse is true for the sub-investment

	1 Year Horizon					2 Year Horizon				
	RWA	Loss-Av	Loss-99%	MtM-Av	MtM-99%	RWA	Loss-Av	Loss-99%	MtM-Av	MtM-99%
January 2004										
0% Asset Correl.										
Base										
0 Starting Dur.	128,952	741	1,322	-532	-144	127,343	1,263	1,983	3,622	4,090
1Y Starting Dur.	127,309	1,277	1,983	-257	244	125,767	1,818	2,644	3,879	4,437
Prop. Drift										
0 Starting Dur.	0.03	1.58	0.00	-0.73	-7.56	-0.11	4.89	3.33	-0.06	1.04
1Y Starting Dur.	0.01	3.63	3.33	-0.54	11.34	-0.37	8.00	7.50	0.63	2.14
Non-Prop. Drift										
0 Starting Dur.	-0.12	4.81	0.00	-1.79	-60.90	-0.30	8.44	6.67	1.63	2.73
1Y Starting Dur.	-0.18	5.35	3.33	-0.87	13.61	-0.41	7.66	5.00	-1.43	0.15
50% Asset Correl.										
Base										
0 Starting Dur.	128,952	741	8,527	-532	4,875	127,343	1,263	12,295	3,622	10,608
1Y Starting Dur.	127,309	1,277	12,493	-257	7,620	125,767	1,818	15,600	3,879	12,716
Prop. Drift										
0 Starting Dur.	0.03	1.58	0.78	-0.73	1.00	-0.11	4.89	0.54	-0.06	0.43
1Y Starting Dur.	0.01	3.63	0.00	-0.54	0.56	-0.37	8.00	1.27	0.63	1.48
Non-Prop. Drift										
0 Starting Dur.	-0.12	4.81	1.55	-1.79	2.29	-0.30	8.44	1.61	1.63	1.52
1Y Starting Dur.	-0.18	5.35	0.53	-0.87	1.02	-0.41	7.66	1.27	-1.43	1.01
January 2008										
0% Asset Correl.										
Base										
0 Starting Dur.	126,749	1,587	2,380	20,895	21,159	123,188	2,980	4,032	2,087	2,697
1Y Starting Dur.	125,604	1,974	2,842	21,203	21,533	122,102	3,407	4,495	2,458	3,125
Prop. Drift										
0 Starting Dur.	-0.01	3.51	2.78	1.44	1.59	-0.30	7.56	4.92	28.22	22.73
1Y Starting Dur.	-0.18	8.96	6.98	2.95	3.20	-0.86	15.02	11.76	40.96	33.84
Non-Prop. Drift										
0 Starting Dur.	-0.59	11.79	8.33	3.73	4.05	-1.47	18.61	14.75	53.71	43.85
1Y Starting Dur.	-0.44	11.03	9.30	2.51	2.86	-0.89	13.25	11.76	28.23	24.24
50% Asset Correl.										
Base										
0 Starting Dur.	126,749	1,587	12,493	20,895	24,014	123,188	2,980	17,913	2,087	10,876
1Y Starting Dur.	125,604	1,974	14,872	21,203	24,499	122,102	3,407	19,962	2,458	12,156
Prop. Drift										
0 Starting Dur.	-0.01	3.51	0.53	1.44	2.15	-0.30	7.56	0.18	28.22	4.28
1Y Starting Dur.	-0.18	8.96	0.44	2.95	2.74	-0.86	15.02	1.66	40.96	7.03
Non-Prop. Drift										
0 Starting Dur.	-0.59	11.79	0.53	3.73	6.20	-1.47	18.61	2.21	53.71	11.43
1Y Starting Dur.	-0.44	11.03	0.89	2.51	4.53	-0.89	13.25	1.99	28.23	7.47

Table 4.6: Credit portfolio loss simulation results - Speculative Grade Portfolio

Simulation results for regulatory capital, actual loss and MtM loss for the hypothetical portfolio consisting of 20%-80% investment and sub-investment grade bonds respectively. Portfolio results are provided for the periods January 2004-January 2006, and January 2008-January 2010. The first period corresponds to benign economic conditions, while the second period is representative of a stress economic environment. To properly capture the impact of both baseline and rating momentum duration, I differentiate between all bonds starting from 0 duration and all bonds having spent 1 year in their current grade. For the "Base" case I only use the CFNAI index as a covariate. For the "Prop. Drift" case I add the proportional rating momentum effect to the baseline business cycle fluctuations of the "Base" case. Finally, the "Non-Prop. Drift" case refers to the full model, which allows the rating momentum to vary with time spent in current grade and business cycle movements. All results correspond to loss estimates and, therefore, negative figures indicate a gain. All rating momentum results are reported as differences from Base case, expressed in %.

grade portfolio. This difference in behaviour between investment and sub-investment grade portfolios can be explained but the non-linearity of the RWA function. The PD level that maximises the RWA formula in (4.15) is close to the TTC PD for the Ba3 grade. As the average portfolio grade shifts to the B and Caa-C grades, the RWA figure will decrease. Given the severity of the economic conditions during 2008-2009, the non-linearity of the RWA formula causes the sub-investment grade portfolio to have lower regulatory capital requirements. This behaviour is counterbalanced by the significantly higher provisions that financial institutions would have to raise during the period 2008-2009 to cover the impaired loans. Finally, losses are higher under the assumption that

all the bonds in the portfolio have maintained their rating for 1 year at the beginning of the observation window. This observation is consistent with the baseline hazard curves in figure 4.2; migration and default intensities peak between 1 and 2 years since rating assignment.

Examining the results for the rating momentum clearly shows that the drift in rating migrations causes a shift in the loss distribution, with significantly higher average and tail losses. This increase in losses is more pronounced during periods of severe economic conditions; using the 2008 downturn as a proxy for stress conditions, the 1 year average losses for the rating momentum can be more than 11% higher than the base case (close to 20% for the 2 year cumulative losses), while during the 2004-2005 benign economic period, the average loss for rating momentum can be approximately 5% higher than the base case (close to 8% for the 2 year cumulative losses). When the rating momentum effect is constrained to be proportional, the above mentioned figures can be 1.5-4 times lower, and this difference is significantly higher when starting from a 0 duration. This finding is consistent with the shape of the time profiles in Figure 4.3; the migration intensity between 0 and 365 days is much higher than the average intensity across the entire time horizon, and therefore, the proportionality assumption would lead to materially different results over that timeframe. This extremely high difference between proportional and non-proportional rating drift highlights the importance of correctly specifying the momentum effect, especially for portfolios of recently re-rated bonds.

Due to the relative scarcity of default events and the much higher base case loss estimates, the percentage increase of the tail loss under the presence of rating momentum is less pronounced than the increase in average losses. When assuming 0% asset correlation, the tail loss increase over the base case is 20%-25% less than the equivalent increase in average loss (this percentage is calculated by averaging the reported results for all cases). For the 50% asset correlation, the magnitude of the tail loss increase can be 10 times lower than the increase in average losses; across the different periods and the 2 portfolio structures used in the analysis, the Gaussian copula based tail losses are between <1bp and 200bp higher than the base case, with the average deviation in the 80bp range. This range of percentage increase in 99% losses under the presence of rating momentum, is broadly consistent with the 107bp average increase in VaR

reported in Güttler and Raupach (2010) across the period 1996-2005. Nevertheless, a few differences should be highlighted when comparing results across the two studies: a) Güttler and Raupach (2010) use S&P ratings from 1996 to 2005, unlike the present study that uses Moody's ratings over the period 1983-2008 and therefore includes two additional periods of stress (early '90s recession, 2008 credit crisis), b) the base case for Güttler and Raupach (2010) is the time homogeneous Markov chain transition matrix, unlike the present study that assumes for the base case a transition matrix depended on the macroeconomic environment, c) Güttler and Raupach (2010) use 99.9% confidence level for the VaR, unlike the 99% level chosen in the current study, d) the portfolio composition across the 2 studies is different, e) the methodology of deriving momentum sensitive transition matrices is different across the 2 studies. Furthermore, despite the lower impact relative to the average losses, the absolute difference from the base case can be marginally more pronounced for low quality portfolios. With 0% asset correlation, the increase in tail losses due to the presence of momentum is \$50 higher than the equivalent increase in average losses over the 1 year horizon (the absolute amount impact of the rating momentum for tail losses is approximately \$265), and \$75 higher over the 2 year horizon (the absolute amount impact of the rating momentum for tail losses is approximately \$530).

The MtM estimates include credit spread fluctuations in addition to defaults, and, therefore, losses can be more extreme than the pure jump-to-default case, depending on the direction corporate bond yields move over the observation period. Exploring the base case results for the stress period January 2008-January 2009, it is very apparent that the extreme increase in corporate borrowing premiums resulted in average MtM losses 35 times higher than purely default related write-downs for the low risk portfolio (across the 1 year horizon, approximately \$14,000 across 0 and 1 year initial duration as compared to approximately \$400); the high risk portfolio has a substantial jump-to-default component and therefore the equivalent increase for this portfolio is approximately tenfold. Given the extreme jump in base MtM losses, the average 3.5% increase due to the presence of rating momentum is substantially lower than the equivalent increase in jump-to-default losses (approximately 11% across the 2 portfolios). Nevertheless, in absolute figures, the momentum-adjusted MtM loss estimates are higher than the equivalent default only estimates; for the extreme stress period starting on January 2008, the absolute impact of rating momentum on average MtM losses is

more than \$700 over the 1 year horizon, while the equivalent increase in jump-to-default losses is close to \$200 (over the 2 year horizon, the absolute impact is approximately \$900 for the MtM losses and \$500 for the jump-to-default losses). This remark highlights that the combination of increased migrations caused by the rating drift and high credit spreads can severely hurt investors that risk manage their fixed-income positions based purely on the state of the business cycle. It is worth noting that, consistently with the jump-to-default losses, the time varying momentum leads to 1.5-4 times higher MtM losses across the 2 portfolios if the portfolio is re-rated at the beginning of the observation period. Nevertheless, due to the non-monotonic relative movement of credit spreads across the rating scale, when starting from 1 year duration time, the time varying momentum results are generally lower than the proportional momentum.

4.6 Conclusion

In this chapter I explore the dynamics of the momentum effect in credit rating migrations. I confirm the results of earlier studies, by reporting a very strong effect of previous rating movement on subsequent migration direction, with the, so-called, rating momentum effect quantified as an approximately 5 times increase in downgrade intensity and 3 times increase in upgrade intensity on average. Unlike previous studies that use pre-credit crunch data, the results of this study are based on an extended dataset including the recent economic downturn and, therefore, the reported impact of rating momentum effect reflects the recent increase in credit risk.

I further show that there is a strong dependence of this momentum effect on time spent in a given rating grade and that this dependence is highly non-linear. The shape of the time profile for the rating momentum can vary from an exponential decay for investment grade downgrades to a concave increasing function for investment grade upgrades. For sub-investment grade firms, the momentum effect for both downgrades and upgrades monotonically increases to reach a maximum and then slowly decays as firms stay longer in their current rating grade. These findings provide strong evidence that the time-invariant momentum effects used in earlier studies are inadequate to properly capture the dynamics of credit transition matrices. Furthermore, failing to account for the time-varying nature of the rating momentum effect questions the robustness of

any statistical testing based on proportional hazard models, since any dependence on survival time violates the proportionality assumption of these models.

In addition to the dependence on time spent in the current grade, I also test whether the rating momentum effect varies across the business cycle, possibly non-proportionally. The results indicate that the rating momentum effect for investment grade firms remains largely unaffected across the different phases of the economic cycle, even though for investment grade upgrades there are very few migrations during periods of stress in order to draw robust conclusions. For speculative grade firms, the effect of the business cycle on the rating momentum effect is very pronounced, with a clear non-proportional increase in downward momentum and an approximately proportional decrease in upward momentum during periods of stress. Furthermore, the downward momentum can be more than twice as strong in stress conditions as compared to the long run average, shedding some light into the amplification of economic shocks via the domino effect that corporate downgrades have on credit conditions.

The impact of each of the above effects is quantified across different portfolio structures and economic conditions. Basel II/III regulatory capital remains largely unaffected by the presence of rating momentum, with noticeable differences mainly during stress conditions for high credit quality portfolios. On the contrary, the impact of rating momentum on jump-to-default losses is proved to be highly significant, with an increase of 11% and 20% in 1-year and 2-year average losses during periods of stress. These significant uplifts can be attributed mainly to time and business cycle dependence of the rating momentum, since removing those non-proportional effects leads to 1.5-4 times lower loss figures, especially if the portfolio under consideration is recently re-rated. I report results for the tail losses consistent with the study of Güttler and Raupach (2010), despite the differences in methodology and data used. I choose the 99% confidence level to define the tail losses and, in percentage terms, the impact of rating momentum is smaller than the average losses. The absolute amount impact strongly depends on the underlying portfolio's composition. For high credit quality portfolios with extremely low default probabilities, the presence of momentum adds very little as compared to the average loss estimates. For low credit quality portfolios, with a significant proportion defaulting, rating momentum adjusted tail losses increase substantially more than the average losses, especially as defaults accumulate over a 2

year horizon. Finally, due to the extreme movements of credit spreads during periods of stress, MtM losses can be more pronounced than jump-to-default losses; for the 2008-2009 period, the absolute MtM loss amount is 3.5 times higher than the equivalent amount for jump-to-default losses over the 1-year horizon and 1.8 times over the 2-year horizon.

Due to the focus of this chapter on capturing the duration and business cycle dependence of the rating drift effect, the statistical specification is restricted to ensure the robustness of results. Therefore, the specification can be extended in various ways to facilitate more complex dynamics. A straightforward extension is to make the time varying momentum effect rating grade and industry specific. Nevertheless, credit migrations are relatively infrequent events and the sample of rated companies is not large enough to allow for a robust analysis of the time varying rating momentum dynamics at such a granular level. Therefore, differentiating the time dependent rating momentum dynamics by rating grade and industry is left for future research as more data become available.

Another extension can be the inclusion of latent systematic variables. My specification assumes a single source of systematic fluctuations in credit migrations, which takes the form of the CFNAI diffusion index. This variable does not correspond to an actual economic measure, but it is constructed so that it summarises the co-movements in a very large panel of macroeconomic and financial measures. Despite being a calculated measure, it is treated as observable for modelling purposes, since it is not estimated jointly with the credit migrations. A number of studies have reported a strong residual effect in rating migrations after generic macroeconomic and financial movements are taken into account. This unobservable clustering in rating migrations is usually attributed to economy-wide or industry-wide frailty, and is captured by a set of autocorrelated factors jointly estimated with the observable part of the model, see Koopman et al. (2008) for more details. To include such effects parametric assumptions are needed, therefore relaxing the semi-parametric nature of the Cox model used in this study. Furthermore, the likelihood evaluation for such type of models requires multi-dimensional integration, leading to simulation based methods for the estimation that can dramatically increase the computational time.

Chapter 5

Conclusion

5.1 Summary

In the aftermath of the recent financial crisis, the thesis aims at exploring the interlinkages between credit risk and the macro-financial environment. The research themes revolve around 3 major topics in today's risk management and policy making. First, conditional on the macroeconomic conditions how well can be forecast corporate defaults? Default forecasting is at the heart of future loss projections for any financial institution while it also has implications concerning the pricing of fixed-income portfolios. Second, what is the impact of structural shocks on credit losses? Structural shocks and their impact on credit portfolios are very important in policy making, but can be extremely helpful when forming hypothetical scenarios for stress testing of financial institutions. Third, what are the dynamics of serial dependence in rating migrations and what are the implications of using the Markovian assumption when projecting transition matrices? The non-Markovian characteristics in corporate transition matrices can affect both investors (that MtM their corporate bond portfolios) and financial institutions (that base their capital requirements on rating grades).

The first question deals with the forecasting power of the macroeconomic environment on corporate bond defaults. The 2008 credit crisis highlighted the importance of a systematic view on credit risk; nevertheless, empirical evidence is inconclusive concerning which economy-wide variables help to robustly predict the default cycles. To

avoid the individual variable selection bias, I develop a high dimensionality dynamic factor model from a panel of 103 macroeconomic and financial indicators. Unlike the existing literature on dynamic factor modelling for credit risk prediction, the proposed specification provides a clearly interpretable decomposition of the economic environment fluctuations into a small set of factors. I propose a novel MCMC algorithm based on Gibbs Sampling to combine the economy-wide dynamic factors with US corporate bond specific latent factors and forecast default probabilities at 3-month, 6-month and 1-year horizons. I provide evidence that the dynamic factors can lead to significant gains in default probability forecasting performance of the model. The improvement in forecasting performance over simplified specifications having Industrial Production as the only covariate can be as high as 90% in terms of RMSE and more than 80% in terms of MAE. Nevertheless, the gains from using dynamic and latent factors becomes less pronounced the longer the default forecasting horizon.

The second question deals with the effects of monetary policy and macro-financial shocks on corporate bond defaults and recoveries. I use a VAR model with a novel semi-structural identification scheme to isolate the various shocks at the macro level. In order to quantify the impact of the shocks on Moody's default and recovery data, I use a partially non-Gaussian specification. The econometric specification is sufficiently flexible to separate the macroeconomic effects from the credit-specific systematic dynamics and to isolate the different shock transmission mechanisms. I report intuitive impulse response functions for all macroeconomic variables in the analysis and the results highlight the importance of the cost/working capital and balance sheet channels in the monetary transmission process. Corporate default likelihood is strongly affected by balance sheet and real economy shocks for the cyclical industry sectors, while the effects of monetary policy shocks typically take up to one year to materialise. In contrast, recovery rates tend to be more sensitive to asset price shocks, while real economy shocks mainly affect secured debt recovery values. Finally, I present strong evidence that semi-structural shocks account for approximately 45% of the long-run forecast variance for defaults and approximately 34% for recovery rates.

The third question explores the stability of the autocorrelation in credit rating migrations and their, possibly non-linear, dependence on the business cycle. I provide robust evidence that there is a strong correlation between successive credit rating move-

ments. I show that this so-called rating momentum effect is not constant through the cycle; its magnitude exhibits a non-linear dependence to the time spent in a given rating grade and changes across the different states of the economy. Furthermore, the rating momentum is more pronounced for downgrades as compared to upgrades, confirming the existing evidence of a predominately downwards drift in credit ratings. Empirical evidence from Moody's Default & Recovery database suggests that the presence of this directional dependence has a significant impact on portfolio losses and regulatory capital. In extreme periods of stress, average losses can be 20% higher over a 2-year horizon and this increase can be twice as high for MtM losses. The Basel II/III regulatory capital is less affected by the presence of momentum, but capital requirement for high quality portfolios can be underestimated by more than 3% in times of stress.

5.2 Contribution to knowledge

The results of the thesis provide significant extensions to the existing literature and further enhance the understanding of the macroeconomic impact on corporate credit risk.

First, I provide further evidence that the use of dynamic factors can lead to significant improvements in forecasting performance. This is in line with the work of Koopman et al. (2011) that use dynamic factors based on principal components. While Koopman et al. (2011) benchmark the forecasting performance using a constant-only base specification, I provide a more robust comparison by using a base specification with Industrial Production as the sole covariate. I show that the improvement in forecasting performance strongly depends on the industrial sector and the forecasting horizon, but can be as high as 90% in terms of RMSE for the 3-month horizon. Unlike prior work in the literature, I introduce a semi-structurally identified set of factors that provides an intuitive interpretation of the different macro-financial dynamics, in addition to generating independent covariates for use in corporate default regressions. Furthermore, unlike the contemporaneous nature of the existing work on dynamic factors and the credit side of the economy, I assess the forecasting power using different horizons for the default rates. That allows for a better understanding of the forward looking performance of the model and the corresponding implications for risk management.

Second, I provide the first structural decomposition of corporate bond defaults and recoveries in terms of fundamental macroeconomic shocks. I introduce a new semi-structural identification scheme for a small scale VAR model, that combines a set of short- and long-run restrictions to isolate theory-based economic shocks. This is a significant addition to the existing literature that typically identifies shocks based on recursive restrictions. I show that the identified shocks have a meaningful interpretation and the impulse-response analysis ensures that the responses of all the input variables are intuitive. Being the first structural analysis of corporate bond credit risk metrics, I show that this newly introduced semi-structural model generates macroeconomic shocks that account for approximately 45% of the default forecast variance and approximately 34% of the recovery forecast variance. Finally, I find empirical support for the work of Barth and Ramey (2002), Chowdhury et al. (2006), and Ravenna and Walsh (2006) by reporting robust and strong evidence for the presence of a working capital/cost channel transmission mechanism in monetary policy. This is of particular importance for policy makers as it implies that corporates are able to partially pass on increases in their cost of debt to consumers. The presence of a cost channel causes some inflationary pressures when corporates struggle to attract funding due to balance sheet shocks, or face increasing yields in their debt as a result of monetary policy shocks.

Third, I provide a thorough analysis of the so-called momentum effect on corporate bond ratings. The analysis is an extension of the work by Lando and Skødeberg (2002) and Güttler and Raupach (2010) and explores a number of themes not considered before. While rating momentum in the existing literature is treated as time-invariant with a proportional impact on rating migrations, I provide robust evidence that the serial dependence in rating migrations depends non-linearly on the business cycle and the time a given issuer spends in a rating grade. These highly non-proportional extensions of the basic Cox proportional hazard model used in prior studies, also help to remove the bias associated with testing non-proportional effects in a proportional hazard setting. The work presented in this thesis enhances our knowledge on the behaviour of rating agencies and shows that the proportionality typically assumed in time-to-event studies of corporate migrations can lead to misleading inference. Furthermore, I provide a detailed impact analysis of the rating momentum effect on credit loss and regulatory capital metrics. These results are very important for risk management as they show that the Markovian assumption typically associated with transition matrices can

lead to a severe underestimation of rating migration risk. While Güttler and Raupach (2010) also explore the impact of serial dependence on VaR, the results reported in this thesis constitute the first attempt to quantify the impact of non-proportional rating momentum on a number of credit loss metrics, ranging from actual and MtM losses to regulatory capital for financial institutions.

5.3 Implications for Risk Management and Policy Making

The analysis and results presented in chapters 2-4 touch a number of topics in the broader field of credit risk modelling and have implications for investor and financial institution risk management, micro-prudential regulation, financial stability and monetary policy. While the chapters are independent and each one addresses a distinct research question, the results can be seen as a unified framework to assess the sensitivity of the different aspects of credit risk to a changing economic environment and to quantify potential losses that might occur as the results of adverse market movements.

The effects of credit risk on the banking system, in the aftermath of the recent credit crisis, have been the focus of regulators on both sides of the Atlantic. Regulatory actions, in the form of stress testing exercises, aim to assess the solvency of the banking sector and identify the need for additional capital requirements. The analysis and results of this thesis provide useful insights into two complimentary approaches to the scenario building and assessment process. First, the dynamic factor approach of chapter 2 can lead to internally consistent scenarios across many macroeconomic variables. As the dynamic factors in chapter 2 have an economic interpretation, stress scenarios can be defined directly in the factor space. For example a deterioration in business cycle conditions by 2 standard deviations coupled with a sharp increase in the cost of debt by 1 standard deviation can form the scenario narrative; based on this scenario narrative, observed macroeconomic series paths can be derived directly from the model. A second approach is to use the structural view of the economy followed in chapter 3 to form extreme but plausible shocks to drive the stress scenarios. The structural view would enable one to express forward looking views on the economy in terms of independent, fundamental sources of macroeconomic activity. For example the scenario narrative might be defined in terms of drop in aggregate demand, combined with a corporate

balance sheet shock and tightening of monetary policy. Under both approaches, the reduced-form econometric specifications for the default and recovery rates can be used to quantify the effect of the defined macro-financial scenario on credit losses.

In addition to scenario analysis and stress testing, the methodologies presented in chapter 2 can also be used for PiT forecasting of credit portfolios losses, both from an investor's and a financial intermediary's perspective. For financial intermediaries loss forecasting is essential in provision setting and pricing. Collective (or general) impairment provisioning reflects the estimated amount of losses incurred on a collective basis, without having been individually identified. To set collective provisions financial institutions typically use aggregate loss forecasting based on appropriate portfolio segments (sector and/or credit quality). Therefore, robust forecasting tools that make best use of the large amount of macroeconomic data available are very important in ensuring that default rate forecasts over various horizons do not materially over- or under-shoot the actual realisations. In addition to provisioning, accurate default probabilities are essential for the pricing of private or public debt instruments. Pricing of loans should reflect the counterparty's probability of default over the lending duration. Similarly for investors in the corporate bond market, accurate predictions of default occurrence can drive investment decisions (if the investor thinks the bond is over/under-priced based on his/her belief of future evolution of default rates) or can feed in the risk management process to determine the maximum potential loss a given portfolio might incur.

In addition to the PiT credit measures, corporate credit ratings play a major role in today's risk management. Investors in the corporate bond market often base their decisions on the rating of issuer (which is often link to the bond's yield). Potential future changes of the issuer's credit grade can therefore have a material impact on the value of a portfolio that is MtM. Furthermore, the capital requirements for banking book positions of financial institutions are based on long-run default probability assessments, typically mimicking the behaviour of external rating agency credit grades. Stress testing of capital positions required a forward looking assessment of the possible change in the credit quality of a portfolio of loans. Following the results of chapter 4, corporate ratings exhibit a high degree of serial dependence that can greatly affect the prediction of MtM losses and capital requirements alike. Furthermore, I show that the serial dependence is strongly influenced by the state of the business cycle. If not accounted

properly, both capital requirements and market losses can be severely underestimated, especially in periods of stress. This is particularly evident for MtM losses, where the loss estimate during periods of stress is a combination of rating downgrades and credit spread widening.

Except from the micro view of individual portfolio credit losses, the results presented in chapter 3 have important implications for monetary policy and financial stability. Inflation targeting monetary policy without a complete assessment of the macro-financial linkages could miss the credit and cost channel transmission mechanisms. As I show, monetary policy tightening shocks can lead to temporary inflationary pressures 3-4 quarters post policy changes, and central banks need to adjust their expectations accordingly. Furthermore, the sensitivity of default and recovery rates to monetary policy and balance sheet shocks, indicate that regulatory oversight should be adjusted accordingly when assessing capital requirements of financial intermediaries in low interest rate and high leverage regimes; rise in interest rates and over-leveraging can increase substantially write-offs in corporate and consumer loans, shrinking the available capital to absorb further losses.

5.4 Limitations and Extensions

Chapters 2-4 explore 3 key issues in today's credit risk and macroeconomic modelling. No analysis can cover a research topic from all the possible angles or provide answers to all the relevant questions. Rather, the results provided in this thesis aim at stimulating some discussion around how credit risk relates to the macroeconomic environment and at highlighting the advantages of advanced econometric techniques in credit risk modelling. Therefore, some research questions and approaches have been left out scope and could be addressed as part of future research projects.

The thesis approaches the 3 main research questions from a macro perspective. Even chapter 4, that uses firm specific momentum indicators and duration times, tries to infer aggregate relationships and effects. Due to the macro focus of the analysis, no firm-specific financial data are used. Therefore, an obvious extension of the approach followed in this thesis is to answer the same questions by adding financial variables for the firms in the sample. As Duffie et al. (2007) show, while macroeconomic data are very

useful in explaining part of the time series aspect of defaults, for cross-sectional analysis financial ratios and most importantly the Distance-to-Default are essential. That has implications for the analysis of chapter 2; if differentiating the default probability across bond issuers is the focus then forecasting defaults will need to include firm-specific information in addition to macroeconomic factors. Furthermore, the structural analysis of chapter 3 can greatly benefit from the use of corporate balance sheet information. Macroeconomic shocks have complex transmission mechanisms and granular financial data can help to disentangle the impact on both asset and liability side of the corporate balance sheet. Finally, for the rating momentum analysis of chapter 4 adding balance sheet information could help explain the quantitative element of the serial dependence in credit ratings. Rating momentum is partially based on the change in the balance sheet strength following a rating action and partially on behavioural aspects of rating assignment process. While financial data do not explain the latter, they can help trace the potential contagion following rating changes.

The dynamic factor and structural approaches of chapter 2 and 3 respectively can be unified under a structural dynamic factor model that is based on economic and financial theory to identify the different factors/shocks. The dynamic factor model of chapter 2 is based on a number of 0 restrictions on the matrix of factor loadings. This particular structural of the factor loadings matrix essentially restricts the short-run responses to the dynamic factors for a subset of macroeconomic variables. Chapter 3 uses a number of short- and long-run restrictions to make the shock identification more in line with traditional economic theory views. In principle, since a VAR can be written in a dynamic factor form, these two approaches could be unified in one model. Nevertheless, the presence of long-run restrictions complicates the estimation process for a dynamic factor model, and the existing literature has not provided any estimation scheme to address this class of problems. The MCMC sampling scheme used in chapter 2 provides a good starting point for the development of an efficient estimation algorithm, as it decomposes the different sampling steps into manageable sub-groups.

The use of recovery rates in this thesis is limited to the impact analysis of structural macroeconomic shocks on credit risk presented in chapter 3. The forecasting analysis of chapter 2 could also be extended to cover the effect of the dynamic factors on

recovery rates. Forecasting recoveries can be more complicated as compared to defaults. Recovery rates are only applicable for defaulted debt and therefore the overall sample size reduces significantly in size. Furthermore, recoveries are calculated by discounting a number of post-default cash flows that do not necessarily correspond to a single period. Therefore it is difficult to associate a single realisation of a macroeconomic variable or factor with the overall recovery rate. Chapter 3 simplifies the analysis by using post-default trading prices as a proxy for ultimate recoveries. Further work on structural macroeconomic decomposition or forecasting of credit losses could also consider the realised recoveries, building on the work of Khieu et al. (2012).

Recoveries and defaults need not be treated as independent. Multiple studies show that there is a strong correlation between default clustering and low levels of recoveries, see Altman (2008) and the references therein. As structural credit models imply, both defaults and recovery values depend on the asset value of a firm and the underlying drivers of this asset value process; the difference is the timing, as defaults depend on the pre-default dynamics of the asset value, while recoveries depend on the post-default dynamics. Therefore, a straightforward extension of the analysis presented in chapter 3 is to jointly assess the impact of structural macroeconomic shocks on defaults and recoveries. Creal et al. (2013) introduce a set-up that allows the joint modelling of default and recovery time series, albeit restricting the dimensionality to a small set of macroeconomic variables and a simplified identification scheme. Joint modelling of the credit risk constituents can also be applied to forecasting, therefore extending the analysis of chapter 2. While joint modelling the default and recovery dynamics offers clear benefits when assessing tail losses (the co-dependence between defaults and recoveries becomes very important when focusing on extreme quantiles of the loss distribution), the benefits for average loss forecasting are less clear.

Building on the potential interdependence between defaults and recoveries, credit metrics and macro-financial environment could also be jointly modelled, along the lines of Creal et al. (2013). Joint modelling of defaults, recoveries and macroeconomic conditions can help incorporating feedback loops between the credit and the real sides of the economy. At the same time though, increasing the dimensionality of the system to be estimated, increases substantially the computational burden. This poses constraints on the structural and forecasting specifications that can be feasibly estimated. For

example the semi-structural identification restrictions of chapter 3 are extremely difficult to be implemented in a large scale dynamic and non-Gaussian state space model. Furthermore, lag optimisation for forecasting purposes is very time consuming and only contemporaneous correlations are allowed in the existing literature.

When moving to the continuous time setup of chapter 4, I employ a semi-parametric Cox proportional specification to test the rating momentum effect. While the semi-parametric nature removes the restrictions concerning the baseline intensities in a rating transition matrix, at the same time it limits the dynamic effects that can be incorporated in the model. It is well documented that rating migrations, just like defaults, exhibit clustering in time in addition to the observed business cycle movements. This is usually captured by economy-wide or industry-wide unobserved factors, as in Koopman et al. (2008). Therefore, an extension of the analysis presented in chapter 4 could incorporate unobserved frailty factors by adding a parametric form for the baseline intensities and random effects through calendar time.

Finally, due to data constraints I only use default and recovery information provided by Moody's and I limit the analysis to the US market. Moody's does not rate the entire universe of corporates, especially the small and medium-sized companies that typically do not issue public debt. A complete assessment of corporate default forecasts and the sensitivity of corporate defaults and recoveries to macroeconomic shocks would need to include defaults and bankruptcies from the entire corporate sector. This means that the rated corporate universe would need to also include data from the other major rating agencies and also private loan defaults. While the rated corporate universe can be complemented with data by S&P and Fitch, private debt defaults and recoveries are very difficult to obtain and require proprietary data from financial institutions. Extending the analysis to non-US geographical jurisdictions also required a vast amount of data. Corporate bond defaults are typically very rare and the time series length is typically very short. For public companies, default information can be obtained from cross-referencing databases; see Duan and Van Laere (2012) that describe their approach for building a global default probability model for public companies. For private firms the same limitations as in the US apply.

Technical Appendix

A Markov chain Monte Carlo

A.1 Gibbs Sampling

Gibbs sampling is a stochastic integration frequently used to sample from posterior distribution in Bayesian inference. The method generates random draws from the posterior conditional distributions for each of the parameters in succession, each time conditioning on fixed values of all the other parameters and the observed data. For a set of observed data points \mathbf{Y} and a vector of model specific parameters $\boldsymbol{\vartheta}$, the desired posterior distribution of the parameter vector is denoted by $p(\boldsymbol{\vartheta}|\mathbf{Y})$. Assuming that the full posterior for each parameter is not of closed-form, Gibbs sampling partitions the full parameter set $\boldsymbol{\vartheta}$ into G groups, $\boldsymbol{\vartheta} = \{\boldsymbol{\vartheta}_1, \dots, \boldsymbol{\vartheta}_G\}$, and samples each set in succession.

The algorithm uses full conditional distributions for each parameter group g . Therefore, to sample each $\boldsymbol{\vartheta}_g$, Gibbs sampling treats the rest of the parameter set $\boldsymbol{\vartheta}_{-g}$ and the data \mathbf{Y} as fixed and generates random draws from $p(\boldsymbol{\vartheta}_g|\boldsymbol{\vartheta}_{-g}, \mathbf{Y})$. Starting from initial values $\boldsymbol{\vartheta}^0 = \{\boldsymbol{\vartheta}_1^0, \dots, \boldsymbol{\vartheta}_G^0\}$, the j -th iteration of the algorithm generates the set of random draws $\boldsymbol{\vartheta}_1^j, \dots, \boldsymbol{\vartheta}_G^j$ as follows:

1. Sample $\boldsymbol{\vartheta}_1^j$ from $p(\boldsymbol{\vartheta}_1^j|\boldsymbol{\vartheta}_2^{j-1}, \boldsymbol{\vartheta}_3^{j-1}, \dots, \boldsymbol{\vartheta}_G^{j-1}, \mathbf{Y})$
2. Sample $\boldsymbol{\vartheta}_2^j$ from $p(\boldsymbol{\vartheta}_2^j|\boldsymbol{\vartheta}_1^j, \boldsymbol{\vartheta}_3^{j-1}, \dots, \boldsymbol{\vartheta}_G^{j-1}, \mathbf{Y})$
- ⋮
3. Sample $\boldsymbol{\vartheta}_G^j$ from $p(\boldsymbol{\vartheta}_G^j|\boldsymbol{\vartheta}_1^j, \boldsymbol{\vartheta}_2^j, \boldsymbol{\vartheta}_3^j, \dots, \boldsymbol{\vartheta}_{G-1}^j, \mathbf{Y})$,

assuming the full conditionals for all parameter sets are of known form.

It has been shown that under mild conditions the Gibbs sampling full conditional random draws are guaranteed to converge to posterior marginal draws. More specifically, the random draws from the full conditionals converge in distribution to the true parameter values

$$\{\boldsymbol{\vartheta}_1^j, \dots, \boldsymbol{\vartheta}_G^j\} \xrightarrow{d} \{\boldsymbol{\vartheta}_1, \dots, \boldsymbol{\vartheta}_G\}$$

as $j \rightarrow \infty$. Furthermore, the joint posterior of $\{\boldsymbol{\vartheta}_1^j, \dots, \boldsymbol{\vartheta}_G^j\}$ converges to the true joint posterior distribution of $\{\boldsymbol{\vartheta}_1, \dots, \boldsymbol{\vartheta}_G\}$ at a geometric rate in j . Finally, for any measurable function $f(\boldsymbol{\vartheta}_1, \dots, \boldsymbol{\vartheta}_G)$ whose expectation exists the following applies

$$\lim_{J \rightarrow \infty} \frac{1}{J} \sum_{j=1}^J f(\boldsymbol{\vartheta}_1^j, \dots, \boldsymbol{\vartheta}_G^j) \xrightarrow{a.s.} E[f(\boldsymbol{\vartheta}_1, \dots, \boldsymbol{\vartheta}_G)],$$

where $\xrightarrow{a.s.}$ denotes that the left-hand side converges almost surely to the right-hand side. Therefore, inference based on the Gibbs sampling is guaranteed convergence.

This section is not intended to be a thorough exposition of the Gibbs sampling, but rather to give a high level overview of the algorithm and some basic results. For more details, see the classic textbooks of Zellner (1971) and Robert and Casella (1999).

A.2 Slice Sampling

The Gibbs sampling scheme described in section A.1 is straightforward to implement as long as the full conditionals for each parameter group are of closed form. When this is not the case, other sampling schemes would need to complement Gibbs sampling, such as rejection sampling, Metropolis-Hastings, adaptive rejection sampling, or slice sampling to name a few. For the analysis in chapter 3, I opt for the slice sampling algorithm of Neal (2003) since it is relatively easy to implement, is generic without the need to tailor it to the individual likelihood function to sample from and exhibits very good convergence rates.

Slice sampling exploits the fact that random draws from a distribution $p(x)$ can be obtained by uniformly sampling from the points underneath the curve of such a distribution or, alternatively, points (x, u) that satisfy $0 \leq u \leq p(x)$. For ease of

exposition I only cover the univariate case; for the multivariate case Gibbs sampling steps can be used to sample each variable conditional on the remaining variables. In that case, starting from an initial point x_0 , the algorithm involves 3 main steps:

1. Draw u from $U(0, p(x_0))$, defining the slice $S = \{x : u \leq p(x)\}$.
2. Find an interval $I = (L, R)$ around x_0 that contains all of the slice S . There is a trade-off between using large sampling regions and allowing large moves in the distribution space, and using a simpler sampling region to maximise efficiency. To address this problem local "stepping-out" and "shrinkage" procedures are used:
 - "Stepping-out" procedure: After defining a random interval of width w around x_0 , expand the interval in steps of size w until both ends are outside the slice S .
 - "Shrinkage" procedure: Generate draws from a uniform distribution on I until a point inside the slice is found. Points outside the slice are used to shrink the interval.
3. Draw a new point x_1 from the part of the slice within this interval

B Bayesian Analysis of Multivariate Regression Models

I assume a generic Gaussian linear regression of the form:

$$y_i = \mathbf{x}_i \boldsymbol{\beta} + \varepsilon_i, \quad \varepsilon_i \sim N(0, \sigma_\varepsilon^2), \quad (1)$$

where y_i is an observed data point, \mathbf{x}_i is a vector of K covariates, $\boldsymbol{\beta}$ is a vector of K coefficients and ε_i is the error term, which is normally distributed with 0 mean and σ_ε^2 variance. The regression can be re-written in matrix form:

$$\mathbf{y} = \mathbf{X} \boldsymbol{\beta} + \boldsymbol{\varepsilon}, \quad \boldsymbol{\varepsilon} \sim N(\mathbf{0}, \mathbf{R}), \quad (2)$$

where $\mathbf{y} = [y_1, \dots, y_n]'$ is the stacked vector of the n observations, $\boldsymbol{\varepsilon} = [\varepsilon_1, \dots, \varepsilon_n]'$ is the stacked vector of residuals, $\mathbf{R} = \sigma_\varepsilon^2 I_n$ is the expanded error variance-covariance matrix

and \mathbf{X} is the matrix of k covariates:

$$\mathbf{X} = \begin{bmatrix} \mathbf{x}_{1,1} & \dots & \mathbf{x}_{1,k} \\ \vdots & \ddots & \vdots \\ \mathbf{x}_{n,1} & \dots & \mathbf{x}_{n,k} \end{bmatrix}.$$

The specification in (2) is quite generic and can incorporate panel data where each observation y_i is tracked through time t . Under the Bayesian paradigm, inference is based on the joint posterior of $\boldsymbol{\beta}$ and \mathbf{R} conditional on the observed data, $p(\boldsymbol{\beta}, \mathbf{R}_\varepsilon | \mathbf{y})$. Using Bayes theorem this quantity can be expressed as:

$$p(\boldsymbol{\beta}, \mathbf{R} | \mathbf{y}) = p(\boldsymbol{\beta}, \mathbf{R})p(\mathbf{y} | \boldsymbol{\beta}, \mathbf{R}), \quad (3)$$

where $p(\boldsymbol{\beta}, \mathbf{R})$ is the joint prior distribution for the parameters and $p(\mathbf{y} | \boldsymbol{\beta}, \mathbf{R})$ is the likelihood function. As \mathbf{R} in this case is a diagonal matrix that depends on the single parameter σ_ε^2 , (6) can be re-written in terms of only σ_ε^2 . Nevertheless, the posterior expression (6) is also applicable to cases with a generic variance-covariance matrix \mathbf{R} . The likelihood function $p(\mathbf{y} | \boldsymbol{\beta}, \mathbf{R})$ for the regression model in (2) takes the form:

$$p(\mathbf{y} | \boldsymbol{\beta}, \mathbf{R}) \propto |\mathbf{R}|^{-\frac{n}{2}} e^{-\frac{1}{2}[(\mathbf{y} - \mathbf{X}\boldsymbol{\beta})' \mathbf{R}^{-1}(\mathbf{y} - \mathbf{X}\boldsymbol{\beta})]}. \quad (4)$$

An integral part of Bayesian inference is the choice of priors for the parameters. The conjugate priors for the regression model (2) correspond to:

$$\begin{aligned} p(\boldsymbol{\beta}, \sigma_\varepsilon^2) &= p(\boldsymbol{\beta} | \sigma_\varepsilon^2) p(\sigma_\varepsilon^2), \\ \boldsymbol{\beta} | \sigma_\varepsilon^2 &\sim N(\mathbf{b}_0, \sigma_\varepsilon^2 \mathbf{B}_0), \\ \sigma_\varepsilon^2 &\sim \text{invGamma}(c_0, C_0), \end{aligned} \quad (5)$$

where \mathbf{b}_0 , \mathbf{B}_0 , c_0 and C_0 are hyper-parameters defining the moments of the prior distributions. Under this choice of priors, the posteriors are of closed form and of the same Normal-Inverse Gamma family as the priors. The joint posterior distribution $p(\boldsymbol{\beta}, \sigma_\varepsilon^2 | \mathbf{y})$ is given by:

$$p(\boldsymbol{\beta}, \sigma_\varepsilon^2 | \mathbf{y}) \propto p(\boldsymbol{\beta} | \sigma_\varepsilon^2, \mathbf{y}) p(\sigma_\varepsilon^2 | \mathbf{y}), \quad (6)$$

where $p(\boldsymbol{\beta}|\sigma_\varepsilon^2, \mathbf{y})$ corresponds to a Normal density with moments:

$$\begin{aligned}\boldsymbol{\beta}|\sigma_\varepsilon^2, \mathbf{y} &\sim N(\mathbf{b}_1, \mathbf{B}_1), \\ \mathbf{B}_1 &= (\mathbf{B}_0^{-1} + \mathbf{X}'\mathbf{X})^{-1}, \\ \mathbf{b}_1 &= \mathbf{B}_1(\mathbf{B}_0^{-1}\mathbf{b}_0 + \mathbf{X}'\mathbf{y})\end{aligned}\tag{7}$$

and $p(\sigma_\varepsilon^2|\mathbf{y})$ corresponds to an Inverse-Gamma density with parameters:

$$\begin{aligned}\sigma_\varepsilon^2|\mathbf{y} &\sim \text{InvGamma}(c_1, C_1), \\ c_1 &= c_0 + n/2, \\ C_1 &= C_0 + (\mathbf{y}'\mathbf{y} + \mathbf{b}_0'\mathbf{B}_0^{-1}\mathbf{b}_0 - \mathbf{b}_1'\mathbf{B}_1^{-1}\mathbf{b}_1)/2.\end{aligned}\tag{8}$$

The expression for the regression coefficients posteriors in (7) corresponds to the OLS, adjusted for the prior beliefs for the coefficient values (the mean of the prior \mathbf{b}_0) and the relative strength of the prior prior beliefs (the variance of the prior \mathbf{B}_0).

Unfortunately, this choice of priors implies that the prior for the regression coefficients $\boldsymbol{\beta}$ gets noisier if the data do, which is a somewhat unrealistic feature. Instead, in this thesis I choose to work with two alternative sets of independent priors. The first is a variation of the conjugate prior that assumes the same Normal-Inverse Gamma form for the priors, but removes the dependence of $p(\boldsymbol{\beta})$ on σ_ε^2 :

$$\begin{aligned}p(\boldsymbol{\beta}, \sigma_\varepsilon^2) &= p(\boldsymbol{\beta})p(\sigma_\varepsilon^2), \\ \boldsymbol{\beta}|\sigma_\varepsilon^2 &\sim N(\mathbf{b}_0, \mathbf{B}_0), \\ \sigma_\varepsilon^2 &\sim \text{invGamma}(c_0, C_0).\end{aligned}\tag{9}$$

Under the Independent Normal-Inverse Gamma prior, posterior marginals for $\boldsymbol{\beta}$ and σ_ε^2 are not of closed form. To approximate the joint density Gibbs sampling can be used. For Gibbs sampling, only the full conditionals need to be of closed form. It can be proven that the full conditionals for this choice of priors is of the Normal-Inverse Gamma form. For the regression coefficients $\boldsymbol{\beta}$, the full conditionals take the form:

$$\begin{aligned}\boldsymbol{\beta}|\mathbf{R}, \mathbf{y} &\sim N(\mathbf{b}_1, \mathbf{B}_1), \\ \mathbf{B}_1 &= (\mathbf{B}_0^{-1} + \mathbf{X}'\mathbf{R}^{-1}\mathbf{X})^{-1}, \\ \mathbf{b}_1 &= \mathbf{B}_1(\mathbf{B}_0^{-1}\mathbf{b}_0 + \mathbf{X}'\mathbf{R}^{-1}\mathbf{y})\end{aligned}\tag{10}$$

where I have substituted the expanded matrix \mathbf{R} instead of σ_ε^2 . The full conditional mean \mathbf{b}_1 in (10) is equal to the generalised least squares estimator $(\mathbf{X}'\mathbf{R}^{-1}\mathbf{X})^{-1}\mathbf{X}'\mathbf{R}^{-1}\mathbf{y}$, adjusted for the prior mean \mathbf{b}_0 according to its strength \mathbf{B}_0 . The posterior moments

in (10) can be re-expressed in the following prediction-correction filter form:

$$\begin{aligned}
\mathbf{b}_1 &= \mathbf{b}_0 + \mathbf{K}_1(\mathbf{y} - \mathbf{X}\mathbf{b}_0), \\
\mathbf{B}_1 &= (\mathbf{I} - \mathbf{K}_1\mathbf{X})\mathbf{B}_0, \\
\mathbf{K}_1 &= \mathbf{B}_0\mathbf{X}'\mathbf{C}^{-1}, \\
\mathbf{C} &= \mathbf{X}\mathbf{B}_0\mathbf{X}' + \mathbf{R}.
\end{aligned}
\tag{11}$$

This alternative, but numerically equivalent, form expresses the posterior mean \mathbf{b}_1 as a correction to the prior mean \mathbf{b}_0 . This correction is based on the residuals $\mathbf{y} - \mathbf{X}\mathbf{b}_0$, which is the prediction error when using the prior mean \mathbf{b}_0 as an estimator of $\boldsymbol{\beta}$. This representation is helpful for the derivation of the Kalman filter in section C.2. For the same set of priors, the full conditional for σ_ε^2 is given by:

$$\begin{aligned}
\sigma_\varepsilon^2 | \boldsymbol{\beta}, \mathbf{Y} &\sim \text{invGamma}(c_1, C_1) \\
c_1 &= c_0 + n/2, \\
C_1 &= C_0 + (\mathbf{y} - \mathbf{X}\boldsymbol{\beta})'(\mathbf{y} - \mathbf{X}\boldsymbol{\beta})/2
\end{aligned}
\tag{12}$$

The second choice of priors assumes no prior knowledge on either $\boldsymbol{\beta}$ or σ_ε^2 . This corresponds to the so-called Jeffreys (or diffuse) prior:

$$\begin{aligned}
p(\boldsymbol{\beta}, \sigma_\varepsilon^2) &= p(\boldsymbol{\beta} | \sigma_\varepsilon^2) p(\sigma_\varepsilon^2), \\
\boldsymbol{\beta} | \sigma_\varepsilon^2 &\propto 1, \\
\sigma_\varepsilon^2 &\propto 1/\sigma_\varepsilon^2.
\end{aligned}
\tag{13}$$

For the diffuse choice of priors, the posteriors (just like the conjugate priors) are of closed form. The posterior for $\boldsymbol{\beta}$ is again a normal density with moments:

$$\begin{aligned}
\boldsymbol{\beta} | \sigma_\varepsilon^2, \mathbf{y} &\sim N(\mathbf{b}_1, \mathbf{B}_1), \\
\mathbf{B}_1 &= (\mathbf{X}'\mathbf{X})^{-1}, \\
\mathbf{b}_1 &= \mathbf{B}_1\mathbf{X}'\mathbf{y}
\end{aligned}
\tag{14}$$

which correspond to the normal OLS quantities. The posterior of σ_ε^2 is again an Inverse-Gamma density with moments:

$$\begin{aligned}
\sigma_\varepsilon^2 | \mathbf{y} &\sim \text{InvGamma}(c_1, C_1), \\
c_1 &= n/2, \\
C_1 &= (\mathbf{y} - \mathbf{X}\boldsymbol{\beta})'(\mathbf{y} - \mathbf{X}\boldsymbol{\beta})/2.
\end{aligned}
\tag{15}$$

C State Space Models

C.1 State Space Models and Bayesian Inference

State space models consist of an equation describing the dependence of the observed quantities on the system's state variables (*measurement equation*), and a second equation describing the stochastic evolution of the state variables (*state equation*). The measurement equation for a Gaussian state space model takes the form:

$$\mathbf{y}_t = \mathbf{Z}\mathbf{F}_t + \boldsymbol{\epsilon}_t, \quad \boldsymbol{\epsilon}_t \sim N(0, \boldsymbol{\Sigma}_\epsilon), \quad (16)$$

for a vector of observed time series \mathbf{y}_t and a vector of state variables \mathbf{F}_t . \mathbf{Z} is the matrix of the sensitivities of the observed variables to the system's states and for simplicity it is assumed time-invariant; in more generic state space models, the matrix is allowed to vary with t . The state equation dictating the stochastic evolution of \mathbf{F}_t in the same Gaussian state space model can be expressed as:

$$\mathbf{F}_t = \boldsymbol{\Phi}\mathbf{F}_{t-1} + \boldsymbol{\eta}_t, \quad \boldsymbol{\eta}_t \sim N(0, \boldsymbol{\Sigma}_\eta), \quad (17)$$

where $\boldsymbol{\Phi}$ is the transition matrix defining how future values of \mathbf{F}_t depend on the process's past values.

Bayesian inference in state space models is very appealing, since the parameter space provides a natural decomposition and efficient algorithms with Gibbs sampling steps can be employed to integrate the joint posterior. More specifically, Bayesian inference is based on deriving the joint posterior density

$$p(\mathbf{F}, \mathbf{Z}, \boldsymbol{\Phi}, \boldsymbol{\Sigma}_\epsilon, \boldsymbol{\Sigma}_\eta | \mathbf{y}) \propto p(\mathbf{y} | \mathbf{F}, \mathbf{Z}, \boldsymbol{\Phi}, \boldsymbol{\Sigma}_\epsilon, \boldsymbol{\Sigma}_\eta) p(\mathbf{F} | \boldsymbol{\Phi}, \boldsymbol{\Sigma}_\eta) p(\mathbf{Z}, \boldsymbol{\Phi}, \boldsymbol{\Sigma}_\epsilon, \boldsymbol{\Sigma}_\eta), \quad (18)$$

where $\mathbf{y} = [\mathbf{y}_1, \dots, \mathbf{y}_T]'$ and $\mathbf{F} = [\mathbf{F}_1, \dots, \mathbf{F}_T]'$ are the stacked observation and state vectors respectively. Due to the recursive nature of a state space model (18) can be re-written as:

$$p(\mathbf{F}, \boldsymbol{\vartheta} | \mathbf{y}) \propto p(\mathbf{F}_0 | \boldsymbol{\vartheta}) p(\boldsymbol{\vartheta}) \prod_{t=1}^T p(\mathbf{y}_t | \mathbf{F}_t, \boldsymbol{\vartheta}) p(\mathbf{F}_t | \mathbf{F}_{t-1}, \boldsymbol{\vartheta}), \quad (19)$$

where for brevity $\boldsymbol{\vartheta} = \{\mathbf{Z}, \boldsymbol{\Phi}, \boldsymbol{\Sigma}_\epsilon, \boldsymbol{\Sigma}_\eta\}$. The densities $p(\mathbf{y}_t|\mathbf{F}_t, \boldsymbol{\vartheta})$ and $p(\mathbf{F}_t|\mathbf{F}_{t-1}, \boldsymbol{\vartheta})$ are defined in (16) and (17) respectively. $p(\mathbf{F}_0|\boldsymbol{\vartheta})$ is the initial value for the state vector, while $p(\boldsymbol{\vartheta})$ is the prior density for all model parameters.

The posterior density (19) is not of closed form and therefore MCMC techniques can be used to evaluate the multidimensional integration. A typical Gibbs sampling scheme involves 2 main steps:

- Sample the states conditionally on the data and the model parameters: This is translated to sampling \mathbf{F} from $p(\mathbf{F}|\mathbf{y}, \boldsymbol{\vartheta})$. Even though this step can be implemented by sampling \mathbf{F}_t sequentially at each t using a set of second layers Gibbs sampling steps, it is more efficient to obtain random draws for the entire sequence of T states in one step using some form of simulation smoothing. I follow Durbin and Koopman (2002) and proceed in 5 steps:
 1. Obtain the smoothed estimates for the states $\hat{\mathbf{F}}_{t|T}$ by a forward pass of the Kalman filter and a backwards pass of the Kalman smoother. Kalman filter and smoother are described in section C.2.
 2. Randomly generate T observation and state innovations. The random innovations at time t are denoted by $\boldsymbol{\epsilon}_t^*$ and $\boldsymbol{\eta}_t^*$.
 3. Create random observations and states by plugging $\boldsymbol{\epsilon}_t^*$ and $\boldsymbol{\eta}_t^*$ into (16) and (17) respectively. The pseudo-observations and states at time t are denoted by \mathbf{y}_t^* and \mathbf{F}_t^* respectively.
 4. Obtain smoothed states $\hat{\mathbf{F}}_{t|T}^*$ by applying the Kalman filter and smoother on the pseudo-observations \mathbf{y}_t^* .
 5. Obtain the posterior draws as $\tilde{\mathbf{F}}_{t|T} = \hat{\mathbf{F}}_{t|T} - \hat{\mathbf{F}}_{t|T}^* + \mathbf{F}_t^*$.
- Sample the model parameters conditionally on the data and the states: This is translated to sampling all the elements of $\boldsymbol{\vartheta}$ from the complete-data posterior density $p(\boldsymbol{\vartheta}|\mathbf{F}, \mathbf{y})$. It is clear that conditional on the states being known, both (16) and (17) can be treated as linear regressions and therefore all the elements of $\boldsymbol{\vartheta}$ can be sampled using the results of section B.

C.2 The Kalman Filter and Smoother

The Kalman filter is a recursive algorithm that operates on the series of measurements observed over time in (16) and produces a statistically optimal estimate of the underlying system state in (17). The algorithm works in a two-step process. In the prediction step, the Kalman filter produces estimates of the current state variables, along with the corresponding estimation uncertainty. Once the next measurement is observed (with some noise), these estimates are updated using a weighted average: estimates with higher certainty (lower variance) get a higher weight.

The equations forming the Kalman filter can be derived in multiple ways, see for example Harvey (1990). Here, I base the derivation on the Bayesian information filter as it provides a more intuitive way to understand the mechanics of the algorithm. Furthermore, it links naturally to the Bayesian estimation procedure used to estimate the parameters in (16)-(17). To understand the filtering logic, I make use of the results in section B. To derive the prediction and update steps of the Kalman filter, in addition to the formula provided in section B, the explicit form of the marginal likelihood $p(\mathbf{y}|\mathbf{R})$ for the regression model (2) is also needed. Using Bayes' theorem, the marginal likelihood is obtained by evaluating:

$$p(\mathbf{y}|\mathbf{R}) = \frac{p(\mathbf{y}|\boldsymbol{\beta}, \mathbf{R})p(\boldsymbol{\beta})}{p(\boldsymbol{\beta}|\mathbf{R}, \mathbf{y})}. \quad (20)$$

Evaluating the marginal likelihood at $\boldsymbol{\beta} = \mathbf{b}_0$, from the joint normality of data and prior results in:

$$\mathbf{y}_t|\mathbf{R} \sim N(\mathbf{X}\mathbf{b}_0, \mathbf{C}), \quad (21)$$

where $\mathbf{C} = \mathbf{X}\mathbf{B}_0\mathbf{X}' + \mathbf{R}$ is defined in (11).

Turning now to the state space model setup in (16) and (17), the Kalman filter provides the estimates for \mathbf{F}_t based on the posterior density $p(\mathbf{F}_t|\mathbf{y}_t)$ (density of the state vector \mathbf{F} at time t , given the information available up to and including t). Using Bayes' theorem this posterior density can be obtained as:

$$p(\mathbf{F}_t|\mathbf{y}_t) = \frac{p(\mathbf{y}_t|\mathbf{F}_t)p(\mathbf{F}_t|\mathbf{y}_{t-1})}{p(\mathbf{y}_t|\mathbf{y}_{t-1})}, \quad (22)$$

where $p(\mathbf{y}_t|\mathbf{F}_t)$ is the likelihood of \mathbf{y}_t and $p(\mathbf{y}_t|\mathbf{y}_{t-1})$ is the marginal likelihood after

integrating out the state vector \mathbf{F}_t . (22) essentially states that the posterior density of \mathbf{F}_t is a combination of the observed model's likelihood, $p(\mathbf{y}_t|\mathbf{F}_t)$, when the state vector follows the prior $p(\mathbf{F}_t|\mathbf{y}_t)$. When the state space model is Gaussian (as it is the case here), all the above densities are normal, with moments given by the Kalman filter: the prediction step provides the prior $p(\mathbf{F}_t|\mathbf{y}_{t-1})$ for the new estimate of \mathbf{F}_t , while the update step combines that prior with the regression model (16) for the observations \mathbf{y}_t to derive the best linear estimate for \mathbf{F}_t .

The prediction-correction form (11) and the marginal likelihood function in (21) serve as the basis for the derivation of the Kalman filter. To estimate the densities in (22), the Kalman filter iterates the following steps assuming that at the beginning of each iteration the moments of the filtered density $\mathbf{F}_{t-1}|\mathbf{y}_{t-1} \sim N(\hat{\mathbf{F}}_{t-1|t-1}, \mathbf{P}_{t-1|t-1})$ are given by:

1. *State Prediction Step*: This step determines the state prediction density $p(\mathbf{F}_t|\mathbf{y}_{t-1})$ in (22). Using Bayes' theorem this density can be written as

$$p(\mathbf{F}_t|\mathbf{y}_{t-1}) = \frac{p(\mathbf{F}_t|\mathbf{F}_{t-1})p(\mathbf{F}_{t-1}|\mathbf{y}_{t-1})}{p(\mathbf{F}_{t-1}|\mathbf{F}_t, \mathbf{y}_{t-1})},$$

which is exactly the marginal likelihood in (20) for a linear regression with \mathbf{F}_t as the dependent variable, and \mathbf{F}_{t-1} as the unknown regression parameter, having $p(\mathbf{F}_{t-1}|\mathbf{y}_{t-1})$ as its prior. Based on the normal marginal likelihood form in (21) and the state equation (17), the moments of the normal density $p(\mathbf{F}_t|\mathbf{y}_{t-1})$ are:

$$\begin{aligned} \mathbf{F}_t|\mathbf{y}_{t-1} &\sim N(\hat{\mathbf{F}}_{t|t-1}, \mathbf{P}_{t|t-1}), \\ \hat{\mathbf{F}}_{t|t-1} &= \mathbf{\Phi}\hat{\mathbf{F}}_{t-1|t-1}, \\ \mathbf{P}_{t|t-1} &= \mathbf{\Phi}\mathbf{P}_{t-1|t-1}\mathbf{\Phi}' + \Sigma_\eta. \end{aligned} \tag{23}$$

2. *State Correction Step*: This step corrects the state prediction from step I. to incorporate the information from a new data point \mathbf{y}_t .

- (a) *Observation Prediction Step*: The corresponding predictive density for the observations $p(\mathbf{y}_t|\mathbf{y}_{t-1})$ is obtained based on the predicted state value from step I. Using Bayes' theorem this density can be written as:

$$p(\mathbf{y}_t|\mathbf{y}_{t-1}) = \frac{p(\mathbf{y}_t|\mathbf{F}_t)p(\mathbf{F}_t|\mathbf{y}_{t-1})}{p(\mathbf{F}_t|\mathbf{y}_t, \mathbf{y}_{t-1})},$$

which is exactly the marginal likelihood in (20) for a linear regression of with \mathbf{y}_t as the dependent variable, and \mathbf{F}_t as the unknown regression parameter, having $p(\mathbf{F}_t|\mathbf{y}_{t-1})$ as its prior. Based on the normal marginal likelihood form in (21) and the measurement equation (16), the moments of the normal density $p(\mathbf{y}_t|\mathbf{y}_{t-1})$ are:

$$\begin{aligned}\mathbf{y}_t|\mathbf{y}_{t-1} &\sim N(\hat{\mathbf{y}}_{t|t-1}, \mathbf{C}_{t|t-1}), \\ \hat{\mathbf{y}}_{t|t-1} &= \mathbf{Z}\hat{\mathbf{F}}_{t|t-1}, \\ \mathbf{C}_{t|t-1} &= \mathbf{Z}\mathbf{P}_{t|t-1}\mathbf{Z}' + \Sigma_\epsilon.\end{aligned}\tag{24}$$

(b) *State Update Step*: The filter density $p(\mathbf{F}_t|\mathbf{y}_t)$ is based on the state and observation predictions, $p(\mathbf{F}_t|\mathbf{y}_{t-1})$ and $p(\mathbf{y}_t|\mathbf{y}_{t-1})$ respectively. From (22) this density is equivalent to the posterior of the regression parameter $\boldsymbol{\beta}$ in (11), when $\boldsymbol{\beta} = \mathbf{F}_t|\mathbf{y}_t$ with $p(\mathbf{F}_t|\mathbf{y}_{t-1})$ as the prior. Therefore, using the prediction-correction form for the regression coefficient in (11):

$$\begin{aligned}\mathbf{F}_t|\mathbf{y}_t &\sim N(\hat{\mathbf{F}}_{t|t}, \mathbf{P}_{t|t}), \\ \hat{\mathbf{F}}_{t|t} &= \boldsymbol{\Phi}\hat{\mathbf{F}}_{t|t-1} + \mathbf{K}_t(\mathbf{y}_t - \hat{\mathbf{y}}_{t|t-1}), \\ \mathbf{P}_{t|t} &= (1 - \mathbf{K}_t\mathbf{Z})\mathbf{P}_{t|t-1}, \\ \mathbf{K}_t &= \mathbf{P}_{t|t-1}\mathbf{Z}'\mathbf{C}_{t|t-1}^{-1},\end{aligned}\tag{25}$$

which concludes the Kalman filter iteration for time t . For each new t , the filter uses the density $\mathbf{F}_{t-1}|\mathbf{y}_{t-1} \sim N(\hat{\mathbf{F}}_{t-1|t-1}, \mathbf{P}_{t-1|t-1})$ from the previous iteration and goes back to step I.

Grouping the Kalman filter recursions in a concise form:

$$\begin{aligned}\text{(Prior) Predicted state estimate:} & \quad \hat{\mathbf{F}}_{t|t-1} = \boldsymbol{\Phi}\hat{\mathbf{F}}_{t-1|t-1} \\ \text{(Prior) Predicted state variance:} & \quad \mathbf{P}_{t|t-1} = \boldsymbol{\Phi}\mathbf{P}_{t-1|t-1}\boldsymbol{\Phi}' + \Sigma_\eta \\ \text{Observation Predicted estimate:} & \quad \hat{\mathbf{y}}_{t|t-1} = \mathbf{Z}\hat{\mathbf{F}}_{t|t-1} \\ \text{Observation Predicted variance:} & \quad \mathbf{C}_{t|t-1} = \mathbf{Z}\mathbf{P}_{t|t-1}\mathbf{Z}' + \Sigma_\epsilon \\ \text{Optimal Kalman gain:} & \quad \mathbf{K}_t = \mathbf{P}_{t|t-1}\mathbf{Z}'\mathbf{C}_{t|t-1}^{-1} \\ \text{(Posterior) Updated state estimate:} & \quad \hat{\mathbf{F}}_{t|t} = \boldsymbol{\Phi}\hat{\mathbf{F}}_{t|t-1} + \mathbf{K}_t(\mathbf{y}_t - \hat{\mathbf{y}}_{t|t-1}) \\ \text{(Posterior) Updated state variance:} & \quad \mathbf{P}_{t|t} = (\mathbf{I} - \mathbf{K}_t\mathbf{Z})\mathbf{P}_{t|t-1}\end{aligned}\tag{26}$$

To start the Kalman filter, one has to choose the initial normal prior $\mathbf{F}_0|\mathbf{y}_0 \sim N(\hat{\mathbf{F}}_{0|0}, \mathbf{P}_{0|0})$. Since for this chapter I assume that the state vector \mathbf{F}_t is stationary (implying that Φ does not have roots on unit circle), it can be shown that the optimal moments for the initial distribution are the unconditional mean and variance of the \mathbf{F}_t process. The unconditional mean is by construction equal to $\mathbf{0}$. The ergodic variance of \mathbf{F}_t is the solution to:

$$\Phi \mathbf{P}_{0|0} \Phi' + \Sigma_\eta = \mathbf{P}_{0|0}, \quad (27)$$

where I assume that \mathbf{F}_t follows a VAR(1) process. (27) can be generalised to VARs of higher order if the VAR is re-written as a VAR(1) process with Φ being the equivalent VAR(1) matrix. Solving (27) for $\mathbf{P}_{0|0}$ and using the $\mathbf{0}$ mean assumption leads to the initial distribution for the filter, given by:

$$\mathbf{F}_0|\mathbf{y}_0 \sim N(\mathbf{0}, \text{devec} [(I_{K^2} - \Phi \otimes \Phi)^{-1} \cdot \text{vec}(\Sigma_\eta)]_{K,K}), \quad (28)$$

where \otimes denotes the Kronecker product, $\text{vec}(\Sigma_\eta)$ represents the vectorised version of Σ_η , and $\text{devec}[\cdot]_{K,K}$ represents the de-vectorised version of the expression in brackets, transformed into a K -by- K matrix. For a complete description of the filter initialisation problem, for both stationary and non-stationary cases, see Durbin and Koopman (2001).

The Kalman filter delivers best estimates of the state vector \mathbf{F}_t at each t conditional on all available information at time t . For the simulation smoother algorithm described in section C.1, random draws conditional on the entire sample T are needed. The backwards passing Kalman smoother is designed to provide optimal estimates for $\mathbf{F}_{t|T}$. The algorithm can be written in many forms, but here I follow Durbin and Koopman (2001) and defines the smoother recursions as follows:

$$\begin{aligned} \text{Smoothed estimate of state:} \quad & \hat{\mathbf{F}}_{t|T} = \hat{\mathbf{F}}_{t|t-1} + \mathbf{P}_{t|t-1} \mathbf{r}_t \\ & \mathbf{r}_t = \mathbf{Z}' \mathbf{C}_{t|t-1}^{-1} (\mathbf{y}_t - \hat{\mathbf{y}}_{t|t-1}) + \mathbf{L}_t' \mathbf{r}_{t+1} \\ & \mathbf{L}_t = \Phi - \Phi \mathbf{K}_t \mathbf{Z} \quad (29) \\ \text{Smoothed state variance:} \quad & \mathbf{P}_{t|T} = \mathbf{P}_{t|t-1} - \mathbf{P}_{t|t-1} \mathbf{N}_t \mathbf{P}_{t|t-1}' \\ & \mathbf{N}_t = \mathbf{Z}' \mathbf{C}_{t|t-1}^{-1} \mathbf{Z} + \mathbf{L}_t' \mathbf{N}_{t+1} \mathbf{L}_t \end{aligned}$$

where the recursion runs from time $T - 1$ backwards and $\mathbf{r}_t/\mathbf{N}_t$ are initialised at time T with $\mathbf{r}_T = \mathbf{0}$ and $\mathbf{N}_T = \mathbf{0}$. At time T the smoothed estimates are the Kalman filtered estimates.

References

- Aalen, O., Ø. Borgan, H. Gjessing, and S. Gjessing (2008). *Survival and Event History Analysis: A Process Point of View*. Springer.
- Acharya, V. V., S. T. Bharath, and A. Srinivasan (2007). “Does industry-wide distress affect defaulted firms? Evidence from creditor recoveries”. In: *Journal of Financial Economics* 85.3, pp. 787–821.
- Agarwal, R. and M. Gort (2002). “Firm and Product Life Cycles and Firm Survival”. In: *The American Economic Review* 92.2, pp. 184–190.
- Altman, E. (2008). “Default recovery rates and LGD in credit risk modeling and practice: an updated review of the literature and empirical evidence”. In: *Advances in Credit Risk Modeling and Corporate Bankruptcy Prediction*. Ed. by S. Jones. Cambridge University Press, pp. 175–206.
- Altman, E. I. (1968). “Financial Ratios, Discriminant Analysis and the Prediction of Corporate Bankruptcy”. In: *The Journal of Finance* 23.4, pp. 589–609.
- Altman, E. I. and E. A. Kalotay (2014). “Ultimate recovery mixtures”. In: *Journal of Banking & Finance* 40.0, pp. 116–129.
- Altman, E. I. and D. L. Kao (1992). “The Implications of Corporate Bond Ratings Drift”. In: *Financial Analysts Journal* 48.3, pp. 64–75.
- Altman, E. I. and H. A. Rijken (2004). “How rating agencies achieve rating stability”. In: *Journal of Banking & Finance* 28.11, pp. 2679–2714.
- Altman, E. I., B. Brady, A. Resti, and A. Sironi (2005). “The Link between Default and Recovery Rates: Theory, Empirical Evidence, and Implications”. In: *The Journal of Business* 78.6, pp. 2203–2228.
- Amato, J. D. and C. H. Furfine (2004). “Are credit ratings procyclical?” In: *Journal of Banking & Finance* 28.11, pp. 2641–2677.

- Andersen, P. K., Ø. Borgan, R. D. Gill, and N. Keiding (1993). *Statistical Models Based on Counting Processes*. Springer.
- Avramov, D., T. Chordia, G. Jostova, and A. Philipov (2007). “Momentum and Credit Rating”. In: *The Journal of Finance* 62.5, pp. 2503–2520.
- Bai, J. and S. Ng (2002). “Determining the Number of Factors in Approximate Factor Models”. In: *Econometrica* 70.1, pp. 191–221.
- (2013). “Principal components estimation and identification of static factors”. In: *Journal of Econometrics* 176.1, pp. 18–29.
- Bangia, A., F. X. Diebold, A. Kronimus, C. Schagen, and T. Schuermann (2002). “Ratings migration and the business cycle, with application to credit portfolio stress testing”. In: *Journal of Banking & Finance* 26.2–3, pp. 445–474.
- Bannier, C. E. and C. W. Hirsch (2010). “The economic function of credit rating agencies – What does the watchlist tell us?” In: *Journal of Banking & Finance* 34.12, pp. 3037–3049.
- Barth, M. J. and V. A. Ramey (2002). “The Cost Channel of Monetary Transmission”. In: *NBER Macroeconomics Annual 2001, Volume 16*, pp. 199–256.
- Bernanke, B. and M. Gertler (1989). “Agency Costs, Net Worth, and Business Fluctuations”. In: *The American Economic Review* 79.1, pp. 14–31.
- Bernanke, B., M. Gertler, and G. Simon (1996). “The Financial Accelerator and the Flight to Quality”. In: *The Review of Economics and Statistics* 78.1, pp. 1–15.
- Bernanke, B. S. and M. Gertler (1995). “Inside the Black Box: The Credit Channel of Monetary Policy Transmission”. In: *Journal of Economic Perspectives* 9.4, pp. 27–48.
- Bernanke, B. S. and I. Mihov (1998). “The liquidity effect and long-run neutrality”. In: *Carnegie-Rochester Conference Series on Public Policy* 49.0, pp. 149–194.
- Bernanke, B. S., J. Boivin, and P. Eliasziw (2005). “Measuring the Effects of Monetary Policy: A Factor-Augmented Vector Autoregressive (FAVAR) Approach”. In: *The Quarterly Journal of Economics* 120.1, pp. 387–422.
- Bjørnland, H. C. and K. Leitemo (2009). “Identifying the interdependence between US monetary policy and the stock market”. In: *Journal of Monetary Economics* 56.2, pp. 275–282.
- Black, F. and J. C. Cox (1976). “Valuing Corporate Securities: Some Effects of Bond Indenture Provisions”. In: *The Journal of Finance* 31.2, pp. 351–367.

- Blanchard, O. J. and D. Quah (1989). “The Dynamic Effects of Aggregate Demand and Supply Disturbances”. In: *The American Economic Review* 79.4, pp. 655–673.
- Boivin, J., M. P. Giannoni, and D. Stevanovic (2013). “Dynamic Effects of Credit Shocks in a Data-Rich Environment”. In: *Federal Reserve Bank of New York Staff Reports*, no. 615.
- Bräuning, F. and S. J. Koopman (2014). “Forecasting macroeconomic variables using collapsed dynamic factor analysis”. In: *International Journal of Forecasting* 30.3, pp. 572–584.
- Breslow, N. (1974). “Covariance Analysis of Censored Survival Data”. In: *Biometrics* 30.1, pp. 89–99.
- Campbell, J. Y. and G. B. Taksler (2003). “Equity Volatility and Corporate Bond Yields”. In: *The Journal of Finance* 58.6, pp. 2321–2350.
- Campbell, J. Y., J. Hilscher, and J. A. N. Szilagyi (2008). “In Search of Distress Risk”. In: *The Journal of Finance* 63.6, pp. 2899–2939.
- Cantor, R. (2001). “Moody’s investors service response to the consultative paper issued by the Basel Committee on Bank Supervision “A new capital adequacy framework””. In: *Journal of Banking & Finance* 25.1, pp. 171–185.
- Carter, C. K. and R. Kohn (1994). “On Gibbs Sampling for State Space Models”. In: *Biometrika* 81.3, pp. 541–553.
- Chen, H. U. I. (2010). “Macroeconomic Conditions and the Puzzles of Credit Spreads and Capital Structure”. In: *The Journal of Finance* 65.6, pp. 2171–2212.
- Chen, L., P. Collin-Dufresne, and R. S. Goldstein (2009). “On the Relation Between the Credit Spread Puzzle and the Equity Premium Puzzle”. In: *Review of Financial Studies* 22.9, pp. 3367–3409.
- Chowdhury, I., M. Hoffmann, and A. Schabert (2006). “Inflation dynamics and the cost channel of monetary transmission”. In: *European Economic Review* 50.4, pp. 995–1016.
- Christensen, J. H. E., E. Hansen, and D. Lando (2004). “Confidence sets for continuous-time rating transition probabilities”. In: *Journal of Banking & Finance* 28.11, pp. 2575–2602.
- Collin-Dufresne, P. and R. S. Goldstein (2001). “Do Credit Spreads Reflect Stationary Leverage Ratios?” In: *The Journal of Finance* 56.5, pp. 1929–1957.

- Collin-Dufresne, P., R. S. Goldstein, and J. S. Martin (2001). “The Determinants of Credit Spread Changes”. In: *The Journal of Finance* 56.6, pp. 2177–2207.
- Cox, D. R. (1972). “Regression Models and Life-Tables”. In: *Journal of the Royal Statistical Society. Series B (Methodological)* 34.2, pp. 187–220.
- (1975). “Partial Likelihood”. In: *Biometrika* 62.2, pp. 269–276.
- Creal, D., B. Schwaab, S. J. Koopman, and A. Lucas (2013). “Observation Driven Mixed-Measurement Dynamic Factor Models with an Application to Credit Risk”. In: *Review of Economics and Statistics*, forthcoming.
- CreditMetrics (2007). *Technical Document*. Report. RiskMetrics Group.
- Das, S. R., D. Duffie, N. Kapadia, and L. Saita (2007). “Common Failings: How Corporate Defaults Are Correlated”. In: *The Journal of Finance* 62.1, pp. 93–117.
- Diebold, F. X. and C. Li (2006). “Forecasting the term structure of government bond yields”. In: *Journal of Econometrics* 130.2, pp. 337–364.
- Diebold, F. X., G. D. Rudebusch, and S. Boragan Aruoba (2006). “The macroeconomy and the yield curve: a dynamic latent factor approach”. In: *Journal of Econometrics* 131.1–2, pp. 309–338.
- Duan, J.-C. and E. Van Laere (2012). “A public good approach to credit ratings – From concept to reality”. In: *Journal of Banking & Finance* 36.12, pp. 3239–3247.
- Duan, J.-C., J. Sun, and T. Wang (2012). “Multiperiod corporate default prediction—A forward intensity approach”. In: *Journal of Econometrics* 170.1, pp. 191–209.
- Duffie, D. and D. Lando (2001). “Term Structures of Credit Spreads with Incomplete Accounting Information”. In: *Econometrica* 69.3, pp. 633–664.
- Duffie, D., L. Saita, and K. Wang (2007). “Multi-period corporate default prediction with stochastic covariates”. In: *Journal of Financial Economics* 83.3, pp. 635–665.
- Duffie, D., A. Eckner, G. Horel, and L. Saita (2009). “Frailty Correlated Default”. In: *The Journal of Finance* 64.5, pp. 2089–2123.
- Durbin, J. and S. J. Koopman (2002). “A Simple and Efficient Simulation Smoother for State Space Time Series Analysis”. In: *Biometrika* 89.3, pp. 603–615.
- Durbin, J. and S. Koopman (2001). *Time Series Analysis by State Space Methods*. Clarendon Press.
- Efron, B. (1977). “The Efficiency of Cox’s Likelihood Function for Censored Data”. In: *Journal of the American Statistical Association* 72.359, pp. 557–565.

- Eichenbaum, M. (1992). “Interpreting the macroeconomic time series facts: The effects of monetary policy’: by Christopher Sims”. In: *European Economic Review* 36.5, pp. 1001–1011.
- Embrechts, P., F. Lindskog, and A. Mcneil (2003). “Chapter 8 - Modelling dependence with copulas and applications to risk management”. In: *Handbook of Heavy Tailed Distributions in Finance*. Ed. by S. T. Rachev. North-Holland, pp. 329–384.
- Engle, R. and M. Watson (1981). “A One-Factor Multivariate Time Series Model of Metropolitan Wage Rates”. In: *Journal of the American Statistical Association* 76.376, pp. 774–781.
- Fei, F., A.-M. Fuertes, and E. Kalotychou (2012). “Credit Rating Migration Risk and Business Cycles”. In: *Journal of Business Finance & Accounting* 39.1-2, pp. 229–263.
- Figlewski, S., H. Frydman, and W. Liang (2012). “Modeling the effect of macroeconomic factors on corporate default and credit rating transitions”. In: *International Review of Economics & Finance* 21.1, pp. 87–105.
- Flannery, M. J. and K. P. Rangan (2006). “Partial adjustment toward target capital structures”. In: *Journal of Financial Economics* 79.3, pp. 469–506.
- Flannery, M. J., S. Nikolova, and Ö. Öztekin (2012). “Leverage Expectations and Bond Credit Spreads”. In: *Journal of Financial and Quantitative Analysis* 47.04. (Stas), pp. 689–714.
- Fleming, T. and D. Harrington (2011). *Counting Processes and Survival Analysis*. Wiley.
- Frühwirth-Schnatter, S. and R. Frühwirth (2007). “Auxiliary mixture sampling with applications to logistic models”. In: *Computational Statistics & Data Analysis* 51.7, pp. 3509–3528.
- Frühwirth-Schnatter, S. and H. Wagner (2006). “Auxiliary mixture sampling for parameter-driven models of time series of counts with applications to state space modelling”. In: *Biometrika* 93.4, pp. 827–841.
- (2008). “Marginal likelihoods for non-Gaussian models using auxiliary mixture sampling”. In: *Computational Statistics & Data Analysis* 52.10, pp. 4608–4624.
- Frühwirth-Schnatter, S., R. Frühwirth, L. Held, and H. Rue (2009). “Improved auxiliary mixture sampling for hierarchical models of non-Gaussian data”. In: *Statistics and Computing* 19.4, pp. 479–492.

- Frye, J. (2000a). “Collateral Damage”. In: *Risk magazine* 13.4, pp. 91–94.
- (2000b). “Depressing Recoveries”. In: *Risk Magazine* 13.11, pp. 106–111.
- Geanakoplos, J. (2010). “The Leverage Cycle”. In: *NBER Macroeconomics Annual 2009, Volume 24*. Vol. 24. National Bureau of Economic Research, University of Chicago Press, pp. 1–65.
- Giesecke, K. (2004). “Correlated default with incomplete information”. In: *Journal of Banking & Finance* 28.7, pp. 1521–1545.
- Gilchrist, S. and E. Zakrajšek (2012). “Credit Spreads and Business Cycle Fluctuations”. In: *American Economic Review* 102.4, pp. 1692–1720.
- Gilchrist, S., V. Yankov, and E. Zakrajšek (2009). “Credit market shocks and economic fluctuations: Evidence from corporate bond and stock markets”. In: *Journal of Monetary Economics* 56.4, pp. 471–493.
- Güttler, A. and P. Raupach (2010). “The Impact of Downward Rating Momentum”. In: *Journal of Financial Services Research* 37.1, pp. 1–23.
- Güttler, A. and M. Wahrenburg (2007). “The adjustment of credit ratings in advance of defaults”. In: *Journal of Banking & Finance* 31.3, pp. 751–767.
- Hamilton, D. D. and R. Cantor (2004). *Rating Transitions and Defaults Conditional on Watchlist, Outlook and Rating History*. Special Comment. Moody’s Investors Service.
- Harvey, A. (1990). *Forecasting, Structural Time Series Models and the Kalman Filter*. Cambridge University Press.
- Hertz-Picciotto, I. and B. Rockhill (1997). “Validity and Efficiency of Approximation Methods for Tied Survival Times in Cox Regression”. In: *Biometrics* 53.3, pp. 1151–1156.
- Hu, Y. and W. Perraudin (2002). “The dependence of recovery rates and defaults”. In: *Mimeo, BirkBeck College*.
- Hull, J., M. Predescu, and A. White (2004). “The relationship between credit default swap spreads, bond yields, and credit rating announcements”. In: *Journal of Banking & Finance* 28.11, pp. 2789–2811.
- Jacobson, T., J. Lindé, and K. Roszbach (2006). “Internal ratings systems, implied credit risk and the consistency of banks’ risk classification policies”. In: *Journal of Banking & Finance* 30.7, pp. 1899–1926.

- Jankowitsch, R., F. Nagler, and M. G. Subrahmanyam (2014). “The determinants of recovery rates in the US corporate bond market”. In: *Journal of Financial Economics* 114.1, pp. 155–177.
- Jarrow, R. (2001). “Default Parameter Estimation Using Market Prices”. In: *Financial Analysts Journal* 57.5, pp. 75–92.
- Jarrow, R. A. and S. M. Turnbull (1995). “Pricing Derivatives on Financial Securities Subject to Credit Risk”. In: *The Journal of Finance* 50.1, pp. 53–85.
- Jorion, P. and G. Zhang (2007). “Good and bad credit contagion: Evidence from credit default swaps”. In: *Journal of Financial Economics* 84.3, pp. 860–883.
- (2009). “Credit Contagion from Counterparty Risk”. In: *The Journal of Finance* 64.5, pp. 2053–2087.
- Jungbacker, B. and S. J. Koopman (2014). “Likelihood-based dynamic factor analysis for measurement and forecasting”. In: *The Econometrics Journal*, n/a–n/a.
- Kalbfleisch, J. D. and R. L. Prentice (1973). “Marginal likelihoods based on Cox’s regression and life model”. In: *Biometrika* 60.2, pp. 267–278.
- Kalbfleisch, J. and R. Prentice (2002). *The Statistical Analysis of Failure Time Data*. Wiley Series in Probability and Statistics. Wiley.
- Kavvathas, D. (2000). *Estimating Credit Rating Transition Probabilities for Corporate Bonds*. Working Paper. University of Chicago.
- Khieu, H. D., D. J. Mullineaux, and H.-C. Yi (2012). “The determinants of bank loan recovery rates”. In: *Journal of Banking & Finance* 36.4, pp. 923–933.
- Kiyotaki, N. and J. Moore (1997). “Credit Cycles”. In: *Journal of Political Economy* 105.2, pp. 211–248.
- Koopman, S. J. and A. Lucas (2008). “A Non-Gaussian Panel Time Series Model for Estimating and Decomposing Default Risk”. In: *Journal of Business & Economic Statistics* 26.4, pp. 510–525.
- Koopman, S. J., A. Lucas, and A. Monteiro (2008). “The multi-state latent factor intensity model for credit rating transitions”. In: *Journal of Econometrics* 142.1, pp. 399–424.
- Koopman, S. J., R. Kräussl, A. Lucas, and A. B. Monteiro (2009). “Credit cycles and macro fundamentals”. In: *Journal of Empirical Finance* 16.1, pp. 42–54.

- Koopman, S. J., A. Lucas, and B. Schwaab (2011). “Modeling frailty-correlated defaults using many macroeconomic covariates”. In: *Journal of Econometrics* 162.2, pp. 312–325.
- (2012). “Dynamic Factor Models With Macro, Frailty, and Industry Effects for U.S. Default Counts: The Credit Crisis of 2008”. In: *Journal of Business & Economic Statistics* 30.4, pp. 521–532.
- Lando, D. and M. S. Nielsen (2010). “Correlation in corporate defaults: Contagion or conditional independence?” In: *Journal of Financial Intermediation* 19.3, pp. 355–372.
- Lando, D. and T. M. Skødeberg (2002). “Analyzing rating transitions and rating drift with continuous observations”. In: *Journal of Banking & Finance* 26.2–3, pp. 423–444.
- Löffler, G. (2004). “An anatomy of rating through the cycle”. In: *Journal of Banking & Finance* 28.3, pp. 695–720.
- (2013). “Can rating agencies look through the cycle?” In: *Review of Quantitative Finance and Accounting* 40.4, pp. 623–646.
- Longstaff, F. A. and E. S. Schwartz (1995). “A Simple Approach to Valuing Risky Fixed and Floating Rate Debt”. In: *The Journal of Finance* 50.3, pp. 789–819.
- McNeil, A. J. and J. P. Wendin (2006). “Dependent credit migrations”. In: *Journal of Credit Risk* 2.3, pp. 87–114.
- (2007). “Bayesian inference for generalized linear mixed models of portfolio credit risk”. In: *Journal of Empirical Finance* 14.2, pp. 131–149.
- Merton, R. C. (1974). “On the Pricing of Corporate Debt: The Risk Structure of Interest Rates”. In: *The Journal of Finance* 29.2, pp. 449–470.
- Metz, A., S. Sorensen, D. Keisman, and D. Chiu (2012). “Trading Prices as Predictors of Ultimate Corporate Recovery Rates”. In: *Moody’s Investor Service, Special Comment*.
- Neal, R. M. (2003). “Slice sampling”. In: pp. 705–767.
- Nickell, P., W. Perraudin, and S. Varotto (2000). “Stability of rating transitions”. In: *Journal of Banking & Finance* 24.1–2, pp. 203–227.
- Pesaran, M. H., T. Schuermann, B.-J. Treutler, and S. M. Weiner (2006). “Macroeconomic Dynamics and Credit Risk: A Global Perspective”. In: *Journal of Money, Credit, and Banking* 38.5. Volume 38, Number 5, August 2006, pp. 1211–1261.

- Pykhtin, M. (2003). “Unexpected Recovery Risk”. In: *Risk Magazine* 16.8, pp. 74–78.
- Ravenna, F. and C. E. Walsh (2006). “Optimal monetary policy with the cost channel”. In: *Journal of Monetary Economics* 53.2, pp. 199–216.
- Robert, C. and G. Casella (1999). *Monte Carlo Statistical Methods*. Springer New York.
- Rosenberg, J. V. and T. Schuermann (2006). “A general approach to integrated risk management with skewed, fat-tailed risks”. In: *Journal of Financial Economics* 79.3, pp. 569–614.
- Royston, P. and D. G. Altman (1994). “Regression Using Fractional Polynomials of Continuous Covariates: Parsimonious Parametric Modelling”. In: *Journal of the Royal Statistical Society. Series C (Applied Statistics)* 43.3, pp. 429–467.
- Rubio-Ramírez, J. F., D. F. Waggoner, and T. Zha (2010). “Structural Vector Autoregressions: Theory of Identification and Algorithms for Inference”. In: *The Review of Economic Studies* 77.2, pp. 665–696.
- Sargent, T. J. (1989). “Two Models of Measurements and the Investment Accelerator”. In: *Journal of Political Economy* 97.2, pp. 251–287.
- Shephard, N. (1994). “Partial non-Gaussian state space”. In: *Biometrika* 81.1, pp. 115–131.
- Sims, C. A. (1992). “Interpreting the macroeconomic time series facts: The effects of monetary policy”. In: *European Economic Review* 36.5, pp. 975–1000.
- Stefanescu, C., R. Tunaru, and S. Turnbull (2009). “The credit rating process and estimation of transition probabilities: A Bayesian approach”. In: *Journal of Empirical Finance* 16.2, pp. 216–234.
- Stock, J. and M. Watson (2010). “Dynamic Factor Models”. In: *Clements MP, Henry DF Oxford Handbook of Economic Forecasting*. Oxford University Press.
- Stock, J. and M. W. Watson (1991). “A Probability Model of the Coincident Economic Indicators”. In: *The Leading Economic Indicators: New Approaches and Forecasting Records*. Ed. by G. Moore and K. Lahiri. Cambridge University Press, pp. 63–90.
- Stock, J. H. and M. W. Watson (1989). “New Indexes of Coincident and Leading Economic Indicators”. In: *NBER Macroeconomics Annual 1989, Volume 4*. Ed. by O. J. Blanchard and S. Fischer. National Bureau of Economic Research, Inc, pp. 351–409.
- (2002a). “Forecasting Using Principal Components from a Large Number of Predictors”. In: *Journal of the American Statistical Association* 97.460, pp. 1167–1179.

- Stock, J. H. and M. W. Watson (2002b). “Macroeconomic Forecasting Using Diffusion Indexes”. In: *Journal of Business & Economic Statistics* 20.2, pp. 147–162.
- (2005). *Implications of Dynamic Factor Models for VAR Analysis*. Working Paper 11467. National Bureau of Economic Research.
- Varotto, S. (2012). “Stress testing credit risk: The Great Depression scenario”. In: *Journal of Banking & Finance* 36.12, pp. 3133–3149.
- Vasicek, O. (2002). “Loan portfolio value”. In: *Risk Magazine*, pp. 160–162.
- Wand, M. and C. Jones (1995). *Kernel Smoothing*. Monographs on Statistics & Applied Probability. Chapman & Hall.
- Zellner, A. (1971). *An introduction to Bayesian inference in econometrics*. J. Wiley.