UNIVERSITY OF AMSTERDAM

UvA-DARE (Digital Academic Repository)

Essays on macro-financial risks

Sperna Weiland, R.C.

Link to publication

Citation for published version (APA):
Sperna Weiland, R. C. (2019). Essays on macro-financial risks. Amsterdam: Tinbergen Institute.

General rights
It is not permitted to download or to forward/distribute the text or part of it without the consent of the author(s) and/or copyright holder(s), other than for strictly personal, individual use, unless the work is under an open content license (like Creative Commons).

Disclaimer/Complaints regulations
If you believe that digital publication of certain material infringes any of your rights or (privacy) interests, please let the Library know, stating your reasons. In case of a legitimate complaint, the Library will make the material inaccessible and/or remove it from the website. Please Ask the Library: https://uba.uva.nl/en/contact, or a letter to: Library of the University of Amsterdam, Secretariat, Singel 425, 1012 WP Amsterdam, The Netherlands. You will be contacted as soon as possible.

UvA-DARE is a service provided by the library of the University of Amsterdam (http://dare.uva.nl)

Download date: 09 Oct 2019

Robertus C. Sperna Weiland

Essays on Macro-Financial Risks

This dissertation consists of three main chapters that explore the pricing implications of several 'macro-financial' risks in different contexts. More specifically, the first chapter investigates the dynamic interplay between credit and liquidity risk in the US corporate bond market. The second chapter considers the modeling and pricing of sovereign credit risk. The third chapter examines the relationship between labor and financial markets and studies the impact of aggregate labor income risk on stock returns.

Robertus C. (Rob) Sperna Weiland (1990) holds a M.Sc. degree in Stochastics and Financial Mathematics (2014, Cum Laude) from VU Amsterdam. He received the first prize in the "EY TopQuants Quantitative Finance Thesis Awards" and the second prize in the "CFA Society European Quant Awards" competitions for his M.Sc. thesis. In 2014, Rob started his Ph.D. research at the Finance Department of the University of Amsterdam and worked under the supervision of Prof. dr. R.J.A. Laeven, Prof. dr. F.C.J.M. De Jong, and Prof. dr. P.J.C. Spreij. During his Ph.D., Rob obtained a research qualification in Econometrics from the Tinbergen Institute, was Junior Fellow at the Amsterdam Center for Excellence in Risk and Macro Finance, and was a visiting researcher at the Bendheim Center for Finance at Princeton University.

Essays on Macro-Financial Risks Robertus C. Sperna Weiland

744

Universiteit van Amsterdam

ROBERTUS C. SPERNA WEILAND

ESSAYS ON MACRO-FINANCIAL RISKS

ISBN 978 90 361 0566 8

© R.C. Sperna Weiland: *Essays on Macro-Financial Risks,* Academisch proefschrift, July 2019

Cover design: Crasborn Graphic Designers bno, Valkenburg a.d. Geul

This book is no. 744 of the Tinbergen Institute Research Series, established through cooperation between Rozenberg Publishers and the Tinbergen Institute. A list of books which already appeared in the series can be found in the back.

ESSAYS ON MACRO-FINANCIAL RISKS

ACADEMISCH PROEFSCHRIFT

ter verkrijging van de graad van doctor
aan de Universiteit van Amsterdam
op gezag van de Rector Magnificus
prof. dr. ir. K.I.J. Maex
ten overstaan van een door het College voor Promoties ingestelde
commissie, in het openbaar te verdedigen in de Agnietenkapel
op vrijdag 27 september 2019, te 10:00 uur
door

Robertus Cornelis Sperna Weiland
geboren te Leusden

PROMOTIECOMMISSIE

Promotores:	
prof. dr. R.J.A. Laeven	Universiteit van Amsterdam
prof. dr. F.C.J.M. de Jong	Tilburg University
Copromotor:	
prof. dr. P.J.C. Spreij	Universiteit van Amsterdam
Overige leden:	
prof. dr. F.R. Kleibergen	Universiteit van Amsterdam
prof. dr. H.P. Boswijk	Universiteit van Amsterdam
prof. dr. J.J.A.G. Driessen	Tilburg University
prof. dr. P. Feldhütter	Copenhagen Business School
dr. E. Eiling	Universiteit van Amsterdam

Faculteit Economie en Bedrijfskunde

Non nobis solum nati sumus

— Marcus Tullius Cicero

ACKNOWLEDGEMENTS

For a long time I have been trying to picture this moment: My dissertation is handed in and under review by the committee, and the only thing left to do is writing these acknowledgements. A perfect moment to reflect on everything that has happened on this journey of pursuing a PhD. Some highlights that immediately come to mind are my time at the Tinbergen Institute, the many PhD student outings, the conferences I attended, and my research visit at Princeton University. However, much more than these specific events, I realize that all these experiences would not have been meaningful (or even possible) without the help, advice, and friendship of many persons. In this section, perhaps the most important section of this whole dissertation, I therefore would like (trying) to thank all of you that supported me during these last few years.

First, I would like to thank my supervisors Roger Laeven, Frank de Jong, and Peter Spreij for their guidance and mentorship from which I benefitted a lot the past few years. It goes without saying that without your numerous comments, discussions, and general involvement this dissertation would not be what it is today. However, apart from all your intellectual input, I want to thank you foremost for the persons you are and for creating such a pleasant working environment in which I experienced a lot of freedom, trust, and support.

Roger, you are truly someone to look up to: smart, ambitious, extremely knowledgeable, but next to that a real family man and a people's person. You somehow seem to have found a certain balance in your life which I find very admirable and I hope to approach myself one day. I want to thank you for all the opportunities, support, and trust you have given me the past few years. It is safe to say that without you I would probably not have done a PhD.

Frank, it is hard to imagine a more involved and caring supervisor and I was (and am) very happy and honored you wanted to continue supervising me even after you formally left the UvA. I benefitted a lot from your deep knowledge of many subjects, but also from your direct, pragmatic and caring personality. It says enough that on multiple occasions

people told me how lucky I was to have you as a supervisor. I cannot but concur.

Peter, apart from your undisputed and vast knowledge on (financial) mathematical topics, I respect you a lot for your unique character. You are open-minded, easy to approach, and you judge people for who they are instead of looking at their status or background. If only more people would do that. I hold good memories of the Winter schools you organized, the drinks we had, and the many personal stories you shared. Your 'bike-stealing' story is already a true classic.

Next, I would like to thank all the members of the defense committee for their time and effort of going through my dissertation. Joost Driessen and Frank Kleibergen, many thanks for your valuable comments and useful discussions on earlier versions of some of the chapters. Peter Boswijk, thanks for the inspiring courses in which you taught me most of the econometrics that I know. Peter Feldhütter, I am really honored you accepted the invitation of being part of my defense committee. Some of your work has been very inspirational for me. And lastly, thanks to Esther Eiling, not only for being part of the committee, but also for being one of my (very appreciated) co-authors, your open and warm personality, and for your role as head of the Finance department this last year.

This brings me to thanking all of my colleagues at the Finance department and the Actuarial Sciences and Mathematical Finance department. I had a great time being part of these two sections and I look back on many interesting seminars and research days and many fun social activities. In particular, I would like to mention and thank Rafael Perez Ribas, Philippe Versijp, Esther Eiling, Aleksandar Andonov, Evgenia Zhivotova, and Aerdt Houben for our collaborations in teaching the courses Fixed Income Risk Management, Investment and Portfolio Theory 2, and Banking Risk Management. I would like to thank Jeroen Ligterink for his role as head of the Finance department during the first years I was at UvA. I always felt that I could knock on your door for anything, which is much appreciated. I would also like to thank Michel Vellekoop for a similar role in the Actuarial Sciences and Mathematical Finance department and for his role in organizing the annual Winter school in Lunteren to which I am always very much looking forward. Thanks to Servaas van Bilsen for the many coffee breaks, lunches, and interesting conversations we had over the past few years, I will miss those for sure. Patrick Tuijp, thanks as well for all our good conversations. Furthermore,

thanks to Enrico Perotti for making your house available for some of the department outings and for the very nice dinner you once invited me to. I would also very much like to thank everyone of the administrative staff for making it possible to do our jobs. So thanks to Lorena Zevedei, Deyanira Gonzalez Alvarez, Welmoed van Hoogen, Sidonie Rademaker, Mariëlle Lameijer, and Bas Bouten for all your support and quick replies to the many important (and less important) questions I had for you these years.

One of the highlights during my PhD was my research visit at the Bendheim Center for Finance at Princeton University in the Fall semester of 2017. I met many smart, kind, and inspiring people there and it would be impossible to thank everyone here who made my stay there so great. But at least I would like to thank Yacine Aït-Sahalia for supervising me during this visit. I really learned a lot from all our conversations and the course you were teaching during that time. I would also like to thank my two PhD office mates and friends, Chris "Christoffer" Wolf and "Fresh" Tyler Abbot. I do miss all the office banter with you guys and you really made my stay a blast. Lastly, I would like to thank Nyla and Jawad for being the best landlords and friends I could wish for. You really made your place my home far away from home.

The past few years would not have been so enjoyable without the support and friendship of many of my PhD colleagues and friends. Thanks to all the current and former inhabitants of office M404, in particular, Adam, Annemijn, Dorinth, Felipe, Fleur, Guilherme, Ieva R., Ieva S., Joris, Magda D., Magda R., Mark, Merve, Oscar, Pascal, Richard, Robin, Ruobing, Sin Man, Yumei (and everyone else I did not explicitly mention here) for making it the best office room in the whole UvA. Dorinth, thanks for being my office-bro, eilandgenoot, 50% of #dreamteam, collaborator on the "Working 50 project", travel companion, and for becoming one of my best friends these past few years. Next to my M404 colleagues, I would also like to thank all the other ABS PhD students for the many fun PhD events. Special thanks to Eloisa and Hannah (a.k.a. team Phoenix) for (losing) the many foosball games, for helping me conquering my fear of rollercoasters, and for all the other epic epicness. Also thanks to my (former) PhD colleagues from the Actuarial Sciences and Mathematical Finance department, Andrei, Frank, Jan, Jitze, Merrick, Robert, Servaas, Xiye, Yuan, and Zhenzhen, for making me feel like a full member of the ASMF group even though I did not see you all on a daily basis. Jan, thanks for sharing your latex template (used

to typeset this thesis) with me. Thanks as well to Dieter, Huyen, Jamie, Magda, Timo and all my other TI friends and classmates for many unforgettable moments. Thanks to Bas (and Renee) and Nick for the many laughs, your support, the awesome trip to Bali, and your true friendship. Same goes for my other JvO friends (Rob V., Mart, Steven, Kasper, Dennis, Paul, Jort). Martine, thanks for all our walks and talks. Erik and Misha, thanks for the occasional drinks, dinners, and banter. Erik and Guusje, I am very much looking forward to working with you again.

Finally, thanks to my family[1] for all your love and support. Koen, extraordinary gentleman, I could not have wished for a better housemate and I really value us living together (and everything around it) these past few years. Ruud, you always give the wisest advices and more than once you showed me a different perspective of things. Thanks for all the good talks, laughs, and exploring the world together. Christa, the best sister in the world of who I am extremely proud. I know you are always very involved with everyone's well-being and really try to make time for everyone despite you being extremely busy, know that it is much appreciated. Dr. Niek, you already gave the right example of how to become a doctor, just as you have been an example for me in many things throughout my life. Thanks for paving the way for me. Oma, de lunches op zaterdag zijn ook voor mij een hoogtepunt van de week. Mom, it is nice to have a soulmate in this world, and even nicer if this soulmate happens to be your mother. Dad, apart from being the best father in the world, you will soon also become my best colleague at my new job, how awesome is that? Sin Man, your resilience, even in the toughest of times, is amazing and I respect you deeply for that. Thanks for your love, support, chocolate boosts, GEKness, and your overall positive attitude towards life.

Finishing this PhD means finishing a big chapter in my life. And even though I am very much looking forward to taking up new challenges and getting to know new colleagues in my new job, I would like to end with the following quote that perfectly summarizes my feeling at the end of this moment of reflection:

"How lucky I am to have something that makes saying goodbye so hard."
— A.A. Milne, The Complete Tales of Winnie-the-Pooh

Rob Sperna Weiland, June 2019

1 Including all aunts, uncles, cousins, boyfriends, girlfriends, and Dirkxen.

CONTENTS

List of Tables xv
List of Figures xvii
1 INTRODUCTION 1
1.1 On part I: Credit risk 1
1.2 On part II: Labor income risk 6

I CREDIT RISK 9
2 FEEDBACK BETWEEN CREDIT AND LIQUIDITY RISK IN THE US CORPORATE BOND MARKET 11
2.1 Introduction 11
2.2 The model 17
2.2.1 Dynamic credit-liquidity interactions 17
2.2.2 Bond pricing 19
2.3 Data 21
2.3.1 Bond transaction data and bond-specific information 21
2.3.2 Construction of bid-ask spreads and measurement equations 23
2.4 Estimation methodology 26
2.4.1 Joint posterior density 26
2.4.2 Joint conditional posterior density and MCMC sampling 29
2.4.3 MCMC algorithm 30
2.4.4 Parameter assumptions, practical considerations and finite sample behavior 32
2.5 Estimation results 33
2.5.1 Ford Motor Company 33
2.5.2 Rating portfolios 38
2.5.3 Risk premia 42
2.6 Model implications 43
2.6.1 Credit-liquidity decompositions 43
2.6.2 Implications for Value-at-Risk 50
2.7 Conclusions 54
2.A Construction of the liquidity jump indicator 55
2.B Monte Carlo simulation study 57

2.C Closed-form pricing formulas 59
2.D Price corrections for the rating portfolios 64
2.E Details of the Metropolis-Hastings steps 66
2.E.1 Metropolis steps for drawing the states 66
2.E.2 Metropolis steps for drawing the parameters 71
2.F Bayes factors 72
3 CREDIT RISK PREMIA IN SOVEREIGN CREDIT DEFAULT SWAPS 75
3.1 Introduction 75
3.2 Data 79
3.3 The model 83
3.3.1 CDS pricing 88
3.4 Estimation methodology 89
3.4.1 Estimation risk-neutral intensities $\lambda^{Q}_{j,t}$ 89
3.4.2 Estimation default event risk premium parameter 97
3.5 CDS spread decomposition 101
3.6 Conclusions 109
3.A Closed-form pricing formulas 109
3.B MCMC details 112
3.B.1 Metropolis steps for drawing the states 115
3.B.2 Metropolis steps for drawing the parameters 117
3.B.3 Estimation country-specific components 118
II LABOR INCOME RISK 121
4 LABOR INCOME RISK AND STOCK RETURNS: THE ROLE OF HORIZON EFFECTS 123
4.1 Introduction 123
4.2 Theoretical framework 127
4.3 Empirical methodology and data 132
4.3.1 Hedging demand due to labor income risk across horizons 132
4.3.2 Cross-sectional regression methodology 134
4.3.3 Data and summary statistics 140
4.4 Empirical results 144
4.4.1 Horizon-specific hedging demands due to labor income risk 144
4.4.2 Asset pricing tests across different horizons 151
4.4.3 Discussion of the results 163
4.5 Robustness checks 164
4.5.1 Other test assets 164

4.5.2 Maximum scale specification 170
4.5.3 Multivariate versus univariate betas 172
4.5.4 Nominal versus real labor income growth 172
4.6 Conclusions 175
4.A Model derivations 176
4.A.1 Log-linearization of the wealth dynamics 177
4.A.2 Optimal portfolio 178
4.A.3 Equilibrium pricing 180
4.B Cross-sectional standard errors 182
4.B.1 Multivariate betas 182
4.B.2 Univariate betas 184

BIBLIOGRAPHY 187
Summary 197
Samenvatting (summary in Dutch) 201
List of co-authors 205

LIST OF TABLES

Table 2.1 Descriptive statistics: transaction data 23
Table 2.2 Descriptive statistics: relative bid-ask spreads 25
Table 2.3 Parameter estimates 33
Table 2.4 Average yield spread decomposition 47
Table 2.5 VaR capital requirements per dollar invested 54
Table 2.6 Monte Carlo results univariate model 58
Table 2.7 Monte Carlo results multivariate model 59
Table 3.1 Summary statistics sovereign CDS spreads 81
Table 3.2 Historical sovereign cumulative average default rates 83
Table 3.3 Parameter estimates risk-neutral common and country-specific intensity processes 92
Table 3.4 Estimates of the default event risk premium parameter μ 99
Table 3.5 Decomposition of five-year CDS spreads into expected risk and risk premia components 104
Table 3.6 Decomposition of CDS spreads into expected risk and risk premia components across maturities 105
Table 3.7 Decomposition of five-year CDS spreads into country-specific and systemic risk components 107
Table 3.8 Decomposition of five-year CDS spreads into country-specific and systemic expected risk and risk premia components 108
Table 4.1 Mapping between time-scales and time spans 137
Table 4.2 Summary statistics of quarterly labor income growth rates 141
Table 4.3 Labor income risk induced portfolio hedging weights for different horizons 147
Table 4.4 "Keeping-up-with-the-Joneses" induced portfolio hedging weights for different horizons 151
Table 4.5 Exposures of 25 size-BM portfolios to labor income risk across different horizons 153
Table 4.6 Cross-sectional regressions for 25 size-BM portfolios (multivariate betas) 158

Table 4.7 Cross-sectional regressions for 25 size-BM portfolios - comparison with alternative asset pricing models 161
Table 4.8 Cross-sectional regressions for 50 combined 25 size-BM and 25 size-investment portfolios (multivariate betas 166
Table 4.9 Cross-sectional regressions for 55 combined 25 size-BM and 30 industry portfolios (multivariate betas) 167
Table 4.10 Cross-sectional regressions for 50 size-BM and size-investment, and 55 size-BM and industry portfolios - comparison with alternative asset pricing models 168
Table 4.11 Cross-sectional regressions for 25 size-BM portfolios - different maximum scales J 171
Table 4.12 Cross-sectional regressions for 25 size-BM portfolios (univariate betas) 173
Table 4.13 Cross-sectional regressions for 25 size-BM portfolios - real aggregate labor income growth (multivariate betas) 174

LIST OF FIGURES

Figure 2.1 Representative Ford bond price path 24
Figure 2.2 Price decay credit and liquidity shocks Ford 37
Figure 2.3 Estimated credit intensities and credit jump probabilities Ford 38
Figure 2.4 Estimated liquidity intensities and liquidity jump probabilities Ford 39
Figure 2.5 Liquidity identification 40
Figure 2.6 Credit identification 41
Figure 2.7 Estimated prices representative Ford bond 42
Figure 2.8 Decomposition estimated intensities Ford 46
Figure 2.9 Decomposition of 10y Ford bond yield spread 48
Figure 2.10 Decomposition of 10y B and lower bond yield spread 49
Figure 2.11 Absolute term structure per component for Ford 51
Figure 2.12 Absolute term structure per component for B and lower-rated bonds 52
Figure 2.13 Simulated 10-day bond return distributions B and lower-rated bonds 53
Figure 3.1 Five-year sovereign CDS spreads 80
Figure 3.2 Estimated common factor intensities and jump times 95
Figure 3.3 Model-fit and estimated country-specific factor intensities and jump times 96
Figure 3.4 Historical and model-implied conditional default probabilities per rating class 100
Figure 4.1 Aggregate labor income growth scales at different frequencies 143
Figure 4.2 Coefficients of determination of portfolio hedging weight regressions across different horizons 146

1 INTRODUCTION

One of the central themes in finance is to explain how financial risks affect asset prices and returns. A good understanding of the pricing implications of risks is important for several reasons. First, prices might contain information that could be useful for risk analyses and being able to distill such information could be beneficial for risk management purposes. Second, it helps in understanding why there are cross-sectional differences in average returns across assets: Investors are typically risk averse and therefore not only price in their risk expectations, but also demand compensation for being exposed to unpredictable deviations of underlying risk factors from these expectations. In principle, these so-called risk premia are larger for riskier assets and, hence, cross-sectional differences in risk exposures can explain why some assets tend to earn higher average returns than others.

In this dissertation, the pricing implications of several 'macro-financial' risks are analyzed in different contexts. In particular, Part I, consisting of Chapters 2 and 3, centers around issues related to credit risk. More specifically, Chapter 2 investigates the dynamic interplay between credit and liquidity risk at the corporate bond market level, whereas Chapter 3 focuses on the modeling and pricing of credit risk at the sovereign level. Part II, consisting of Chapter 4, examines the relationship between labor and financial markets and studies the impact of aggregate labor income risk on stock returns. This first chapter serves as a brief introduction to the subsequent chapters and discusses their motivation and position in the literature.

1.1 ON PART I: CREDIT RISK

Credit risk refers to the risk associated with any credit-linked events such as rating changes, variations of credit spreads, and ultimately default events. A default is said to occur if a counterparty in a financial contract does not fulfil its financial obligations as specified by the contract in a timely manner. Default risk can thus be thought of as a combination of three components, namely the value of the contract that is

exposed to default risk, the fraction of this value that would be recovered in the case of a default event, and the probability that the default event occurs. Credit risk models typically address all of these components by establishing a link between a statistical model describing default risk and a formal economic pricing model. Such a link is important since it facilitates the construction of prices from estimated default event distributions or, conversely, the estimation of market-implied default probabilities from observed market prices.

In general, two streams of modeling approaches can be identified in the credit risk literature. The class of structural models considers credit risk at the level of the firm and views debt and equity as contingent claims on the value of the firm's assets. In these models, a default event occurs if the value of the firm's assets falls below a certain threshold that is either specified exogenously or endogenously. Structural models focus on modeling the firm's value and capital structure and are thus able to link default events to the economic fundamentals of the firm (see, e.g., Merton, 1974; Black and Cox, 1976; Leland, 1994; Longstaff and Schwartz, 1995; Leland and Toft, 1996). The class of reduced-form models, and in particular intensity-based models, on the other hand, treats default events as jumps in an underlying stochastic default process and does not link default occurrences to economic fundamentals (see, e.g., Jarrow and Turnbull, 1995; Lando, 1998; Duffee, 1999; Duffie and Singleton, 1999). The advantages of these models over structural models are, however, that they accommodate more flexible pricing formulas, more formal estimation procedures, and typically fit the data much better.

Historically, the credit risk literature has mainly been concerned with the pricing of defaultable corporate bonds. Originally, this literature focused solely on credit risk to explain the existence of yield spreads over benchmark Treasury or swap rates. However, many empirical studies have pointed out that credit risk alone cannot account for the observed credit spreads (Huang and Huang, 2012), and argue that liquidity, which refers to the ease with which an asset can be traded, is one of the main candidates in explaining this so-called "credit spread puzzle".[1] Longstaff et al. (2005), for example, show that the non-default component in corporate bond spreads is substantial and that it is strongly

1 Other components that have been proposed to (partly) explain credit spreads are tax premia and systemic risk premia (Elton et al., 2001; Huang and Huang, 2012), jump risk premia (Cremers et al., 2008) and equity volatility effects (Campbell and Taksler, 2003).

related to measures of bond liquidity. Many other empirical studies confirm that liquidity is being priced into bonds (see, e.g., Houweling et al., 2005; Chen et al., 2007; Lin et al., 2011; De Jong and Driessen, 2012), and that the price impact of liquidity is significantly more severe during times of financial crisis and for speculative grade bonds (see Edward et al., 2007; Nashikkar et al., 2011; Bao et al., 2011; Friewald et al., 2012; Dick-Nielsen et al., 2012).

The corporate bond literature mainly addresses the interplay between credit and liquidity risk in a 'static' way by, for example, comparing liquidity measures across rating classes or by comparing tranquil and distressed periods. Several structural credit risk studies argue, however, that credit-liquidity interactions are 'dynamic' and that these interactions can have a substantial impact on yield spreads. For example, in the theoretical models of He and Milbradt (2014) and Chen et al. (2017) a dynamic default-liquidity feedback loop arises through a debt rollover channel (see also He and Xiong, 2012; Nagler, 2019): By assuming a relatively illiquid post-default secondary bond market, higher default risk worsens a bond's pre-default secondary market liquidity because investors anticipate the higher likelihood of higher holding costs in the case of default. This, in turn, amplifies equity holders' rollover losses and increases the endogenous probability of default. In empirical applications, they explicitly quantify the effects of credit-liquidity interactions on yield spreads and find that these interaction terms account for approximately 25% of total credit spread and for 16% to 46% of the changes in spreads over time. Similar to most of the structural credit risk literature, however, these empirical results are based on model calibrations that seem to have difficulties with capturing the data, especially in crisis periods. A more formal empirical analysis of credit-liquidity interactions is therefore outside the scope of their research and is the topic of Chapter 2.

In particular, Chapter 2 proposes a new parsimonious reduced-form framework that allows for a dynamic feedback loop between credit and liquidity risk as alluded to in the structural literature. The novelty of the modeling approach is that credit and liquidity intensities are mod-

eled as bivariate self- and cross-exciting processes.[2] [3] The self-excitation in the credit and liquidity intensities capture the empirically observed tendency of credit events to cluster in time (Aït-Sahalia et al., 2014) and the persistence of liquidity (see, e.g., Acharya and Pedersen, 2005). At the same time, the cross-excitation between the credit and liquidity intensities gives rise to an explicit dynamic feedback mechanism in which credit and liquidity shocks interact in a possibly asymmetric fashion and tend to cluster in time. Compared to other reduced-form models, and in accordance with the structural models of He and Milbradt (2014) and Chen et al. (2017), the model accommodates a more detailed decomposition of yield spreads in which the price impact of credit risk, liquidity risk, and their interactions terms can explicitly be quantified. The model is estimated using a newly developed Bayesian estimation procedure and US corporate bond transaction data. The empirical analysis centers around an illustrative case study on Ford Motor Company and an analysis of systematic credit and liquidity interactions based on bond portfolios sorted by credit ratings. The results show that the impact of liquidity shocks on credit risk is mostly negligible, whereas the impact of credit shocks on liquidity, on the other hand, can be substantial. This so-called credit-induced liquidity component is found to be particularly large during the most turbulent times and for bonds with lower credit ratings.

Chapter 3 considers sovereign credit risk instead of corporate credit risk and investigates credit risk premia embedded in sovereign credit default swaps (CDS). A sovereign CDS can be viewed as an insurance contract in which the buyer seeks protection against defaults of the underlying reference sovereign specified in the contract. The protection buyer makes

2 The intensity-based literature typically models credit-liquidity interactions by correlating stochastic credit and liquidity intensity processes (see, e.g., Janosi et al., 2002; Duffie et al., 2003; Driessen, 2005; Liu et al., 2006; and Buhler and Trapp, 2010). By modeling dependencies with correlations, however, these models only consider symmetric and contemporaneous interactions and are not able to address asymmetric feedback effects.

3 Self- and cross-exciting processes were first introduced by Hawkes (1971b) (see also Hawkes, 1971c; Hawkes and Oakes, 1974) and are also often referred to as Hawkes processes. They have mainly been used in disciplines other than finance and economics. Ogata and Akaike (1982), for example, used these processes for the modeling of earthquake occurrences. Other applications of these processes can be found in genome analysis (see Reynaud-Bouret and Schbath, 2010) and in models representing the transmission of contagious diseases in epidemiology. Some economic applications of these processes can be found in Errais et al. (2010), Aït-Sahalia et al. (2014), and Aït-Sahalia et al. (2015).

periodic premium payments to the protection seller, who in return pays a certain amount to the protection buyer in case the reference sovereign defaults before the maturity of the contract. The 'price' of a CDS contract is called the CDS spread and refers to annual amount that the protection buyer needs to pay to the protection seller and is quoted as the number of basis points of the notional value of the CDS contract. In principle, CDS spreads are higher for riskier countries and are thus useful for assessing the relative riskiness of countries. In absolute terms, however, the market-implied default probabilities that can be backed out from the CDS spreads are typically much higher than default probability estimates based on historical default data. This indicates that sovereign CDS spreads are not pure measures of expected default risk and that other components are priced in as well.

Chapter 3 investigates the pricing effects of two types of risk premia related to sovereign credit risk. The first type of risk premium that is considered is referred to as the 'distress risk premium' and captures the compensation (risk averse) investors require for being exposed to unexpected variations in credit spreads. Other studies also analyze distress risk premia embedded in sovereign CDS spreads and find their price effects to be substantial and show that they are closely related to global and macroeconomic factors (Pan and Singleton, 2008; Longstaff et al., 2011; Zinna, 2013). The second type of risk premium that is considered is referred to as the 'default event risk premium'. The default event risk premium captures the compensation investors demand for bearing the risk of the unpredictable default event itself. In principle, such an event risk premium is priced if a default has a systemic impact. Evidence for default event risk premia being priced has been found in the corporate credit risk literature (see e.g., Driessen, 2005; Berndt et al., 2008). The sovereign credit risk literature, however, has largely ignored this component.

To estimate the credit risk premia, Chapter 3 introduces a new model for the term structure of sovereign credit risk. In the model, a country can default either due to a shock in a common factor, or due to a shock in a country-specific factor. Both these factors are modeled as self-exciting processes, making them able to capture the commonality and clustering of jump-like increases in CDS spreads observed in the data. The estimated model facilitates a decomposition of CDS spreads and shows that both distress and default event risk premia are substantial. In relative terms, distress risk premia seem to be more heavily priced in CDS

spreads of countries with higher credit ratings, whereas default event risk premia are more heavily priced in CDS spreads of countries with lower credit ratings.

1.2 ON PART II: LABOR INCOME RISK

For many investors, most of their wealth is made up by the value of their risky human capital, which is defined as the present discounted value of all future labor income. Indeed, it is estimated that human capital accounts for 90 to 93 percent of overall wealth in the economy (Lustig et al., 2013; Palacios, 2015). At the same time, labor is also an important input for many firms and the employees' human capital forms an important asset for the firm. As such, labor and financial markets are fundamentally interrelated and the question arises whether labor income risk affects expected stock returns.

The growing literature on labor and finance has developed substantial evidence for a variety of possible channels on how labor markets can affect portfolio allocation and expected stock returns. From the firm's perspective, for example, the fact that labor is not owned by the firm, but merely rented from its employees, exposes the firm to the risk that valuable human capital can leave at any time. Especially for firms operating in sectors where employees' skill sets are easily transferable to other industries, employees might be more inclined to search for higher salaries elsewhere when the industry in which they are currently employed is hit by adverse shocks. That is, labor mobility could amplify a firm's exposure to systematic risk (Donangelo et al., 2010, and Donangelo, 2014). From an investor's perspective, one way in which labor income risk can affect stock prices is through a portfolio channel. As human capital makes up most of the typical investor's wealth, a non-zero correlation between human capital returns and stock returns induces the investor to be exposed to equity market risk, even in the absence of direct stock holdings. On the individual level, this could have implications for the investor's optimal portfolio holdings as he wants to hedge this indirect exposure to equity market risk. On the aggregate level, these hedging demands could lead to labor income risk affecting equilibrium stock market prices and returns (see, e.g., Mayers, 1972). Next to these direct hedging demands, evidence has also been found for a "Keeping-up-with-the-Joneses" (KUJ) effect, in which investors care about their

wealth relative to that of their 'peers' (Abel, 1990). If this KUJ effect is in place, investors particularly demand stocks that correlate highly with the labor income of their peers, since these stocks provide a good hedge against labor income shocks of their peers. Again, in the aggregate these demand effects may affect expected stock returns (Gomez et al., 2009, and Gomez et al., 2016).

Chapter 4 further analyzes the relation between labor and equity markets, thereby focusing in particular on possible horizon effects that might play a role. There are several reasons to expect differences in labor-equity market linkages across different horizons. First, wages are sticky and set infrequently. As a result wages are smooth and therefore the correlation between stock returns and labor income growth at short horizons is typically found to be low (see, e.g., Fama and Schwert, 1977, and Cocco et al., 2005). At the same time, labor income growth and stock returns are not expected to diverge forever, and indeed several studies find stronger comovements between the two at longer horizons (see, e.g., Baxter and Jermann, 1997; Benzoni et al., 2007). Furthermore, labor income earning investors may also have different investment horizons, for example, due to differences in their remaining pre-retirement years in which they still earn wages. These heterogeneous investment horizons may also affect the frequency at which systematic risk is priced (Kamara et al., 2016).

The aim of Chapter 4 is thus to test whether the impact of labor income risk on returns displays any horizon effects and, if so, to identify which horizon(s) are most pronounced. Chapter 4 first introduces a stylized model in which optimal portfolio allocation decisions and expected stock returns explicitly depend on the covariances between one-period equity returns and labor income growth rates across different horizons. Using the model predictions as a basis for the empirical tests, a flexible empirical methodology that simultaneously includes labor income risk across different horizons is employed. The results provide robust evidence that the two- to four-year horizon strongly dominates as this is the only horizon for which labor income risk is significantly priced. Furthermore, a simple two-factor model that includes an equity market risk factor and a labor income risk factor at this medium horizon outperforms many traditional asset pricing models in explaining the cross-section of stock returns. Portfolio adjustments are also found to be most significant at this horizon. In all, the results are in line with evidence of wage rigidity in which wages adjust to changes in the marginal product of labor every two to four years due to infrequent wage setting.

Part I

CREDIT RISK

2 FEEDBACK BETWEEN CREDIT AND LIQUIDITY RISK IN THE US CORPORATE BOND MARKET

This chapter is based on:

Sperna Weiland, R. C., Laeven, R. J. A., De Jong, F. (2019). Feedback between Credit and Liquidity Risk in the US Corporate Bond Market, Unpublished working paper. University of Amsterdam, Amsterdam, and Tilburg University, Tilburg.[1]

2.1 INTRODUCTION

It is well documented that credit and liquidity risk in the corporate bond market are interrelated and that liquidity conditions tend to worsen when credit risk manifests itself. Dick-Nielsen et al. (2012) and Friewald et al. (2012), for example, show that corporate bond market liquidity dried up substantially during the sub-prime crisis, and especially so for speculative grade bonds. In these studies, as is the case in most of the bond literature, the interactions between credit and liquidity risk are mainly addressed in a 'static' way, e.g., by comparing liquidity measures across rating classes or by comparing tranquil periods with periods of distress. As argued by the recent structural credit risk literature, however, it is plausible that credit-liquidity interactions are 'dynamic' and that these interactions contribute for a non-trivial part to bond yield

1 We are very grateful to Jean-Paul Renne (discussant), Cynthia Wu (discussant), Yacine Aït-Sahalia, Frank Kleibergen, David Lando, Thomas Philippon, Marti Subrahmanyam, Wei Xiong, Izhakian Yehuda, and conference and seminar participants at the 10th Annual Society for Financial Econometrics (SoFiE) Conference in New York, the GRETA 16th International Conference on Credit Risk Evaluation in Venice, the 5th Conference on Credit Analysis and Risk Management in Basel, the 30th Australasian Finance and Banking Conference in Sydney, the 17th Winter School on Mathematical Finance in Lunteren, the Dutch Central Bank, the National Bank of Belgium, Princeton University, the Tinbergen Institute, the University of New South Wales, and the University of Amsterdam for their comments and suggestions. This research was funded in part by the Netherlands Organization for Scientific Research (NWO) under grant NWO Vidi 2009 (Laeven).

spreads (see, e.g., Ericsson and Renault, 2006; He and Milbradt, 2014; Chen et al., 2017).

The structural credit risk literature usually relies on calibration to empirically analyze its models, which often provide a limited fit to the data especially in crisis episodes. Reduced-form, most noticeably intensity-based, credit risk models typically capture the data better and allow for formal estimation and testing procedures. In important work, Chen et al. (2017) state, however, that "such dynamic interactions [between credit and liquidity risk] are not easy to capture using reduced-form models". In this chapter, we introduce a new parsimonious reduced-form framework in which dynamic feedback loops between credit and liquidity risk are possible. Our framework allows for a dynamic feedback mechanism as alluded to in structural models while preserving the specific advantages of a reduced-form model. We develop a corresponding Bayesian estimation procedure, enabling us to conduct a formal econometric investigation of the dynamic interactions between bond credit and liquidity risk and their effects on yield spreads and bond investment risk.

The novelty of our modeling approach, compared to other reduced-form models, is that we model the default and liquidity intensities as bivariate self- and cross-exciting processes. Self-excitation in the credit and liquidity intensities, where adverse credit (liquidity) shocks increase the probability of further credit (liquidity) shocks, captures the empirical observation that credit events tend to cluster in time (see e.g., Lando, 1998, Duffie and Singleton, 1999, Aït-Sahalia et al., 2014), and that liquidity risk is persistent (see, e.g., Acharya and Pedersen, 2005, and references therein). The cross-excitation between credit risk build-up and liquidity dry-up gives rise to an explicit dynamic feedback mechanism over time that generates positive default-liquidity loops in which credit and liquidity shocks cluster in a flexible, potentially asymmetric, fashion.

We derive (semi-)closed-form bond price formulae, which depend on both intensity processes. The liquidity intensity appears through a liquidity discount factor, reflecting that higher illiquidity drives down prices (see, e.g., Duffie and Singleton, 1999; Longstaff et al., 2005). As a consequence, credit (liquidity) shocks affect bond prices in two ways, namely by a self-exciting (cross-exciting) jump in the credit intensity, reflecting the increased likelihood of going into default, and by a cross-exciting (self-exciting) jump in the liquidity intensity, which increases the illiquidity discount. Compared to other intensity-based models, and in line with the structural models of He and Milbradt (2014) and Chen et al. (2017),

our model facilitates a more refined decomposition of yield spreads, allowing us to explicitly quantify the price impact of and risk associated with pure credit risk, pure liquidity risk, and the two credit-liquidity interaction terms.

The main challenges in estimating our model stem from the fact that both the credit and liquidity intensity processes, as well as the credit and liquidity shocks, are inherently latent, and the fact that we work in a bond context in which data sampling is irregular. Because of this irregular sampling, we cannot rely on existing estimation procedures dealing with self- and mutual excitation (see, e.g., Aït-Sahalia et al., 2015; Boswijk et al., 2016). Instead, building on Eraker (2004), we propose a Bayesian Markov Chain Monte Carlo (MCMC) procedure, based on a discretization of the intensity processes. We add an additional measurement equation to our model that explicitly incorporates liquidity information embedded in estimated bid-ask spreads. This improves identification and facilitates interpretation of the liquidity factor. We show in a Monte Carlo study that our estimation methodology yields very good performance.

For our empirical analysis we use corporate bond transaction data from the Trade Reporting and Compliance Engine (TRACE) database, which contains price information on practically all secondary market corporate bond transactions. We focus on the crisis period starting from July 2007 up to July 2009, as credit-liquidity interactions are likely to be particularly prevailing in distressed periods. Since most structural feedback mechanisms between credit and liquidity risk identified by the literature are of a firm-specific nature, we first illustrate our methodology in a firm-specific analysis of Ford Motor Company, which is among the companies with the most, and most actively traded, bonds outstanding. Next, we consider bond portfolios sorted by credit rating to investigate the interaction between systematic credit and liquidity shocks.

Our empirical results show that our model captures the data very well. For example, our model yields average relative pricing errors between 4% and 10% in the crisis period we analyze. Chen et al. (2017) in a structural model calibrated to a longer time period report relative pricing errors in credit spreads of over 100% in the crisis period. Our model identifies multiple clusters of credit and liquidity shocks and we find strong evidence of both self- and cross-exciting behavior of credit and liquidity risk. Adverse credit shocks significantly affect bond market liquidity and these cross-excitation effects tend to be stronger for bonds

with a lower credit rating. The reverse feedback from liquidity to credit is found to be statistically significant for below investment grade bonds, but economically limited for all rating classes.

Our model allows for a decomposition of the bond yield spreads into credit, liquidity, and two interaction components. We find that the widening of yield spreads during the financial crisis can for a large part be explained by worsening liquidity conditions. The decomposition reveals, however, that this illiquidity is in fact for a substantial part driven by deteriorating credit conditions. This so-called 'credit-induced liquidity' component accounts for 0.50 (AAA/AA rating) to 0.73 (B and lower rating) percentage points of total 10-year spreads on average, but in the most turbulent times, it accounts for up to 2.05 percentage points (B and lower rating). In the case study on Ford Motor Company, the credit-induced liquidity component makes up over 60% in relative terms of total yield spreads at the peak of the crisis. The 'liquidity-induced credit' component, on the other hand, is typically small and thus remains economically limited.

We also analyze the yield spread decompositions along the term structure dimension. We show that the credit-liquidity interaction terms influence bonds of different maturities in different ways. We find in particular a pronounced downwards sloping term structure for the credit-induced liquidity interaction term, most similar in shape to the pure liquidity component. For example, shortly after the default of Lehman Brothers, the credit-induced liquidity component contributed for over 10 percentage points to the 3-year yield spreads of Ford, whereas this component contributed for 3 percentage points to Ford's 20-year yield spreads. Furthermore, the liquidity-induced credit component, while economically limited for most maturities, becomes relevant for short maturities of lower-rated bonds at the peak of the crisis.

Finally, we show that our results have important implications for standard risk management procedures such as Value-at-Risk computations. The self- and cross-excitation mechanisms induce feedback effects and non-trivial impulse responses, resulting in a fat tail in the loss domain of the bond return distribution. Therefore, ignoring credit-liquidity interactions may result in a considerable underestimation of capital requirements. We show that a substantial part of the capital requirements is due to credit-liquidity interactions, and this effect typically becomes more pronounced for bonds with lower credit ratings. We find, for example, that the 10-day 99% VaR of a 10-year Ford bond decreases from

5.7 to 4.6 cents per dollar invested if one ignores only the credit-liquidity interactions.

Our research is part of the literature on the interactions between credit and liquidity risk in bond markets.[2] Ericsson and Renault (2006) introduce a structural model in which the illiquidity of the market for distressed debt influences renegotiations under financial distress. He and Milbradt (2014) develop a structural model in which a default-liquidity loop arises through a debt rollover channel (see also He and Xiong, 2012). Assuming a relatively illiquid post-default secondary bond market, they show that higher default risk worsens a bond's pre-default secondary market liquidity. This, in turn, amplifies equity holders' rollover losses and increases the endogenous probability of default. Nagler (2019) provides some empirical evidence that the rollover channel from bond liquidity to credit risk exists and has a substantial price impact around liquidity shocks. Chen et al. (2017) study the default-liquidity spiral across the business cycle. Our reduced-form framework allows to conduct a formal econometric analysis of the dynamic credit-liquidity interactions.

Within the intensity-based credit risk literature, the credit and liquidity intensities are typically modeled by stochastic diffusive processes (see, e.g., Janosi et al., 2002; Duffie et al., 2003; Driessen, 2005; Liu et al., 2006; Buhler and Trapp, 2010). By using possibly correlated diffusive components, these models only consider contemporaneous and sym-

2 There are several reasons to expect dynamic interactions between credit and liquidity risk in the corporate bond market. As documented by Dick-Nielsen et al. (2012) and Bongaerts et al. (2017), credit spreads increase while corporate bond turnover decreases in times of distress. The decrease in corporate bond turnover can partly be explained by a flight-to-liquidity type of behavior, where investors refrain from trading relatively illiquid corporate bonds and instead use more liquid assets for rebalancing and risk-shifting purposes. Due to the over-the-counter nature of the corporate bond market, the lower turnover increases dealers' search costs of finding counterparties and therefore induces them to quote higher bid-ask spreads, i.e., transaction costs increase when credit quality deteriorates. This supports a dynamic interaction from credit to liquidity. Another reason why higher credit risk can imply lower liquidity is the cost of market making. If higher credit risk is associated with higher bond price volatility, bond market makers will charge wider bid-ask spreads to cover their inventory costs (Ho and Stoll, 1981). Higher credit risk and higher volatility may also lead to increased information asymmetry among corporate bond traders and for this reason have a detrimental effect on liquidity, see, e.g., the standard microstructure model of Glosten and Milgrom (1985). He and Milbradt (2014) argue that there is also a channel from corporate bond market liquidity to credit risk. Starting from the observation that bond transaction costs increase in times of distress, they show that a decrease in bond market liquidity results in rollover losses, which in turn increases default risk.

metric interactions,[3] and are not able to address the (potentially asymmetric) lead-and-lag relationships of the interaction effects. One notable exception to this modeling approach is Monfort and Renne (2014), who model default and liquidity interactions through credit- and liquidity-related regimes which influence the process parameters of the default and liquidity intensities. Our model preserves the specific advantages of intensity-based models, but features a dynamic feedback mechanism through mutual excitation between credit and liquidity risk that does not assume symmetry between the prevalence of credit-liquidity interactions. This enables us to explicitly quantify the impact of and risk associated with the dynamic credit-liquidity interactions over time and across maturities.

Our research is also related to the literature on MCMC methods in a financial context (see, e.g., Johannes et al., 1999; Eraker et al., 2003; Eraker, 2004; Johannes and Polson, 2009). More specifically, we develop an estimation methodology that builds upon and significantly extends the procedure of Eraker (2004) for estimating different classes of jump diffusion models of stock price dynamics using joint option and stock price data. The main differences between our setting and his are that we consider irregularly sampled bond price data instead of option and stock data, that our state processes are latent instead of partially observable, and that our model set-up is more general, since we allow for self- and mutual excitation between the state variables instead of perfectly correlated jumps as in Eraker (2004). The result of these differences is that identification and estimation are more challenging in our model. To accommodate these challenges we propose a different order of sampling in the MCMC algorithm and construct an additional measurement equation.

The remainder of this chapter is structured as follows: Section 2.2 introduces the model set-up. Section 2.3 provides a data description and Section 2.4 describes our estimation methodology. Section 2.5 presents the estimation results of our empirical analysis. Section 2.6 illustrates some of the implications of our findings, and Section 2.7 concludes. Proofs and some auxiliary details of the estimation methodology are relegated to six Appendices [2.A, 2.B, 2.C, 2.D, 2.E, 2.F].

3 Buhler and Trapp (2010) assume that default and liquidity intensities are linear combinations of two underlying independent factors, which they label the pure default and pure liquidity factors. In this way, they allow for asymmetry between the contemporaneous effects of credit on liquidity, and vice versa.

2.2 THE MODEL

2.2.1 *Dynamic credit-liquidity interactions*

Throughout we consider a filtered probability space $\left(\Omega, \mathcal{F}, \{\mathcal{F}\}_{t\geqslant 0}, \mathbb{Q}\right)$[4] satisfying the usual conditions of right-continuity and completeness. We propose to model credit and liquidity interactions through mutually exciting jump processes defined on this probability space. That is, we define two point processes N_t^c and N_t^l, denoting the credit and liquidity shock arrival processes, respectively, with corresponding stochastic intensity processes λ_t^c and λ_t^l, respectively. The jump intensities describe the $\mathcal{F}_t$-conditional probabilities of credit and liquidity shocks occurring in a time interval of length Δ as follows:

$$\begin{cases} \mathbb{Q}\left[N_{t+\Delta}^c - N_t^c = 0 \middle| \mathcal{F}_t\right] = 1 - \lambda_t^c \Delta + o(\Delta), \\ \mathbb{Q}\left[N_{t+\Delta}^c - N_t^c = 1 \middle| \mathcal{F}_t\right] = \lambda_t^c \Delta + o(\Delta), \\ \mathbb{Q}\left[N_{t+\Delta}^l - N_t^l = 0 \middle| \mathcal{F}_t\right] = 1 - \lambda_t^l \Delta + o(\Delta), \\ \mathbb{Q}\left[N_{t+\Delta}^l - N_t^l = 1 \middle| \mathcal{F}_t\right] = \lambda_t^l \Delta + o(\Delta), \end{cases} \tag{2.1}$$

and $\mathbb{Q}\left[N_{t+\Delta}^k - N_t^k > 1 | \mathcal{F}_t\right] = o(\Delta)$, for $k = \{c, l\}$.

The dynamic feedback between the two point processes comes from the fact that both intensities respond to jumps in both processes, supposed to occur in the following mean-reverting way:

$$\begin{aligned} d\lambda_t^c &= \alpha_c(\lambda_\infty^c - \lambda_t^c)dt + \sigma_c\sqrt{\lambda_t^c}dW_t^c + \beta_{1,1}dN_t^c + \beta_{1,2}dN_t^l, \\ d\lambda_t^l &= \alpha_l(\lambda_\infty^l - \lambda_t^l)dt + \sigma_l\sqrt{\lambda_t^l}dW_t^l + \beta_{2,1}dN_t^c + \beta_{2,2}dN_t^l, \end{aligned} \tag{2.2}$$

where W^c and W^l are independent Brownian motions capturing small fluctuations in credit and liquidity risk. In the case of a large credit shock (i.e., $dN_t^c = 1$), the credit intensity jumps up by $\beta_{1,1}$ and the liquidity intensity by $\beta_{2,1}$. Similarly, in the case of a large liquidity shock (i.e., $dN_t^l = 1$), the liquidity intensity increases by $\beta_{2,2}$ and the credit intensity by $\beta_{1,2}$. Hence, occurrences of large credit and liquidity shocks instantaneously increase the probabilities of future credit and liquidity

4 In this chapter, we consider primarily bond-implied measurements of credit and liquidity event arrival processes, except in Section 2.5.3. This means that we focus primarily on obtaining information on these arrival processes under an equivalent martingale measure ($\mathbb{Q}$), instead of under the physical probability measure ($\mathbb{P}$).

shocks via their direct impacts on the intensity processes. Following a jump there is mean reversion in the intensity processes at exponential rate α_k to a base-line intensity λ_∞^k, $k = \{c, l\}$. Self-excitation in the credit intensity captures the empirical observation that credit events tend to cluster. Similarly, the self-excitation in liquidity captures the persistence in liquidity risk. The mutual excitation between credit and liquidity provides a parsimonious way of modeling feedback between the two. Unlike contemporaneously correlated diffusive intensities, which are typically employed in the literature, this feedback mechanism is dynamic in nature, allowing for lead-and-lag relationships. Furthermore, it does not superimpose any symmetry (i.e., we do not impose $\beta_{1,2} = \beta_{2,1}$) between the effects of credit risk on liquidity and vice versa. Finally, the feedback between the counting processes N^c and N^l is probabilistic rather than certain.[5]

It can be shown that the quadruple $\left(N^c, N^l, \lambda^c, \lambda^l\right)$ is jointly Markov (see Hawkes, 1971a; Aït-Sahalia et al., 2015), and, under suitable parameter restrictions, stationary and nonnegative. More specifically, to assure nonnegativity of the processes, we assume that $\lambda_\infty^k \geqslant 0$, $\alpha_k \geqslant 0$, $\sigma_k \geqslant 0$ $(k = c, l)$, $\beta_{i,j} \geqslant 0$ $(i, j = 1, 2)$, and that the Feller-condition (Feller, 1951) in the intensity processes is satisfied (i.e., $2\alpha_k\lambda_\infty^k \geqslant \sigma_c^k$, $k = c, l$). In order to assure stationarity, we assume sufficiently strong mean reversion and we impose the sufficient conditions $\alpha_c > \beta_{1,1} + \beta_{1,2}$ and $\alpha_l > \beta_{2,1} + \beta_{2,2}$.

The nature of our jump model, which aims to capture the subtle features of the credit and liquidity intensity dynamics, requires a sample covering a sufficiently turbulent time period and a rich panel of bond prices to properly estimate the jump intensity dynamics and the jump process parameters. This follows as a result of the Peso problem. In our empirical analysis, we therefore focus on the crisis period from July 2007 until July 2009. It would be tenuous, however, to accurately estimate risk premia based on this relatively short sample period. Since the main aim of this chapter is to quantify credit-liquidity interactions, and not the evaluation of risk premia, we therefore choose to abstract away from risk

5 It would be feasible to let the diffusive parts of our model also be correlated. We choose, however, to only consider interactions between liquidity and credit risk through the cross-excitation mechanism dictated by the parameters $\beta_{1,2}$ and $\beta_{2,1}$, as we focus on the dynamic feedback that arises in episodes of distress (i.e., large credit or liquidity shocks), rather than on the contemporaneous correlation between the two in tranquil periods. Furthermore, we find that the benefits of adding a correlation parameter do not outweigh the econometric costs of having to estimate this additional parameter.

premia for most of this chapter, although we will briefly touch upon this topic in Section 2.5.3.

2.2.2 *Bond pricing*

We derive our bond price formulae by extending the results of Duffie and Singleton (1999), Lando (1998), and others. These authors derive pricing formulae for defaultable coupon bonds using the doubly stochastic Poisson framework to model default risk. They first show that within intensity-based credit risk models, the time t value P_t of an n-year defaultable bond with semi-annual coupon c and face value 1, under the assumption of recovery of face value with recovery rate $(1-L)$, is given by

$$\begin{aligned} P_t = c\sum_{j=1}^{2n} \mathbb{E}^Q\left[e^{-\int_t^{t+0.5j}(r_s+\lambda_s)ds}\middle|\mathcal{F}_t\right] + \mathbb{E}^Q\left[e^{-\int_t^{t+n}(r_s+\lambda_s)ds}\middle|\mathcal{F}_t\right] \\ + \int_t^{t+n}(1-L)\mathbb{E}^Q\left[\lambda_s e^{-\int_t^{t+s}(r_u+\lambda_u)du}\middle|\mathcal{F}_t\right]ds, \end{aligned} \tag{2.3}$$

where r is the risk-free interest rate process, λ the default intensity process, and the expectation $\mathbb{E}^Q\left[e^{-\int_t^T \lambda_s ds}\middle|\mathcal{F}_t\right]$ the (risk-neutral) survival probability up to time $T > t$ where default occurs at the first jump of the underlying point process. The first part of Eqn. (2.3) represents the current value of all future coupon payments; the second part of (2.3) represents the current value of the principal repayment; and the third part represents the value of the recovery payment in case of default. Typically a recovery rate of 40% is assumed (i.e., $L = 0.6$).

Eqn. (2.3) is derived under the standard assumption that the company defaults at the first jump of the underlying point process. We assume, however, that, more generally, upon the occurrence of a credit event, a default occurs with probability $0 < \gamma \leqslant 1$. Letting τ denote the time-to-default, the $\mathcal{F}_t$-probability of survival to time $T > t$, $p(t,T)$, is then given by

$$p(t,T) = \mathbb{E}^Q\left[\mathbb{Q}\left[\tau > T|\mathcal{F}_T\right]|\mathcal{F}_t\right] = \mathbb{E}^Q\left[(1-\gamma)^{N_T^c - N_t^c}\middle|\mathcal{F}_t\right], \tag{2.4}$$

where we will use the convention $0^0 = 1$ and $0^c = 0$ for $c > 0$ as this makes formula (2.4) also valid for $\gamma = 1$.

Furthermore, Eqn. (2.3) does not yet take liquidity effects into account and, therefore, in line with the intensity-based credit risk modeling literature, we introduce an additional liquidity discount factor which can be seen as a convenience yield that investors require for holding an illiquid (defaultable) asset (see, e.g., Duffie and Singleton, 1999, and Longstaff et al., 2005). Finally, we assume that the risk-free rate process is independent of the credit and liquidity intensities, and we denote the time t price of a risk-free zero-coupon bond with maturity T by $D(t,T) = \mathbb{E}^{\mathbb{Q}}\left[\exp\left(-\int_t^T r_s ds\right)\middle|\mathcal{F}_t\right]$. All-in-all, we obtain the following general bond pricing formula, which depends on both credit and liquidity risk:

$$\begin{aligned} P_t = c\sum_{j=1}^{2n} D(t,t+0.5j)\mathbb{E}^{\mathbb{Q}}\left[e^{-\int_t^{t+0.5j}\lambda_s^l ds}(1-\gamma)^{N^c_{t+0.5j}-N^c_t}\middle|\mathcal{F}_t\right] \\ + D(t,t+n)\mathbb{E}^{\mathbb{Q}}\left[e^{-\int_t^{t+n}\lambda_s^l ds}(1-\gamma)^{N^c_{t+n}-N^c_t}\middle|\mathcal{F}_t\right] \\ + \int_t^{t+n}(1-L)D(t,s)\mathbb{E}^{\mathbb{Q}}\left[e^{-\int_t^s\lambda_u^l du}\lambda_s^c\gamma(1-\gamma)^{N^c_s-N^c_t}\middle|\mathcal{F}_t\right]ds. \end{aligned} \tag{2.5}$$

Although our setting with mutual excitation implies that we are outside the doubly stochastic Poisson framework that is conventionally adopted in the intensity-based credit risk literature, we show that the expectations appearing in (2.5) can be calculated explicitly (up to a solution of the system of ODEs), by exploiting the affine jump-diffusion framework of Duffie et al. (2000) in its generalized version as defined in their Appendix B. The explicit expressions and proofs are given in Appendix 2.C.

From (2.5), we see that credit (liquidity) shocks affect bond prices in two ways, namely by a self-exciting (cross-exciting) jump in the credit intensity, reflecting the increased likelihood of going into default, and by a cross-exciting (self-exciting) jump in the liquidity intensity, which increases the illiquidity discount. Bond prices compatible with our model thus contain information on the dynamics of both credit and liquidity shocks, and, therefore, provide an appropriate basis for our empirical analysis of the dynamic credit-liquidity interactions.

2.3 DATA

2.3.1 *Bond transaction data and bond-specific information*

In our empirical analysis, we consider US bond portfolios sorted by rating and conduct a case study on Ford Motor Company, both over the period from July 2007 until July 2009. We focus on Ford because it is among the companies with the most, and most actively traded, bonds outstanding. Furthermore, as we show below, bid-ask spreads of Ford's bonds rose to very high levels at the peak of the crisis, and many of Ford's bonds were subject to multiple rating changes during our sample period, making it an interesting case to analyze credit-liquidity interactions.

Our empirical study is based on bond transaction data, which we obtain from the TRACE (Trade Reporting and Compliance Engine) database. The TRACE system was launched on July 1, 2002, and contains the transaction reports of secondary over-the-counter corporate bond transactions by members of the Financial Industry Regulatory Authority. Initially, not all transaction reports were disseminated. From October 2004 onwards the *standard* TRACE database covers practically all secondary market corporate bond transactions. We use the *Enhanced* TRACE data, which differs from the standard TRACE data in a few respects. First, the Enhanced TRACE data contains all transactions back to July 1, 2002 (including the previously non-disseminated transactions). Second, the Enhanced TRACE data does not cap the transaction volumes.[6] Third, the Enhanced TRACE data contains historical buy and sell information.[7] We filter out erroneous trades using the methods proposed by Dick-Nielsen (2014). Furthermore, we use a price sequence filter that filters out price outliers, transactions with prices that are more than 20 percent away from the median transaction price that day, and transactions with prices that are more than 20 percent away from the previous trading price.

We focus on fixed coupon bonds and we exclude bonds that are putable, convertible, denominated in other currencies than US dollars, or have a maturity below 1 or beyond 30 years. We obtain this bond-specific information from the Mergent Fixed Income Securities Database (FISD). This

6 In standard TRACE data, transactions are capped at $5,000,000 for investment grade bonds and at $1,000,000 for speculative grade bonds.

7 These data improvements, compared to the standard TRACE data, come at the cost of a delay in the availability of the data.

database also contains bond credit rating histories and bond information relevant for pricing, such as the coupon rates, coupon frequencies, principal values, and day count conventions. We normalize all transaction data to have a principal value of $1,000. Finally, we discard retail-sized transactions with volumes below $10,000 for the Ford case study, and below $100,000 for the rating portfolios.

The resulting Ford sample consists of 72 bonds, traded on 506 days, with a total of 40,174 transactions. We use this full set of transactions to estimate daily bid-ask spreads (see Section 2.3.2). To estimate our model, however, we employ only the closing prices of the bonds on each day they are traded.[8] We have therefore an unbalanced panel dataset consisting of 10,903 end-of-day transactions. The corresponding summary statistics are reported in Table 2.1. The bonds of Ford were all rated B or lower throughout our sample period. Figure 2.1 illustrates the price path of a typical bond over the period 2007 to 2009 and highlights several major credit events that occurred during that period.[9] This figure already provides suggestive evidence that credit events tend to cluster in time and that bond prices typically drop around the time of a credit event.

The rating portfolios are subdivided into AAA/AA, A, BBB, BB, and B and lower-rated bonds, and are updated every quarter. A bond is included in the rating class portfolio corresponding to its rating prevailing at its first transaction in that quarter. In order to prevent the portfolios from containing many infrequently traded bonds, we include in every quarter only bonds of firms that have two or more bonds that trade on at least 20 different days, and subsequently only consider these relatively frequently traded bonds. Since not all bonds within the same rating class are equally credit risky, we correct the transaction prices for bond-specific fixed effects which measure the differences in baseline credit-worthiness. This procedure is explained in Appendix 2.D. The summary statistics of the resulting rating portfolios are given in Table 2.1.

8 The use of the full set of transactions would result in excessive computing time.

9 We consider a credit event to be any major adverse information shock regarding creditworthiness of a company. Concrete examples are credit rating downgrades or negative press releases.

Table 2.1. Descriptive statistics: transaction data. This table reports descriptive statistics of the transaction prices, maturities, number of transactions per day, and coupon rates for the rating portfolios and Ford; see Section 2.3.1. The sample period is July 2007 until July 2009.

		AAA/ AA	A	BBB	BB	B and lower	Ford
#Firms		219	377	353	236	298	1
#Transactions		12,215	22,149	26,982	26,240	32,807	10,903
	Mean	1001.5	982.3	958.2	896.0	799.4	600.9
Price	Std. dev.	43.8	82.2	84.7	91.4	151.0	213.2
($)	Min	490.2	17.0	326.4	488.2	136.0	93.3
	Max	1261.8	1434.7	1307.3	1218.8	1387.2	952.2
	Mean	8.7	14.9	14.0	7.5	6.3	11.0
Maturity	Std. dev.	9.1	10.9	10.3	4.5	3.1	6.6
(years)	Min	1.0	1.0	1.0	1.0	1.0	3.1
	Max	30.0	30.0	30.0	29.8	26.6	25.1
	Mean	24.6	44.8	54.5	52.8	66.0	21.6
#Obs.	Std. dev.	10.4	21.3	21.5	16.2	22.3	6.6
per day	Min	2	5	2	1	1	3
	Max	90	131	160	130	196	64
	Mean	5.03	5.93	6.45	7.10	8.26	6.64
Coupon rate	Std. dev.	1.33	0.94	1.03	1.02	1.60	0.83
(%)	Min	1.625	3.625	3.6	4.8	5.0	5.65
	Max	9.375	10.35	10.5	12	13.5	9.38

2.3.2 *Construction of bid-ask spreads and measurement equations*

In order to gain insight in how the bond market liquidity of the rating portfolios and of Ford has evolved over the sample period, we construct estimates of bid-ask spreads using the buy- and sell-indicators provided by TRACE. For this purpose, we first sort the data by bond and transaction time on each transaction date. Next, we construct pairs of consecutive trades on each date and for every bond i, and run the following linear regression over all transaction pairs $k = 1, 2, \ldots$ on a certain day $t = 1, 2, \ldots, N$:

$$Y^i_{t_k} - Y^i_{t_{k-1}} = \alpha_t + \beta_t \left(d^i_{t_k} - d^i_{t_{k-1}} \right) + \epsilon^i_{t_k}. \tag{2.6}$$

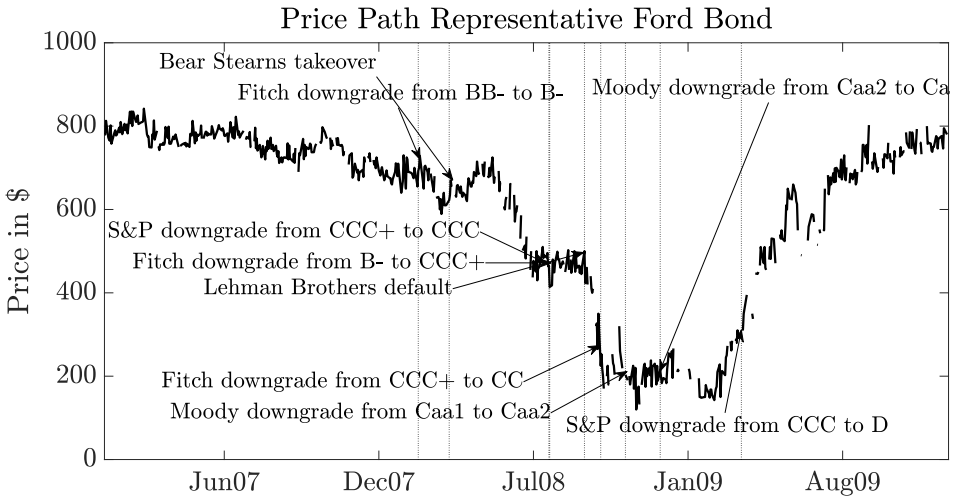


Figure 2.1. Representative Ford bond price path. The figure plots the price path of a representative bond of Ford, and indicates several credit events that Ford was exposed to. The reported prices are the final transaction prices at each date the bond was traded. Sources: TRACE and Mergent FISD.

Here $Y^i_{t_k}$ is bond i's log-price at transaction time t_k, $d^i_{t_k}$ is a dummy that takes the value $+1$ if the transaction was a buy order and -1 if the transaction was a sell order, and $\epsilon^i_{t_k}$ is an i.i.d. normally distributed error term. In this regression equation, the parameter β_t represents the average relative bid-ask spread at day t.[10] Table 2.2 displays the summary statistics of the estimated relative bid-ask spreads for Ford and the rating portfolios. Both for Ford and the rating portfolios, we observe strong increases of relative bid-ask spreads at the peak of the crisis. Furthermore, the overall bid-ask spreads of Ford bonds were particularly large compared to other firms (of the same credit rating) during our sample period.

Since we explicitly model, and wish to estimate, liquidity jumps, it is insightful to analyze the bid-ask spread data, to see when potential liquidity jumps have occurred. In principle, days on which the bid-ask spread is much higher compared to the preceding day, are days on which it is likely that a liquidity jump has occurred. However, because we do not

10 If we would use actual prices rather than log-prices in (2.6), we would obtain estimates of the bid-ask spreads in dollars, but because prices drop sharply during our sample period, relative bid-ask spreads are arguably more informative.

Table 2.2. Descriptive statistics: relative bid-ask spreads. This table reports descriptive statistics of the estimated relative (to transaction prices) bid-ask spreads of Ford and the rating portfolios over the whole sample period (panel A) and the period September 15, 2008 until March 31, 2009, reflecting the peak of the crisis (panel B); see Section 2.3.2. We report the mean, median, and standard deviation of the relative bid-ask spread over the corresponding period.

Panel A: Full sample period			
	Mean	Median	Std. Dev.
AAA/AA	0.71%	0.62%	0.32%
A	0.78%	0.72%	0.30%
BBB	0.83%	0.76%	0.32%
BB	0.69%	0.63%	0.29%
B and lower	0.88%	0.69%	0.53%
Ford	2.44%	1.95%	1.65%
Panel B: Peak of crisis			
AAA/AA	1.09%	1.00%	0.33%
A	1.13%	1.03%	0.30%
BBB	1.19%	1.15%	0.29%
BB	0.94%	0.84%	0.37%
B and lower	1.41%	1.18%	0.71%
Ford	4.00%	3.91%	2.02%

have true bid-ask spread data but only estimated bid-ask spread data, we have to account for estimation error before we can use these data to indicate potential ('likely') liquidity jumps. For this purpose, we assume that β_t is the true, but unobserved, bid-ask spread at day t, and follows an AR(1) process:

$$\beta_t = \phi\beta_{t-1} + \eta_t,$$

where $0 \leqslant \phi \leqslant 1$ and where η_t is an i.i.d. error term with zero mean and finite variance. The bid-ask spread at day t is, however, estimated with error u_t according to $\hat{\beta}_t = \beta_t + u_t$.

At a daily frequency, we can assume ϕ to be close to one. Therefore, innovations in the bid-ask spread are dictated by the error term η_t. Hence, by filtering the η_t series, we can indicate potential liquidity jumps, through identifying those dates on which (the filtered) η_t is relatively large. In particular, we indicate potential liquidity jumps by days on which the filtered level of η_t is more than 2.25 standard deviations above its mean. The details of the filtering procedure are described in Appendix 2.A.[11] In our estimation procedure detailed in the next section, we will exploit the potential liquidity jumps as indicated by the procedure above through additional measurement equations. We will assume that the potential liquidity jumps are noisy indicators of actual liquidity jumps, and utilize this information in identifying actual liquidity jumps.

2.4 ESTIMATION METHODOLOGY

In this section, we develop our estimation procedure for identifying the latent jump times, latent intensity processes, and model parameters based on a Markov Chain Monte Carlo (MCMC) method. Our estimation methodology builds upon Eraker (2004) and extends it mainly to allow for self- and cross-excitation in the latent jump intensities. As we will show below, the key to the challenge of dealing with mutual excitation is to include both the latent jump counters and the latent jump intensities into the state vector, inducing a joint Markovian structure, and then to properly discretize and orderly sample (in a specific way that accommodates their latency) the intensity processes defined in (2.2). Furthermore, we add an additional measurement equation to explicitly incorporate liquidity information embedded in the estimated bid-ask spreads.

2.4.1 *Joint posterior density*

We first develop an expression for the joint density of the observed data, model parameters, and state variables. This joint density is important in constructing the MCMC sampler and therefore we will show explicitly how it is derived. We write the (model-based) bond prices P_t as a func-

11 In total, we indicate 11 potential liquidity shocks for Ford. In a similar way, we indicate 5, 4, 6, 8, and 7 potential liquidity shocks for the AAA/AA, A, BBB, BB, and B and lower-rated portfolios, respectively.

tion F of the state variables $X_t = \{N_t^c, N_t^l, \lambda_t^c, \lambda_t^l\}$, the parameter vector Θ, and other arguments such as time to maturity and coupon frequency collected in a vector χ. That is, $P_t = F(X_t, \chi, \Theta)$, where F is given by the bond price formula (2.5), and
$\Theta = \{\alpha_c, \lambda_\infty^c, \sigma_c^2, \alpha_l, \lambda_\infty^l, \sigma_l^2, \beta_{1,1}, \beta_{1,2}, \beta_{2,1}, \beta_{2,2}, \gamma\}$. We have an unbalanced panel dataset with daily closing transaction price observations and we let $S(t)$ denote the set of bonds for which we have an observation at day t, $t = 1, \ldots, N$. We assume that the log-bond transaction prices $Y_{t,j}$, $j \in S(t)$, are observed with normally distributed pricing errors,[12] that is,

$$Y_{t,j} = \log\big(F(X_t, \chi_j, \Theta)\big) + \xi_{t,j}, \qquad j \in S(t), t = 1, \ldots, N,$$

with $\xi_{t,j} \sim \mathcal{N}(0, h^2)$.

In addition to the (log-)bond prices, we also employ the liquidity jump indicators described in Section 2.3.2. We let

$$I_t^l = \mathbb{1}(\Delta\beta_t > c), \tag{2.7}$$

where $\Delta\beta_t$ is the increment in the (relative) bid-ask spread between days $t-1$ and t. We take $c > 0$ to be a fixed threshold value, which we set at the mean change in the relative bid-ask spread plus 2.25 times the standard deviation of the change in the bid-ask spread. The indicator function I_t^l defined in (2.7) is viewed as a noisy measure of the latent liquidity jump occurrences, and we specify the corresponding measurement equation as follows:

$$I_t^l = \begin{cases} \text{Bernoulli}(\varphi_{\text{high}}) & \text{if } \Delta N_t^l = 1, \\ \text{Bernoulli}(\varphi_{\text{low}}) & \text{if } \Delta N_t^l = 0, \end{cases} \tag{2.8}$$

where $\text{Bernoulli}(p)$ is a Bernoulli random variable with success probability p, and $0 < \varphi_{\text{low}} < \varphi_{\text{high}} < 1$. The rationale underlying this measurement equation is that it helps identifying actual liquidity jumps based on bid-ask spread changes, which aids in overall state and parameter identification. Note that we do not impose or pre-fix liquidity jump times, nor do we restrict them to be necessarily in the set of dates identified by the bid-ask spread data. We merely assume that, on a day for which the liquidity jump indicator I_t^l equals 1, it is more likely to identify an actual liquidity jump ($\Delta N_t^l = 1$) because $\varphi_{\text{high}} > \varphi_{\text{low}}$.

12 For the rating portfolios we assume that fixed effects corrected log-prices are observed with normally distributed pricing errors.

Let $Y = \{Y_{t,j} : t = 1,\dots,N, j \in S(t)\}$ denote the vector of all closing transaction price observations, $I^l = \{I^l_t | t = 1,\dots,N\}$ the vector of all liquidity jump indicators, $\bar{Y} = \{Y, I^l\}$ the vector with all measurement observations, $X = \{X_t : t = 1,\dots,N\}$ the vector with all states, and $\bar{\Theta} = \{\Theta, \varphi_{high}, \varphi_{low}, h^2\}$ the vector with all parameters. Then the specification of the pricing errors and the liquidity jump indicator functions allows us to write the conditional density of the joint observations as

$$\begin{aligned} p(\bar{Y}|X,\bar{\Theta}) &\propto \prod_{t=1}^{N} p(I^l_t|X_t,\bar{\Theta}) \prod_{j\in S(t)} \phi(Y_{t,j}; F(X_t,\chi_j,\Theta), h^2) \\ &=: \prod_{t=1}^{N} p(I^l_t|X_t,\bar{\Theta}) p(Y_t|X_t,\Theta), \end{aligned}$$

where $\phi(x; m, s^2)$ denotes a normal density with mean m and variance s^2 evaluated at x, and $p(I^l_t|X_t,\bar{\Theta})$ is a Bernoulli density with the relevant success probability parameter (depending on the value of ΔN^l_t).

The full *joint posterior density* $p(\bar{Y}, X, \bar{\Theta})$ is then given by

$$p(\bar{Y}, X, \bar{\Theta}) \propto p(\bar{Y}|X,\bar{\Theta}) p(X|\bar{\Theta}) p(\bar{\Theta}), \tag{2.9}$$

where $p(\bar{\Theta})$ is the prior for $\bar{\Theta}$. We choose our priors to be proper but in such a way that they impose little information. More details can be found in Appendix 2.E. Using the Markovian property of the joint intensities and jump processes (see Aït-Sahalia et al., 2014), we can rewrite (2.9) as a product over the observation times. That is,

$$p(\bar{Y}, X, \bar{\Theta}) \propto p(\bar{\Theta}) \prod_{t=1}^{N} p(I^l_t|X_t,\bar{\Theta}) p(Y_t|X_t,\bar{\Theta}) p(X_t|X_{t-1},\bar{\Theta}), \tag{2.10}$$

where $p(X_t|X_{t-1},\bar{\Theta})$ is the transition density of the state processes.

Since the transition density of the state processes is unknown, we consider, with slight abuse of notation, the following discretized version of (2.2):

$$\begin{aligned} \lambda^c_{t+1} - \lambda^c_t &= \alpha_c \lambda^c_\infty \Delta - \alpha_c \lambda^c_t \Delta + \sigma_c \sqrt{\lambda^c_t \Delta} \epsilon^c_{t+1} + \beta_{1,1} \Delta N^c_{t+1} + \beta_{1,2} \Delta N^l_{t+1}, \\ \lambda^l_{t+1} - \lambda^l_t &= \alpha_l \lambda^l_\infty \Delta - \alpha_l \lambda^l_t \Delta + \sigma_l \sqrt{\lambda^l_t \Delta} \epsilon^l_{t+1} + \beta_{2,1} \Delta N^c_{t+1} + \beta_{2,2} \Delta N^l_{t+1}, \end{aligned}$$

where Δ is the length of the time interval between times t and $t+1$ (a business day), ϵ^c_{t+1} and ϵ^l_{t+1} are independent standard normal random variables, and $\Delta N^k_{t+1} = 1$ ($k = c, l$) indicates a jump arrival. The

jump counters ΔN^k_{t+1} are Bernoulli random variables with non-constant intensities $\lambda^k_t \Delta$ ($k = c, l$). For this discretization to be appropriate, the observation intervals Δ should not be too long. Eraker et al. (2003) show in a simulation study that at a daily frequency such a discretization bias is negligible. Using this discretization, the transition density $p(X_t|X_{t-1}, \bar{\Theta})$ can be decomposed as

$$\begin{aligned} p(X_t|X_{t-1}, \bar{\Theta}) &= p(\lambda^c_t, \lambda^l_t|N^c_t, N^l_t, X_{t-1}, \bar{\Theta}) p(N^c_t, N^l_t|X_{t-1}, \bar{\Theta}) \\ &= p(\lambda^c_t, \lambda^l_t|N^c_t, N^l_t, X_{t-1}, \bar{\Theta}) p(N^c_t|X_{t-1}, \bar{\Theta}) p(N^l_t|X_{t-1}, \bar{\Theta}), \end{aligned} \tag{2.11}$$

where $p(\lambda^c_t, \lambda^l_t|N^c_t, N^l_t, X_{t-1}, \Theta)$ is bivariate Gaussian, and $p(N^k_t|X_{t-1}, \Theta)$ ($k = c, l$) are Bernoulli. That is, upon discretization, the transition density is a mixture of normal densities. The decomposition in (2.11) also enables us to deal with the self- and cross-excitation, since we can sequentially draw N^c_t and N^l_t from the Bernoulli densities $p(N^c_t|X_{t-1}, \Theta)$ and $p(N^l_t|X_{t-1}, \Theta)$, respectively, and then draw λ^c_t and λ^l_t from $p(\lambda^c_t, \lambda^l_t|N^c_t, N^l_t, X_{t-1}, \Theta)$ using the newly drawn N^c_t and N^l_t in the conditioning information.

2.4.2 *Joint conditional posterior density and MCMC sampling*

Ultimately, we are interested in the *joint conditional posterior density* $p(\bar{\Theta}, X|\bar{Y})$. The MCMC algorithm, which we describe below, will draw from this conditional distribution, and with these draws we can estimate the marginal posterior density of the parameters $p(\bar{\Theta}|\bar{Y})$, since, by Bayes theorem, $p(\bar{\Theta}|\bar{Y}) \propto p(\bar{\Theta}, X|\bar{Y})$. Hence, we can take the sample average of $\bar{\Theta}^{(1)}, \bar{\Theta}^{(2)}, \ldots, \bar{\Theta}^{(G)}$, $\hat{\bar{\Theta}} = \frac{1}{G} \sum_{g=1}^{G} \bar{\Theta}^{(g)}$, with G the number of simulations, as our estimate of $\bar{\Theta}$. In a similar way, we can estimate the latent jump intensities by considering sample averages. Furthermore, to estimate the jump times, i.e., to decide whether a jump occurred at time t, $t = 1, \ldots, N$, we define a threshold, $\omega > 0$, and say that a jump occurred at time t if $\frac{1}{G} \sum_{g=1}^{G} \Delta N^{k(g)}_t > \omega$ ($k = c, l$) (see Johannes et al., 1999).

Because the joint conditional posterior density is high-dimensional and non-standard, we cannot sample from this density directly. To overcome these problems, we employ a Gibbs sampler, allowing us to sequentially draw all random variables from the joint posterior density. Furthermore, by using appropriate Metropolis-Hastings algorithms, we are able to

deal with the problem of non-standard densities. The general idea of our MCMC algorithm is outlined next, the details can be found in Appendix 2.E.

2.4.3 *MCMC algorithm*

The Gibbs sampler consists of the following steps, initialized by an appropriate set of starting values for X and $\bar{\Theta}$ when $g = 0$:

For each $g = 0, \ldots, G-1$ simulate for each $t = 1, \ldots, N$

1. $X_t^{(g+1)}$ from $p(X_t | X_{1:N_{\setminus t}}^{(g)}, \bar{\Theta}^{(g)}, \bar{Y})$,

and next simulate

2. $\bar{\Theta}^{(g+1)}$ from $p(\bar{\Theta} | X^{(g+1)}, \bar{Y})$,

where $X_{1:N_{\setminus t}}$ denotes the collection of state vectors X_s at all $s = 1, \ldots, N$ except at $s = t$. Note that we first sample both the jump intensities and jump counters at each given transaction date before considering the next transaction date. Eraker (2004) first samples all volatilities for all transaction dates and then, given these newly drawn volatilities, samples the jump counters for all transaction dates, and so on. This difference in the order of sampling is in principle asymptotically irrelevant, as the Markov chain converges regardless of the order of drawing in each iteration. However, as we will explain in the next subsection, the order matters in our case of fully unobservable states, because one might get stuck in a situation where jumps are never sampled, despite the information content embedded in (observable) bond prices. Sampling all states day-by-day will make it possible to alternate between two different sampling schemes, enabling us to circumvent this problem.

2.4.3.1 *Drawing intensities and jump counters*

Let us first consider simulating $X_t^{(g+1)}$ from $p(X_t | X_{1:N_{\setminus t}}^{(g)}, \bar{\Theta}^{(g)}, \bar{Y})$. We take advantage of the fact that $p(X_t | X_{1:N_{\setminus t}}^{(g)}, \bar{\Theta}^{(g)}, \bar{Y})$ is again characterized by its complete set of conditional distributions. Therefore, we can split the drawing of $X_t^{(g+1)}$ into the following three steps:

1. draw $N_t^{c(g+1)}$ from $p(N_t^c | \lambda_t^{c(g)}, \lambda_t^{l(g)}, N_t^{l(g)}, X_{1:N_{\setminus t}}^{(g)}, \bar{\Theta}^{(g)}, \bar{Y})$;

2. draw $N_t^{l(g+1)}$ from $p(N_t^l|\lambda_t^{c(g)}, \lambda_t^{l(g)}, N_t^{c(g+1)}, X_{1:N_{\setminus t}}^{(g)}, \bar{\Theta}^{(g)}, \bar{Y})$;

3. draw $\lambda_t^{c(g+1)}$ and $\lambda_t^{l(g+1)}$ from
$p(\lambda_t^c, \lambda_t^l|N_t^{c(g+1)}, N_t^{l(g+1)}, X_{1:N_{\setminus t}}^{(g)}, \bar{\Theta}^{(g)}, \bar{Y})$.

In Appendix 2.E we show that, under this simulation scheme, the probability of drawing jumps depends on the (previously drawn) paths of the intensities and the liquidity jump indicators but not on the observed bond prices. As a consequence, we may never sample jumps even when the bond data would clearly indicate large price drops, since we decoupled the drawing of jumps from the data likelihood. This problem is directly linked to the fact that our intensity processes are latent, unlike in Eraker (2004), where one of the state processes is observable.

To overcome this problem, we alternate between this simulation scheme and one in which we draw the whole state vector X_t jointly. That is, under the latter scheme, we do not split drawing X_t into three substeps as above, but draw the whole vector in one go, at each given transaction date. Because the λ^k ($k = c, l$) impact bond prices, the acceptance of a new draw for X_t then depends on both the transition densities as well as on the bond price data, and this increases the likelihood of drawing an initial jump, if the data suggest so. A disadvantage of this simulation scheme is, however, that the measurement equations induced by the liquidity jump indicators do not play a direct role in the proposal of new liquidity jumps, which effectively diminishes the information contained in these measurement equations. Therefore, we alternate between these two simulation schemes.

2.4.3.2 *Drawing parameters*

Next we consider drawing $\bar{\Theta}^{(g+1)}$ from $p(\bar{\Theta}|X^{(g+1)}, \bar{Y})$. Where possible, we draw parameters in a Gibbs step using conjugate priors. This can be done for h^2, φ_{high}, and φ_{low}, for which we specify an inverse Gamma prior, and Beta priors, respectively. Because the other parameters appear in both the transition densities as well as in the bond price formulae, their conditional densities are unknown, and therefore we have to rely on Metropolis steps. From a computational perspective, it is better to update these parameters in blocks. Indeed, each time one of these parameters gets updated, the ODEs of the bond price formulae have to be recomputed. On the other hand, we do not want the blocks to be too

large, because drawing from a high-dimensional density can result in low acceptance probabilities. We therefore draw the parameters in the following blocks: $\{\alpha_c, \lambda_\infty^c, \sigma_c^2, \gamma\}$, $\{\alpha_l, \lambda_\infty^l, \sigma_l^2\}$, and $\{\beta_{1,1}, \beta_{1,2}, \beta_{2,1}, \beta_{2,2}\}$. We use, for each block, a random walk Metropolis step with a Gaussian proposal density having as mean vector the previous draw and having a diagonal covariance matrix, similar to Eraker (2004). Details can again be found in Appendix 2.E.

2.4.4 *Parameter assumptions, practical considerations and finite sample behavior*

The credit and liquidity jump intensity processes and the associated counting processes are latent. The paths and parameters of these processes are estimated as part of the estimation methodology we develop. Although they are theoretically identified, it can be tenuous in practice to accurately estimate all parameters in our multivariate model. This applies in particular to the long-term average liquidity parameter, λ_∞^l. Therefore, we pre-identify the value of this parameter and treat it as fixed when estimating the model. The parameter λ_∞^l reflects the liquidity discount in relatively tranquil times, since for mean-reverted values of the intensity process this discount is given by $e^{-\lambda_\infty(T-t)} \approx 1-\lambda_\infty(T-t)$. According to Amihud and Mendelson (1986), the liquidity discount is given by the product of transaction costs and the trading frequency. In the relatively tranquil first year of our sample, we find the average relative bid-ask spreads to be approximately 50 and 150 basis points for all rating classes and Ford, respectively, similar to Bongaerts et al. (2017). As documented by Bongaerts et al. (2017), the average trading frequency of corporate bonds is close to once per year. This implies that the liquidity discount in tranquil times is approximately 0.5% and 1.5% for the rating portfolios and Ford. Hence, we fix $\lambda_\infty^l = 0.005$ and $\lambda_\infty^l = 0.015$ for the rating portfolios and Ford, respectively.

In order to test the performance of our estimation methodology, we conduct a Monte Carlo study, first for a restricted version of the model in which the liquidity intensity is absent, and next for the full model specification. The results, detailed in Appendix 2.B, show that our estimation methodology performs very well in a realistic setting and that the parameters are usually estimated with high precision.

2.5 ESTIMATION RESULTS

Equipped with these Monte Carlo results we now apply our estimation methodology to the data.

2.5.1 *Ford Motor Company*

We consider the Ford case study first. The last column of Table 2.3 reports the posterior means, posterior standard deviations, and 99 percent highest posterior density (HPD) intervals[13] of the Bayesian MCMC parameter estimates for Ford. Table 2.3 reveals that the parameters are usually estimated with high precision, as indicated by the relatively low posterior standard deviations compared to the (positive-valued) posterior means. Furthermore, the 99 percent HPD intervals are away from zero for all parameters, suggesting that all parameters are positive.

Table 2.3. Parameter estimates. This table reports the posterior means, standard deviations (in parenthesis), and 99 percent HPD intervals (in square brackets) for the parameters of the mutually exciting credit-liquidity model for the rating portfolios and Ford in the period July 2007 until July 2009. The last row reports the average relative pricing errors (ARPE). Throughout, we keep $\lambda_\infty^l = 0.005$ (rating portfolios) and $\lambda_\infty^l = 0.015$ (Ford) fixed.

	AAA/AA	A	BBB	BB	B and lower	Ford
α_c	1.372	2.264	2.663	3.219	2.335	3.897
	(0.988)	(0.512)	(0.252)	(0.068)	(0.379)	(0.174)
	[0.089, 2.679]	[1.327, 3.167]	[2.313, 3.115]	[3.112, 3.431]	[1.486, 2.792]	[3.629, 4.290]
λ_∞^c	0.086	0.090	0.077	0.184	0.150	0.132
	(0.025)	(0.012)	(0.012)	(0.027)	(0.030)	(0.026)
	[0.062, 0.140]	[0.075, 0.120]	[0.063, 0.109]	[0.138, 0.232]	[0.119, 0.235]	[0.109, 0.228]
σ_c^2	0.172	0.395	0.384	1.163	0.698	0.977
	(0.112)	(0.147)	(0.110)	(0.174)	(0.245)	(0.253)
	[0.021, 0.350]	[0.199, 0.730]	[0.266,0.626]	[0.835, 1.447]	[0.383, 1.300]	[0.740, 1.720]
γ	0.282	0.291	0.303	0.264	0.639	0.366
	(0.037)	(0.050)	(0.057)	(0.044)	(0.138)	(0.146)
	[0.215, 0.329]	[0.194, 0.362]	[0.185, 0.382]	[0.190, 0.350]	[0.367, 0.813]	[0.005, 0.523]

13 The 99 percent HPD interval is the smallest possible interval that contains 99% of the mass of the posterior distribution. It is analogous (but not identical) to a confidence interval in frequentist statistics.

Continuation of Table 2.3

	AAA/AA	A	BBB	BB	B and lower	Ford
α_l	4.904	4.838	5.117	5.78	2.487	3.720
	(0.865)	(0.194)	(0.295)	(0.446)	(0.522)	(0.101)
	[3.950, 6.609]	[4.564, 5.273]	[4.384, 5.572]	[4.893, 6.416]	[1.764, 3.652]	[3.507, 4.017]
σ_l^2	0.522	0.962	0.951	0.905	2.457	2.503
	(0.125)	(0.242)	(0.068)	(0.100)	(0.45)	(0.217)
	[0.305, 0.765]	[0.698, 1.614]	[0.807, 1.137]	[0.715, 1.143]	[1.852, 3.082]	[2.044, 2.937]
$\beta_{1,1}$	0.155	0.241	0.403	0.408	0.277	0.405
	(0.089)	(0.081)	(0.081)	(0.078)	(0.057)	(0.162)
	[0.000, 0.307]	[0.101, 0.414]	[0.211, 0.542]	[0.273, 0.574]	[0.187, 0.410]	[0.026, 0.689]
$\beta_{1,2}$	0.063	0.149	0.189	0.158	0.155	0.210
	(0.039)	(0.115)	(0.113)	(0.045)	(0.067)	(0.063)
	[0.000, 0.160]	[0.000, 0.418]	[0.000, 0.387]	[0.061, 0.315]	[0.051, 0.298]	[0.044, 0.359]
$\beta_{2,1}$	0.073	0.070	0.085	0.142	0.202	1.074
	(0.039)	(0.045)	(0.028)	(0.030)	(0.086)	(0.074)
	[0.000, 0.143]	[0.000, 0.140]	[0.028, 0.135]	[0.070, 0.197]	[0.094, 0.359]	[0.930, 1.203]
$\beta_{2,2}$	0.339	0.245	0.233	0.211	0.325	0.384
	(0.050)	(0.060)	(0.088)	(0.017)	(0.044)	(0.077)
	[0.230, 0.483]	[0.098, 0.329]	[0.067, 0.380]	[0.150, 0.249]	[0.254, 0.414]	[0.263, 0.528]
φ_{high}	0.424	0.418	0.453	0.509	0.526	0.405
	(0.141)	(0.130)	(0.150)	(0.147)	(0.140)	(0.128)
	[0.112, 0.781]	[0.129, 0.746]	[0.134, 0.818]	[0.162, 0.837]	[0.198, 0.852]	[0.122, 0.734]
φ_{low}	0.011	0.009	0.012	0.012	0.012	0.019
	(0.005)	(0.004)	(0.005)	(0.005)	(0.005)	(0.007)
	[0.003, 0.025]	[0.001, 0.021]	[0.003, 0.026]	[0.002, 0.026]	[0.002, 0.026]	[0.004, 0.037]
h	0.047	0.104	0.065	0.062	0.130	0.109
	(0.002)	(0.001)	(0.000)	(0.000)	(0.002)	(0.002)
	[0.043, 0.050]	[0.102, 0.105]	[0.064, 0.066]	[0.061, 0.063]	[0.128, 0.134]	[0.107, 0.117]
ARPE	3.36%	5.59%	5.18%	4.87%	9.85%	6.57%

Central to our analysis are the mutual excitation parameters $\beta_{ij}, i, j \in \{1, 2\}$. We find relatively large estimated values for the self-excitation parameters $\beta_{1,1}$ and $\beta_{2,2}$, which supports the existence of self-excitation in both credit and liquidity risk. Furthermore, the cross-excitation parameter $\beta_{2,1}$ is estimated to be particularly large, indicating a strong effect of credit shocks on liquidity risk. The reverse effect of liquidity shocks on credit risk is found to be smaller. This is in line with what can be expected from the structural credit risk literature, in which the main channel through which (idiosyncratic) bond market liquidity affects credit risk is a debt rollover mechanism (see, e.g., He and Xiong, 2012, He and Milbradt, 2014; Nagler 2019; Chen et al., 2017). This chan-

nel was unlikely to be in play for Ford. Indeed, within our sample period, none of Ford's bonds matured, and there were only two new bond issuances (one on April 28, 2008, and one on May 28, 2009).

To address more formally the significance of the mutual excitation parameters, we compare our model with nested model specifications using Bayes factors.[14] In particular, we compare our full model with (i) a model without mutual excitation (i.e., $\beta_{1,1} = \beta_{1,2} = \beta_{2,1} = \beta_{2,2} = 0$), (ii) a model without self-excitation (i.e., $\beta_{1,1} = \beta_{2,2} = 0$), and (iii) a model without cross-excitation (i.e., $\beta_{1,2} = \beta_{2,1} = 0$). In all cases the Bayes factors are essentially zero, and, hence, favor the full model specification with mutual excitation over the nested model specifications.

The upper panels of Figure 2.2 plot the model-based evolution of the bond prices in absolute (upper left panel) and relative (upper right panel) terms after either a liquidity or a credit shock.[15] The figures reveal that the initial price effect of a credit shock is estimated to be larger than the initial price effect of a liquidity shock. However, the price effect of the credit shock decays faster than that of a liquidity shock. The differences in impulse response and persistence that these patterns reflect are what makes possible to disentangle credit from liquidity jumps and to econometrically identify the shocks. Furthermore, the average daily volatilities of the continuous components of the credit and liquidity intensities are estimated to be $\sqrt{\sigma_c^2 \times \lambda_\infty^c \times \frac{1}{252}} \approx 0.02$ and $\sqrt{\sigma_l^2 \times \lambda_\infty^l \times \frac{1}{252}} \approx 0.01$, respectively. These average daily volatilities are much smaller than the self- and cross-excitation parameters, which makes it relatively easy to distinguish a credit or liquidity jump from a typical fluctuation caused by the Brownian components. The long-term mean of the credit intensity process, λ_∞^c, is estimated to be 0.132, hence the average 1-year (risk-neutral) default probability is $\lambda_\infty^c \times \gamma = 0.048$.

The estimated values of the liquidity jump indicator parameters φ_{high} and φ_{low} indicate that there is a reasonable degree of correspondence between the liquidity shocks identified by the estimation algorithm and those suggested by the bid-ask measurement equation. $\varphi_{\text{low}} = 0.019$, for

14 The details of how the Bayes factors are computed are given in Appendix 2.F.

15 We take the long-term average model-implied price of a 10-year bond and consider the price effect (i.e., impulse response) of a single credit or liquidity jump. That is, we compute the 10-year bond price using the estimated parameters and setting the credit and liquidity intensities equal to their long-term averages $\lambda_\infty^c = 0.132$ and $\lambda_\infty^l = 0.015$, respectively, and then analyze the price change when either one credit or one liquidity shock occurs.

example, implies that on days where our estimation algorithm did not detect a liquidity jump, the probability that the bid-ask data would have suggested a liquidity jump is only approximately 1.9 percent. Furthermore, $\varphi_{\text{high}} = 0.405$ implies that on days where our algorithm detects a liquidity jump, the bid-ask spread data also suggests a liquidity jump with probability 0.405. This confirms that the measurement equation provides valuable information on the occurrences of liquidity jumps, but also that the algorithm still has the freedom to identify liquidity jumps on dates different from the measurement equation.

The estimated credit and liquidity intensity paths are displayed in the upper panels of Figures 2.3 and 2.4. Both the credit and liquidity intensities increase at the peak of the crisis, reflecting an increase in credit risk and a worsening market liquidity. The lower panels of Figures 2.3 and 2.4 display the estimated credit and liquidity jump probabilities. Specifically, we compute, for each day in the sample, the fraction of MCMC iterations in which a credit or liquidity jump was drawn.[16]

To verify whether the estimated credit and liquidity intensity paths indeed represent adequate measures of credit and liquidity risk, we compare them to other measures of credit and liquidity risk. In the upper left panel of Figure 2.5, we plot both the (normalized) estimated liquidity intensities and the (normalized) estimated relative bid-ask spreads of Ford. Spearman's correlation between the two is 0.72, suggesting that the estimated liquidity intensity indeed reflects liquidity risk. In a similar fashion, the upper panel of Figure 2.6 plots the estimated credit intensities and the (normalized) 10-year CDS spreads of Ford.[17] Spearman's correlation between the two is 0.88, confirming that the estimated credit intensity process is indeed a meaningful representation of credit risk.

16 Because the MCMC algorithm sometimes draws jumps randomly, we define actual jumps to be those days on which the estimated jump probability is larger than 0.5. In total we identify 5 credit jumps and 3 liquidity jumps.

17 This maturity is chosen because it is closest to the average maturity of bonds used in our sample (see Table 2.1).

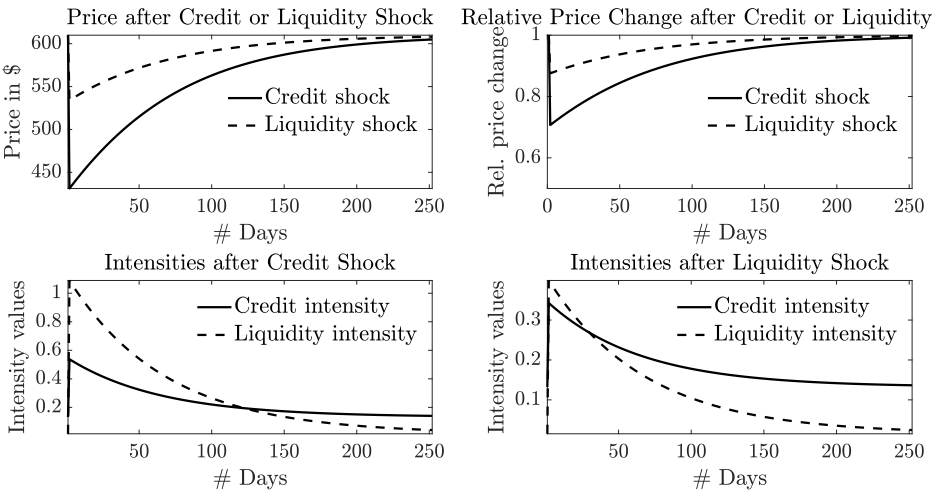


Figure 2.2. Price decay credit and liquidity shocks Ford. The figure plots the effects of credit and liquidity shocks on bond prices and the credit and liquidity intensities. In every panel, the credit and liquidity intensities start at their long-term means $\lambda^c_\infty = 0.132$ and $\lambda^l_\infty = 0.015$, respectively, and we assume that either a credit shock or a liquidity shock occurs (at Day 0). The upper left panel plots the price patterns after the credit or liquidity shock. The upper right panel plots the evolution of the relative price differences (compared to the initial price) after the credit or liquidity shock. The bottom left panel plots the credit and liquidity intensity paths after the credit jump. The bottom right panel plots the credit and liquidity intensity paths after a liquidity jump.

Finally, regarding model fit diagnostics, we find that, even though our model is parsimonious and the data spans a very turbulent time period, bond prices are fitted surprisingly well. The average relative pricing errors in our Ford sample range between 4.1% and 11.4% with an average of 6.6% across all bonds in the sample. Compared to the univariate model we analyzed as part of our Monte Carlo study, this relative pricing error is a great improvement, because that model produced an average relative pricing error of around 19%. This provides evidence that taking into account liquidity risk and the dynamic interactions between credit and liquidity risk greatly improves the model's capability of fitting the observed price paths in this turbulent period. Figure 2.7 displays the true and model-implied price paths of a representative bond in our sample. It shows that a substantial part of our (limited) pricing errors comes from the smoothing nature of our estimators (i.e.,

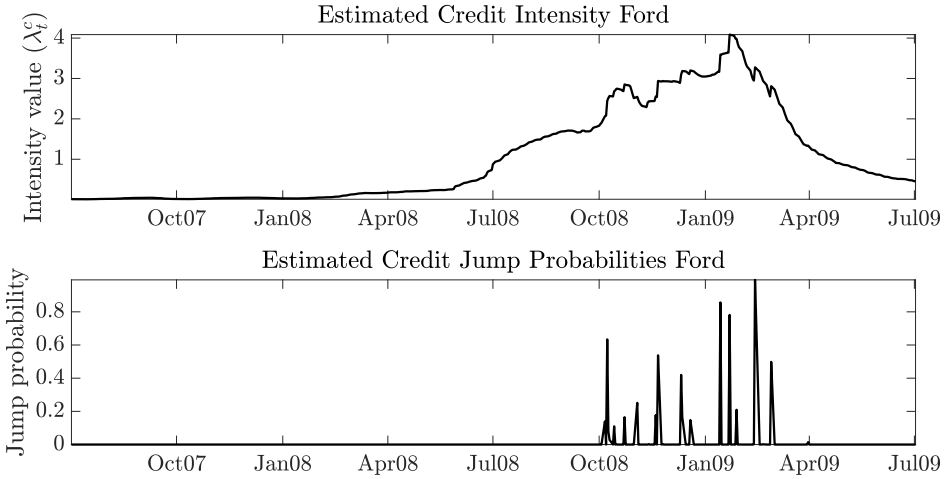


Figure 2.3. Estimated credit intensities and credit jump probabilities Ford. The upper panel plots the estimated credit intensities over the period July 2007 until July 2009. The bottom panel displays the estimated probabilities of credit jump occurrences.

the posterior means are conditional on all observations), which induces somewhat smoother model-implied price paths compared to somewhat noisier true price paths.

2.5.2 *Rating portfolios*

We next consider bond portfolios sorted by rating to analyze the dynamic interactions between systematic credit and liquidity shocks. Table 2.3 presents the Bayesian MCMC parameter estimates. We find clear evidence of self-excitation in the credit and liquidity intensities in nearly all rating portfolios. Only for the highest AAA/AA-rated bonds the HPD interval of the credit self-excitation parameter $\beta_{1,1}$ contains 0. The two cross-excitation effects, and in particular the credit-induced liquidity effect, become more pronounced as the rating worsens. Only for the highest two rating classes both HPD intervals include 0. We also find that the estimates of the long-run average 1-year default probability, $\lambda_\infty^c \times \gamma$, typ-

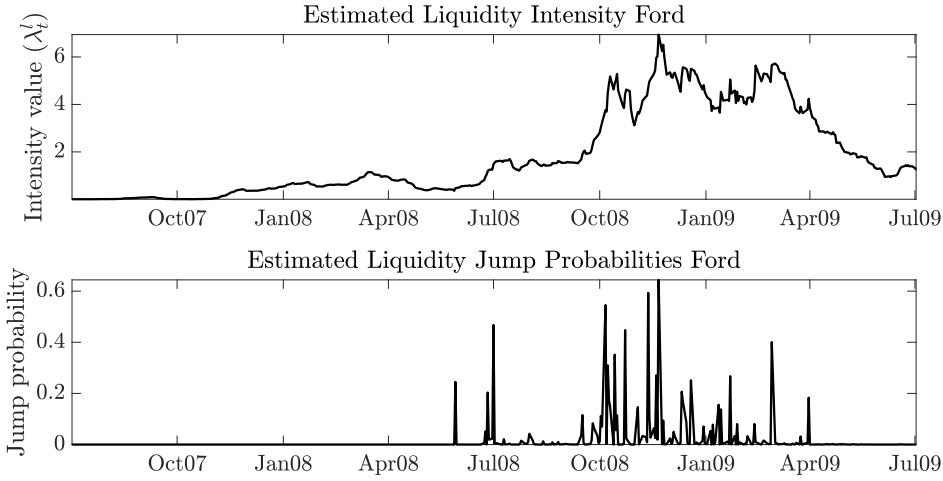


Figure 2.4. Estimated liquidity intensities and liquidity jump probabilities Ford. The upper panel plots the estimated liquidity intensities over the period July 2007 until July 2009. The bottom panel displays the estimated probabilities of liquidity jump occurrences.

ically increase as the credit rating worsens.[18] Furthermore, the estimates of σ_c^2 and σ_l^2 indicate that the credit and liquidity jump intensities are more volatile for bonds with a lower credit rating. Finally, the estimated values of the parameters φ_{high} and φ_{low} are very similar across the different rating portfolios and also relate well to the estimates in the Ford case study. This suggests that the potential liquidity jumps indicated by our bid-ask spread estimates are roughly as informative for the rating portfolios as in the Ford case study.

Figures 2.5 and 2.6 compare the estimated liquidity and credit intensities for the different rating classes with the respective credit default swap indices and estimated relative bid-ask spreads.[19] Just like in the Ford case study, in view of their match, the estimated credit and liquidity inten-

18 Since we estimate the model separately for each rating class, we do not explicitly take into account the possibility of future rating migrations in estimating our parameters. As a consequence, the reported parameter estimates for λ_∞^c need to be interpreted with care, because they might pick up some of the price effects of possible future rating changes.

19 In comparing the credit intensities, we only focus on the BB and B and lower rating classes, because these are the only rating classes for which adequate credit default swap indices are available.

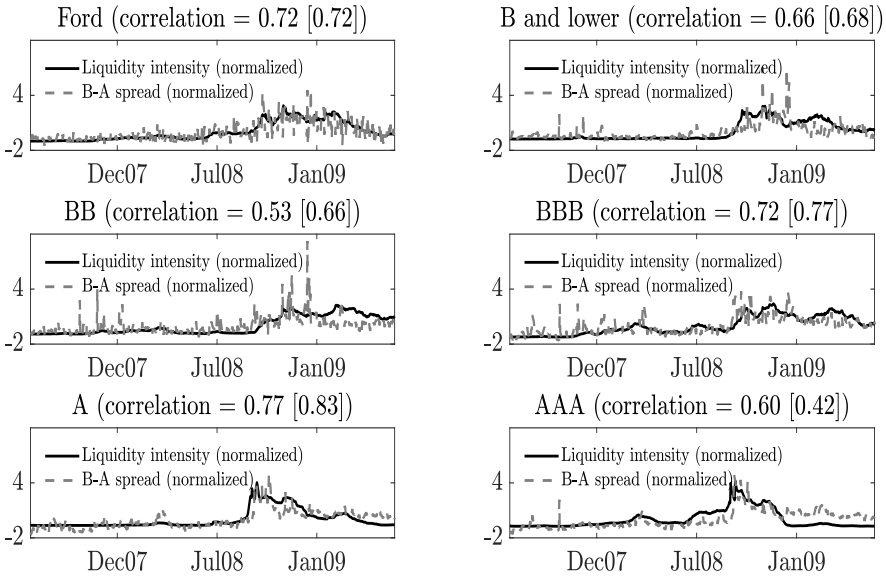


Figure 2.5. Liquidity identification. The figure plots the estimated liquidity intensities and corresponding estimated relative bid-ask spreads for Ford and all rating classes. To enhance comparability, the original series are normalized to series with zero mean and unit standard deviation by first subtracting the series' mean from the observations, and subsequently dividing by the original series' standard deviation. For each panel, the corresponding linear and Spearman's rank (in square brackets) correlations between the (normalized) liquidity intensities and (normalized) relative bid-ask spreads are reported.

sities appear to be meaningful representations of credit and liquidity risk. In estimating the model for the rating portfolios, we are not assessing firm-specific credit and liquidity risk, but systematic credit and liquidity risk for bonds in a certain rating class. Of course, not all firms are affected similarly by systematic shocks.[20] Furthermore, differences in jump times across rating classes are likely to be caused by flight-to-quality and flight-to-liquidity effects.

20 Therefore, jump identification is more tenuous. This is why we take slightly lower boundaries for estimating credit and liquidity jumps than in the Ford case study. More specifically, we set the credit and liquidity jump boundaries to 0.2. These boundaries are chosen such that we identify at least one credit and liquidity jump for all rating classes. For the B and lower rating class we still use a jump boundary of 0.5 for the credit jumps, as previously.

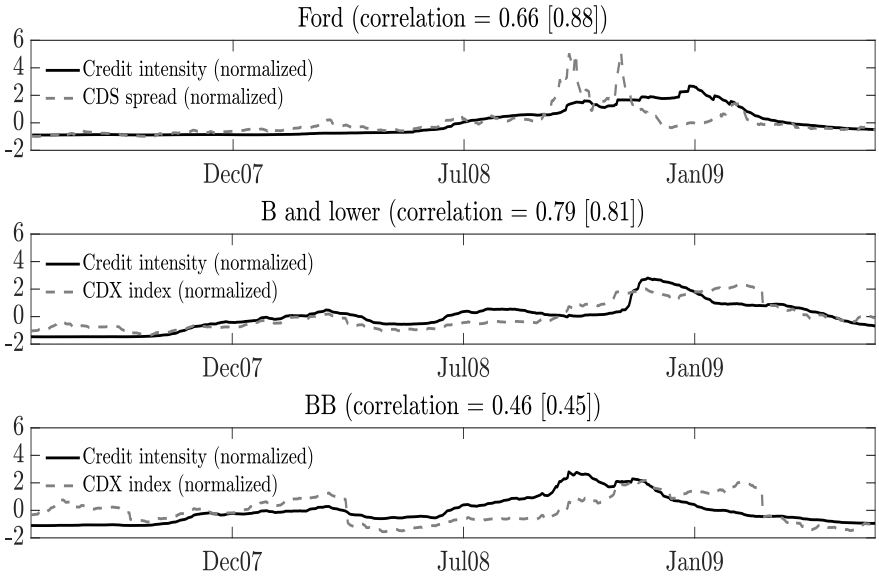


Figure 2.6. Credit identification. The figure compares the estimated credit intensities with market-implied measures of credit risk. To enhance comparability, all the original series are normalized to series with zero mean and unit standard deviation by first subtracting the series' mean from the observations, and subsequently dividing by the original series' standard deviation. For Ford (upper panel) we use the (normalized) 10-year CDS spread as a market-implied measure of credit risk. For the B and lower (middle panel) and BB (bottom panel) rating classes, we use the (normalized) 5-year maturity Markit CDS indices (CDX) of high yield firms with a B and BB rating, respectively, as market-implied measures of credit risk. For each panel, the corresponding linear and Spearman's rank (in square brackets) correlations between the (normalized) credit intensities and (normalized) CDS/CDX data are reported.

To assess the fit of our model we next compute relative pricing error diagnostics per rating portfolio. We find that our model captures the bond price data very well. The excellent model fit highlights one of the major advantages of our reduced-form model of dynamic credit-liquidity interactions. In particular, we obtain relative pricing errors of only 5.18% and 4.87% for BBB- and BB-rated bonds, respectively, over the crisis period that we analyze. Chen et al. (2017) in a structural credit-liquidity risk model that is calibrated over a longer time period report

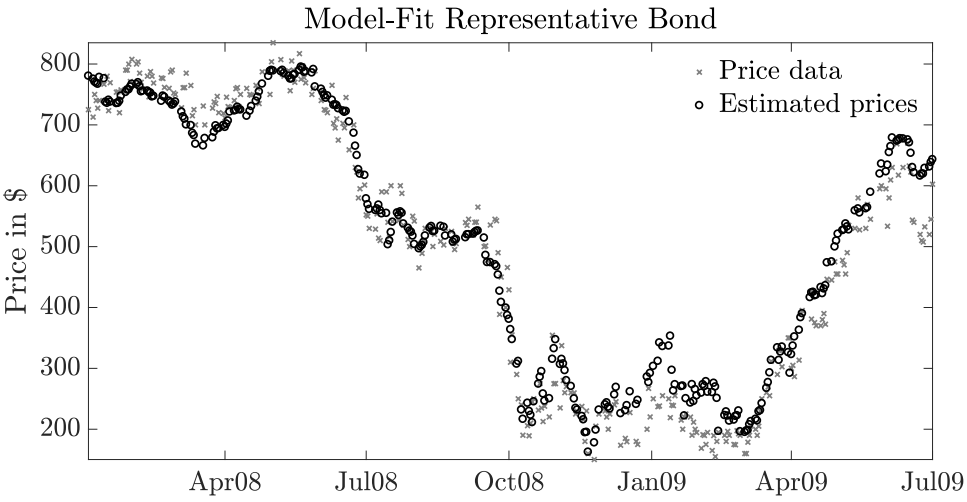


Figure 2.7. Estimated prices representative Ford bond. The figure plots the actual and model-implied price paths of a representative bond of Ford.

relative pricing errors of over 100% for Baa and Ba-rated bonds in the crisis period.

2.5.3 *Risk premia*

While the focus of this chapter is not on risk premia, we briefly discuss them in this subsection. Our framework allows for a general model specification in which risk premia related to changes in credit and liquidity risk can in principle be included. This would go by first specifying the $\mathbb{P}$-dynamics of the credit and liquidity intensity processes; next using these dynamics as state transition equations in our estimation procedure; and finally calculating the risk premia from the differences between the resulting $\mathbb{P}$- and $\mathbb{Q}$-parameters. However, given the relatively short sample period we focus on, it is impossible to estimate these risk premia with good precision.

While the default risk premia alluded to above are beyond our scope, in light of our jump model, it is interesting to briefly explore default *event* risk premia related to the (timing of the) credit events themselves. Following Jarrow et al. (2005), we assume $\lambda_t^{c,\mathbb{Q}} = \mu\lambda_t^{c,\mathbb{P}}$, with μ the default event risk premium and $\lambda_t^{c,\mathbb{P}}$ the credit intensity process under the

physical measure $\mathbb{P}$. We estimate μ by minimizing the distance between model-implied and historical conditional default probabilities, across all rating classes.[21] Using historical default rates per rating class obtained from the S&P rating agency, we find an estimated value of μ equal to 5.81. This value is roughly consistent with, but slightly larger than, estimates typically reported in the literature (see, e.g., Driessen, 2005; Berndt et al., 2018); the slightly larger estimate we find can be explained by the fact that, compared to these papers, we consider only a crisis period.

2.6 MODEL IMPLICATIONS

In this section, we illustrate some implications of our model and estimation results. First, we show how our model gives rise to a natural decomposition of yield spreads into a pure credit, a pure liquidity, a credit-induced liquidity, and a liquidity-induced credit component. Next, we analyze implications of incorporating dynamic credit-liquidity interactions in standard risk management procedures based on Value-at-Risk computations.

2.6.1 *Credit-liquidity decompositions*

The bond literature typically decomposes yield spreads into a credit and a liquidity component, in an additive way. Such a decomposition ignores the important role played by credit and liquidity interactions.[22] We propose a yield spread decomposition that nests the common additive credit-liquidity decomposition, but further subdivides the credit and liquidity components into pure credit, liquidity-induced credit, pure liquidity, and credit-induced liquidity components, similar to He and Milbradt (2014) and Chen et al. (2017). Contrary to these structural modeling approaches, which rely on calibration, our reduced-form model and

21 More specifically, we employ the estimation methodology described in Driessen (2005).

22 Policies aimed at improving market liquidity, for example, might not only affect yield spreads in a direct way by improving liquidity for a given amount of default risk, but also in an indirect way by lowering default risk. Similarly, policies aimed at improving a firm's solvency, might influence yield spreads not only by lowering default risk, but also by a corresponding improvement in market liquidity. Taking into account the indirect effects is thus crucial in designing effective policy measures and risk management strategies.

associated Bayesian MCMC estimation methodology enable us to provide a more formal econometric investigation of the contribution of the interaction terms to yield spreads.

Throughout, we consider (model-based) bond yield spreads, ys, with respect to the corresponding Treasury yields. The decomposition scheme is then given by

$$ys = \underbrace{ys_{pureCRED} + ys_{LIQ\to CRED}}_{\text{Credit Component } ys_{CRED}} + \underbrace{ys_{pureLIQ} + ys_{CRED\to LIQ}}_{\text{Liquidity Component } ys_{LIQ}}. \quad (2.12)$$

To achieve (2.12), we first decompose yield spreads into a credit and a liquidity component. Specifically, we first evaluate the credit component, ys_{CRED}, by computing yield spreads that would occur if there was no liquidity discounting. We can compute these yield spreads by setting the parameters governing the liquidity intensity process equal to zero, and by using the estimated credit intensities and credit parameters of the full model. Next, we define the liquidity component as the difference between the total yield spread and the credit component, i.e., $ys_{LIQ} = ys - ys_{CRED}$, as is standard in the literature (see, e.g., Longstaff et al., 2005; Chen et al., 2017).

We note that, even though there is no liquidity discount in computing ys_{CRED}, the credit intensities within our model, and hence ys_{CRED}, do contain liquidity influences through the cross-excitation parameter $\beta_{1,2}$. Similarly, in computing ys_{LIQ}, there is no default component, but the liquidity intensities within our model do contain the spillover effects from credit risk through the cross-excitation parameter $\beta_{2,1}$. Precisely these features make possible a further decomposition of the credit (liquidity) component into pure credit (liquidity) and liquidity-induced credit (credit-induced liquidity) subcomponents.

To this end, we first decompose the credit intensity into a component attributable to pure credit risk and a component attributable to liquidity risk, i.e., $\lambda^c_t = \lambda^c_{pureCRED,t} + \lambda^c_{LIQ\to CRED,t}$. We compute $\lambda^c_{LIQ\to CRED,t}$ using the fact that, within our model, each time there is a liquidity jump, the credit intensity jumps up by $\beta_{1,2}$. Hence, we can compute at any point in time the total cumulative effect of liquidity jumps on the credit intensity, based on the estimated liquidity jump times and taking into account the decay over time due to the mean-reverting behavior of the credit intensity process. Next, the pure credit component of the credit intensity is defined as $\lambda^c_{pureCRED,t} = \lambda^c_t - \lambda^c_{LIQ\to CRED,t}$. The first panel

of Figure 2.8 illustrates the decomposition of λ_t^c into $\lambda^c_{\text{pureCRED},t}$ and $\lambda^c_{\text{LIQ}\to\text{CRED},t}$ in the case of Ford. We observe that, since $\beta_{1,2}$ is relatively small and there is only a limited number of liquidity jumps, the liquidity-induced credit component is small relative to the pure credit component.

Having decomposed the credit intensity into a pure credit and a liquidity-induced credit component, we define the pure credit yield spread component to be the yield spread induced by the restricted model driven only by the pure credit intensity part $\lambda^c_{\text{pureCRED},t}$, and ignoring the liquidity intensity and liquidity intensity parameters. Furthermore, the liquidity-driven credit component is defined as the difference between the total credit component and the pure credit component, i.e., $\text{ys}_{\text{LIQ}\to\text{CRED}} = \text{ys}_{\text{CRED}} - \text{ys}_{\text{pureCRED}}$. Similarly, we derive the pure and credit-induced liquidity components of the yield spread by first decomposing the liquidity intensity into a pure liquidity part and a credit-induced liquidity part. The second panel of Figure 2.8 plots the decomposition of λ_t^l into the pure liquidity part, $\lambda^l_{\text{pureLIQ},t}$, and the credit-induced liquidity part, $\lambda^l_{\text{CRED}\to\text{LIQ},t}$, in the case of Ford. Both components appear to be prominent. Next, to compute the credit-induced liquidity component of the yield spread, we use only the credit-induced liquidity intensity values in computing the liquidity discount. Finally, the pure liquidity yield spread is given by the difference between the total liquidity component and the credit-induced liquidity component, i.e., $\text{ys}_{\text{pureLIQ}} = \text{ys}_{\text{LIQ}} - \text{ys}_{\text{CRED}\to\text{LIQ}}$.

2.6.1.1 *Decompositions over the cross-section and over time*

We first explore the yield spread decompositions across the different rating classes and analyze their evolution over time. This sheds light on what drove yield spread changes during the financial crisis. We consider model-based prices of 10-year zero coupon bonds and compute, for each day in our sample, the yield spread decomposition by implementing the decomposition scheme described above. Table 2.4 summarizes the cross-sectional decomposition results and displays the average contribution of the different components to the total yield spreads, both in absolute and relative terms, over the whole sample period and over the period September 15, 2008 to March 31, 2009, covering the peak of the crisis.

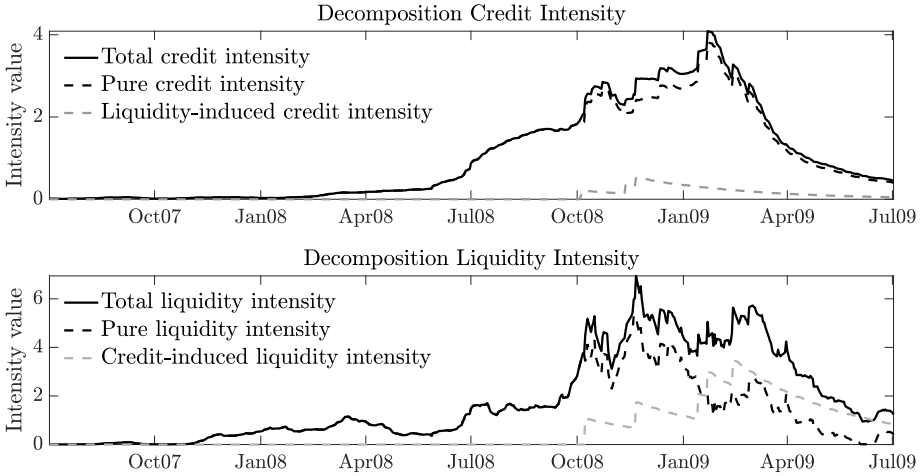


Figure 2.8. Decomposition estimated intensities Ford. The top panel plots the estimated credit intensities, the pure credit intensities, and the liquidity-induced credit intensities for the Ford sample. The bottom panel plots the estimated liquidity intensities, the pure liquidity intensities, and the credit-induced liquidity intensities for the Ford sample.

We find that, in absolute terms, the pure credit, pure liquidity, and credit-induced liquidity components are all increasing as the credit rating worsens. The liquidity-induced credit component is small, nearly constant across rating classes, and economically limited. Furthermore, we find that the yield spread contributions of all components increase in absolute terms at the peak of the crisis, when compared to the average decomposition over the full sample period. This effect is again particularly pronounced for the lower credit rating classes. Table 2.4 also shows that the increase in yield spreads at the peak of the crisis was mainly the result of deteriorating liquidity.[23] Our decomposition reveals, however, that a substantial part of this effect is actually credit-driven. Indeed, the credit-induced liquidity component is already large on average for the BB and B and lower rating classes, and becomes even more pronounced at the peak of the crisis. For example, for B and lower-rated bonds, the absolute credit-induced liquidity component is 0.73 percentage points

23 This is in line with the empirical literature; see, e.g., Bao et al., 2011; Dick-Nielsen et al., 2012; Friewald et al., 2012.

Table 2.4. Average yield spread decomposition. This table reports the average contribution of the different components to the 10-year bond yield spreads of Ford and the rating portfolios over the whole sample period (panel A) and the period September 15, 2008 until March 31, 2009, reflecting the peak of the crisis (panel B). We report both the average absolute contribution in percentage points as well as the relative contribution to the total credit spread (in parentheses).

Panel A: Full sample period					
	$ys_{pureCRED}$	$ys_{LIQ \to CRED}$	$ys_{pureLIQ}$	$ys_{CRED \to LIQ}$	ys
AAA/AA	1.20 (60.4%)	0.02 (1.0%)	0.30 (13.3%)	0.50 (25.3%)	2.04
A	1.38 (60.2%)	0.05 (1.6%)	0.49 (15.1%)	0.53 (23.1%)	2.46
BBB	1.40 (52.3%)	0.03 (0.7%)	0.85 (26.2%)	0.54 (20.7%)	2.81
BB	2.42 (63.4%)	0.02 (0.3%)	1.00 (21.3%)	0.60 (15.1%)	4.03
B and lower	3.93 (74.2%)	0.03 (0.4%)	0.91 (12.8%)	0.73 (12.6%)	5.60
Ford	2.43 (28.6%)	0.02 (0.1%)	4.10 (41.6%)	3.10 (29.7%)	9.64
Panel B: Peak of crisis					
AAA/AA	1.33 (55.7%)	0.07 (3.0%)	0.52 (18.9%)	0.54 (22.5%)	2.47
A	1.45 (41.8%)	0.16 (4.5%)	1.42 (37.0%)	0.59 (16.7%)	3.62
BBB	1.76 (45.0%)	0.06 (1.5%)	1.52 (37.6%)	0.62 (16.0%)	3.96
BB	2.70 (51.6%)	0.05 (0.8%)	1.71 (31.3%)	0.86 (16.3%)	5.35
B and lower	4.37 (55.3%)	0.07 (0.9%)	2.43 (29.5%)	1.14 (14.3%)	8.02
Ford	2.89 (18.7%)	0.05 (0.3%)	7.04 (45.3%)	5.74 (35.7%)	15.73

on average over the whole sample period, but 1.14 percentage points at the peak of the crisis.

The decomposition results in Table 2.4 make explicit that Ford is not a typical firm within the B and lower rating class. The total yield spreads for Ford are much larger than for the average B and lower-rated bonds. In particular, the liquidity component is much larger, indicating that the liquidity risk of Ford bonds was larger than that of the average B and lower firm. This is consistent with our bid-ask spread estimates (see Table 2.2). Table 2.4 also shows that the liquidity-induced credit risk component is particularly small for Ford, which is in line with the fact that Ford did not roll over any debt during our sample period. The credit-induced liquidity component, on the other hand, is very large making up over 30% of the yield spreads.

Next to the average decompositions, our model set-up also allows us to consider the evolution of the decompositions over time. To illustrate, we plot this evolution, in absolute and relative terms, for 10-year bonds of Ford and B and lower-rated firms in Figures 2.9 and 2.10. For both Ford and B and lower-rated firms, we see that liquidity dried up at the peak of the crisis. In particular, the credit-induced liquidity component became large, and at some days accounted for over 60% of Ford's total yield spreads. This illustrates that in turbulent times the interaction terms can be even much more pronounced than the average decomposition results of Table 2.4 would suggest.

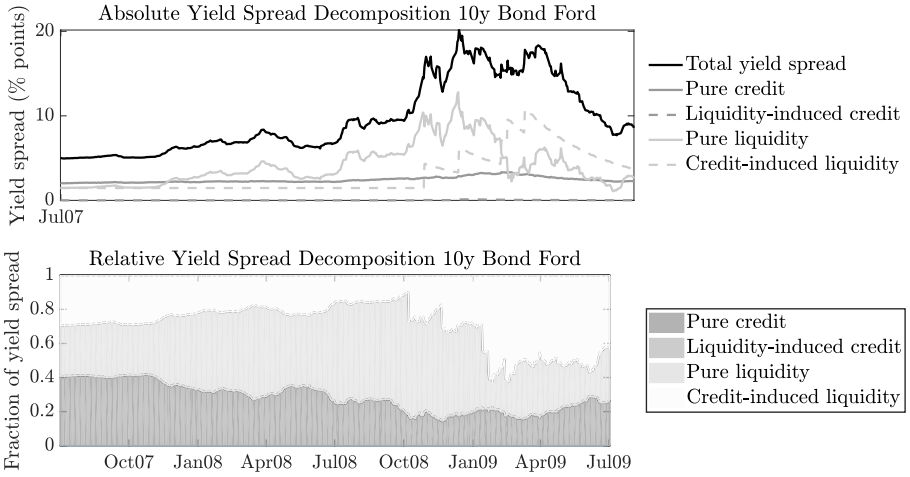


Figure 2.9. Decomposition of 10y Ford bond yield spread. The upper panel of the figure plots the absolute contribution of the pure credit, pure liquidity, credit-driven liquidity, and liquidity-driven credit components to the 10-year Ford bond yield spread over time. The lower panel displays the relative size of the different components in terms of the total yield spread over time. Note that the liquidity-driven credit component is small and therefore hardly visible.

It is interesting to compare our decomposition results with those of Chen et al. (2017).[24] We both find the liquidity-induced credit compo-

24 We compare our results with Chen et al. (2017) instead of He and Milbradt (2014), because the former extend the analysis of the latter, and claim to have more reliable results. We focus on their results in "bad states" of the economy, since they label our complete sample period as a period in which the economy is in a bad state. In comparing the results, we consider the results for bonds with a maturity of 10 years, as Chen et al. (2017) do not consider the decomposition in the term structure dimension.

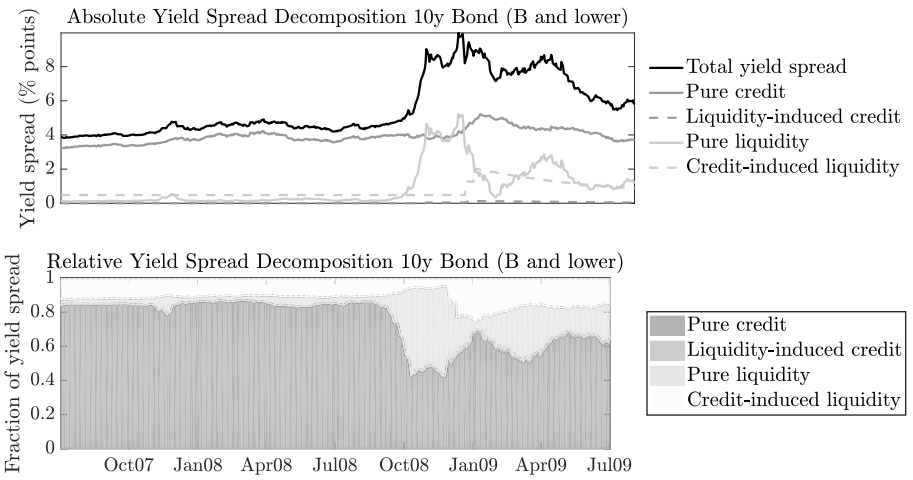


Figure 2.10. Decomposition of 10y B and lower bond yield spread. The upper panel of the figure plots the absolute contribution of the pure credit, pure liquidity, credit-driven liquidity, and liquidity-driven credit components to the 10-year B and lower-rated bond yield spread over time. The lower panel displays the relative size of the different components in terms of the total yield spread over time. Note that the liquidity-driven credit component is small and therefore hardly visible.

nent to be by far the smallest of the four components. Compared to Chen et al. (2017), we typically find this component to be even slightly smaller, but the differences with their results are only a few basis points. The differences between our decomposition results and theirs are more pronounced in the liquidity components. First, we typically find larger credit-induced liquidity components. Second, our results reveal that the pure liquidity component is highly variable over time and differs substantially across rating classes, whereas Chen et al. (2017) keep this component fixed both over time and across ratings.

2.6.1.2 *Decomposition across maturities*

In addition to the evolution of the decompositions over time for given maturities, we also investigate the term structures of the different components that make up the decompositions. Focusing again on Ford and B and lower-rated bonds, Figures 2.11 and 2.12 plot the term structures

of the absolute decompositions on different dates in the sample. More specifically, we display the term structures of the different components on July 2, 2007, representing the start of our sample period, and on three dates during the period that we labeled as 'the peak of the financial crisis': September 15, 2008, which is the default date of Lehman Brothers, December 31, 2008, and March 31, 2009.

Figures 2.11 and 2.12 show that at the start of our sample period the term structures of yield spread components are relatively flat for Ford and slightly decreasing for B and lower-rated bonds. At the peak of the crisis, however, all components are much larger and have clearly downwards sloping term structures. The negative slope of the term structures comes from the mean-reverting behavior of the intensity processes, which dampens the effects of shocks in the long run. Therefore, the yield impact (in absolute terms) of shocks on long maturity bonds is smaller than for short maturity bonds. We clearly see from the figures that credit and liquidity conditions worsen as the crisis develops, and that especially the credit-induced liquidity component becomes more prevalent. The liquidity-induced credit component is negligible for Ford. For the B and lower-rated bonds, the liquidity-induced credit component is economically relevant only for short maturity bonds at the peak of the crisis. For example, this component contributes for approximately 1 percentage point to 3-year bond yield spreads on December 31st, 2008. In relative terms, this amounts to 5% of the total yield spread. In general, the credit-liquidity interaction terms matter most at the peak of the crisis and for short maturity bonds.

2.6.2 *Implications for Value-at-Risk*

Our results also have important implications for standard risk management procedures such as Value-at-Risk (VaR) computations. The self- and cross-excitation mechanisms induce a bond return distribution that has a fatter tail in the loss domain compared to distributions implied by conventional diffusion models or jump-diffusion models with i.i.d. shocks. This is so because not only do both credit and liquidity shocks in our model have an adverse price impact, but their occurrences also make it more likely that additional price shocks will follow. Thus, our dynamic feedback mechanism generates pronounced clusters where credit and

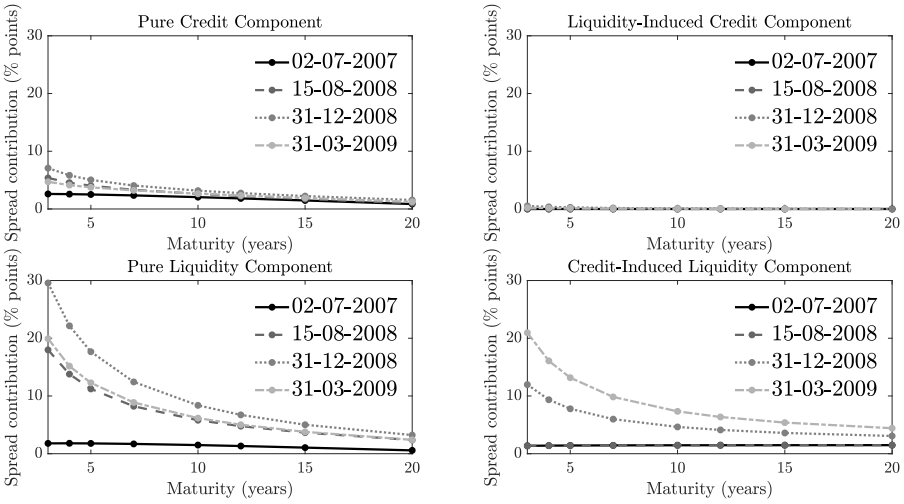


Figure 2.11. Absolute term structure per component for Ford. The figure plots the absolute term structures for the pure credit, pure liquidity, credit-driven liquidity, and liquidity-driven credit components of Ford bond yield spreads on the start date of our sample (02-07-2007) and on different dates in the period of the sample that we defined as the peak of the crisis (i.e., September 15, 2008 until March 31, 2009).

liquidity shocks accumulate, leading to fat and skew tails of the loss distribution.

To illustrate the effects of the self- and cross-excitation components on VaR capital requirements, we compute the 10-day 99% VaR of a hypothetical 10-year semi-annual coupon bond with a coupon rate of 6.64% (which agrees with the Ford sample average) under three different model specifications: (i) the full model, (ii) the model without cross-excitation ($\beta_{1,2} = \beta_{2,1} = 0$), and (iii) the model without self- and cross-excitation (i.e., $\beta_{1,1} = \beta_{1,2} = \beta_{2,1} = \beta_{2,2} = 0$). For each model specification, we simulate 10,000 10-day sample paths of the intensity and counting processes using the long-term intensity values λ^c_∞ and λ^l_∞ as starting points, and compute the implied 10-day return distribution of the bond. We then take the first percentile of the simulated return distribution to obtain the 99% VaR.

Figure 2.13 displays the resulting return distributions and 99% VaR levels for B and lower-rated bonds. It is clearly visibly that the full model has a much fatter tail in the loss domain than the restricted models. The

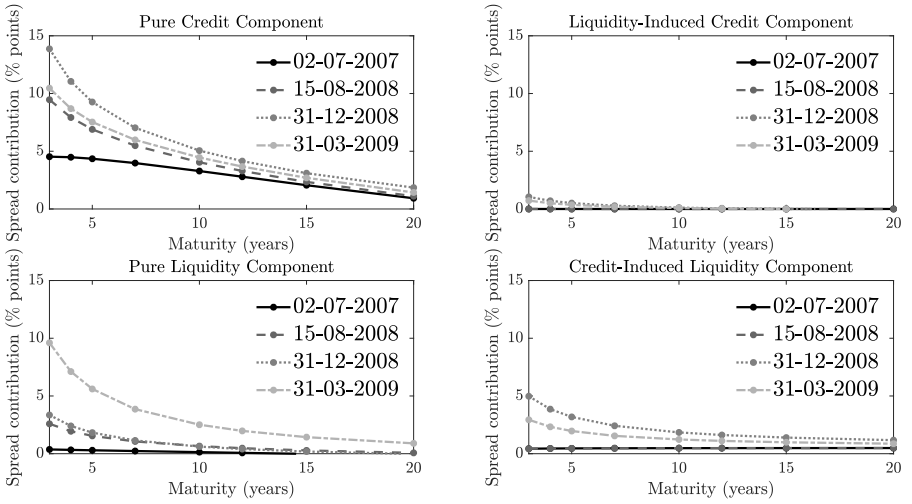


Figure 2.12. Absolute term structure per component for B and lower-rated bonds. The figure plots the absolute term structures for the pure credit, pure liquidity, credit-driven liquidity, and liquidity-driven credit components of B and lower-rated bond yield spreads on the start date of our sample (02-07-2007) and on different dates in the period of the sample that we defined as the peak of the crisis (i.e., September 15, 2008 until March 31, 2009).

model without cross-excitation in turn has a fatter tail than the model without excitation. Table 2.5 reports the 10-day 99% VaR capital requirements under the different model specifications for both Ford and the rating portfolios. We compute the amount of required capital in cents per dollar invested in the respective bonds. For example, the VaR of B and lower-rated bonds implied by the full model is -6.2%, while the VaR implied by the model without cross-excitation is -5.2%. Therefore, if one would ignore the interactions between credit and liquidity risk, the VaR capital requirement would be $6.2 - 5.2 = 1.0$ cents per dollar too low. An alternative interpretation is that 16% of the VaR can be attributed to credit-liquidity interactions. In a similar fashion, the capital requirement resulting from the model without excitation is 4.6 cents per dollar invested. This implies that ignoring the self-exciting components results in an additional underestimation of capital of $5.2 - 4.6 = 0.6$ cents per dollar invested. As can be seen from Table 2.5, the cross-excitation channel affects the VaR capital requirements the most for lower-rated bonds, because the differences between the full model's VaR and the

VaR resulting from the model without cross-excitation become larger as the credit rating deteriorates.

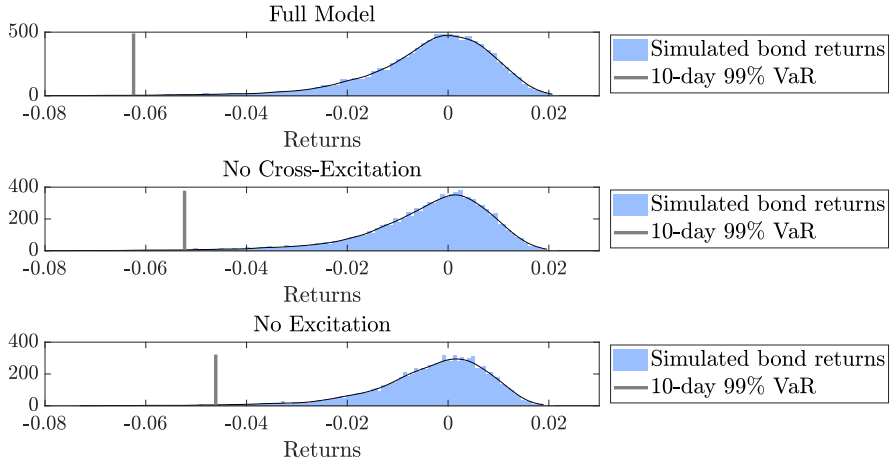


Figure 2.13. Simulated 10-day bond return distributions B and lower-rated bonds. This figure plots the simulated 10-day bond returns of B and lower-rated bonds using the full model (upper panel), the model without cross-excitation (middle panel), and the model without cross- and self-excitation (bottom panel). In every panel, the first percentile (i.e., the 10-day 99% VaR) of the simulated return distribution is indicated by a vertical line.

Table 2.5. VaR capital requirements per dollar invested. This table reports the VaR capital requirement in cents per dollar invested in a 10-year semi-annual coupon bond based on different model specifications. For every rating class/-Ford, we simulate 10,000 10-day paths of the intensity and jump processes using the parameter estimates of the corresponding rating class/Ford and taking the long-run intensity values λ_∞^c and λ_∞^l as starting points.

	Full model	No cross-excitation	No excitation
AAA/AA	1.3	1.2	1.0
A	1.7	1.6	1.3
BBB	1.4	1.3	1.2
BB	1.9	1.7	1.4
B and lower	6.2	5.2	4.6
Ford	5.7	4.6	3.9

2.7 CONCLUSIONS

We have designed a framework to analyze the dynamic interactions between credit and liquidity risk in the corporate bond market. We have proposed a novel way of modeling these interactions through self- and cross-exciting processes. The interaction mechanism gives rise to dynamic feedback loops between credit and liquidity risk. We have developed a corresponding Bayesian estimation procedure, based on bond price formulae that we have derived in closed form, and using bond transaction data.

In our empirical analysis, we have considered US bond portfolios sorted by rating and a case study on Ford Motor Company. Our results have shown that our model captures the data very well. We have documented strong evidence of both self- and cross-exciting behavior of credit and liquidity risk, have shown that the cross-excitation effects are asymmetric, and tend to be stronger for bonds with a lower credit rating. Our model makes possible a decomposition of bond yield spreads into a pure credit, a pure liquidity, a credit-induced liquidity, and a liquidity-induced credit component. We have revealed that the widening of bond yield spreads during the financial crisis is mainly caused by worsening liquidity conditions, which, in turn, are for a large part the consequence of deteriorating credit conditions. The credit-induced liquidity compo-

nent accounts for 0.50 (AAA/AA rating) to 0.73 (B and lower rating) percentage points of total 10-year spreads on average, but in the most turbulent times, it accounts for up to 2.05 percentage points. The reverse effect of liquidity on credit risk is typically smaller, but still makes up a non-trivial part of the yield spreads for lower-rated short maturity bonds at the peak of the crisis. We have illustrated that our results also have important implications for risk management. In particular, ignoring the dynamic credit-liquidity interactions results in a significant underestimation of risk, leading to capital requirements that are substantially too low especially for bonds with lower credit ratings.

2.A CONSTRUCTION OF THE LIQUIDITY JUMP INDICATOR

As in the main text, let β_t be the true, but unobserved, bid-ask spread at day t which is assumed to follow an AR(1) process:

$$\beta_t = \phi\beta_{t-1} + \eta_t,$$

with $0 \leqslant \phi \leqslant 1$ and where η_t is an i.i.d. error term with zero mean and finite variance. In the following, we assume that ϕ is known. In typical higher frequency applications, for example with daily data, ϕ will usually be close to one. The bid-ask spread is estimated with error u_t according to

$$\hat{\beta}_t = \beta_t + u_t.$$

Now let $y_t = \hat{\beta}_t - \phi\hat{\beta}_{t-1}$, so that

$$y_t = \eta_t + u_t - \phi u_{t-1}.$$

The variance and first-order auto-covariance of y_t are given by

$$\mathrm{Var}(y_t) = \sigma_\eta^2 + (1+\phi^2)\sigma_u^2, \quad \mathrm{Cov}(y_t, y_{t-1}) = -\phi\sigma_u^2,$$

and let us denote $\rho = -\mathrm{Cov}(y_t, y_{t-1})/\mathrm{Var}(y_t)$.

We aim to obtain a best estimate of η_t given observations of the series $\hat{\beta}_t$ or, equivalently, of y_t. Leads and lags of y_t also contain information about η_t as they are correlated with y_t. We can use a regression setup to estimate η_t, using two leads and lags of y_t, as follows. Let $\mathcal{Y}_t = (y_{t-2}, y_{t-1}, y_t, y_{t+1}, y_{t+2})$. Then, the best estimate of η_t given $\mathcal{Y}_t$ is

$$\mathbb{E}[\eta_t|\mathcal{Y}_t] = \mathrm{Cov}(\eta_t, \mathcal{Y}_t)\mathrm{Var}(\mathcal{Y}_t)^{-1}\mathcal{Y}_t'.$$

Notice that

$$\mathrm{Var}(y_t) = \begin{pmatrix} 1 & -\rho & 0 & 0 & 0 \\ -\rho & 1 & -\rho & 0 & 0 \\ 0 & -\rho & 1 & -\rho & 0 \\ 0 & 0 & -\rho & 1 & -\rho \\ 0 & 0 & 0 & -\rho & 1 \end{pmatrix} \mathrm{Var}(y_t),$$

and

$$\mathrm{Cov}(\eta_t, y_t) = (0, 0, \sigma_\eta^2, 0, 0).$$

The inverse of the variance matrix is given by

$$\mathrm{Var}(y_t)^{-1} = \frac{1}{\mathrm{Var}(y_t)(1-2\rho^2)} \begin{pmatrix} 1-\rho^2 & \rho & \rho^2 & \rho^3 & \rho^4 \\ \rho & 1 & \rho & \rho^2 & \rho^3 \\ \rho^2 & \rho & 1 & \rho & \rho^2 \\ \rho^3 & \rho^2 & \rho & 1 & \rho \\ \rho^4 & \rho^3 & \rho^2 & \rho & 1-\rho^2 \end{pmatrix}.$$

Hence, we find the best estimate of η_t as

$$\mathbb{E}[\eta_t | y_t] = \frac{\sigma_\eta^2}{\mathrm{Var}(y_t)(1-2\rho^2)} (\rho^2 y_{t-2} + \rho y_{t-1} + y_t + \rho y_{t+1} + \rho^2 y_{t+2}).$$

Finally, we need an estimator of σ_η^2. Following De Jong and Schotman (2010), we use the non-parametric estimator

$$\hat{\sigma}_\eta^2 = \sum_{k=-2}^{2} \mathrm{Cov}(y_t, y_{t-k}).$$

Collecting terms and denoting $\rho_k = \mathrm{Cov}(y_t, y_{t-k})/\mathrm{Var}(y_t)$, we finally obtain

$$\hat{\eta}_t = \frac{\sum_{k=-2}^{2} \rho_i}{1-2\rho^2} (\rho^2 y_{t-2} + \rho y_{t-1} + y_t + \rho y_{t+1} + \rho^2 y_{t+2}),$$

with $\rho = -\rho_1$.

2.B MONTE CARLO SIMULATION STUDY

To assess the performance of our estimation methodology, we have conducted a Monte Carlo simulation study, which we describe in more detail in this Appendix. In each simulation run, we construct a dataset of bond prices by first simulating the jump intensities and counting processes, next applying bond price formula (2.5) to these series, and finally adding normally distributed measurement errors to the resulting (log) prices. In simulating the state vector series, we use realistic parameter values that mimic the values obtained in the empirical analysis, and we approximate the continuous-time dynamics with an Euler discretization scheme using a very fine time grid. We choose the bond characteristics such that each simulated sample closely reflects the actual Ford data. More specifically, each simulated sample consists of daily observations of 23 semi-annual coupon bonds with coupon rates equal to the sample average of 6.64% and maturities of $3, 4, 5, \ldots, 25$ years.

In our Monte Carlo analysis, we focus on two nested model specifications, which we believe to provide good insight in the performance of our estimation procedure. The first model specification we consider, is one in which we abstract away from liquidity issues and assume there is only credit risk. In this univariate model specification, we do not have a liquidity discount factor, and, hence, no liquidity jumps. In line with our full model, we assume that the credit intensity process contains self-exciting jumps. We first take this univariate model specification to the empirical data and use the thus obtained parameter estimates as input in the Monte Carlo experiment. The univariate model is only used as an auxiliary model in this Monte Carlo study, and we therefore do not discuss its empirical results in much detail; however, we do want to mention that its model fit is poor (relative pricing errors are around 20%) compared to the multivariate model, which supports the use of a multivariate model that includes liquidity risk, and credit-liquidity interactions. In the second model specification, we employ the full model, but we assume that the liquidity jump indicator is fully informative about liquidity jumps. That is, we assume that liquidity jump times are known, and, hence, we do not need to simulate the bid-ask spread data explicitly. We use the model parameters as obtained in the empirical analysis described in Section 2.5 as input in the Monte Carlo study.

For both model specifications, we conduct 100 simulations consisting of 252 daily observations (i.e., one business year) of our bond panel and

apply our estimation procedure to each of the constructed datasets. The results for the univariate model can be found in Table 2.6 and the results for the multivariate model are in Table 2.7. The results confirm that our estimation methodology yields good performance in a realistic setting designed to mimic the features of the empirical analysis. Using our estimation methodology, the population parameters of our model can be recovered with high degree of precision. This applies in particular also to the mutual excitation parameters $\beta_{ij}, i, j \in \{1, 2\}$, which are central to our analysis.

Table 2.6. Monte Carlo results univariate model. This table reports the results of the Monte Carlo study using the univariate model specification with self-exciting credit intensity process following the dynamics $d\lambda_t^c = \alpha_c(\lambda_\infty^c - \lambda_t^c)dt + \sigma_c\sqrt{\lambda_t^c}dW_t^c + \beta_{1,1}dN_t^c$. The study consists of 100 simulated bond panels with maturities 3,4,...,25 years, each containing 252 observations (i.e., 1 year of simulated data at daily frequency assuming 252 business days per year). The true parameter values, Monte Carlo sample means, standard deviations, and minimum and maximum values are reported in separate columns.

	True	MC mean	MC std. dev.	Min	Max
α_c	1.75	1.81	0.08	1.60	1.94
λ_∞^c	0.55	0.56	0.03	0.46	0.62
σ_c^2	1.85	1.81	0.10	1.67	2.02
$\beta_{1,1}$	1.70	1.73	0.08	1.55	1.90
γ	0.10	0.10	0.01	0.08	0.15
h	0.01	0.01	0.00	0.01	0.01

Table 2.7. Monte Carlo results multivariate model. This table reports the results of the Monte Carlo study using the bivariate model specification with mutually exciting credit and liquidity intensity processes with the following dynamics: $d\lambda_t^c = \alpha_c(\lambda_\infty^c - \lambda_t^c)dt + \sigma_c\sqrt{\lambda_t^c}dW_t^c + \beta_{1,1}dN_t^c + \beta_{1,2}dN_t^l$ and $d\lambda_t^l = \alpha_l(\lambda_\infty^l - \lambda_t^l)dt + \sigma_l\sqrt{\lambda_t^l}dW_t^l + \beta_{2,1}dN_t^c + \beta_{2,2}dN_t^l$. The study consists of 100 simulated bond panels with maturities 3,4,...,25 years, each containing 252 observations (i.e., 1 year of simulated data at daily frequency assuming 252 business days per year). Throughout, we assume that liquidity jump times are known, and we keep $\lambda_\infty^l = 0.015$ fixed. The true parameter values, Monte Carlo sample means, standard deviations, and minimum and maximum values are reported in separate columns.

	True	MC mean	MC std. dev.	Min	Max
α_c	3.90	4.10	0.30	3.19	4.78
λ_∞^c	0.13	0.12	0.01	0.10	0.16
σ_c^2	0.98	0.90	0.10	0.66	1.20
γ	0.37	0.46	0.07	0.17	0.60
α_l	3.72	2.73	0.77	1.23	5.66
σ_l^2	2.50	2.61	0.29	1.88	3.39
$\beta_{1,1}$	0.41	0.39	0.13	0.04	0.73
$\beta_{1,2}$	0.21	0.26	0.12	0.05	0.63
$\beta_{2,1}$	1.07	1.06	0.09	0.81	1.29
$\beta_{2,2}$	0.38	0.37	0.07	0.11	0.55
h	0.05	0.05	0.00	0.05	0.06

2.C CLOSED-FORM PRICING FORMULAS

In this Appendix, we show how to compute the expectations appearing in (2.5) in closed-form, up to the solution of a system of ODEs, by exploiting the fact that the state vector consisting of the jump counting and jump intensity processes falls into the class of affine jump-diffusions of Duffie et al. (2000) in its generalized version as defined in their Appendix B (see also Aït-Sahalia et al., 2014; Errais et al., 2010). A gener-

alized affine jump-diffusion process X_t in a state space $D \subset \mathbb{R}^{2\times m}$ is defined as the solution to the stochastic differential equation

$$dX_t = \mu^X(X_t)dt + \sigma^X(X_t)dW_t^X + \sum_{j=1}^{m} dJ_{j,t}, \tag{2.C.1}$$

where $\mu^X : D \to \mathbb{R}^{2\times m}$, $\sigma^X : D \to \mathbb{R}^{(2\times m)\times(2\times m)}$, W^X is a Brownian motion in $\mathbb{R}^{2\times m}$, J_i, $i = 1, \ldots, m$, are pure jump processes with jump intensities $\lambda_{i,t}^X = \lambda_i^X(X_t)$ for some $\lambda_i^X : D \to [0, \infty)$ and with fixed jump size distribution ν^i with jump transforms $\theta^i(c) = \int_{\mathbb{R}^n} \exp(c \cdot z)\, d\nu(z)$ for $c \in \mathbb{C}^n$, and where μ^X, $\sigma^X\sigma^{X\prime}$ and λ_i^X are affine on D:

$$\begin{aligned}
\mu^X &= K_0 + K_1 x, \text{ for } K = (K_0, K_1) \in \mathbb{R}^{2\times m} \times \mathbb{R}^{(2\times m)\times(2\times m)},\\
\left(\sigma^X\sigma^{X\prime}\right)_{ij} &= (H_0)_{ij} + (H_1)_{ij} \cdot x, \quad \text{for } H_0 \in \mathbb{R}^{(2\times m)\times(2\times m)}\\
&\qquad \text{and } H_1 \in \mathbb{R}^{(2\times m)\times(2\times m)\times(2\times m)}\\
\lambda_i^X &= l_0 + l_1 \cdot x, \text{ for } l = (l_0, l_1) \in \mathbb{R} \times \mathbb{R}^{(2\times m)}.
\end{aligned}$$

Now let $X_t = (N_t^c, N_t^l, \lambda_t^c, \lambda_t^l)'$. Then we have that de dynamics of X are given by

$$\begin{aligned}
dX_t &= d\begin{pmatrix} N_t^c \\ N_t^l \\ \lambda_t^c \\ \lambda_t^l \end{pmatrix}\\
&= \begin{pmatrix} 0 \\ 0 \\ \alpha_c(\lambda_\infty^c - \lambda_t^c) \\ \alpha_l(\lambda_\infty^l - \lambda_t^l) \end{pmatrix} dt + \begin{pmatrix} 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & \sigma_c\sqrt{\lambda_t^c} & 0 \\ 0 & 0 & 0 & \sigma_l\sqrt{\lambda_t^l} \end{pmatrix} d\begin{pmatrix} 0 \\ 0 \\ W_t^c \\ W_t^l \end{pmatrix}\\
&\quad + \begin{pmatrix} 1 \\ 0 \\ \beta_{1,1} \\ \beta_{2,1} \end{pmatrix} dN_t^c + \begin{pmatrix} 0 \\ 1 \\ \beta_{1,2} \\ \beta_{2,2} \end{pmatrix} dN_t^l.
\end{aligned} \tag{2.C.2}$$

From this specification it is clear that the joint state vector X is affine, since we can write $n = 2 \times 2 = 4$, $m = 2$, $K_0 = (0, 0, \alpha_c \lambda_\infty^c, \alpha_l \lambda_\infty^l)^\top$,

$$K_1 = \begin{pmatrix} 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & -\alpha_c & 0 \\ 0 & 0 & 0 & -\alpha_l \end{pmatrix},$$

$H_0 = \mathbf{0}$,

$$H_1^1 = H_1^2 = \begin{pmatrix} 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix},$$

$$H_1^3 = \begin{pmatrix} 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & \sigma_c^2 & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix}, \quad H_1^4 = \begin{pmatrix} 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & \sigma_l^2 \end{pmatrix},$$

$l_0^c = l_0^l = 0$, $l_1^c = (0, 0, 1, 0)^\top$, and $l_1^l = (0, 0, 0, 1)^\top$. Furthermore, because N_t^c and N_t^l are counting processes and the coefficients $\beta_{i,j}$ for $i, j = 1, 2$ are constants, we have that $\theta^c(c) = \exp\left(c_1 + \beta_{1,1} c_3 + \beta_{2,1} c_4\right)$ and $\theta^l(c) = \exp\left(c_2 + \beta_{1,2} c_3 + \beta_{2,2} c_4\right)$.

Using results of Duffie et al. (2000) for the class of affine jump-diffusions in its generalized version defined in their Appendix B, we can now find closed-form expressions for the transform $u \mapsto \mathbb{E}\left[e^{-\int_t^T R(X_s)ds} e^{u \cdot X_T} \middle| \mathcal{F}_t\right]$ and extended transform $(u, v) \mapsto \mathbb{E}\left[e^{-\int_t^T R(X_s)ds} (v \cdot X_T) e^{u \cdot X_T} \middle| \mathcal{F}_t\right]$, where $R(x) = \rho_0 + \rho_1 \cdot x$ for $(\rho_0, \rho_1) \in \mathbb{R} \times \mathbb{R}^{2 \times m}$. We state the following proposition:

Proposition 1.
$\mathbb{E}^{\mathbb{Q}}\left[e^{-\int_t^T \lambda_s^l ds}(1-\gamma)^{N_T^c}\middle|\mathcal{F}_t\right] = e^{\alpha(t)+\beta_1(t)N_t^c+\beta_2(t)N_t^l+\beta_3(t)\lambda_t^c+\beta_4(t)\lambda_t^l}$,

with

$$\begin{aligned}
\dot{\alpha}(t) &= -\alpha_c\lambda_\infty^c\beta_3(t) - \alpha_l\lambda_\infty^l\beta_4(t),\\
\dot{\beta}_1(t) &= \dot{\beta}_2(t) = 0,\\
\dot{\beta}_3(t) &= \alpha_c\beta_3(t) - \frac{1}{2}(\sigma_c^2\beta_3^2(t)) - \left(e^{\beta_1(t)+\beta_{1,1}\beta_3(t)+\beta_{2,1}\beta_4(t)} - 1\right),\\
\dot{\beta}_4(t) &= 1 + \alpha_l\beta_4(t) - \frac{1}{2}(\sigma_l^2\beta_4^2(t)) - \left(e^{\beta_2(t)+\beta_{1,2}\beta_3(t)+\beta_{2,2}\beta_4(t)} - 1\right),\\
\alpha(T) &= \beta_2(T) = \beta_2(s) = \beta_3(T) = \beta_4(T) = 0,\\
\beta_1(T) &= \log(1-\gamma) = \beta_1(s), \quad t \leqslant s \leqslant T.
\end{aligned}$$

Proof. Consider an affine jump-diffusion process X in some state space $D \subset \mathbb{R}^n$ solving the stochastic differential equation

$$dX_t = \mu(X_t)dt + \sigma(X_t)dW_t + \sum_{i=1}^{m} dZ_t^i,$$

where Z^i are pure jump processes whose jumps have a fixed probability distribution ν^i on $\mathbb{R}^n$ and arrive with intensity $\lambda^i(X_t)$ for some $\lambda^i : D \to [0,\infty)$. Let us fix an affine process $R : D \to \mathbb{R}$. Then we have that the complete affine structure of the model is captured by:

$$\begin{aligned}
\mu(x) &= K_0 + K_1 x, \text{ for } K = (K_0, K_1) \in \mathbb{R}^n \times \mathbb{R}^{n\times n}.\\
\sigma(x)\sigma(x)^\top &= H_0 + \sum_{k=1}^{n} H_1^{(k)} x_k, \text{ for } H = (H_0, H_1) \in \mathbb{R}^{n\times n} \times \mathbb{R}^{n\times n\times n}.\\
\lambda^i(x) &= l_0^i + l_1^i \cdot x, \text{ for } l = (l_0, l_1) \in \mathbb{R} \times \mathbb{R}^n.\\
R(x) &= \rho_0 + \rho_1 \cdot x, \text{ for } \rho = (\rho_0, \rho_1) \in \mathbb{R} \times \mathbb{R}^n.
\end{aligned}$$

Let us furthermore denote the jump-transforms, which determine the jump-size distributions, as $\theta^i(c) = \int_{\mathbb{R}^n} \exp(c \cdot z)\, d\nu^i(z)$ for $c \in \mathbb{C}^n$.

We want to compute an expression of the form

$$\phi^X(u, X, t, T) = \mathbb{E}^X\left[\exp\left(-\int_t^T R(X_s)ds\right) e^{u\cdot X_T}\middle|\mathcal{F}_t\right].$$

According to Proposition 1 of Duffie et al. (2000), we have, under some technical assumptions on the processes being well-behaved, that we can write

$$\phi^X(u, x, t, T) = e^{\alpha(t)+\beta(t)\cdot x},$$

where β and α satisfy the following (complex-valued) ODEs:

$$\dot{\beta}(t) = \rho_1 - K_1^\top \beta(t) - \frac{1}{2}\beta(t)^\top H_1 \beta(t) - \sum_{i=1}^{m} l_1^i (\theta^i(\beta(t)) - 1)$$

$$\dot{\alpha}(t) = \rho_0 - K_0 \cdot \beta_t - \frac{1}{2}\beta(t)^\top H_0 \beta(t) - \sum_{i=1}^{m} l_0^i (\theta^i(\beta(t)) - 1),$$

with boundary conditions $\beta(T) = u$ and $\alpha(T) = 0$.

Since we aim at computing $\mathbb{E}^Q\left[e^{-\int_t^T \lambda_s^l ds}(1-\gamma)^{N_T^c}\middle|\mathcal{F}_t\right]$, we have that $\rho_0 = 0$, $\rho_1 = (0,0,0,1)^\top$ and $u = (\log(1-\gamma),0,0,0)^\top$. Together with the system matrices, K_0, K_1, H_0, H_1, l_0^c, l_0^l, l_1^c, l_1^l, $\theta^c(c)$, and $\theta^l(c)$ as specified above, the desired result thus follows by applying Proposition 1 of Duffie et al. (2000). □

In a similar way, we can find a system of ODEs for $\mathbb{E}^Q\left[\lambda_T^c \gamma e^{-\int_t^T \lambda_u^l du}(1-\gamma)^{N_T^c}\middle|\mathcal{F}_t\right]$:

Proposition 2.

$$\mathbb{E}^{\mathbb{Q}}\left[\lambda_T^c \gamma e^{-\int_t^T \lambda_u^l du}(1-\gamma)^{N_T^c}\middle|\mathcal{F}_t\right] = e^{\alpha(t)+\log(1-\gamma)N_t^c+\beta_3(t)\lambda_t^c+\beta_4(t)\lambda_t^l} \times \left(A(t)+B_3(t)\lambda_t^c+B_4(t)\lambda_t^l\right),$$

with

$$\begin{aligned}
-\dot{A}(t) &= \alpha_c\lambda_\infty^c B_3(t)+\alpha_l\lambda_\infty^l B_4(t),\\
-\dot{B}_3(t) &= -\alpha_c B_3(t)+\beta_3(t)\sigma_c^2 B_3(t)\\
&\quad +\beta_{1,1}e^{\beta_1(t)+\beta_{1,1}\beta_3(t)+\beta_{2,1}\beta_4(t)}B_3(t)\\
&\quad +\beta_{2,1}e^{\beta_1(t)+\beta_{1,1}\beta_3(t)+\beta_{2,1}\beta_4(t)}B_4(t),\\
-\dot{B}_4(t) &= -\alpha_l B_4(t)+\beta_4(t)\sigma_l^2 B_4(t)\\
&\quad +\beta_{1,2}e^{\beta_2(t)+\beta_{1,2}\beta_3(t)+\beta_{2,2}\beta_4(t)}B_3(t)\\
&\quad +\beta_{2,2}e^{\beta_2(t)+\beta_{1,2}\beta_3(t)+\beta_{2,2}\beta_4(t)}B_4(t),\\
\alpha(T) &= \beta_2(T)=\beta_2(s)=\beta_3(T)=\beta_4(T)=A(T)=B_1(s)\\
&= B_2(s)=B_4(T)=0, \quad t\leqslant s\leqslant T,\\
\beta_1(T) &= \log(1-\gamma)=\beta_1(s), \quad t\leqslant s\leqslant T,\\
B_3(T) &= \gamma,
\end{aligned}$$

where $\alpha(t)$ *and* $\beta(t)$ *satisfy the ODEs presented in Proposition 1.*

Proof. Proposition 3 of Duffie et al. (2000) with $u=(\log(1-\gamma),0,0,0)^\top$ and $v=(0,0,\gamma,0)^\top$ yields the desired result. □

2.D PRICE CORRECTIONS FOR THE RATING PORTFOLIOS

Since not all bonds within the same rating class are equally credit risky, we correct for bond-specific fixed effects, which measure the differences in baseline credit-worthiness. A simple fixed effects correction would, however, also correct for coupon and maturity differences between bonds, because these differences are also constant over time. Fortunately, our pricing formulae use maturity and coupon rate as input, and are thus able to deal with these differences without correcting for them. Therefore, we propose to correct prices for bond-specific fixed effects within groups obtained by double-sorting on maturity and coupon rates. That is, we collect bonds with similar maturity and coupon rates in the same

bucket and apply a fixed effects correction per bucket. Since we update the rating portfolios every quarter, we apply the price corrections per rating class, per quarter.

For each rating class and quarter, we create nine different buckets, consisting of different combinations of maturity and coupon rates (low, middle, high, based on lower, middle, and upper 33th percentiles), and apply the following fixed effects price correction per bucket: Let the log-price of bond j be given by

$$Y^0_{t,j} = \alpha_j + f(X_t, \chi_j, \Theta) + \epsilon_{t,j}, \tag{2.D.3}$$

where α_j is the bond-specific fixed effect, f our log-price function depending on the states X_t at time t, the bond characteristics χ_j (maturity, coupon rate, etc.), and the parameters Θ, and $\epsilon_{t,j}$ an i.i.d. normally distributed measurement error. We have that the time series average of the log-price of bond j is given by

$$\frac{1}{T}\sum_{t=1}^{T} Y^0_{t,j} \approx \alpha_j + \frac{1}{T}\sum_{t=1}^{T} f(X_t, \chi_j, \Theta). \tag{2.D.4}$$

Denoting $\overline{Y^0_{t,j}} = \frac{1}{T}\sum_{t=1}^{T} Y^0_{t,j}$ and $\overline{f(X_t, \chi_j, \Theta)} = \frac{1}{T}\sum_{t=1}^{T} f(X_t, \chi_j, \Theta)$, we then get

$$\frac{1}{N}\sum_{j=1}^{N} \overline{Y^0_{t,j}} \approx \frac{1}{N}\sum_{j=1}^{N} \alpha_j + \overline{f(X_t, \chi_j, \Theta)} = \overline{f(X_t, \chi_j, \Theta)}, \tag{2.D.5}$$

where N is the number of bonds in the same maturity-coupon bucket. The last equality follows from $\sum_{j=1}^{N} \alpha_j = 0$ and the fact that we assume $\chi_i = \chi_j \; \forall i, j$ in this bucket. This gives the following estimator for α_j:

$$\hat{\alpha_j} = \overline{Y^0_{t,j}} - \frac{1}{N}\sum_{j=1}^{N} \overline{Y^0_{t,j}}.$$

All-in-all, we obtain fixed effects corrected prices $Y^0_{t,j} - \hat{\alpha_j}$, and employ these prices to estimate our model.

The above procedure sometimes produces an overly strong price correction when the price development of a bond in a certain quarter is not comparable to the other bonds in the same maturity-coupon bucket. For example, if a bond defaults or gets downgraded, its prices for the remainder of the quarter will typically be much lower than those of the

average bond in the same bucket. In such cases, the fixed effects correction produces adapted prices that are much higher than the original price, since we subtract the relatively low bond's time series average and add the relatively high cross-sectional bucket average to the original transaction prices. In these cases, the price correction may be so large that we should not consider these transactions representative for that rating class and maturity-coupon bucket. Therefore, when the price correction is more than 50% away from the original price, we discard these observations. This results in discarding 0, 140, 340, 630, and 2459 transactions for the AAA/AA, A, BBB, BB, and B and lower-rated portfolios, respectively.

2.E DETAILS OF THE METROPOLIS-HASTINGS STEPS

In this Appendix, we provide more details on the Metropolis-Hastings steps conducted in our estimation methodology. We first discuss the drawing of new states and next consider the drawing of new parameters.

2.E.1 *Metropolis steps for drawing the states*

As explained in the main text, we alternate between two schemes for simulating the jump intensities and jump counters. In the first scheme, we draw at each transaction date t the credit jump, the liquidity jump, and the intensities in consecutive steps. In the second scheme, we draw at each transaction date t the whole state vector X_t in one step. Below we provide the details of the Metropolis steps used in both schemes.

State simulation scheme 1

In the first scheme, we take advantage of the fact that $p(X_t|X^{(g)}_{1:N_{\setminus t}}, \bar{\Theta}^{(g)}, \bar{Y})$ is characterized by its full conditionals. Therefore, we can split the drawing of $X_t^{(g+1)}$ into the following three steps:

1. draw $N_t^{c(g+1)}$ from $p(N_t^c|\lambda_t^{c(g)}, \lambda_t^{l(g)}, N_t^{l(g)}, X^{(g)}_{1:N_{\setminus t}}, \bar{\Theta}^{(g)}, \bar{Y})$;
2. draw $N_t^{l(g+1)}$ from $p(N_t^l|\lambda_t^{c(g)}, \lambda_t^{l(g)}, N_t^{c(g+1)}, X^{(g)}_{1:N_{\setminus t}}, \bar{\Theta}^{(g)}, \bar{Y})$;

3. draw $\lambda_t^{c(g+1)}$ and $\lambda_t^{l(g+1)}$ from
$p(\lambda_t^c, \lambda_t^l | N_t^{c(g+1)}, N_t^{l(g+1)}, X_{1:N_{\setminus t}}^{(g)}, \bar{\Theta}^{(g)}, \bar{Y})$.

Under the adopted discretization scheme, we have that the full posterior of N_t^c is a Bernoulli density with success probability

$$p(\Delta N_t^c = 1 | \lambda_t^{c(g)}, \lambda_t^{l(g)}, N_t^{l(g)}, X_{1:N_{\setminus t}}^{(g)}, \bar{\Theta}^{(g)}, \bar{Y}) = \frac{p(\lambda_t^{c(g)}, \lambda_t^{l(g)} | \Delta N_t^c = 1, N_t^{l(g)}, X_{t-1}^{(g+1)}, \bar{\Theta}^{(g)}) p(\Delta N_t^c = 1 | X_{t-1}^{(g+1)})}{\sum_{s=0,1} p(\lambda_t^{c(g)}, \lambda_t^{l(g)} | \Delta N_t^c = s, N_t^{l(g)}, X_{t-1}^{(g+1)}, \bar{\Theta}^{(g)}) p(\Delta N_t^c = s | X_{t-1}^{(g+1)})},$$

which is easy to compute, because
$p(\lambda_t^{c(g)}, \lambda_t^{l(g)} | \Delta N_t^c = s, N_t^{l(g)}, X_{t-1}^{(g+1)}, \bar{\Theta}^{(g)})$ is a bivariate normal density and $p(\Delta N_t^c = s | X_{t-1}^{(g+1)})$ is Bernoulli with success probability $\lambda_{t-1}^{c(g+1)} \Delta$. The credit jump probability does not depend on $\bar{Y}$, since the bond prices only depend on the intensities and parameter vector, and the liquidity jump indicator I_t^l only depends on N_t^l and the parameters and not on N_t^c.

Similarly, the full posterior of $N_{t_i}^l$ is Bernoulli with success probability given by

$$p(\Delta N_t^l = 1 | \lambda_t^{c(g)}, \lambda_t^{l(g)}, N_t^{c(g+1)}, X_{1:N_{\setminus t}}^{(g)}, \bar{\Theta}^{(g)}, \bar{Y}) = \frac{p(I_t^l | \Delta N_t^l = 1, \bar{\Theta}) p(\lambda_t^{c(g)}, \lambda_t^{l(g)} | \Delta N_t^l = 1, N_t^{c(g+1)}, X_{t-1}^{(g+1)}, \bar{\Theta}^{(g)}) p(\Delta N_t^l = 1 | X_{t-1}^{(g+1)})}{\sum_{s=0,1} p(I_t^l | \Delta N_t^l = s, \bar{\Theta}) p(\lambda_t^{c(g)}, \lambda_t^{l(g)} | \Delta N_t^l = s, N_t^{c(g+1)}, X_{t-1}^{(g+1)}, \bar{\Theta}^{(g)}) p(\Delta N_t^l = s | X_{t-1}^{(g+1)})},$$

which is again easy to compute, because
$p(\lambda_t^{c(g)}, \lambda_t^{l(g)} | \Delta N_t^l = s, N_t^{c(g+1)}, X_{t-1}^{(g+1)}, \bar{\Theta}^{(g)})$ is bivariate normal,
$p(I_t^l | \Delta N_t^l = s, \bar{\Theta})$ is Bernoulli with success probability φ_{high} if $\Delta N_t^l = 1$ and φ_{low} if $\Delta N_t^l = 0$, and $p(\Delta N_t^l = s | X_{t-1}^{(g+1)})$ is Bernoulli with success probability $\lambda_{t-1}^{l(g+1)} \Delta$. Note that the posterior liquidity jump probability depends on $\bar{Y}$, since N_t^l appears in the density of I_t^l.

After having drawn $N_t^{c(g+1)}$ and $N_t^{l(g+1)}$, we draw the new intensities $\lambda_t^{c(g+1)}$ and $\lambda_t^{l(g+1)}$ from the density

$$\begin{aligned}
&p(\lambda_t^c, \lambda_t^l | N_t^{c(g+1)}, N_t^{l(g+1)}, X_{1:N_{\setminus t}}^{(g)}, \bar{\Theta}^{(g)}, \bar{Y}) \\
&\propto p(Y_t | \lambda_t^c, \lambda_t^l, N_t^{c(g+1)}, N_t^{l(g+1)}, \bar{\Theta}^{(g)}) \\
&\times p(\lambda_{t+1}^{c(g)}, \lambda_{t+1}^{l(g)} | \lambda_t^c, \lambda_t^l, N_{t+1}^{c(g)}, N_{t+1}^{l(g)}, \bar{\Theta}^{(g)}) \\
&\qquad \times p(\lambda_t^c, \lambda_t^l | \lambda_{t-1}^{c(g+1)}, \lambda_{t-1}^{l(g+1)}, N_t^{c(g+1)}, N_t^{l(g+1)}, \bar{\Theta}^{(g)}),
\end{aligned}$$

which is the product of a multivariate normal density and two bivariate normal densities. In this density, I^l and φ_{high} and φ_{low} do not play a role, since they do not depend on the intensity values. This density is non-standard and therefore we use a Metropolis step with proposal density

$$\begin{aligned}
&q(\lambda_t^c, \lambda_t^l | N_t^{c(g+1)}, N_t^{l(g+1)}, X_{1:N_{\setminus t}}^{(g)}, \bar{\Theta}^{(g)}, \bar{Y}) \\
&\qquad = p(\lambda_t^c, \lambda_t^l | \lambda_{t-1}^{c(g+1)}, \lambda_{t-1}^{l(g+1)}, N_t^{c(g+1)}, N_t^{l(g+1)}, \bar{\Theta}^{(g)}),
\end{aligned}$$

which corresponds to a bivariate normal distribution with mean

$$\begin{pmatrix} \lambda_{t-1}^{c(g+1)} + \alpha_c^{(g)} \lambda_\infty^{c(g)} \Delta - \alpha_c^{(g)} \lambda_{t-1}^{c(g+1)} \Delta + \beta_{1,1}^{(g)} \Delta N_t^{c(g+1)} + \beta_{1,2}^{(g)} \Delta N_t^{l(g+1)} \\ \lambda_{t-1}^{l(g+1)} + \alpha_l^{(g)} \lambda_\infty^{l(g)} \Delta - \alpha_l^{(g)} \lambda_{t-1}^{l(g+1)} \Delta + \beta_{2,1}^{(g)} \Delta N_t^{c(g+1)} + \beta_{2,2}^{(g)} \Delta N_t^{l(g+1)} \end{pmatrix} \tag{2.E.6}$$

and variance matrix

$$\begin{pmatrix} \sigma_c^{2(g)} \lambda_{t-1}^{c(g+1)} \Delta & 0 \\ 0 & \sigma_l^{2(g)} \lambda_{t-1}^{l(g+1)} \Delta \end{pmatrix}. \tag{2.E.7}$$

Using this proposal density, the acceptance criterion becomes

$$\min \left(\frac{p(Y_t | X_t^{(g+1)}, \bar{\Theta}^{(g)}) p(\lambda_{t+1}^{c(g)}, \lambda_{t+1}^{l(g)} | N_{t+1}^{c(g)}, N_{t+1}^{l(g)}, X_t^{(g+1)}, \bar{\Theta}^{(g)})}{p(Y_t | X_t^{(g)}, \bar{\Theta}^{(g)}) p(\lambda_{t+1}^{c(g)}, \lambda_{t+1}^{l(g)} | N_{t+1}^{c(g)}, N_{t+1}^{l(g)}, X_t^{(g)}, \bar{\Theta}^{(g)})}, 1 \right).$$

For the end point $t = N$ the acceptance criterion simplifies, since the terms $p(\lambda_{t+1}^{c(g)}, \lambda_{t+1}^{l(g)} | N_{t+1}^{c(g)}, N_{t+1}^{l(g)}, X_t^{(g+1)}, \bar{\Theta}^{(g)})$ and

$p(\lambda_{t+1}^{c(g)}, \lambda_{t+1}^{l(g)} | N_{t+1}^{c(g)}, N_{t+1}^{l(g)}, X_t^{(g)}, \bar{\Theta}^{(g)})$ do not appear anymore in the numerator and denominator, respectively. For the starting point $t = 1$, we use a slightly different proposal density, since we cannot condition on X_{t-1}. We therefore draw $\lambda_1^{c(g+1)}$ and $\lambda_1^{l(g+1)}$ from a bivariate normal density with mean vector $\left(\lambda_2^{c(g)}, \lambda_2^{l(g)}\right)^\top$ and variances $\sigma\sqrt{\lambda_2^{c(g)}\Delta}$ and $\sigma\sqrt{\lambda_2^{l(g)}\Delta}$, respectively.

State simulation scheme 2

In the second simulation scheme we sample the complete vector X_t at once from $p(X_t | X_{1:N_{\setminus t}}, \bar{\Theta}, \bar{Y})$. From (2.10) we have

$$p(X_t | X_{1:N_{\setminus t}}^{(g-1)}, \bar{\Theta}, \bar{Y}) \propto p(I_t^l | X_t, \bar{\Theta}) p(Y_t | X_t, \bar{\Theta}) p(X_t | X_{t-1}^{(g)}, \bar{\Theta}) p(X_{t+1}^{(g-1)} | X_t, \bar{\Theta}), \tag{2.E.8}$$

where we exploit the Markovian property of the state processes, and, hence, only need to consider the terms in (2.10) where (parts of) X_t enters directly. We observe that in this density, both the likelihoods as well as the transition densities play a role and this allows us to circumvent the problems encountered in the sampling procedure above.

We showed in Eqn. (2.11) that the transition density can be written as

$$\begin{aligned} p(X_t | X_{t-1}, \bar{\Theta}) &= p(\lambda_t^c, \lambda_t^l | N_t^c, N_t^l, X_{t-1}, \bar{\Theta}) p(N_t^c, N_t^l | X_{t-1}, \bar{\Theta}) \\ &= p(\lambda_t^c, \lambda_t^l | N_t^c, N_t^l, X_{t-1}, \bar{\Theta}) p(N_t^c | X_{t-1}, \bar{\Theta}) p(N_t^l | X_{t-1}, \bar{\Theta}). \end{aligned} \tag{2.E.9}$$

We have that $p(Y_t | X_t, \bar{\Theta})$ is multivariate normal with dimension equal to the number of observations at time t. More specifically, let there be k observations at time t. Then,

$$p(Y_t | X_t, \bar{\Theta}) \sim \mathcal{N}\left(\begin{pmatrix} F(X_t, \chi_1, \Theta) \\ F(X_t, \chi_2, \Theta) \\ \vdots \\ F(X_t, \chi_k, \Theta) \end{pmatrix}, \begin{pmatrix} h^2 & 0 & 0 & \cdots \\ 0 & h^2 & 0 & \cdots \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & h^2 \end{pmatrix} \right).$$

Furthermore, $p(I_t^l | N_t^l, \bar{\Theta})$ is Bernoulli with success probability φ_{high} if $\Delta N_t^l = 1$ and φ_{low} if $\Delta N_t^l = 0$.

In order to draw $X_t^{(g+1)}$ from (2.E.8), we use the following proposal density

$$\begin{aligned} q(X_t|X_{1:N_{\setminus t}}^{(g)}, \bar{\Theta}^{(g)}) &= p(X_t|X_{t-1}^{(g+1)}, \bar{\Theta}^{(g)}) \\ &= p(\lambda_t^c, \lambda_t^l|N_t^c, N_t^l, X_{t-1}^{(g+1)}, \bar{\Theta}^{(g)}) \\ &\times p(N_t^c|X_{t-1}^{(g+1)}, \bar{\Theta}^{(g)}) p(N_t^l|X_{t-1}^{(g+1)}, \bar{\Theta}^{(g)}). \end{aligned}$$

From the discretization of the intensity processes, we get that this is a mixture of bivariate normal distributions. To see this, note that $p(\lambda_t^c, \lambda_t^l|N_t^c, N_t^l, X_{t-1}^{(g+1)}, \bar{\Theta}^{(g)})$ is bivariate normal with mean

$$\begin{pmatrix} \lambda_{t-1}^{c(g+1)} + \alpha_c^{(g)} \lambda_\infty^{c(g)} \Delta - \alpha_c^{(g)} \lambda_{t-1}^{c(g+1)} \Delta + \beta_{1,1}^{(g)} \Delta N_t^c + \beta_{1,2}^{(g)} \Delta N_t^l \\ \lambda_{t-1}^{l(g+1)} + \alpha_l^{(g)} \lambda_\infty^{l(g)} \Delta - \alpha_l^{(g)} \lambda_{t-1}^{l(g+1)} \Delta + \beta_{2,1}^{(g)} \Delta N_t^c + \beta_{2,2}^{(g)} \Delta N_t^l \end{pmatrix}$$

and variance matrix

$$\begin{pmatrix} \sigma_c^{2(g)} \lambda_{t-1}^{c(g+1)} \Delta & 0 \\ 0 & \sigma_l^{2(g)} \lambda_{t-1}^{l(g+1)} \Delta \end{pmatrix}.$$

Now, depending on the four possible combinations of ΔN_t^c and ΔN_t^l (i.e., these are both 1, both 0 or one of them is 1 and the other 0), we have that
$p(\lambda_t^c, \lambda_t^l|N_t^c, N_t^l, X_{t-1}^{(g+1)}, \bar{\Theta}^{(g)}) p(N_t^c|X_{t-1}^{(g+1)}, \bar{\Theta}^{(g)}) p(N_t^l|X_{t-1}^{(g+1)}, \bar{\Theta}^{(g)})$ is a mixture of 4 (bivariate) normal distributions.

When we draw from this distribution, we can first draw $N_t^{c(g+1)}$ and $N_t^{l(g+1)}$ from independent Bernoulli distributions with success probabilities $\lambda_{t-1}^{c(g+1)} \Delta$ and $\lambda_{t-1}^{l(g+1)} \Delta$, respectively, and then, given their outcomes, draw $p(\lambda_t^c, \lambda_t^l|N_t^{c(g+1)}, N_t^{l(g+1)}, X_{t-1}^{(g+1)}, \Theta^{(g)})$ from a normal distribution with the appropriate mean (depending on outcomes of the draws of $N_t^{c(g+1)}$ and $N_t^{l(g+1)}$).
So let $X_t^{(g+1)} = \left\{ \lambda_t^{c(g+1)}, \lambda_t^{l(g+1)}, N_t^{c(g+1)}, N_t^{l(g+1)} \right\}$ denote the new draw of the intensities and jump processes obtained from

$$\begin{aligned} &p(\lambda_t^{c(g+1)}, \lambda_t^{l(g+1)}|N_t^{c(g+1)}, N_t^{l(g+1)}, X_{t-1}^{(g+1)}, \Theta^{(g)}) \\ &\times p(N_t^{c(g+1)}|X_{t-1}^{(g+1)}, \Theta^{(g)}) p(N_t^{l(g+1)}|X_{t-1}^{(g+1)}, \Theta^{(g)}). \end{aligned}$$

Then we have that the acceptance criterion simplifies, since the proposal density is a part of the true density and therefore multiple terms cancel out. More specifically, we will have the following acceptance criterion:

$$\min\left(\frac{p(I_t^{l}|N_t^{l(g+1)},\bar{\Theta})p(Y_t|X_t^{(g+1)},\Theta^{(g)})p(\lambda_{t+1}^{c(g)},\lambda_{t+1}^{l(g)},N_{t+1}^{c(g)},N_{t+1}^{l(g)}|X_t^{(g+1)},\Theta^{(g)})}{p(I_t^{l}|N_t^{l(g)},\bar{\Theta})p(Y_t|X_t^{(g)},\Theta^{(g)})p(\lambda_{t+1}^{c(g)},\lambda_{t+1}^{l(g)},N_{t+1}^{c(g)},N_{t+1}^{l(g)}|X_t^{(g)},\Theta^{(g)})},1\right),$$

where

$$\begin{aligned}
&p(\lambda_{t+1}^{c(g)},\lambda_{t+1}^{l(g)},N_{t+1}^{c(g)},N_{t+1}^{l(g)}|X_t^{(g+1)},\bar{\Theta}^{(g)})\\
&=p(\lambda_{t+1}^{c(g)},\lambda_{t+1}^{l(g)}|N_{t+1}^{c(g)},N_{t+1}^{l(g)},X_t^{(g+1)},\bar{\Theta}^{(g)})\\
&\qquad\times p(N_{t+1}^{c(g)}|X_t^{(g+1)},\bar{\Theta}^{(g)})p(N_{t+1}^{l(g)}|X_t^{(g+1)},\bar{\Theta}^{(g)}).
\end{aligned}$$

Since all densities are standard (i.e., multivariate normal, bivariate normal or Bernoulli), we can easily evaluate this acceptance criterion. For the end points $t = 1$ and $t = N$ similar comments apply as in simulation scheme 1.

2.E.2 *Metropolis steps for drawing the parameters*

We next explain the Metropolis steps for estimating the parameters in more detail. Where possible, we update the parameters using a Gibbs step with conjugate priors. Let us first consider φ_{high} and φ_{low}, for which we specify the priors to be $\text{Beta}(a_{high}, b_{high})$ and $\text{Beta}(a_{low}, b_{low})$, respectively. Let $S(\Delta N^{l(g+1)} = 1)$ denote the days for which $\Delta N_t^{l(g+1)} = 1$. Then we use standard results that a Beta prior on a Bernoulli success probability, where the Bernoulli trial is repeated several times, is again a Beta distribution with updated parameters. We get $\varphi_{high}^{(g+1)} \sim \text{Beta}(a_{high}^{posterior}, b_{high}^{posterior})$ with $a_{high}^{posterior} = a_{high} + \sum_{t\in S(\Delta N^{l(g+1)}=1)} I_t^{l}$ and $b_{high}^{posterior} = b_{high} + \sum_{t=1}^{N} \Delta N_t^{l(g+1)} - \sum_{t\in S(\Delta N^{l(g+1)}=1)} I_t^{l}$, and $\varphi_{low}^{(g+1)} \sim \text{Beta}(a_{low}^{posterior}, b_{low}^{posterior})$ with $a_{low}^{posterior} = a_{low} + \sum_{t\notin S(\Delta N^{l(g+1)}=1)} I_t^{l}$ and $b_{low}^{posterior} = b_{low} + \left(N - \sum_{t=1}^{N} \Delta N_t^{l(g+1)}\right) - \sum_{t\notin S(\Delta N^{l(g+1)}=1)} I_t^{l}$.

For h^2 we use an inverse gamma prior specification. Since h^2 only appears as variance in the measurement equations for the bonds, we get that the posterior distribution for h^2 is again inverse gamma with updated parameters. This follows from $p(h^2|\bar{\Theta}\backslash\{h^2\}, X, \bar{Y}) \propto p(Y|\bar{\Theta}, X)p(h^2)$, where $p(Y|\bar{\Theta}, X)$ is multivariate normal with diagonal variance matrix

with h^2 as variance and $p(h^2)$ the prior inverse gamma density. Explicit computations are standard and are therefore omitted. All the other parameters appear in both the transition densities as well as the bond price formulae and therefore it is not possible to specify conjugate prior specifications for them. We use Metropolis steps to update these parameters in multiple blocks, since this is computationally less demanding than updating one parameter at the time. Following Eraker (2004), we propose to use a Gaussian proposal density with as mean vector the current draw of the parameters in the block, and with a diagonal covariance matrix.

We choose the priors on all parameters to be proper but uninformative in the sense that the prior variances should be high compared to the estimated posterior variances. We use the following priors: (i) α_c, α_l, λ^c_∞, λ^l_∞, σ^2_c, and σ^2_l: Gamma(0.02,10); (ii) $\beta_{1,1}$, $\beta_{1,2}$, $\beta_{2,1}$, and $\beta_{2,2}$: Normal(0.02,10); (iii) γ: Uniform(0,1); (iv) h^2: IG$(100, 1)$; (v) φ_{high}: Beta$(5, 5)$; and (vi) φ_{low}: Beta$(1, 10)$.

In line with theoretical parameter restrictions for nonnegativity and stationarity of the processes, we impose some parameters to be nonnegative by using gamma priors. For $\beta_{i,j}$ $(i, j = 1, 2)$ we take, however, normally distributed priors. This choice does not restrict them from taking negative values, although we can easily assure their positivity by rejecting all parameter draws for which one of these parameters becomes negative. The reason for choosing normal priors is that in order to compute the Bayes factors for testing model restrictions, we need that the priors are well-defined for zero values. Since γ is the probability of going into default in the case a credit shock arrives, we put a Uniform(0,1) prior on it. In general, the means and variances of all parameters are chosen arbitrarily, but such that the means are small and positive for all parameters and the variances relatively large compared to their means. All-in-all, our results are robust against prior specification, since typically the likelihood contribution of the priors is small compared to the likelihood contribution of the data in the acceptance criteria. Furthermore, the posterior standard deviations are also much smaller than the prior standard deviations, indicating that our priors do not impose much information.

2.F BAYES FACTORS

In order to formally test for self- and cross-excitation, we use Bayes factors to compare the full model with different models in which we

put restrictions on the self- and cross-excitation parameters. We let $\mathcal{M}_1$ denote the full model, $\mathcal{M}_2$ the model with restrictions $\beta_{1,1} = \beta_{1,2} = \beta_{2,1} = \beta_{2,2} = 0$ (no excitation), $\mathcal{M}_3$ the model with restrictions $\beta_{1,2} = \beta_{2,1} = 0$ (no cross-excitation), and $\mathcal{M}_4$ the model with restrictions $\beta_{1,1} = \beta_{2,2} = 0$ (no self-excitation). The Bayes factors are given by the ratios of the two marginal likelihoods of the models under consideration. That is, if we compare model $\mathcal{M}_i$ ($i = 2, 3, 4$) with model $\mathcal{M}_1$, we have

$$\mathrm{BF}_{\mathcal{M}_i|\mathcal{M}_1} = \frac{p_{\mathcal{M}_i}(Y)}{p_{\mathcal{M}_1}(Y)}. \tag{2.F.10}$$

Let us denote the parameter vector of the full model by $\Theta = \{\theta_1, \theta_2\}$, where θ_2 are the parameters that are restricted in model $\mathcal{M}_i$. For example, when comparing model $\mathcal{M}_2$ with the full model, we have that $\theta_2 = \{\beta_{1,1}, \beta_{1,2}, \beta_{2,1}, \beta_{2,2}\}$ and the restriction is that $\theta_2 = \theta_2^*$, with $\theta_2^* = \{0, 0, 0, 0\}$. Since models $\mathcal{M}_2$, $\mathcal{M}_3$, and $\mathcal{M}_4$ are all nested in model $\mathcal{M}_1$, and since the priors on the common parameters, θ_1, are the same and independent of the restricted parameters, θ_2, we can use Savage-Dickey density ratios to compute the Bayes factors. We get

$$\mathrm{BF}_{\mathcal{M}_i|\mathcal{M}_1} = \frac{p_{\mathcal{M}_1}(\theta_2|Y)\big|_{\theta_2=\theta_2^*}}{p_{\mathcal{M}_1}(\theta_2)\big|_{\theta_2=\theta_2^*}}, \tag{2.F.11}$$

that is, the Bayes factor can be computed as the ratio of model $\mathcal{M}_1$'s marginal posterior density of θ_2 and marginal prior density of θ_2, both evaluated in θ_2^*. The computation of the denominator is simple, as we just have to compute the prior marginal density of the full model evaluated in θ_2^*.

In order to compute the numerator of (2.F.11), we follow Chib and Jeliazkov (2001), and we denote the subkernel of the Metropolis-Hastings chain for θ_2 conditional on (θ_1, X) by

$$p(\theta_2, \theta_2^*|Y, \theta_1, X) = \alpha(\theta_2, \theta_2^*|Y, \theta_1, X) q(\theta_2, \theta_2^*|Y, \theta_1, X),$$

where $q(\theta_2, \theta_2'|Y, \theta_1, X)$ is the proposal density for the transition from θ_2 to θ_2' and

$$\alpha(\theta_2, \theta_2'|Y, \theta_1, X) = \min\left(\frac{p(\theta_2'|Y, \theta_1, X) p(\theta_1, \theta_2') q(\theta_2', \theta_2|Y, \theta_1, X)}{p(\theta_2|Y, \theta_1, X) p(\theta_1, \theta_2) q(\theta_2, \theta_2'|Y, \theta_1, X)}, 1\right)$$

denotes the probability of a move. The following (local) reversibility condition holds (see Chib and Jeliazkov, 2001):

$$p(\theta_2, \theta_2^*|Y, \theta_1, X) p(\theta_2|Y, \theta_1, X) = p(\theta_2^*|Y, \theta_1, X) p(\theta_2^*, \theta_2|Y, \theta_1, X). \tag{2.F.12}$$

Multiplying both sides of (2.F.12) by $p(\theta_1, X|Y)$ and integrating over $\psi = (\theta_1, \theta_2, X)$ gives

$$\int p(\theta_2, \theta_2^*|Y, \theta_1, X)p(\theta_2|Y, \theta_1, X)p(\theta_1, X|Y)d\psi = \int p(\theta_2^*|Y, \theta_1, X)p(\theta_2^*, \theta_2|Y, \theta_1, X)p(\theta_1, X|Y)d\psi,$$

or

$$\int p(\theta_2, \theta_2^*|Y, \theta_1, X)p(\theta_1, \theta_2, X|Y)d\psi = \int \pi(\theta_2^*|Y)p(\theta_1, X|Y, \theta_2^*)p(\theta_2^*, \theta_2|Y, \theta_1, X)d\psi.$$

From this last equation and the definitions of $p(\theta_2, \theta_2^*|Y, \theta_1, X)$ and $p(\theta_2^*, \theta_2|Y, \theta_1, X)$ it follows that

$$p(\theta_2^*|Y) = \frac{\mathbb{E}_1\left[\alpha(\theta_2, \theta_2^*|Y, \theta_1, X)q(\theta_2, \theta_2^*|Y, \theta_1, X)\right]}{\mathbb{E}_2\left[\alpha(\theta_2^*, \theta_2|Y, \theta_1, X)\right]}, \tag{2.F.13}$$

where the expectation in the numerator is taken with respect to the distribution $p(\theta_1, \theta_2, X|Y)$, and the expectation in the denominator with respect to $p(\theta_1, X|Y, \theta_2^*)q(\theta_2^*, \theta_2|Y, \theta_1, X)$.

The MCMC algorithm we use to estimate our model provides draws $\{\theta_1^{(g)}, \theta_2^{(g)}, X^{(g)}\}_{g=1}^{G}$ from the full posterior density $p(\theta_1, \theta_2, X|Y)$. Averaging $\alpha(\theta_2, \theta_2^*|Y, \theta_1, X)q(\theta_2, \theta_2^*|Y, \theta_1, X)$ over these draws thus provides an estimate of the numerator of (2.F.13). For the denominator, we fix θ_2^* and continue the MCMC simulation for an additional J iterations with the two full conditional densities $p(\theta_1|Y, \theta_2^*, X)$ and $p(X|Y, \theta_1, \theta_2^*)$. At each iteration we get values $(\theta_1^{(j)}, X^{(j)})$ and we use these to generate $\theta_2^{(j)} \sim q(\theta_2^*, \theta_2|Y, \theta_1^{(j)}, X^{(j)})$. For every iteration we obtain a draw $(\theta_1^{(j)}, \theta_2^{(j)}, X^{(j)})$ from the distribution $p(\theta_1, X|Y, \theta_2^*)q(\theta_2^*, \theta_2|Y, \theta_1, X)$. We can finally estimate $p(\theta_2^*|Y)$ as

$$\hat{p}(\theta_2^*|Y) = \frac{G^{-1}\sum_{g=1}^{G}\alpha(\theta_2^{(g)}, \theta_2^*|Y, \theta_1^{(g)}, X^{(g)})q(\theta_2^{(g)}, \theta_2^*|Y, \theta_1^{(g)}, X^{(g)})}{J^{-1}\sum_{j=1}^{J}\alpha(\theta_2^*, \theta_2^{(j)}|Y, \theta_1^{(j)}, X^{(j)})}. \tag{2.F.14}$$

3 CREDIT RISK PREMIA IN SOVEREIGN CREDIT DEFAULT SWAPS

This chapter is based on:

Sperna Weiland, R. C. (2018). Credit Risk Premia in Sovereign Credit Default Swaps, Unpublished working paper. University of Amsterdam, Amsterdam.[1]

3.1 INTRODUCTION

It is well documented that there is a high degree of commonality and contagion in sovereign credit risk (see, e.g., Longstaff et al., 2011; Aït-Sahalia et al., 2014). This systemic nature of sovereign credit risk makes it plausible that investors do not only demand compensation for the risk of unexpected variations in credit spreads (hereafter referred to as the "distress risk premia"), but also for the risk of default events themselves (hereafter referred to as the "default event risk premia"). Evidence for such default event risk premia has, for example, been found in corporate bond and credit default swaps (CDS) contexts (see, e.g., Jarrow et al., 2005; Driessen, 2005; Saita, 2005; Berndt et al., 2008). In the sovereign credit risk literature, however, default event premia have largely been ignored, and a detailed empirical analysis of both distress and default event premia is still lacking.

The main aim of this chapter is therefore to investigate in more detail distress and default event risk premia embedded in sovereign CDS spreads. In order to do so, I introduce a new model for the term-structure of sovereign credit risk. In particular, I assume that sovereign defaults can be triggered by shocks in either a common factor, or an independent country-specific factor. The novelty of the model is that I specify both the common and country-specific factors to be self-exciting jump processes. In this way, the model can capture, in a parsimonious way, the

1 I am very grateful to David Lando, Yacine Aït-Sahalia, Joost Driessen, Guillaume Roussellet, Frank de Jong, Roger Laeven, and Peter Spreij, as well as numerous conference and seminar participants for useful comments and discussions.

high degree of commonality in sovereign credit risk, the clustering of large credit shocks over time and across countries, and rapid, jump-like increases in sovereign credit spreads observed in the data.

The model gives rise to a natural decomposition of sovereign CDS spreads along two dimensions, allowing for a detailed analysis of the economic importance of the different risk premia. First, I decompose CDS spreads into country-specific and systemic risk components. I find a similar decomposition across rating classes and show that approximately 65% of five-year CDS spreads can be attributed to country-specific risk and 35% of CDS spreads can be attributed to systemic risk. Second, I decompose CDS spreads into distress risk premia, default event risk premia, and expected default risk components. I find that distress risk premia are mainly relevant for countries with high credit ratings, whereas default event risk premia and expected default risk components are more important for countries with lower credit ratings. For example, distress risk premia account on average for 61% and 6% of CDS spreads of A- and B-rated countries, respectively. Default event risk premia (expected default risk components), on the other hand, account for 22% (17%) and 52% (42%) of CDS spreads of A- and B-rated countries, respectively. I show that the differences in decompositions across rating classes are mainly caused by differences in sovereign-specific risk rather than the differences (in exposure to) systemic risk. I also consider the decomposition of CDS spreads in the term-structure dimension and find that the default event and expected default risk components are more important for shorter maturities, whereas distress risk is more important for longer maturities. This suggests that investors care relatively more about default events in the short-term, whereas the uncertainty regarding future default arrival rates is more important in the long-term.

The model facilitates a multi-step estimation procedure. In a first step, I estimate the model parameters using data on the term-structure of sovereign CDS spreads of 28 geographically dispersed countries over the period 01-01-2008 until 30-12-2016. I estimate the common and country-specific factors in an iterative way. In each step, I use a Bayesian MCMC procedure similar to Sperna Weiland et al. (2019) and I find fast convergence of the parameter estimates. CDS spreads can, however, not be used to estimate default event risk premia, since they only contain information about risk-neutral default probabilities, (see, e.g., Yu, 2002). Therefore, in a second step, I estimate the default event risk premium

using historical sovereign default rates per rating class obtained from S&P in a procedure similar to Driessen (2005).

The estimation results of the first step show that the model is able to capture well both systemic and country-specific risk. I find clear clusters of shocks in the common factor during the peaks of the global financial and European sovereign debt crises, which are periods characterized by rapid increases of credit spreads for all countries. On the other hand, the country-specific factors indeed seem to account for country-specific episodes of distress not (directly) affecting other countries. For example, I show in an illustration that the country-specific factors of Brazil and Russia clearly capture the economic recessions and political unrest these countries experienced in the period 2015-2016.

In the second estimation step, I find a default event risk premium parameter that is closely in line with results found previously in the corporate credit risk literature (see, e.g., Driessen, 2005; Berndt et al., 2008). In particular, the results indicate that investors tend to "overweigh" instantaneous default probabilities by a factor of more than two when pricing sovereign CDS spreads. I show that distress risk premia alone cannot explain the differences between risk-neutral and historically observed default probabilities and that taking into account default event risk premia leads to a much better fit of historical default probabilities.

Although there is a large literature on sovereign credit risk, this chapter is closest related to Pan and Singleton (2008), Remolona et al. (2008) Longstaff et al. (2011), Ang and Longstaff (2013), Zinna (2013), Aït-Sahalia et al. (2014), and Monfort et al. (2018). Ang and Longstaff (2013) also use a model structure in which countries can default due to shocks in either a common or country-specific factor. The main difference between my model and theirs is that they model the common and country-specific factors as diffusive processes, whereas I use self-exciting jump processes. In this way, my model is explicitly designed to capture, in a parsimonious way, clusters of large jumps in CDS spreads apparent in the data. Aït-Sahalia et al. (2014) also use self-exciting processes to capture the clustering of sovereign credit shocks in time. Their model, however, accounts for commonality between sovereigns in an implicit way, through contagion effects, which are modeled by cross-exciting shocks. As a consequence, the number of parameters they need to estimate grows quadratically in the number of countries under consideration. In my model set-up, on the other hand, I only need to estimate one common factor (and loadings to this factor) to capture the commonality

of sovereign credit risk. As a result, my model is more parsimonious and can be estimated on a broader cross-section of countries than theirs. This chapter also differs from both Ang and Longstaff (2013) and Aït-Sahalia et al. (2014) in that they do not take into account any risk premia. I show, however, that risk premia make up a substantial fraction of sovereign CDS spreads.

Pan and Singleton (2008), Longstaff et al. (2011), and Zinna (2013) all consider distress risk premia embedded in sovereign CDS spreads, and show that these have a high level of commonality and are closely related to global and macroeconomic factors. These studies do, however, not consider default event risk premia. To the best of my knowledge, the only two papers that do take into account default event risk premia in a sovereign context are Remolona et al. (2008) and Monfort et al. (2018). Remolona et al. (2008) use credit rating data to extract actual default arrival intensities, and construct a measure of expected loss in case of default. They define the difference between the CDS spread and their expected loss measure to be the risk premium embedded in the CDS spread and show that it can be substantial. Since they do not use a formal pricing model, their risk premium measure essentially captures the total risk premium component, but is not able to distinguish between distress and default risk premia. My model, however, gives rise to an explicit decomposition of CDS spreads in both distress and default risk premia allowing me to study these separately. Monfort et al. (2018) develop a discrete-time pricing framework in which they also explicitly allow for commonality, default event risk premia, and contagion. They illustrate their framework in several applications, one of which is sovereign credit risk. The main focus of their paper is on the development of their pricing framework, which is considerably different from this chapter, and not on a detailed analysis of sovereign credit risk.

The remainder of this chapter is structured as follows: Sections 3.2 and 3.3 describe the data and model set-up, respectively. The estimation methodology and results are discussed in Section 3.4. Section 3.5 considers a decomposition of CDS spreads into country-specific and systemic risk premia components, and Section 3.6 concludes. Some proofs and details on the estimation procedure are given in Appendices 3.A and 3.B.

3.2 DATA

In the empirical analysis, I consider daily sovereign CDS data of 28 geographically dispersed countries over the period 01-01-2008 until 30-12-2016.[2] In particular, I consider for every country the term-structure of CDS spreads and obtain the daily two-, three-, five-, and ten-year CDS spreads from Datastream. All CDS contracts are denominated in dollars.

Table 3.1 presents summary statistics. For every country, the average S&P credit rating is determined by mapping prevailing credit rating grades to a numerical scale and taking the average over the sample period. As expected, countries with higher credit ratings have, in general, lower CDS spreads than countries with lower credit ratings. Furthermore, all countries, except Venezuela, have, on average, upwards-sloping term structures of CDS spreads.[3] The standard deviations show that there is substantial time-series variation in the CDS spreads. Table 3.1 also reports the variation of CDS spreads explained by the first principal component across different maturities within each country. For most countries, the first principal component explains well over 90% of the variation across different maturities.

Figure 3.1 plots the five-year CDS spreads grouped by average rating for all countries in our sample and illustrates two eminent features of the data: 1) A high degree of commonality across countries, and 2) the occurrences of clusters of large jumps in CDS spreads (see, e.g., the distressed period 2008-2009). A principal component analysis on the correlation matrix of the five-year CDS levels reveals that over 60% of the daily variation in CDS levels is explained by the first principal component. When restricting the sample period to the distressed period 2008-2009, however, the first principal component explains over 90% of the daily variation in CDS levels, suggesting that the commonality is larger in cri-

2 For some countries, CDS data was only available from a date later than 01-01-2008. The exact start date of the sample for every country is reported in Table 3.1.

3 Augustin (2018) shows that the slope of the term structure of CDS spreads contains information on the relative importance of global and domestic risk factors. He finds that country-specific factors influence spreads mainly when there is a negative slope of the term structure. Indeed, throughout our sample period Venezuela was prone to many country-specific risk factors such as high inflation and political unrest. This is also reflected in the very large CDS spreads throughout our sample period.

sis periods (see also Rigobon and Fernandez-Arias, 2000). Similar results hold when looking at a monthly frequency.[4]

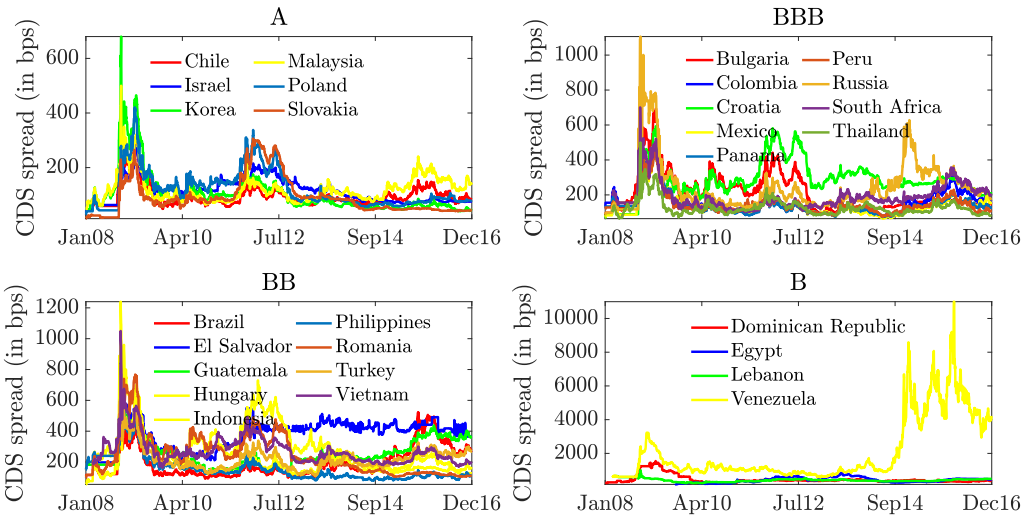


Figure 3.1. Five-year sovereign CDS spreads. This figure plots the five-year CDS spreads for A-rated (upper left panel), BBB-rated (upper right panel), BB-rated (bottom left panel), and B-rated (bottom right panel) countries. Source: Datastream.

In addition to sovereign CDS data, I use historical sovereign default data for estimating the default event risk premium. In particular, I use the average cumulative default rates per rating category as provided by Standard & Poor's (S&P). These default rates are averages of default rates of cohorts of countries that are formed each year. More specifically, each cohort starts on a specific start date and consists of all countries with a similar rating on that start date. The countries of each cohort are followed from the start date onwards, and cumulative default rates are constructed. Finally, in order to filter out time and cohort effects, the average cumulative default rates across cohorts is taken. In the sample, I focus on countries with an average credit rating of A and below. The reason for this is that there are no historical records of sovereign defaults for countries with a credit rating higher than A, and as such the

4 This high degree of commonality in sovereign credit risk is not specific to our sample and has been well-documented and investigated in the literature before (see, e.g., Longstaff et al., 2011; Ang and Longstaff, 2013).

Table 3.1. Summary statistics sovereign CDS spreads. This table reports summary statistics for the CDS spreads of the 28 countries in our sample. The first four columns display the country name, the first date for which observations are available, the total number of (daily) observations, and the average (S&P) rating throughout the sample period, respectively. The average rating is determined by transforming the prevailing ratings to a numerical scale and by taking the average over the sample period. Columns 5 to 8 report the average daily CDS spreads (in basis points) of the two-, three-, five-, and ten-year maturities, respectively. Columns 9 to 12 report the standard deviations of the daily CDS spreads (in basis points) of the two-, three-, five-, and ten-year maturities, respectively. The last column reports the percentage of the variation in the levels of the two-, three-, five-, and ten-year CDS spreads explained by the first principal component. The sample period runs from 01-01-2008 until 12-30-2016. Source: Datastream.

				Average spread				Std. dev. spread				
Country	Start	#Obs	Avg rating	2y	3y	5y	10y	2y	3y	5y	10y	% explained 1st PC
Brazil	1-1-2008	2349	BB	119	146	193	236	76	86	98	104	92.7
Bulgaria	1-2-2008	2348	BBB	160	179	215	243	126	121	110	95	97.9
Chile	2-29-2008	2306	A	55	69	96	124	31	34	40	41	95.3
Colombia	1-2-2008	2348	BBB	102	122	161	203	69	71	73	73	96.0
Croatia	1-2-2008	2348	BBB	215	245	287	317	110	106	103	102	94.3
Dominican Rep.	1-1-2008	2349	B	419	439	454	465	281	274	248	231	98.7
Egypt	8-20-2009	1922	B	368	389	415	425	152	153	150	152	99.0
El Salvador	2-29-2008	2306	BB	318	347	379	403	95	97	95	94	96.8
Guatemala	1-16-2008	2338	BB	200	224	255	288	86	86	89	90	97.8
Hungary	1-1-2008	2349	BB	206	232	272	297	151	144	136	123	97.4

Continuation of Table 3.1

				Average spread				Std. dev. spread				
Country	Start	#Obs	Avg rating	2y	3y	5y	10y	2y	3y	5y	10y	% explained 1st PC
Indonesia	2-29-2008	2306	BB	137	161	216	267	154	145	131	118	98.3
Israel	1-2-2008	2348	A	65	80	115	140	44	45	44	38	89.1
Korea	1-1-2008	2349	A	68	77	104	127	78	77	76	71	98.3
Lebanon	2-29-2008	2306	B	318	358	408	434	99	94	95	101	94.8
Malaysia	1-1-2008	2349	A	65	82	120	154	43	46	49	52	88.6
Mexico	1-2-2008	2348	BBB	86	104	139	177	68	69	68	66	96.5
Panama	1-2-2008	2348	BBB	91	109	145	184	67	69	70	68	97.4
Peru	1-2-2008	2348	BBB	92	111	148	188	63	65	64	62	97.2
Philippines	1-1-2008	2349	BB	92	112	156	203	93	90	83	72	97.6
Poland	1-2-2008	2348	A	75	91	120	152	68	70	69	61	95.2
Romania	1-2-2008	2348	BB	180	205	245	273	154	145	133	119	97.9
Russia	1-1-2008	2349	BBB	190	210	243	271	191	169	145	127	97.4
Slovakia	1-1-2008	2349	A	60	70	92	117	65	67	66	64	96.3
South Africa	2-29-2008	2306	BBB	120	148	189	226	81	79	76	71	93.6
Thailand	2-29-2008	2306	BBB	71	89	127	162	47	47	45	44	93.5
Turkey	10-8-2008	2148	BB	153	176	223	259	90	86	79	73	93.7
Venezuela	2-29-2008	2306	B	2267	2192	2095	1912	2312	2115	1917	1694	99.4
Vietnam	2-29-2008	2306	BB	191	220	273	318	125	116	90	83	96.5

Table 3.2. Historical sovereign cumulative average default rates. This table reports the historical cumulative average default rates (in percentage points) as reported by the S&P rating agency. Source: "2016 Annual Sovereign Default Study and Rating Transitions", S&P.

	Time horizon (years)									
	1	2	3	4	5	6	7	8	9	10
A	0.0	0.0	0.6	1.2	1.8	2.5	3.3	4.0	4.9	5.8
BBB	0.0	0.6	1.3	1.9	2.6	3.3	3.6	3.6	3.6	3.6
BB	0.5	1.7	2.5	2.9	3.8	4.8	5.9	7.1	8.0	8.6
B	2.6	5.5	8.6	11.8	14.6	17.2	20.0	22.8	24.9	26.6

historical default rates are not very informative for countries with a high credit rating. Table 3.2 shows the historical cumulative average default rates as reported by S&P. In general, lower credit ratings are associated with higher default rates. Only the long-term horizons of A- and BBB-rated countries seem to violate this pattern. The reason for this is that the recent (double) default of Greece significantly affected the default rates of A-rated countries and not of BBB-rated countries, since Greece still had a credit rating of A in 2009.

3.3 THE MODEL

Motivated by the apparent features of the data and recent literature on the modeling of sovereign credit risk, I propose a new model for the term structure of sovereign CDS spreads. Similar to Ang and Longstaff (2013), I assume that defaults can occur via two channels: Both a common shock, affecting all sovereigns, as well as independent country-specific factors can induce a country to default.[5] A key distinction of my model compared to Ang and Longstaff (2013), however, is that it departs from the diffusive setting and uses self-exciting jump processes to model the country-specific and common factors (see also Aït-Sahalia et al., 2014). This feature allows the model to capture the empirical observation that large credit shocks tend to cluster both in time as well

5 In line with Ang and Longstaff (2013), I sometimes refer to the common factor as the systemic factor.

as between countries (see Figure 3.1). A major difference between my model and Aït-Sahalia et al. (2014), on the other hand, is that I explicitly account for the commonality in sovereign credit risk by modeling a common factor. Aït-Sahalia et al. (2014) allow for commonality of sovereign credit risk in an implicit way through contagion effects which they model as cross-exciting jumps. The main advantage of my approach over theirs is that I only need to estimate one common factor (and loadings to this factor) to capture commonality, whereas the number of parameters they need to estimate grows quadratically in the number of countries under consideration. As a result, my model can be estimated on a broader cross-section of countries than theirs. Furthermore, Ang and Longstaff (2013) and Aït-Sahalia et al. (2014) do not consider the estimation of risk premia, which is the core focus of this chapter.

Specifically, every country $i, i = 1, 2, \ldots, K$, can be hit by country-specific shocks, N_i, and by shocks in a common factor, N_c, where the subscript c refers to "common".[6] [7] Every time a country is hit by a country-specific shock there is a probability γ_i that this country defaults. Similarly, when a country is hit by a common shock, there is a probability of γ_i^c of going into default. The probabilities γ_i^c are thus sovereign-specific and can be viewed as loadings to the common factor. Since both common and country-specific shocks can trigger a default event, the CDS spread of country i depends on both the common as well as the ith country-specific factors (see equation (3.6) in Section 3.3.1 below). The $N_j, j \in \{c, 1, 2, \ldots, K\}$, are independent counting processes with each an underlying shock arrival intensity process $\lambda_{j,t}^{\mathbb{P}}$ under the actual probability measure $\mathbb{P}$, and an arrival intensity process $\lambda_{j,t}^{\mathbb{Q}}$ under the risk-neutral measure $\mathbb{Q}$. The ratio between the actual and risk-neutral intensities determines the risk premium related to default event risk.

In particular, assuming absence of arbitrage, it can be shown that there exists the following relation between the default intensity processes un-

6 Throughout, I use the subscript j, $j \in \{c, 1, 2, \ldots, K\}$ to refer to any properties holding for both the common and country-specific factors, and the subscript $i, i \in \{1, 2, \ldots, K\}$ to refer to a specific country.

7 The common factor in principle represents exogenous systematic shocks hitting all countries at the same time. Empirically, however, it may capture contagion effects similar to those discussed in Collin-Dufresne et al. (2010). Note, that the model setup in principle could be extended to include contagion effects more directly by allowing country-specific shocks to cross-excite the default intensities of other countries. However, the main reason I specify a common factor is exactly to avoid this, as this would blow up the number of parameters to be estimated.

der the actual probability measure $\mathbb{P}$ and the risk-neutral measure $\mathbb{Q}$ (see, e.g., Jarrow et al., 2005):

$$\lambda^{\mathbb{Q}}_{j,t} = \mu_{j,t}\lambda^{\mathbb{P}}_{j,t}, \quad j \in \{c, 1, 2, \ldots, K\}. \tag{3.1}$$

Here $\mu_{j,t}$ is the risk premium associated with the (unpredictable) default event itself. More specifically, if $\mu_{j,t} > 1$ default event risk is priced as investors overestimate the (instantaneous) probability of default under the risk-neutral measure. In principle, $\mu_{j,t}$ can be time-varying and different across $j, j \in \{c, 1, 2, \ldots, K\}$. However, to estimate $\mu_{j,t}$, data on real-world sovereign default probabilities is needed (Yu, 2002), and, given the scarcity of sovereign default events, it is not feasible to construct accurate time-varying and/or country-specific estimates of real-world default probabilities. For this reason, I assume $\mu_{j,t} = \mu$ to be constant over time and the same for all country-specific and common factors.[8]

Jarrow et al. (2005) argue that there are in principle two reasons for why default event risk could be priced. First, default event risk is priced when there is a positive probability of countries defaulting at the same time (i.e., conditional on the state vectors driving the default intensities, sovereign defaults are not independent). Second, default event risk is priced when there are only a finite number of entities/assets and the effects of a default event can not be diversified away. It is plausible that, especially in a sovereign context, both these conditions are met, and that default event risk should be taken into account.

In addition to default event risk, captured by the factor μ between the default intensities under $\mathbb{P}$ ($\lambda^{\mathbb{P}}_{j,t}$) and $\mathbb{Q}$ ($\lambda^{\mathbb{Q}}_{j,t}$), another source of risk stems from the fact that the likelihood of default changes over time. In the case that fluctuations in the intensities over time are priced, the dynamics of $\lambda^{\mathbb{P}}_{j,t}$ and $\lambda^{\mathbb{Q}}_{j,t}$ also differ under both measures. Risk premia related to changes of default risk over time have been investigated in a sovereign context before (see, e.g., Pan and Singleton, 2008; Longstaff et al., 2011), and I refer to these risk premia as "distress risk premia". In total there are thus three configurations of default intensity processes

8 In estimating μ, I perform a robustness check and estimate μ per rating class. Except for the A-rated countries, I find little variation across these rating-specific estimates, suggesting that this assumption is reasonable.

and their dynamics: The $\mathbb{P}$- and $\mathbb{Q}$-dynamics of $\lambda^{\mathbb{Q}}_{j,t}$, and the $\mathbb{P}$-dynamics of $\lambda^{\mathbb{P}}_{j,t}$, $j \in \{c, 1, 2, \ldots, K\}$.[9]

I assume that the $\mathbb{P}$-dynamics of $\lambda^{\mathbb{Q}}_{j,t}$ are given by the following self-exciting dynamics:[10]

$$d\lambda^{\mathbb{Q}}_{j,t} = \alpha^{\mathbb{P}}_j(\lambda^{\mathbb{P}}_{j,\infty} - \lambda^{\mathbb{Q}}_{j,t})dt + \sigma_j\sqrt{\lambda^{\mathbb{Q}}_{j,t}}dW^{\mathbb{P}}_{j,t} + Z_j dN_{j,t}, \quad j \in \{c, 1, 2, \ldots, K\}, \tag{3.2}$$

where $W^{\mathbb{P}}_{j,t}$ are independent $\mathbb{P}$-Brownian motions and $N_{j,t}$ are the independent credit shock arrival processes with intensity processes $\lambda^{\mathbb{Q}}_{j,t}$ themselves. Every time the counting process $N_{j,t}$ jumps (i.e., a common or country-specific credit event occurs), $\lambda^{\mathbb{Q}}_{j,t}$ jumps by $Z_j > 0$. This again induces an increase in the probability of another jump in $N_{j,t}$, since this jump process is driven by $\lambda^{\mathbb{Q}}_{j,t}$. This self-exciting specification allows the model to capture the clustering of large credit shocks in time and (through the common factor) across countries.

Note that the clustering of large credit shocks can also be obtained by defining the intensity processes $\lambda^{\mathbb{Q}}_{j,t}$ to be standard jump-diffusions with exogenously specified jump components instead of endogenous, self-exciting jumps. That is, the $N_{j,t}$ in (3.2) can be replaced by exogenous jumps $J_{j,t}$, which are driven by their own, separate intensity processes. Such an exogenous specification of the jump part, however, has two important disadvantages compared to the self-exciting specification employed here. First, the interpretation of exogenous shocks affecting the intensities of the credit event processes is not straightforward. Second, the exogenous jump-diffusion specification is less parsimonious, since it requirers the separate modeling of the intensity processes driving the exogenous shocks.

In principle, (3.2) can be generalized in several ways: First, as mentioned above, the model could also take into account direct spillover effects from one country to another by allowing for cross-excitation effects. That

9 In principle, one could also consider the $\mathbb{Q}$-dynamics of $\lambda^{\mathbb{P}}$. This configuration of probability measures does, however, not bear any economic meaning and therefore I ignore it in the discussion.

10 Throughout this chapter, the actual and risk-neutral default intensities are denoted with superscripts $\mathbb{P}$ and $\mathbb{Q}$, respectively (i.e., $\lambda^{\mathbb{P}}$ and $\lambda^{\mathbb{Q}}$). Similarly, where necessary, the parameters governing the $\mathbb{P}$- and $\mathbb{Q}$-dynamics are also denoted with superscripts $\mathbb{P}$ and $\mathbb{Q}$, respectively.

is, a shock in the country-specific factor of country i (i.e., a jump of $N_{i,t}$) could directly affect the intensity processes $\lambda^{\mathbb{Q}}_{j,t}$ $j \in \{c, 1, 2, \ldots, K\}$, $j \neq i$. In this case, one can explicitly differentiate between direct contagion effects and exogenous common shocks. Second, shocks in the common factor (i.e., jumps in $N_{c,t}$) could directly affect the country-specific intensities. This would capture the notion that systemic shocks could trigger episodes of country-specific distress. In both these extensions, the number of parameters to be estimated would, however, increase and identification would become infeasible. Therefore, I use the common component to capture both direct contagion and exogenous shocks, and let the country-specific factors reflect pure country-specific risk.[11]

Consistent with the literature (see, e.g., Pan and Singleton, 2008; Longstaff et al., 2011), I assume that the market prices of risk underlying the change of measure from $\mathbb{P}$ to $\mathbb{Q}$ are dependent on the current levels of the default intensities and are given by

$$\xi_{j,t} = \frac{\delta_{j,0}}{\sqrt{\lambda^{\mathbb{Q}}_{j,t}}} + \delta_{j,1}\sqrt{\lambda^{\mathbb{Q}}_{j,t}}, \quad j \in \{c, 1, 2, \ldots, K\}. \tag{3.3}$$

These market prices of risk assure that the $\mathbb{Q}$-dynamics of $\lambda^{\mathbb{Q}}_{j,t}$ are of a similar form as the $\mathbb{P}$-dynamics and are given by

$$d\lambda^{\mathbb{Q}}_{j,t} = \alpha^{\mathbb{Q}}_j(\lambda^{\mathbb{Q}}_{j,\infty} - \lambda^{\mathbb{Q}}_{j,t})dt + \sigma_j\sqrt{\lambda^{\mathbb{Q}}_{j,t}}dW^{\mathbb{Q}}_{j,t} + Z_j dN_{j,t}, \tag{3.4}$$

where $\alpha^{\mathbb{Q}}_j = \alpha^{\mathbb{P}}_j + \delta_{j,1}\sigma_j$, $\alpha^{\mathbb{Q}}_j\lambda^{\mathbb{Q}}_{j,\infty} = \alpha^{\mathbb{P}}_j\lambda^{\mathbb{P}}_{j,\infty} - \delta_{j,0}\sigma_j$, and $W^{\mathbb{Q}}_{j,t}$ are independent $\mathbb{Q}$-Brownian motions. Note that the difference between the $\mathbb{P}$- and $\mathbb{Q}$-dynamics of $\lambda^{\mathbb{Q}}_{j,t}$ stem from the change of measure in the Brownian motions. The market price of risk parameters $\xi_{j,t}$ capture the compensations that investors require for being exposed to unexpected changes in default arrival rates.

Since $\lambda^{\mathbb{Q}}$ and $\lambda^{\mathbb{P}}$ are related through the constant parameter μ, the $\mathbb{P}$- and $\mathbb{Q}$-dynamics of $\lambda^{\mathbb{P}}_{j,t}$ are of a similar form as the $\mathbb{P}$- and $\mathbb{Q}$-dynamics of $\lambda^{\mathbb{Q}}_{j,t}$. In the estimation procedure, detailed in Section 3.4, I first use

11 Another possible generalization of the model is to allow for stochastic jump sizes Z_j. The reason I take fixed jump size parameters $Z_j > 0$ instead of stochastic jump sizes is that stochastic jump sizes also lead to additional risk premium parameters to be estimated. Therefore, even in the case where the jump size distribution depends on only one parameter (e.g., exponentially distributed jump sizes), the estimation and identification of the additional parameters would be challenging.

sovereign CDS spread data to estimate the $\mathbb{P}$- and $\mathbb{Q}$-dynamics of the risk-neutral default intensity processes, $\lambda^{\mathbb{Q}}_{j,t}$. After that, I use historical sovereign default rates obtained from S&P, assumed to contain information on the real-world default intensities $\lambda^{\mathbb{P}}_{j,t}$, to estimate the default event risk premium parameter μ.

3.3.1 *CDS pricing*

The time t level of the CDS spread of country i with maturity M, denoted by $CDS_{i,t}(M)$, is determined by equating the payoff value for the protection buyer to the payoff value for the protection seller. I will make the standard simplifying assumption that the risk-free rate is independent from the common and country-specific factors, and denote $D(t,T) = \mathbb{E}^{\mathbb{Q}}\left[e^{-\int_t^T r_s ds}\middle|\mathcal{F}_t\right] = \mathbb{E}^{\mathbb{Q}}_t\left[e^{-\int_t^T r_s ds}\right]$, where $\mathcal{F}_t$ denotes the conditioning information available at time t.[12] I use US Treasury rates to construct the risk-free discount factors $D(t,T)$.

Specifically, I get (see, e.g., Duffie and Singleton, 2003)

$$\begin{aligned}
&\frac{1}{4}CDS_{i,t}(M)\times\\
&\sum_{j=1}^{4M} D(t,t+0.25j)\mathbb{E}^{\mathbb{Q}}_t\left[(1-\gamma_i)^{N_{i,t+0.25j}-N_{i,t}}\right]\mathbb{E}^{\mathbb{Q}}_t\left[(1-\gamma^c_i)^{N_{c,t+0.25j}-N_{c,t}}\right]\\
&= (1-R)\times\\
&\int_t^{t+M} D(t,u)\mathbb{E}^{\mathbb{Q}}_t\left[\left(\gamma_i\lambda^{\mathbb{Q}}_{i,u}+\gamma^c_i\lambda^{\mathbb{Q}}_{c,u}\right)(1-\gamma_i)^{N_{i,u}-N_{i,t}}(1-\gamma^c_i)^{N_{c,u}-N_{c,t}}\right]du.
\end{aligned} \tag{3.5}$$

The left-hand side of (3.5) reflects the present value of the (quarterly) premium payments that the buyer makes to the seller, contingent upon a default event not having occurred. A default can occur either through a country-specific shock, $N_{i,t}$, or through a shock in the common factor, $N_{c,t}$. The right-hand side of (3.5) reflects the present value of the payout that the seller makes in case of default. I assume fractional recovery of

12 In more technical terms, $\{\mathcal{F}_t\}_{t\geqslant 0}$ is the filtration generated by all the relevant stochastic processes.

face value of the underlying bond and let R denote the constant recovery rate.[13]

Solving for $\mathrm{CDS}_{i,t}(M)$ gives the following CDS pricing formula:

$$\mathrm{CDS}_{i,t}(M) = \frac{(1-R)\int_t^{t+M} D(t,u)\mathbb{E}_t^{\mathbb{Q}}\left[\left(\gamma_i\lambda_{i,u}^{\mathbb{Q}}+\gamma_i^c\lambda_{c,u}^{\mathbb{Q}}\right)(1-\gamma_i)^{N_{i,u}-N_{i,t}}(1-\gamma_i^c)^{N_{c,u}-N_{c,t}}\right]du.}{\frac{1}{4}\sum_{j=1}^{4M} D(t,t+0.25j)\mathbb{E}_t^{\mathbb{Q}}\left[(1-\gamma_i)^{N_{i,t+0.25j}-N_{i,t}}\right]\mathbb{E}_t^{\mathbb{Q}}\left[(1-\gamma_i^c)^{N_{c,t+0.25j}-N_{c,t}}\right]}. \quad (3.6)$$

The expectations appearing in (3.6) can be computed in closed-form (up to a system of ODEs) by exploiting the affine structure of the model and using the framework outlined in Duffie et al. (2000). The computations are detailed in Appendix 3.A.

3.4 ESTIMATION METHODOLOGY

Similar to Driessen (2005), the model setup is such that it can be estimated in two steps. In the first step, I estimate the model governing the risk-neutral intensities $\lambda_{j,t}^{\mathbb{Q}}$, $j \in \{c, 1, 2, \ldots, K\}$ using sovereign CDS data. In the second step, I estimate the default event risk premium parameter μ using S&P historical sovereign default data. The global outline and results of the first and second steps of the estimation procedure are described in Section 3.4.1, and Section 3.4.2, respectively.

3.4.1 *Estimation risk-neutral intensities* $\lambda_{j,t}^{\mathbb{Q}}$

In estimating the model of the risk-neutral intensities, I use a Bayesian Markov chain Monte Carlo (MCMC) procedure similar to Sperna Weiland et al. (2019). This procedure makes use of the sovereign CDS spreads and provides estimates for the parameters driving the $\mathbb{P}$- and $\mathbb{Q}$-dynamics of $\lambda_{j,t}^{\mathbb{Q}}$, values of the latent risk-neutral intensity processes, and latent jump times.

A major advantage of the model set-up is that I can estimate the parameters governing the common and country-specific factors in an iterative procedure. First, I estimate the common factor by pooling the CDS data

13 Following the literature, I assume a constant recovery rate of 25% for all countries and abstract away from recovery rate risk premia (see, e.g., Longstaff et al., 2011).

of all countries and ignoring the country-specific factors. In the second step, I estimate the country-specific factors, keeping the common factor results from the first step fixed. In a third step, I re-estimate the common factor, but now fixing the country-specific factors obtained in step two. In this way, the estimation of the common factor explicitly takes into account the presence of country-specific factors. I investigated whether applying more iterations of steps two and three would lead to significant changes in the parameter estimates, but found this not to be the case, indicating that there is fast convergence of the estimation procedure. This iterative procedure also helps in identifying which parts of the CDS spreads can be attributed to the common factor and which parts to country-specific factors, since the first step explicitly targets comovements in the CDS spreads, and the second step explicitly targets country-specific deviations from these comovements.[14]

The main challenges in estimating the (risk-neutral intensity) model are that the intensity processes and jump times are latent, and that, due to self-excitation, their transition densities are not known. The key of dealing with these issues is to properly discretize and orderly sample the intensity processes defined in (3.2). To see this, consider the following discretized version of (3.2):

$$\lambda^{\mathbb{Q}}_{j,t+1} - \lambda^{\mathbb{Q}}_{j,t} = \alpha^{\mathbb{P}}_j \lambda^{\mathbb{P}}_{j,\infty} \Delta_{t+1} - \alpha^{\mathbb{P}}_j \lambda^{\mathbb{P}}_{j,t} \Delta_{t+1} + \sigma_j \sqrt{\lambda^{\mathbb{Q}}_{j,t} \Delta_{t+1}}\, \epsilon_{j,t+1} + Z_j \Delta N_{j,t+1}, \quad j \in \{c, 1, 2, \ldots, K\}, \tag{3.7}$$

where Δ_{t+1} is the length of the time interval between t and $t+1$ (i.e., a business day), $\epsilon_{j,t+1}$ an independent standard normal random variable, and the jump counters $\Delta N_{j,t+1} \in \{0, 1\}$ are Bernoulli random variables with non-constant success probabilities $\lambda^{\mathbb{Q}}_{j,t} \Delta_{t+1}$ ($\Delta N_{j,t+1} = 1$ indicates a jump arrival). The discretization thus assumes that at most one jump can occur in the time-interval Δ_{t+1}, which follows from the small-time property of self-exciting processes stating that $\mathbb{P}\left[N_{j,t+\Delta} - N_{j,t} > 1 \middle| \mathcal{F}_t\right] = o(\Delta)$. Sperna Weiland et al. (2019) show in a Monte Carlo study that this discretization on a daily frequency does not impose notable biases.

14 Underlying this identification strategy lies the assumption that the comovements in CDS spreads across different countries are caused by the common factor.

Using the discretization (3.7) and denoting the states by $X_t^j = \{\Delta N_{j,t}, \lambda_{j,t}^{\mathbb{Q}}\}$ and $\bar{\Theta}$ the vector with parameters, the transition density of the states can, by Bayes rule, be decomposed as

$$p(X_t^j | X_{t-1}^j, \bar{\Theta}) = p(\lambda_{j,t}^{\mathbb{Q}} | \Delta N_{j,t}, X_{t-1}^j, \bar{\Theta}) p(\Delta N_{j,t} | X_{t-1}^j, \bar{\Theta}), \tag{3.8}$$

where $p(\lambda_{j,t}^{\mathbb{Q}} | \Delta N_{j,t}, X_{t-1}^j, \bar{\Theta})$ is Gaussian, and $p(\Delta N_{j,t} | X_{t-1}^j, \bar{\Theta})$ Bernoulli with success probability $\lambda_{j,t-1}^{\mathbb{Q}} \Delta_t$. That is, under the discretization above, the transition density is a mixture of normal densities, allowing me to sequentially draw $\Delta N_{j,t}$ from the Bernoulli probabilities $p(\Delta N_{j,t} | X_{t-1}^j, \Theta)$ and $\lambda_{j,t}^{\mathbb{Q}}$ from $p(\lambda_{j,t}^{\mathbb{Q}}, | \Delta N_{j,t}, X_{t-1}^j, \Theta)$ using the newly drawn $\Delta N_{j,t}$ in the conditioning information. The discretization above thus simplify the transition densities of the states, which play a crucial role in determining the posterior densities necessary for Bayesian inference. The details of the estimation procedure are explained in Appendix 3.B.

3.4.1.1 *Estimation results risk-neutral intensities*

Table 3.3 reports the posterior means and standard deviations of the parameter estimates of the common and country-specific risk-neutral intensities, the number of estimated jumps in each of the factors, and the average relative pricing errors of the CDS spreads per country.

Table 3.3 shows that the speed-of-mean-reversion parameters governing the $\mathbb{P}$-dynamics of the intensity processes ($\alpha_j^{\mathbb{P}}$) are larger than the speed-of-mean-reversion parameters governing the $\mathbb{Q}$-dynamics of the intensity processes ($\alpha_j^{\mathbb{Q}}$) for all common and country-specific factors. That is, under the risk-neutral dynamics, distressed periods are more persistent. Furthermore, backing out the implied long-term average intensity values, $\lambda_{j,\infty}^{\mathbb{P}}$ and $\lambda_{j,\infty}^{\mathbb{Q}}$, reveals that the long-term average intensities are higher under the $\mathbb{Q}$-dynamics than under the $\mathbb{P}$-dynamics for the common factor and most of the country-specific factors. Only for the Dominican Republic, Egypt, and Venezuela the opposite result holds. In principle, the slower speed-of-mean-reversions and higher long-term average default intensities under the risk-neutral dynamics indicate the presence of distress risk premia (i.e., risk premia related to the differences in process dynamics under the actual and risk-neutral measures). In Section 3.5, I study the economic significance of these risk premia in more detail.

Table 3.3. Parameter estimates risk-neutral common and country-specific intensity processes. This table reports the parameter estimates and posterior standard errors (in brackets) of the risk-neutral common (first row) and country-specific intensity process parameters. The last column reports the average relative pricing errors (ARPE) for each country.

	$\alpha^{\mathbb{P}}$	$\alpha^{\mathbb{P}}\lambda_{\infty}^{\mathbb{P}}$	σ^2	$\alpha^{\mathbb{Q}}$	$\alpha^{\mathbb{Q}}\lambda_{\infty}^{\mathbb{Q}}$	z	γ_i	h	γ^c	ARPE
Common factor	1.269	0.241	0.470	0.537	0.237	0.462		0.127		
	(0.158)	(0.007)	(0.012)	(0.003)	(0.002)	(0.004)		(0.000)		
Brazil	0.928	0.305	0.245	0.197	0.228	0.092	0.027	0.171	0.007	9.73%
	(0.522)	(0.029)	(0.033)	(0.071)	(0.050)	(0.018)	(0.001)	(0.052)	(0.00004)	
Bulgaria	0.506	0.016	0.031	0.151	0.015	0.074	0.275	0.118	0.008	6.28%
	(0.283)	(0.001)	(0.002)	(0.063)	(0.001)	(0.032)	(0.013)	(0.048)	(0.00005)	
Chile	1.783	0.126	0.155	0.206	0.098	0.159	0.016	0.132	0.009	8.77%
	(0.553)	(0.028)	(0.020)	(0.061)	(0.026)	(0.014)	(0.001)	(0.023)	(0.00003)	
Colombia	1.258	0.159	0.169	0.072	0.114	0.030	0.024	0.142	0.012	7.23%
	(0.978)	(0.041)	(0.007)	(0.069)	(0.044)	(0.009)	(0.004)	(0.060)	(0.00004)	
Croatia	0.596	0.010	0.015	0.243	0.008	0.239	0.401	0.202	0.023	11.59%
	(0.404)	(0.006)	(0.005)	(0.019)	(0.002)	(0.031)	(0.064)	(0.049)	(0.00008)	
Dominican Rep.	1.860	0.394	0.074	0.332	0.038	0.150	0.159	0.092	0.020	4.59%
	(0.388)	(0.007)	(0.030)	(0.081)	(0.017)	(0.047)	(0.022)	(0.050)	(0.00013)	
Egypt	0.527	0.110	0.052	0.250	0.029	0.197	0.176	0.116	0.012	4.32%
	(0.185)	(0.007)	(0.036)	(0.050)	(0.018)	(0.030)	(0.028)	(0.090)	(0.00015)	

Continuation of Table 3.3

	$\alpha^{\mathbb{P}}$	$\alpha^{\mathbb{P}}\lambda_{\infty}^{\mathbb{P}}$	σ^2	$\alpha^{\mathbb{Q}}$	$\alpha^{\mathbb{Q}}\lambda_{\infty}^{\mathbb{Q}}$	Z	γ_1	h	γ^c	ARPE
El Salvador	0.838	0.073	0.056	0.329	0.028	0.205	0.199	0.077	0.025	5.54%
	(0.219)	(0.010)	(0.007)	(0.054)	(0.004)	(0.013)	(0.008)	(0.005)	(0.00010)	
Guatemala	1.570	0.314	0.060	0.086	0.030	0.059	0.059	0.073	0.017	5.60%
	(0.701)	(0.028)	(0.003)	(0.018)	(0.002)	(0.005)	(0.001)	(0.016)	(0.00009)	
Hungary	0.397	0.010	0.018	0.251	0.009	0.245	0.358	0.178	0.019	9.19%
	(0.109)	(0.005)	(0.008)	(0.009)	(0.004)	(0.028)	(0.045)	(0.072)	(0.00008)	
Indonesia	0.534	0.027	0.040	0.166	0.022	0.166	0.146	0.147	0.012	8.25%
	(0.235)	(0.009)	(0.004)	(0.004)	(0.002)	(0.004)	(0.043)	(0.053)	(0.00005)	
Israel	1.406	0.143	0.158	0.038	0.098	0.026	0.029	0.191	0.001	11.11%
	(0.870)	(0.027)	(0.021)	(0.034)	(0.023)	(0.009)	(0.006)	(0.070)	(0.00001)	
Korea	1.499	0.006	0.011	0.278	0.006	0.278	0.357	0.249	0.004	14.18%
	(0.828)	(0.001)	(0.001)	(0.002)	(0.000)	(0.002)	(0.054)	(0.075)	(0.00002)	
Lebanon	1.118	0.143	0.091	0.470	0.072	0.289	0.128	0.095	0.026	6.07%
	(0.338)	(0.023)	(0.006)	(0.166)	(0.019)	(0.059)	(0.008)	(0.023)	(0.00014)	
Malaysia	1.785	0.163	0.168	0.191	0.113	0.175	0.022	0.181	0.008	12.66%
	(0.770)	(0.021)	(0.038)	(0.026)	(0.027)	(0.010)	(0.002)	(0.024)	(0.00003)	
Mexico	1.844	0.157	0.186	0.147	0.118	0.108	0.016	0.154	0.012	8.32%
	(0.666)	(0.042)	(0.013)	(0.087)	(0.039)	(0.005)	(0.005)	(0.039)	(0.00003)	
Panama	1.347	0.163	0.192	0.129	0.131	0.070	0.018	0.129	0.012	7.07%
	(0.803)	(0.037)	(0.017)	(0.103)	(0.039)	(0.017)	(0.002)	(0.043)	(0.00003)	
Peru	1.827	0.180	0.176	0.145	0.119	0.089	0.017	0.119	0.015	6.81%

Continuation of Table 3.3

	(0.756)	(0.035)	(0.013)	(0.095)	(0.042)	(0.006)	(0.002)	(0.034)	(0.00003)	
Philippines	1.441	0.010	0.020	0.170	0.013	0.169	0.274	0.139	0.006	9.73%
	(0.280)	(0.001)	(0.001)	(0.006)	(0.000)	(0.001)	(0.011)	(0.007)	(0.00005)	
Poland	0.531	0.005	0.009	0.236	0.005	0.236	0.370	0.280	0.008	16.13%
	(0.200)	(0.001)	(0.000)	(0.005)	(0.000)	(0.005)	(0.062)	(0.065)	(0.00003)	
Romania	1.819	0.257	0.110	0.053	0.055	0.040	0.066	0.166	0.003	8.29%
	(1.351)	(0.057)	(0.023)	(0.012)	(0.012)	(0.014)	(0.012)	(0.096)	(0.00007)	
Russia	1.056	0.053	0.084	0.218	0.042	0.065	0.131	0.142	0.011	9.13%
	(0.430)	(0.013)	(0.001)	(0.031)	(0.001)	(0.010)	(0.002)	(0.014)	(0.00005)	
Slovakia	0.695	0.005	0.009	0.187	0.005	0.186	0.404	0.333	0.002	22.26%
	(0.303)	(0.001)	(0.001)	(0.003)	(0.000)	(0.003)	(0.046)	(0.058)	(0.00002)	
South Africa	2.123	0.201	0.206	0.207	0.135	0.109	0.039	0.138	0.009	8.79%
	(0.330)	(0.011)	(0.030)	(0.058)	(0.031)	(0.012)	(0.003)	(0.031)	(0.00004)	
Thailand	1.994	0.159	0.196	0.095	0.143	0.067	0.016	0.154	0.010	9.57%
	(0.424)	(0.039)	(0.035)	(0.053)	(0.039)	(0.033)	(0.003)	(0.030)	(0.00002)	
Turkey	0.474	0.027	0.044	0.210	0.022	0.194	0.106	0.101	0.022	7.45%
	(0.174)	(0.004)	(0.002)	(0.007)	(0.001)	(0.014)	(0.005)	(0.009)	(0.00007)	
Venezuela	0.689	0.135	0.098	0.880	0.049	0.196	0.470	0.320	0.138	6.69%
	(0.283)	(0.013)	(0.025)	(0.099)	(0.013)	(0.091)	(0.086)	(0.184)	(0.00081)	
Vietnam	0.422	0.027	0.048	0.196	0.024	0.145	0.150	0.103	0.024	6.93%
	(0.192)	(0.004)	(0.002)	(0.019)	(0.001)	(0.017)	(0.023)	(0.023)	(0.00008)	

The upper panel of Figure 3.2 plots the estimated systemic default risk intensities. The model seems to capture systemic risk well. For example, the common factor was especially large during the 2008-2009 crisis period in which the CDS spreads of all countries spiked up. Furthermore, there is an increase in the common factor during the second half of 2011, reflecting the peak of the European sovereign debt crisis. The bottom panel of Figure 3.2 shows the estimated (self-exciting) jumps in the common factor.[15] I find a cluster of jumps shortly after the default of Lehman Brothers. Furthermore, I find a jump on September 22nd, 2011. On this date, global stock markets dropped over 3% and the VIX index spiked with 11% as a result of increasing fear of investors regarding spillovers of the European sovereign debt crisis. In total, I find 7 jumps in the common factor.

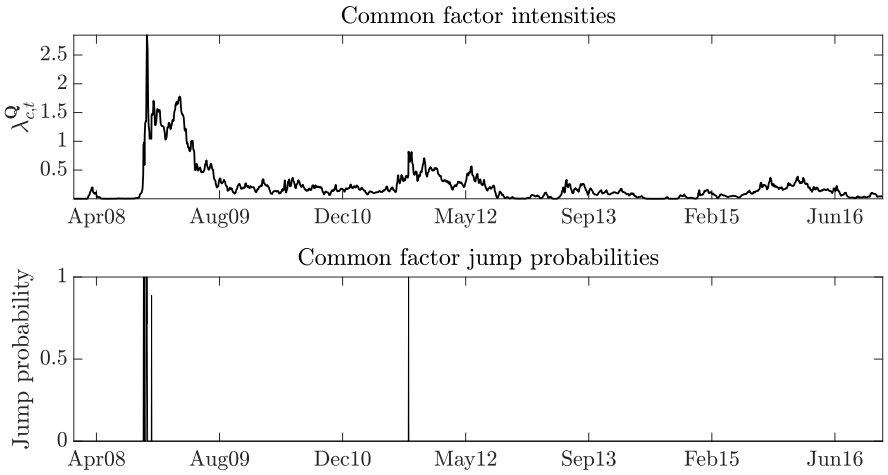


Figure 3.2. Estimated common factor intensities and jump times. This figure plots the estimated common factor intensities (upper panel) and estimated probabilities of arrivals of large shocks in the common factor (bottom panel).

To illustrate the performance of the model regarding country-specific factors, I plot in Figure 3.3 the model-fit of five-year CDS spreads (upper panels), estimated country-specific intensities (middle panels), and estimated country-specific jump probabilities of Brazil (left column) and Russia (right column). I focus on these countries, since their model-fit is

15 For both the common and country-specific factors, I define jumps to occur on those days for which the estimated jump probability is larger than 0.25.

close to the average relative pricing error of 9% over all countries. Furthermore, both Brazil and Russia experienced country-specific distress periods during the sample period, making them appropriate candidates to evaluate the model performance. The middle panels show that the country-specific factors indeed seem to pick up country-specific distress. For Brazil, the intensities spike up from 2015 onwards, coinciding with the start of the economic recession and increased political unrest hitting the country. Similarly, Russia also experienced a recession in 2015-2016 as a result of international sanctions in response to the Ukraine conflict, sharp declines in oil prices, and strong depreciation of the currency. Again, the country-specific intensities seem to capture this episode of distress well.

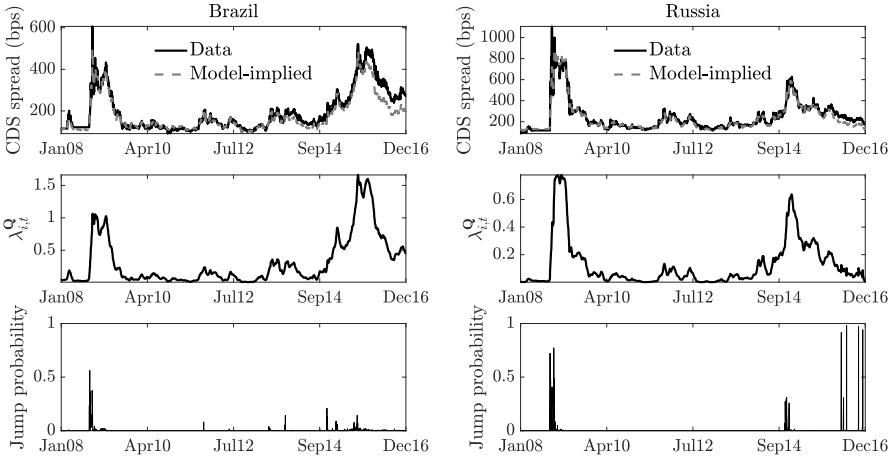


Figure 3.3. Model-fit and estimated country-specific factor intensities and jump times. This figure illustrates the five-year CDS spread model-fit (upper panels), estimated country-specific intensities (middle panels), and estimated jump times (bottom panels) for Brazil (left column) and Russia (right column).

As Table 3.3 indicates, the posterior standard errors of the parameters of, and loadings on (i.e., γ_i^c), the common component are relatively small. The reason for this is that the posterior standard errors of these parameters reported in Table 3.3 are those related to draws in the third substep. During this step, the model-fit did not improve much anymore, resulting in relatively low acceptance rates and, hence, small posterior standard errors. Furthermore, the average relative pricing errors seem to be particularly large for A-rated countries. This is caused by the fact that

CDS spreads of these countries are relatively low, and, therefore, small absolute pricing errors are translated in large relative pricing errors.

3.4.2 *Estimation default event risk premium parameter*

In the second step, I estimate the default event risk premium parameter μ, which defines the ratio between the actual and risk-neutral default intensities (i.e., $\lambda^{\mathbb{Q}}_{j,t} = \mu\lambda^{\mathbb{P}}_{j,t}$, $j \in \{c, 1, 2, \ldots, K\}$). As mentioned in Section 3.3, this parameter can, in principle, be time-varying and different across j, $j \in \{c, 1, 2, \ldots, K\}$, but, given the scarcity of historic sovereign default data, I assume it to be constant and the same across factors. This means that I focus on the average risk premium on default events rather than exploring time-varying aspects of it.

In estimating μ, I follow the procedure proposed by Driessen (2005). That is, I estimate μ by using moment conditions for the conditional default probabilities, which are defined as the probabilities of defaulting in year $t+n$, conditional upon no default between time t and $t+n-1$ (and the average credit rating during the sample period). These moment conditions are given by

$$\mathbb{E}^{\mathbb{P}}_t \left[J_{i,t+n} | R_{i,t} = R, J_{i,t} + J_{i,t+1} + \ldots + J_{i,t+n-1} = 0 \right] = q_{n,R}(\mu, \phi), \\ n = 0, \ldots, 9, R = A, BBB, BB, B, \quad (3.9)$$

where $J_{i,t}$ is a variable that is equal to 1 if country i defaults in the annual time interval $[t, t+1]$, R is the average credit rating of the country during the sample period, and $q_{n,R}(\mu, \phi)$ is the model-implied conditional default probability under the actual probability measure, and ϕ is a parameter vector containing all other parameters of the model.

I use average historical cumulative default rates provided by S&P to consistently estimate the left-hand side of (3.9). In particular, I use the cumulative default rates up to 10 years, since the longest maturity contract in our sample is 10 year. I convert the cumulative default probabilities into yearly conditional default rates $q^{\text{data}}_{n,R}$. Note that I condition on the average rating of country i in the left-hand side of (3.9), since the historical cumulative default rates are categorized by rating class.

The model-implied conditional default rates can be computed explicitly. First, I note that the actual probability that country i defaults within the

next n years, conditional upon that no default has occurred yet, is given by

$$\begin{aligned} p_{i,n,R}(t,\mu,\phi) &= \mathbb{E}^{\mathbb{P}}[J_{i,t} + J_{i,t+1} + \ldots + J_{i,t+n-1} | R_i = R] \\ &= 1 - \mathbb{E}_t^{\mathbb{P}}\left[(1-\gamma_i)^{N_{i,t+n}-N_{i,t}}\right] \mathbb{E}_t^{\mathbb{P}}\left[(1-\gamma_i^c)^{N_{c,t+n}-N_{c,t}}\right]. \end{aligned} \quad (3.10)$$

Because of the affine structure of the model, expression (3.10) can be computed explicitly up to a system of ODEs (see Appendix 3.A).[16] I average out (3.10) over all days in our sample period and denote the obtained probabilities by $p_{i,n,R}(\mu,\phi)$. The yearly conditional default rates are now given by $q_{i,n,R}(\mu,\phi) = 1-(1-p_{i,n+1,R}(\mu,\phi))/(1-p_{i,n,R}(\mu,\phi))$. In a last step, I average the conditional default probabilities over all countries in a given rating category to obtain $q_{n,R}(\mu,\phi)$.

I now estimate μ by using the first step of the generalized method of moments and minimize the sum of squared differences between the model-implied and observed conditional default rates over μ, inserting the estimates for the other parameters $\hat{\phi}$:

$$\min_{\mu}\left[\sum_{R=A,BBB,BB,B}\sum_{n=0}^{9}\left(q_{n,R}(\mu,\hat{\phi}) - q_{n,R}^{data}\right)^2\right]. \quad (3.11)$$

3.4.2.1 *Estimation results default event risk premium parameter*

Using the estimation procedure detailed in the previous section, I find $\hat{\mu} = 2.07$ (see Table 3.4). This implies that investors multiply (instantaneous) default probabilities with a factor of over two when pricing sovereign credit default swaps. This estimated value is in line with values of default event risk premia found previously in the literature on corporate default risk (see, e.g., Driessen, 2005; Berndt et al., 2008). The standard error of the estimate, which takes into account estimation errors in the historical default rates, is 0.61, indicating that μ is statistically significantly larger than 1 at the 5% significance level.[17]

16 Note that the expectations are taken over the distributions implied by the $\mathbb{P}$-dynamics of $\lambda_{j,t}^{\mathbb{P}}$.

17 Unlike Driessen (2005), I do not correct the standard errors for possible biases that might result from the use of estimated process parameters when computing the model-implied default rates $q_{n,R}(\mu,\hat{\phi})$. The reason for this is that the iterative Bayesian setup employed in this chapter does not provide a suitable estimated covariance matrix of all

Table 3.4. Estimates of the default event risk premium parameter μ. This table reports estimates and standard errors (in brackets) of the default event risk premium parameter μ, obtained using S&P historical sovereign cumulative default rates data, Moody's historical sovereign cumulative default rates data, and S&P historical corporate cumulative default rates data. $*$ and $**$ denote significance at the 10% and 5% level, respectively.

	Overall	A	BBB	BB	B
S&P sovereign	2.07**	0.75	2.03	1.59	2.12*
	(0.61)	(0.39)	(1.46)	(0.83)	(0.68)
Moody's sovereign	2.52				
S&P corporate	2.06				

In Figure 3.4, I illustrate the effect of the risk premium parameter μ on default probabilities. For every rating class, the line "Risk-neutral" (solid grey) depicts the risk-neutral model-implied conditional default probabilities (i.e., using the $\mathbb{Q}$-dynamics of $\lambda^{\mathbb{Q}}$ in calculating the default probabilities). The line "$\mu = 1$" (dotted grey) plots the actual model-implied default probabilities, assuming that there is no default event risk premium (i.e., using the $\mathbb{P}$-dynamics of $\lambda^{\mathbb{Q}}$ and assuming $\mu = 1$ in calculating the default probabilities). The difference between these lines is completely caused by the risk premia related to changes in default risk over time (i.e., distress risk premia). Next, the line "S&P historical data" (solid black) presents the empirical conditional default probabilities based on S&P historical default data. For all rating classes except A, the historical default probabilities lie completely below the "Risk-neutral" and "$\mu = 1$" lines, indicating that distress risk premia can not sufficiently explain observed default rates. Finally, the line "$\mu = 2.07$" (striped black) depicts the model-implied actual default probabilities, using $\mu = 2.07$. Taking into account default event risk premia clearly improves the fit of historical default probabilities for the BBB, BB and B ratings.

estimated parameters $\hat{\Phi}$, which is needed in the standard error correction proposed by Driessen (2005). However, given that sovereign default events are quite rare, it is likely that possible biases resulting from misspecifications of the real-world default rates dominate such error-in-variable biases. Following Driessen (2005), I do account for these biases, but nevertheless the statistical significance of the estimation results should be interpreted with caution.

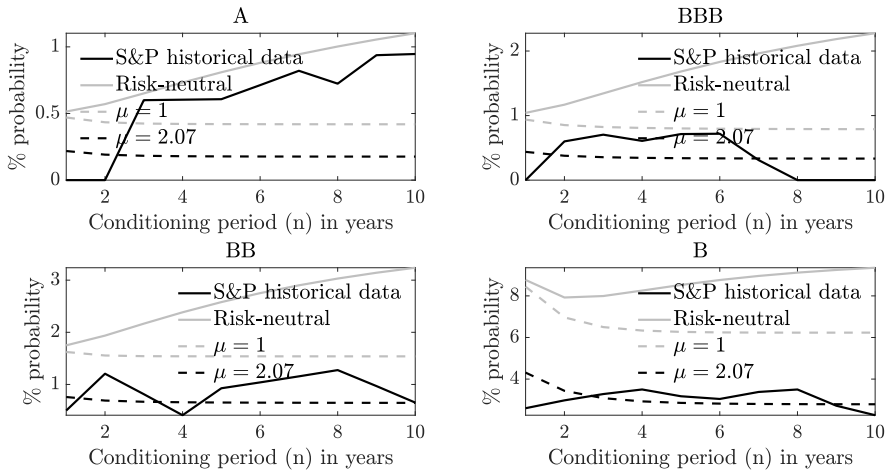


Figure 3.4. Historical and model-implied conditional default probabilities per rating class. This figure plots the conditional default probabilities for A-rated (upper left panel), B-rated (upper right panel), BB-rated (bottom left panel), and B-rated (bottom right panel) countries. The conditional default probability is defined as the probability of a country going into default in period $[n, n+1]$ given no default has occurred before year n. The "S&P historical data" (solid black) line gives the historically estimated conditional default probabilities obtained using S&P data. The line "$\mu = 2.07$" (striped black) represents the model-implied actual conditional default probabilities. The line "$\mu = 1$" (striped grey) represents the model-implied actual conditional default probabilities assuming absence of default event risk premia. The "risk-neutral" (solid grey) line represents model-implied actual conditional default probabilities assuming absence of both distress and default event risk premia.

I perform several robustness checks. First, I also estimate the default event risk premium parameter μ per rating class. For the BBB, BB, and B rating classes I find $\hat{\mu} = 2.03$, $\hat{\mu} = 1.59$, and $\hat{\mu} = 2.12$, respectively (see Table 3.4). Except for the B rating, however, these estimates are not statistically significantly larger than 1. These results are, however, reasonably close to the overall estimate $\hat{\mu} = 2.07$, indicating that there is not much variation in the default event risk premium across countries from these rating classes. For the A rating class, on the other hand, I find $\hat{\mu} = 0.75$, which suggests that there is even a negative default event risk premium for A-rated countries. This result is not statistically significant, but reveals that there are limitations in using historical default rates for higher rated countries. Since sovereign defaults are scarce, especially for

higher rated countries, one or two default events can result in substantial upwards biases in the historical default probability estimates, thereby affecting the default event risk premium estimate. In this case, the historical cumulative default rates of A-rated countries were adjusted upwards as a result of the double default of Greece. Second, I re-estimate μ using 1) average historical sovereign cumulative default rates provided by Moody's, and 2) average historical corporate default rates provided by S&P. I find $\hat{\mu} = 2.52$ ($\hat{\mu} = 2.06$) using Moody's sovereign (S&P corporate) default rates, which are very close to our original estimate using S&P sovereign default rates.[18]

3.5 CDS SPREAD DECOMPOSITION

The differences in the parameters governing the $\mathbb{P}$- and $\mathbb{Q}$-dynamics of $\lambda^{\mathbb{Q}}$ and the estimated value of the default event risk parameter μ indicate the presence of both distress and default event risk premia in sovereign CDS spreads. In this section, I explore the economic significance of these risk premia in more detail, and decompose CDS spreads into distress risk premia, default event risk premia, and expected default risk components.

Similar to Pan and Singleton (2008) and Longstaff et al. (2011), I quantify the magnitude of the distress risk premium by computing the difference in CDS spreads implied by the $\mathbb{P}$- and $\mathbb{Q}$-dynamics of $\lambda^{\mathbb{Q}}_{j,t}$. The CDS spread of country i implied by the $\mathbb{Q}$-dynamics of risk-neutral intensities is given by equation (3.6). This CDS spread includes both the market prices of risk $\xi_{c,t}$ and $\xi_{i,t}$ related to the dynamics of the common and country-specific factors, respectively. The CDS spread of country i implied by the $\mathbb{P}$-dynamics of the risk-neutral intensities, on the other hand, does not include these market prices of risk (i.e., $\xi_{c,t} = \xi_{i,t} = 0$), and the difference between the CDS spreads computed in these ways thus constitutes the distress risk premium embedded in the CDS spread.

I denote the CDS spread of country i implied by the $\mathbb{P}$-dynamics of the risk-neutral intensities by $\text{CDS}^{\mathbb{PQ}}_{i,t}(M)$. Here, the first superscript ($\mathbb{P}$)

18 For the estimates of μ using Moody's sovereign default rates and S&P corporate default rates I cannot provide standard errors, since the computation of these standard errors requires more detailed information on the sizes of the different cohorts used to construct the historical default rate estimates (see Driessen, 2005). Unfortunately, I do not possess this information for these data sets.

refers to the probability measure governing the dynamics of the intensity processes, and the second superscript ($\mathbb{Q}$) refers to the probability measure under which we consider the intensity values (i.e., the superscript $\mathbb{PQ}$ denotes the $\mathbb{P}$-dynamics of $\lambda^{\mathbb{Q}}$). The spread $CDS^{\mathbb{PQ}}_{i,t}(M)$ can thus be computed by using (3.6) and taking expectations with respect to the distributions implied by $\mathbb{P}$-dynamics of the risk-neutral intensities given in (3.2). That is,

$$CDS^{\mathbb{PQ}}_{i,t}(M) = \frac{(1-R)\int_t^{t+M} D(t,u)\mathbb{E}^{\mathbb{P}}_t\left[\left(\gamma_i\lambda^{\mathbb{Q}}_{i,u}+\gamma^c_i\lambda^{\mathbb{Q}}_{c,u}\right)(1-\gamma_i)^{N_{i,u}-N_{i,t}}(1-\gamma^c_i)^{N_{c,u}-N_{c,t}}\right]du.}{\frac{1}{4}\sum_{j=1}^{4M} D(t,t+0.25j)\mathbb{E}^{\mathbb{P}}_t\left[(1-\gamma_i)^{N_{i,t+0.25j}-N_{i,t}}\right]\mathbb{E}^{\mathbb{P}}_t\left[(1-\gamma^c_i)^{N_{c,t+0.25j}-N_{c,t}}\right]}. \quad (3.12)$$

Clearly, if the market prices of risk $\xi_{c,t}$ and $\xi_{i,t}$ are zero, $CDS_{i,t}(M)$ and $CDS^{\mathbb{PQ}}_{i,t}(M)$ are the same, and there is no distress risk premium. If, on the other hand, $\xi_{c,t}$ or $\xi_{i,t}$ are non-zero, $CDS_{i,t}(M)$ and $CDS^{\mathbb{PQ}}_{i,t}(M)$ are not the same and the difference between the two, $[CDS_{i,t}(M) - CDS^{\mathbb{PQ}}_{i,t}(M)]$, constitutes the distress risk premium. I also investigate the distress risk premium in relative terms, which is given by $[CDS_{i,t}(M) - CDS^{\mathbb{PQ}}_{i,t}(M)]/CDS_{i,t}(M)$.

Both $CDS_{i,t}(M)$ and $CDS^{\mathbb{PQ}}_{i,t}(M)$ still contain the default event risk premium parameter μ, since they consider the risk-neutral common and country-specific intensities $\lambda^{\mathbb{Q}}_{c,t}$ and $\lambda^{\mathbb{Q}}_{i,t}$, respectively. To extract the default event risk premium, I can thus go one step further and compute the CDS spreads implied by the $\mathbb{P}$-dynamics of $\lambda^{\mathbb{P}}_{j,t}$, which I denote by $CDS^{\mathbb{PP}}_{i,t}(M)$:

$$CDS^{\mathbb{PP}}_{i,t}(M) = \frac{(1-R)\int_t^{t+M} D(t,u)\mathbb{E}^{\mathbb{P}}_t\left[\left(\gamma_i\lambda^{\mathbb{P}}_{i,u}+\gamma^c_i\lambda^{\mathbb{P}}_{c,u}\right)(1-\gamma_i)^{N_{i,u}-N_{i,t}}(1-\gamma^c_i)^{N_{c,u}-N_{c,t}}\right]du.}{\frac{1}{4}\sum_{j=1}^{4M} D(t,t+0.25j)\mathbb{E}^{\mathbb{P}}_t\left[(1-\gamma_i)^{N_{i,t+0.25j}-N_{i,t}}\right]\mathbb{E}^{\mathbb{P}}_t\left[(1-\gamma^c_i)^{N_{c,t+0.25j}-N_{c,t}}\right]}. \quad (3.13)$$

$CDS^{\mathbb{PP}}_{i,t}(M)$ thus represents the CDS spread absent of any risk premia and is a measure of expected default risk. The difference between $CDS^{\mathbb{PQ}}_{i,t}(M)$ and $CDS^{\mathbb{PP}}_{i,t}(M)$, $[CDS^{\mathbb{PQ}}_{i,t}(M) - CDS^{\mathbb{PP}}_{i,t}(M)]$ is then the default event risk premium embedded in the CDS spread. The relative default event risk premium is given by $[CDS^{\mathbb{PQ}}_{i,t}(M) - CDS^{\mathbb{PP}}_{i,t}(M)]/CDS_{i,t}(M)$.

Table 3.5 reports the average absolute and relative decompositions of five-year CDS spreads per rating class.[19] A few clear patterns emerge: First, I find a strong decreasing pattern in the relative importance of distress risk premia as the rating gets lower. For example, the distress risk premium makes up, on average, 60.8% of CDS spreads of A-rated countries, but only 5.9% of CDS spreads of B-rated countries. Second, I find increasing patterns in both the relative importance of default event risk premia and expected default risk components as the rating gets lower. The default event risk premium and default risk component constitute, on average, 22.0%, and 17.2% of the CDS spread of A-rated countries, respectively. For B-rated countries these relative weights are 52.9%, and 41.3%, respectively. Intuitively, as a default event becomes more likely (i.e., countries with a lower credit rating), investors start caring relatively more about default event risk than distress risk.

The results discussed above focus on the decomposition of five-year CDS spreads. In the term-structure dimension, however, also several interesting patterns emerge. Table 3.6 shows that for all countries the fraction attributable to the distress risk premium increases, whereas the fractions attributable to the default event risk premium and default risk component decrease as the maturity gets longer. This suggests that investors mainly worry about actual default events in short-term horizons. For longer horizons, on the other hand, investors worry more about the uncertainty around future default probabilities.[20]

19 The results on the rating class level are obtained by taking the average of the countries in that rating class. Since Venezuela is quite different from the other countries under consideration, I also report the results for B-rated countries excluding Venezuela. I find that all results stay qualitatively the same when excluding Venezuela from the sample.

20 This result is partly driven by the fact that within the class of jump-diffusion models, the (relative) importance of diffusive (i.e., Brownian) shocks in the intensities builds up for longer maturities because of the persistence in the process dynamics. As a consequence, jump risk is more important for short-term horizons.

Table 3.5. Decomposition of five-year CDS spreads into expected risk and risk premia components. This table reports the average absolute and relative (in brackets) decompositions of five-year CDS spreads into distress risk premia, default event risk premia, and expected default risk components per rating class. For every rating class, the decomposition is computed by taking the average over the individual decompositions of all countries in that rating class. The rating class B* denotes the results of B-rated countries, where Venezuela is excluded.

	Distress premium	Default event premium	Expected default risk
A	54.03	18.49	14.33
	(60.8%)	(22.0%)	(17.2%)
BBB	94.06	36.21	28.12
	(60.2%)	(22.5%)	(17.4%)
BB	100.54	67.29	52.03
	(46.1%)	(30.3%)	(23.6%)
B	-15.84	389.91	284.65
	(5.9%)	(52.9%)	(41.3%)
B*	68.99	169.92	133.39
	(13.1%)	(48.6%)	(38.3%)

Table 3.6. Decomposition of CDS spreads into risk premia components across maturities. This table reports the average relative (panel A) and absolute (panel B) decomposition of CDS spreads into distress risk premia, default event risk premia, and expected default risk components per rating class across different maturities. For every rating class, the decomposition is computed by taking the average over the individual decompositions of all countries in that rating class. The rating class B* denotes the results of B-rated countries, where Venezuela is excluded.

Panel A: Relative decomposition

	Distress premium				Default event premium				Expected default risk			
	2y	3y	5y	10y	2y	3y	5y	10y	2y	3y	5y	10y
A	40.3%	49.8%	60.8%	72.2%	32.6%	27.8%	22.0%	15.8%	27.0%	22.4%	17.2%	12.0%
BBB	40.6%	49.6%	60.2%	71.0%	32.6%	28.0%	22.5%	16.6%	26.8%	22.4%	17.4%	12.5%
BB	27.6%	35.7%	46.1%	57.7%	39.6%	35.6%	30.3%	24.1%	32.9%	28.7%	23.6%	18.2%
B	-0.4%	1.8%	5.9%	11.9%	54.9%	54.4%	52.9%	50.2%	45.5%	43.8%	41.3%	37.8%
B*	4.9%	8.1%	13.1%	20.2%	51.9%	50.8%	48.6%	45.2%	43.2%	41.2%	38.3%	34.7%

panel B: Absolute decomposition

	Distress premium				Default event premium				Expected default risk			
	2y	3y	5y	10y	2y	3y	5y	10y	2y	3y	5y	10y
A	25.5	35.9	54.0	88.2	18.7	18.6	18.5	18.4	15.4	14.9	14.3	13.9
BBB	46.1	64.0	94.1	146.3	37.3	36.8	36.2	35.6	30.8	29.5	28.1	26.8
BB	49.7	68.7	100.5	154.6	66.5	66.8	67.3	67.9	54.9	53.5	52.0	50.7
B	-35.9	-31.6	-15.8	16.2	417.1	404.3	389.9	376.6	330.8	309.5	284.7	260.6
B*	38.7	50.6	69.0	97.5	164.1	166.5	169.9	173.9	136.0	134.5	133.4	132.8

In addition to decomposing CDS spreads in risk premia and default risk components, the model also allows for a decomposition of CDS spreads into systemic risk and country-specific risk components. To compute the systemic risk component, I take (3.6) and ignore the country-specific part. That is,

$$\text{CDS}_{i,t}^{\text{systemic}}(M) = \frac{(1-R)\int_t^{t+M} D(t,u)\mathbb{E}_t^Q\left[\gamma_i^c \lambda_{c,u}^Q (1-\gamma_i^c)^{N_{c,u}-N_{c,t}}\right] du.}{\frac{1}{4}\sum_{j=1}^{4M} D(t,t+0.25j)\mathbb{E}_t^Q\left[(1-\gamma_i^c)^{N_{c,t+0.25j}-N_{c,t}}\right]}. \tag{3.14}$$

Similarly, to compute the country-specific risk component, I take (3.6) and ignore the common factor. That is,

$$\text{CDS}_{i,t}^{\text{country}}(M) = \frac{(1-R)\int_t^{t+M} D(t,u)\mathbb{E}_t^Q\left[\gamma_i \lambda_{i,u}^Q (1-\gamma_i)^{N_{i,u}-N_{i,t}}\right] du.}{\frac{1}{4}\sum_{j=1}^{4M} D(t,t+0.25j)\mathbb{E}_t^Q\left[(1-\gamma_i)^{N_{i,t+0.25j}-N_{i,t}}\right]}. \tag{3.15}$$

Table 3.7 displays the results of the decomposition of five-year CDS spreads in systemic and country-specific risk components. I find a relatively stable decomposition across rating classes, where the country-specific and systemic components account for approximately 65% and 35% of CDS spreads, respectively.

Table 3.7. Decomposition of five-year CDS spreads into country-specific and systemic risk components. This table reports the absolute and relative decomposition of five-year CDS spreads into country-specific and systemic risk components per rating class. For every rating class, the decomposition is computed by taking the average over the individual decompositions of all countries in that rating class. The rating class B* denotes the results of B-rated countries, where Venezuela is excluded.

	Country-specific risk	Country-specific risk fraction	Systemic risk	Systemic risk fraction
A	59.40	67.8%	27.58	32.3%
BBB	94.80	56.8%	63.96	43.5%
BB	143.03	64.4%	77.39	35.9%
B	432.65	66.9%	214.53	32.8%
B*	276.74	73.3%	79.21	26.9%

Lastly, the decompositions outlined above can also be combined. In Table 3.8, I show the results of this two-dimensional decomposition, in which I first decompose the spreads in country-specific and systemic risk components, and then decompose both these parts in the risk premia and expected default risk components. I find that the decomposition of the systemic part (especially in relative terms) is very similar across rating classes: The systemic distress risk premium, systemic default event risk premium, and systemic default risk component account for approximately 20%, 9%, and 6% of CDS spreads across rating classes. The sub-decomposition of the country-specific component, however, differs substantially across rating classes. I find a strong decreasing pattern in the relative importance of country-specific distress risk premia as the rating gets lower. The country-specific default event risk premium and country-specific default risk component, on the other hand, become relatively more important as the rating gets lower. These results thus show that the patterns found in Table 3.5 are mainly due to the differences in country-specific risk components.

Table 3.8. Decomposition of five-year CDS spreads into country-specific and systemic expected risk and risk premia components. This table reports the average relative (panel A) and absolute (panel B) decomposition of five-year CDS spreads into country-specific distress risk premia, country-specific default event risk premia, expected country-specific default risk components, systemic distress risk premia, systemic default event risk premia, and expected systemic default risk per rating class. For every rating class, the decomposition is computed by taking the average over the individual decompositions of all countries in that rating class. The rating class B* denotes the results of B-rated countries, where Venezuela is excluded.

	Panel A: Relative decomposition					
	Country-specific distress premium	Country-specific default event premium	Country-specific expected default risk	Systemic distress premium	Systemic default event premium	Systemic expected default risk
A	41.9%	14.0%	11.9%	19.1%	8.0%	5.3%
BBB	35.0%	11.6%	10.2%	25.5%	10.8%	7.2%
BB	25.4%	21.3%	17.6%	21.0%	9.0%	6.0%
B	-12.5%%	44.1%	35.3%	18.4%	8.5%	5.9%
B*	-2.3%	41.8%	33.8%	15.7%	6.7%	4.5%
	panel B: Absolute decomposition					
A	37.7	11.8	9.9	16.5	6.7	4.4
BBB	56.5	20.5	17.8	38.0	15.7	10.3
BB	55.5	48.1	39.4	45.7	19.1	12.6
B	-137.2	326.8	243.1	117.4	57.1	40.0
B*	13.5	145.8	117.4	56.5	23.9	15.9

3.6 CONCLUSIONS

In this chapter, I investigate credit risk premia embedded in sovereign CDS spreads. In particular, I consider distress risk premia, which are defined as the compensation investors demand for unpredictable variations in future default arrival rates, and default event risk premia, which are defined as the compensation investors demand for the risk of default events themselves.

I propose a new model for the term-structure of sovereign credit risk and assume that the default of a country can be triggered by either a common, systemic factor, or by an independent country-specific factor. The novelty of the model is that I specify both the common and country-specific factors to be self-exciting jump processes. In this way, the model can capture, in a parsimonious way, the commonality in sovereign credit risk and the clustering of large credit shocks over time and across countries observed in the data. I estimate the model using sovereign CDS data and historical sovereign default rates per rating class from S&P.

The model facilitates a decomposition of CDS spreads along two dimensions: First, I decompose CDS spreads in country-specific and systemic risk components. I find a similar decomposition across rating classes and show that approximately 65% of (five-year) CDS spreads can be attributed to country-specific risk and 35% of CDS spreads can be attributed to systemic risk. Second, I decompose CDS spreads into distress risk premia, default event risk premia, and default risk components. I find that the distress risk premium is mainly priced in CDS spreads of countries with higher credit ratings and longer maturities, whereas the default event risk premium and default risk component are relatively more important for countries with lower credit ratings and for shorter maturities. I show that the dissimilarities of the decompositions across rating classes are mainly driven by differences in country-specific risk.

3.A CLOSED-FORM PRICING FORMULAS

In this Appendix, I show how the expectations appearing in the CDS pricing formula (3.6) can be computed explicitly (up to a system of ODEs).

Let $X_{j,t} = (N_{j,t}, \lambda^{Q}_{j,t})', j \in \{c, 1, 2, \ldots, K\}$. The dynamics of $X_{j,t}$ are given by

$$\begin{aligned} dX_{j,t} &= d\begin{pmatrix} N_{j,t} \\ \lambda^{Q}_{j,t} \end{pmatrix} \\ &= \begin{pmatrix} 0 \\ \alpha^{Q}_{j}(\lambda^{Q}_{j,\infty} - \lambda^{Q}_{j,t}) \end{pmatrix} dt + \begin{pmatrix} 0 & 0 \\ 0 & \sigma_j\sqrt{\lambda^{Q}_{j,t}} \end{pmatrix} d\begin{pmatrix} 0 \\ W^{Q}_{j,t} \end{pmatrix} \\ &\quad + \begin{pmatrix} 1 \\ Z_j \end{pmatrix} dN_{j,t}. \end{aligned} \tag{3.A.1}$$

From this specification it is clear that the process $X_{j,t}$ falls into the generalized affine jump-diffusion framework and, therefore, I can use the framework of Duffie et al. (2000) and prove the following Propositions:

Proposition 3.

$$\mathbb{E}^{Q}\left[(1-\gamma_j)^{N_{j,T}}\middle|\mathcal{F}_t\right] = e^{\alpha(t)+\beta_1(t)N_{j,t}+\beta_2(t)\lambda^{Q}_{j,t}},$$

with

$$\begin{aligned} \dot{\alpha}(t) &= -\alpha^{Q}_{j}\lambda^{Q}_{j,\infty}\beta_2(t), \\ \alpha(T) &= 0, \\ \dot{\beta}_1(t) &= 0, \\ \beta_1(T) &= \beta_1(s) = \log(1-\gamma_j) \quad t \leqslant s \leqslant T, \\ \dot{\beta}_2(t) &= \alpha^{Q}_{j}\beta_2(t) - \frac{1}{2}\beta_2^2(t)\sigma_j^2 - \left(e^{\beta_1(t)+Z_j\beta_2(t)} - 1\right), \\ \beta_2(T) &= 0. \end{aligned}$$

Proof. Consider an affine jump-diffusion process X in some state space $D \subset \mathbb{R}^n$ solving the stochastic differential equation

$$dX_t = \mu(X_t)dt + \sigma(X_t)dW_t + \sum_{i=1}^{m} dZ^i_t,$$

where Z^i are pure jump processes whose jumps have a fixed probability distribution ν^i on $\mathbb{R}^n$ and arrive with intensity $\lambda^i(X_t)$ for some

$\lambda^i : D \to [0,\infty)$. Let us fix an affine process $R : D \to \mathbb{R}$. Then we have that the complete affine structure of the model is captured by:

$$\begin{aligned}
\mu(x) &= K_0 + K_1 x, \text{ for } K = (K_0, K_1) \in \mathbb{R}^n \times \mathbb{R}^{n\times n}.\\
\sigma(x)\sigma(x)^\top &= H_0 + \sum_{k=1}^{n} H_1^{(k)} x_k, \text{ for } H = (H_0, H_1) \in \mathbb{R}^{n\times n} \times \mathbb{R}^{n\times n\times n}.\\
\lambda^i(x) &= l_0^i + l_1^i \cdot x, \text{ for } l = (l_0, l_1) \in \mathbb{R} \times \mathbb{R}^n.\\
R(x) &= \rho_0 + \rho_1 \cdot x, \text{ for } \rho = (\rho_0, \rho_1) \in \mathbb{R} \times \mathbb{R}^n.
\end{aligned}$$

Let us furthermore denote the jump-transforms, which determine the jump-size distributions, as $\theta^i(c) = \int_{\mathbb{R}^n} \exp(c \cdot z)\, d\nu^i(z)$ for $c \in \mathbb{C}^n$. We want to compute an expression of the form

$$\phi^X(u, X, t, T) = \mathbb{E}^X\left[\exp\left(-\int_t^T R(X_s)ds\right) e^{u\cdot X_T}\middle|\, \mathcal{F}_t\right].$$

According to Proposition 1 of Duffie et al. (2000), we have, under some technical assumptions on the processes being well-behaved, that we can write

$$\phi^X(u, x, t, T) = e^{\alpha(t)+\beta(t)\cdot x},$$

where β and α satisfy the following (complex-valued) ODEs:

$$\begin{aligned}
\dot\beta(t) &= \rho_1 - K_1^\top \beta(t) - \frac{1}{2}\beta(t)^\top H_1 \beta(t) - \sum_{i=1}^{m} l_1^i(\theta^i(\beta(t)) - 1)\\
\dot\alpha(t) &= \rho_0 - K_0 \cdot \beta_t - \frac{1}{2}\beta(t)^\top H_0 \beta(t) - \sum_{i=1}^{m} l_0^i(\theta^i(\beta(t)) - 1),
\end{aligned}$$

with boundary conditions $\beta(T) = u$ and $\alpha(T) = 0$.

Applying the situation described above to the process defined in (3.A.1), we have $n = 2$, $m = 1$, $K_0 = (0, \alpha_j^Q \lambda_{j,\infty}^Q)^\top$,

$$\begin{aligned}
K_1 &= \begin{pmatrix} 0 & 0 \\ 0 & -\alpha_j^Q \end{pmatrix},\\
H_0 &= \mathbf{0},\\
H_1^1 &= \begin{pmatrix} 0 & 0 \\ 0 & 0 \end{pmatrix}, \quad H_1^2 = \begin{pmatrix} 0 & 0 \\ 0 & \sigma_j^2 \end{pmatrix},
\end{aligned}$$

$l_0^1 = 0$ and $l_1^1 = (0,1)^\top$. Since $N_{j,t}^{\mathbb{Q}}$ is a counting process, and the coefficient Z_j is a constant, we have fixed jump sizes and therefore $\theta^1(c) = \exp(c_1 + Z_j c_2)$. Since we want to compute $\mathbb{E}^{\mathbb{Q}}\left[(1-\gamma_j)^{N_{j,T}} \middle| \mathcal{F}_t\right]$, we have that $\rho_0 = 0$, $\rho_1 = (0,0)^\top$ and $u = (\log(1-\gamma_j), 0)^\top$. The result thus follows by applying Proposition 1 of Duffie et al. (2000). □

Proposition 4.

$$\mathbb{E}^{\mathbb{Q}}\left[\gamma_j \lambda_{j,T}^{\mathbb{Q}} (1-\gamma_j)^{N_{j,T}} \middle| \mathcal{F}_t\right] = e^{\alpha(t) + \log(1-\gamma_j) N_{j,t} + \beta_2(t)\lambda_{j,t}^{\mathbb{Q}}} \times (A(t) + B_2(t)\lambda_{j,t}^{\mathbb{Q}}),$$

with

$$\begin{aligned} -\dot{A}(t) &= \alpha_j^{\mathbb{Q}} \lambda_{j,\infty}^{\mathbb{Q}} B_2(t), \\ -\dot{B}_2(t) &= -\alpha_j^{\mathbb{Q}} B_2(t) + \beta_2(t)\sigma_j^2 B_2(t) + (1-\gamma_j) Z_j e^{Z_j \beta_2(t)}, \\ B_2(T) &= \gamma_j, \\ A(T) &= 0, \end{aligned}$$

and $\alpha(t)$ *and* $\beta(t)$ *satisfy the ODEs in the previous proposition.*

Proof. Proposition 3 of Duffie et al. (2000) with $u = (\log(1-\gamma_j), 0)^\top$ and $v = (0, \gamma_j)^\top$ yields the result. □

3.B MCMC DETAILS

In this Appendix, I provide more details of the estimation methodology used to estimate the model of the risk-neutral intensities. As explained in Section 3.4.1, the estimation of the risk-neutral model is divided into three steps. Since all these steps rely on the same MCMC procedure, I briefly explain the set-up of the first and third steps in which I keep the country-specific factors fixed and only consider the common factor with country-specific loadings γ_i^c to this factor. I denote the fixed country-specific parts of the model with Ψ. That is, Ψ contains all parameters related to the country-specific factors, as well as the country-specific jump times and intensities.[21]

21 $\Psi = \{\alpha_i^{\mathbb{P}}, \alpha_i^{\mathbb{P}}\lambda_{i,\infty}^{\mathbb{P}}, \sigma_i^2, \alpha_i^{\mathbb{Q}}, \alpha_i^{\mathbb{Q}}\lambda_{i,\infty}^{\mathbb{Q}}, Z_i, \gamma_i, h_i, \{\lambda_{i,t}^{\mathbb{Q}}\}_{t=1}^N, \{\Delta N_{i,t}\}_{t=1}^N, i \in \{1,\dots,K\}\}$. In the first substep, I ignore the country-specific factors, which essentially boils down to setting all country-specific parameters and intensities equal to zero. In the third substep, I fix the country-specific components to the estimates found in substep two.

I first denote the right-hand side of the CDS pricing formula (3.6) as a function F of a vector with the common factor parameters and country-specific loadings to this factor, $\Theta_i = \{\alpha_c^{\mathbb{Q}}, \alpha_c^{\mathbb{Q}}\lambda_{c,\infty}^{\mathbb{Q}}, \sigma_c^2, Z_c, \gamma_i^c\}$, $i = 1, \dots, K$, state variables $X_t = \{\Delta N_{c,t}, \lambda_{c,t}^{\mathbb{Q}}\}$, maturity M, and fixed country-specific parts. That is

$$\mathrm{CDS}_{i,t}(M) = F(X_t, \Theta_i, M, \Psi). \tag{3.B.2}$$

Since the data does not start on the same date for all countries, I let $S(t)$ denote the set of CDS contracts for which I have an observation on day $t, t = 1, \dots, N$. I assume that the (log) CDS spreads $Y_{t,k}$, $k \in S(t)$ are observed with normally distributed pricing errors, that is

$$Y_{t,k} = \log(F(X_t, \Theta_k, M, \Psi)) + \xi_{t,k}, \quad \forall k \in S(t), t = 1, \dots, N, \tag{3.B.3}$$

with $\xi_{t,k} \sim \mathcal{N}(0, h_c^2)$.[22]

Let $Y = \{Y_{t,k} : t = 1, \dots, N, k \in S(t)\}$ denote the vector of all CDS price observations, $X = \{X_t : t = 1 \dots, N\}$ the vector with all states, and $\bar{\Theta} = \{h_c^2, \alpha_c^{\mathbb{Q}}, \alpha_c^{\mathbb{Q}}\lambda_{c,\infty}^{\mathbb{Q}}, \sigma_c^2, Z_c, \{\gamma_i^c\}_{i=1}^K, \alpha_c^{\mathbb{P}}, \alpha_c^{\mathbb{P}}\lambda_{c,\infty}^{\mathbb{P}}\}$ the vector with all parameters related to the common factor (i.e., both the parameters governing the $\mathbb{P}$- and $\mathbb{Q}$-dynamics of $\lambda_{c,t}^{\mathbb{Q}}$). From now on, I drop the fixed Ψ from the notation. Then the conditional density of the joint observations can be written as

$$\begin{aligned} p(Y|X, \bar{\Theta}) &\propto \prod_{t=1}^{N} \prod_{k \in S(t)} \phi(Y_{t,k}; F(X_{k,t}, \Theta_k, M), h^2) \\ &=: \prod_{t=1}^{N} p(Y_t|X_t, \Theta), \end{aligned} \tag{3.B.4}$$

where $\phi(x; m, s^2)$ denotes a normal density with mean m and variance s^2 evaluated at x.

The full joint posterior density $p(Y, X, \bar{\Theta})$ is then given by

$$p(Y, X, \bar{\Theta}) \propto p(Y|X, \bar{\Theta}) p(X|\bar{\Theta}) p(\bar{\Theta}), \tag{3.B.5}$$

where $p(\bar{\Theta})$ is the prior for $\bar{\Theta}$. I choose our priors to be proper but in such a way that they impose little information.

22 For every country, the pricing formula (3.6) of the two-, three-, five-, and ten-year maturity CDS spreads depend on the same parameters and state variables. Therefore Θ_k is the same for CDS spreads of the same country across different maturities.

Using the Markovian property of the joint intensity and jump process (see Aït-Sahalia et al., 2014), I can rewrite this as a product over the observations times:

$$p(Y, X, \bar{\Theta}) \propto p(\bar{\Theta}) \prod_{t=1}^{N} p(Y_t|X_t, \bar{\Theta}) p(X_t|X_{t-1}, \bar{\Theta}), \tag{3.B.6}$$

where $p(X_t|X_{t-1}, \bar{\Theta})$ is the transition density of the state process. As explained in Section 3.4.1, I use a discrete-time approximation of the transition densities, making it possible to express the transition density in closed-form.

Ultimately, the goal is to sample from the joint conditional posterior density $p(\bar{\Theta}, X|\bar{Y})$. The reason for this is that, by Bayes theorem, $p(\bar{\Theta}|Y) \propto p(\bar{\Theta}, X|Y)$. Hence, the sample average of $\bar{\Theta}^{(1)}, \bar{\Theta}^{(2)}, \ldots, \bar{\Theta}^{(G)}$, $\hat{\bar{\Theta}} = \frac{1}{G} \sum_{g=1}^{G} \bar{\Theta}^{(g)}$, can be used as the estimate of $\bar{\Theta}$. In a similar way, the latent jump intensities can be estimated by considering the sample averages $\lambda_{j,t}^{Q(g)}$ for all $t = 1, \ldots, N$. To estimate the jump times, i.e., to decide whether a jump occurred at time t, $t = 1, \ldots, N$, I define a threshold, $\omega > 0$, and say that a jump occurred at time t if $\frac{1}{G} \sum_{g=1}^{G} N_{j,t}^{(g)} > \omega$ (see Johannes et al., 1999).

Since the joint conditional posterior density is high-dimensional and nonstandard, it is not possible to sample from this density directly. In order to overcome these problems, I employ a Gibbs sampler, which sequentially draws all random variables from the joint posterior density.

The Gibbs sampler consists of the following steps, initialized by an appropriate set of starting values for X and $\bar{\Theta}$ when $g = 0$:

For $g = 1, \ldots, G$, simulate

1. $X_t^{(g+1)}$ from $p(X_t|X_{1:N\setminus t}^{(g)}, \bar{\Theta}^{(g)}, Y)$, for $t = 1, \ldots, N$, and
2. $\bar{\Theta}^{(g+1)}$ from $p(\bar{\Theta}|X^{(g+1)}, Y)$,

where $X_{1:N\setminus t}$ denotes the collection of state vectors X_s at all $s = 1, \ldots, N$ except at $s = t$. In the sections below, I first explain the details of the steps involved in drawing the new states, and next those of drawing the new parameters.

3.B.1 *Metropolis steps for drawing the states*

The drawing of the state vectors requires alternating between two different sampling schemes. This is to deal with the latency of the states, which makes it challenging for the algorithm to initialize the drawing of jumps. For a more detailed discussion on this, I refer to Sperna Weiland et al. (2019).

3.B.1.1 *State simulation scheme 1*

In the first simulation scheme, I use that $p(X_t|X^{(g)}_{1:N_{\setminus t}}, \bar{\Theta}^{(g)}, Y)$ is characterized by its full conditionals. Therefore, the drawing of $X_t^{(g+1)}$ can be split into the following two steps:

1. draw $\Delta N_{j,t}^{(g+1)}$ from $p(\Delta N_{j,t}|\lambda^{Q(g)}_{j,t}, X^{(g)}_{1:N_{\setminus t}}, \bar{\Theta}^{(g)}, Y)$;
2. draw $\lambda^{Q(g+1)}_{j,t}$ from $p(\lambda^{Q}_{j,t}|\Delta N^{(g+1)}_{j,t}, X^{(g)}_{1:N_{\setminus t}}, \bar{\Theta}^{(g)}, Y)$.

Under the discretization of the state transition dynamics introduced in (3.7), the full posterior of $\Delta N_{j,t}$ is a Bernoulli density with success probability

$$p(\Delta N_{j,t} = 1|\lambda^{Q(g)}_{j,t}, X^{(g)}_{1:N_{\setminus t}}, \bar{\Theta}^{(g)}, Y) = \frac{p(\lambda^{Q(g)}_{j,t}|\Delta N_{j,t} = 1, X^{(g+1)}_{t-1}, \bar{\Theta}^{(g)}) p(\Delta N_{j,t} = 1|X^{(g+1)}_{t-1})}{\sum_{s=0,1} p(\lambda^{Q(g)}_{j,t}|\Delta N_{j,t} = s, X^{(g+1)}_{t-1}, \bar{\Theta}^{(g)}) p(\Delta N_{j,t} = s|X^{(g+1)}_{t-1})}. \tag{3.B.7}$$

(3.B.7) is easy to compute, since $p(\lambda^{Q(g)}_{j,t}|\Delta N_{j,t} = s, X^{(g+1)}_{t-1}, \bar{\Theta}^{(g)})$ is a normal density and $p(\Delta N_{j,t} = s|X^{(g+1)}_{t-1})$ a Bernoulli with success probability $\lambda^{Q(g+1)}_{j,t-1}\Delta_t$. The credit jump probability does not depend on Y, since the CDS prices only depend on the intensities and parameter vector. Therefore, the jump probabilities are determined only by the state transition equations and not by the measurement equations. This makes it hard for the algorithm to draw initial jumps. For this reason, I alternate between this sampling scheme and the one explained below, which does take into account the measurement equations.

After having drawn $\Delta N_{j,t}^{(g+1)}$, the new intensity $\lambda_{j,t}^{\mathbb{Q}(g+1)}$ is drawn from the density

$$p(\lambda_{j,t}^{\mathbb{Q}}|\Delta N_{j,t}^{(g+1)}, X_{1:N_{\setminus t}}^{(g)}, \bar{\Theta}^{(g)}, Y) \propto p(Y_t|\lambda_{j,t}^{\mathbb{Q}}, \Delta N_{j,t}^{(g+1)}, \bar{\Theta}^{(g)}) \times p(\lambda_{j,t+1}^{\mathbb{Q}(g)}|\lambda_{j,t}^{\mathbb{Q}}, \Delta N_{j,t+1}^{(g)}, \bar{\Theta}^{(g)}) p(\lambda_{j,t}^{\mathbb{Q}}|\lambda_{j,t-1}^{\mathbb{Q}(g+1)}, \Delta N_{j,t}^{(g+1)}, \bar{\Theta}^{(g)}),$$

which is the product of a multivariate normal density and two univariate normal densities. This density is non-standard, and, therefore, I use a Metropolis step with proposal density $p(\lambda_{j,t}^{\mathbb{Q}}|\lambda_{j,t-1}^{\mathbb{Q}(g+1)}, \Delta N_{j,t}^{(g+1)}, \bar{\Theta}^{(g)})$, which is a normal distribution with mean $\lambda_{j,t-1}^{\mathbb{Q}(g+1)} + \alpha_j^{\mathbb{P}(g)}\lambda_{j,\infty}^{\mathbb{P}(g)}\Delta_t - \alpha_j^{\mathbb{P}(g)}\lambda_{j,t-1}^{\mathbb{Q}(g+1)}\Delta_t + Z_j^{(g)}\Delta N_{j,t}^{(g+1)}$ and variance $\sigma_j^{2(g)}\lambda_{j,t-1}^{\mathbb{Q}(g+1)}\Delta_t$.

Using this proposal density, the acceptance criterion becomes

$$\min\left(\frac{p(Y_t|X_t^{(g+1)}, \bar{\Theta}^{(g)})p(\lambda_{j,t+1}^{\mathbb{Q}(g)}|\Delta N_{j,t+1}^{(g)}, X_t^{(g+1)}, \bar{\Theta}^{(g)})}{p(Y_t|X_t^{(g)}, \bar{\Theta}^{(g)})p(\lambda_{j,t+1}^{\mathbb{Q}(g)}|\Delta N_{j,t+1}^{(g)}, X_t^{(g)}, \bar{\Theta}^{(g)})}, 1\right).$$

For the end point $t = N$ the acceptance criterion simplifies, since the terms $p(\lambda_{j,t+1}^{\mathbb{Q}(g)}|\Delta N_{j,t+1}^{(g)}, X_t^{(g+1)}, \bar{\Theta}^{(g)})$ and $p(\lambda_{j,t+1}^{\mathbb{Q}(g)}|\Delta N_{j,t+1}^{(g)}, X_t^{(g)}, \bar{\Theta}^{(g)})$ do not appear anymore in the numerator and denominator, respectively. For the starting point $t = 1$, I use a slightly different proposal density, since I cannot condition on X_{t-1}. I therefore draw $\lambda_{j,1}^{\mathbb{Q}(g+1)}$ from a normal density with mean $\lambda_{j,2}^{\mathbb{Q}(g)}$ and variance $\sigma_j^2\lambda_{j,2}^{\mathbb{Q}(g)}\Delta_2$. Denoting this proposal density as $q(\lambda_{j,1}^{\mathbb{Q}(g+1)}|\lambda_{j,2}^{\mathbb{Q}(g)}, \bar{\Theta})$ gives the following acceptance criterion:

$$\min\left(\frac{p(Y_1|X_1^{(g+1)}, \bar{\Theta}^{(g)})p(\lambda_{j,2}^{\mathbb{Q}(g)}|\Delta N_{j,2}^{(g)}, X_1^{(g+1)}, \bar{\Theta}^{(g)})q(\lambda_{j,1}^{\mathbb{Q}(g)}|\lambda_{j,2}^{\mathbb{Q}(g)}, \bar{\Theta}^{(g)})}{p(Y_1|X_1^{(g)}, \bar{\Theta}^{(g)})p(\lambda_{j,2}^{\mathbb{Q}(g)}|\Delta N_{j,2}^{(g)}, X_1^{(g)}, \bar{\Theta}^{(g)})q(\lambda_{j,1}^{\mathbb{Q}(g+1)}|\lambda_{j,2}^{\mathbb{Q}(g)}, \bar{\Theta}^{(g)})}, 1\right).$$

3.B.1.2 *State simulation scheme 2*

In the second simulation scheme, instead of drawing X_t in two steps, I sample the complete vector X_t at once from $p(X_t|X_{1:N_{\setminus t}}, \bar{\Theta}, Y)$. Equation (3.B.6) shows that the Markovian property of the state processes implies that I only need to consider the terms in (3.B.6) where (parts of) X_t enters directly. This gives

$$p(X_t|X_{1:N_{\setminus t}}^{(g-1)}, \bar{\Theta}, Y) \propto p(Y_t|X_t, \bar{\Theta})p(X_t|X_{t-1}^{(g)}, \bar{\Theta})p(X_{t+1}^{(g-1)}|X_t, \bar{\Theta}). \quad (3.B.8)$$

In this density, both the likelihoods as well as the transition densities play a role, and, therefore, the drawing of jumps depends on both the measurement equations as well as the transition densities of the states. By assumption of the normally distributed error terms $p(Y_t|X_t, \bar{\Theta})$ is multivariate normal with dimension equal to the number of observations at time t.

According to (3.8), the transition density can be written as

$$p(X_t|X_{t-1}, \bar{\Theta}) = p(\lambda^Q_{j,t}|\Delta N_{j,t}, X_{t-1}, \bar{\Theta})p(\Delta N_{j,t}|X_{t-1}, \bar{\Theta}), \qquad (3.B.9)$$

In order to draw $X_t^{(g+1)}$ from (3.B.8), I use the following proposal density

$$\begin{aligned} q(X_t|X^{(g)}_{1:N_{\setminus t}}, \bar{\Theta}^{(g)}) &= p(X_t|X_{t-1}^{(g+1)}, \bar{\Theta}^{(g)}) \\ &= p(\lambda^Q_{j,t}|\Delta N_{j,t}, N^l_t, X^{(g+1)}_{t-1}, \bar{\Theta}^{(g)})p(\Delta N_{j,t}|X^{(g+1)}_{t-1}, \bar{\Theta}^{(g)}). \end{aligned}$$

Under the discretization of the intensity processes, this is a mixture of normal distributions. When drawing from this distribution, one can first draw $\Delta N_{j,t}^{(g+1)}$ from an independent Bernoulli distribution with success probability $\lambda^{Q(g+1)}_{j,t-1}\Delta_t$, and then, given the outcome, draw $p(\lambda^Q_{j,t}|\Delta N^{(g+1)}_{j,t}, X^{(g+1)}_{t-1}, \Theta^{(g)})$ from a normal distribution with the appropriate mean (depending on outcome of the draw of $N_{j,t}^{(g+1)}$).

Using this proposal density, the acceptance criterion is as follows:

$$\min\left(\frac{p(Y_t|X_t^{(g+1)}, \Theta^{(g)})p(\lambda^{Q(g)}_{j,t+1}, \Delta N^{(g)}_{j,t+1}|X_t^{(g+1)}, \Theta^{(g)})}{p(Y_t|X_t^{(g)}, \Theta^{(g)})p(\lambda^{Q(g)}_{j,t+1}, \Delta N^{(g)}_{j,t+1}|X_t^{(g)}, \Theta^{(g)})}, 1\right),$$

where $p(\lambda^{Q(g)}_{j,t+1}, \Delta N^{(g)}_{j,t+1}|X_t^{(g+1)}, \bar{\Theta}^{(g)}) = p(\lambda^{Q(g)}_{j,t+1}|\Delta N^{c(g)}_{t+1}, X_t^{(g+1)}, \bar{\Theta}^{(g)}) \times p(\Delta N^{(g)}_{j,t+1}|X_t^{(g+1)}, \bar{\Theta}^{(g)})$. Since all densities are standard (i.e., multivariate normal, univariate normal or bernoulli), evaluating this acceptance criterion is straightforward. For the end points $t = 1$ and $t = N$ similar comments apply as in simulation scheme 1.

3.B.2 *Metropolis steps for drawing the parameters*

Next I explain the Metropolis step for estimating the parameters $\bar{\Theta}^{(g+1)}$ from $p(\bar{\Theta}|X^{(g+1)}, Y)$ in more detail.

For h_j^2, $j \in \{c, 1, 2, \ldots, K\}$, the inverse gamma distribution is a conjugate prior. This follows from $p(h_j^2|\bar{\Theta}\backslash\{h_j^2\}, X, Y) \propto p(Y|\bar{\Theta}, X)p(h_j^2)$, where $p(Y|\bar{\Theta}, X)$ is multivariate normal with diagonal variance matrix with h_j^2 as variance and $p(h_j^2)$ the prior inverse gamma density. Explicit computations are standard and are therefore omitted.

For the other parameters, I use random-walk Metropolis steps with Gaussian proposal densities with as mean vectors the previous draws and with a diagonal covariance matrices. I choose the priors on all parameters to be proper but uninformative in the sense that the prior variances should be high compared to the estimated posterior variances. In particular, I use Gamma$(0.02, 10)$ priors for $\alpha_j^{\mathbb{Q}}$, $\alpha_j^{\mathbb{P}}$, $\alpha_j^{\mathbb{Q}}\lambda_{j,\infty}^{\mathbb{Q}}$, $\alpha_j^{\mathbb{Q}}\lambda_{j,\infty}^{\mathbb{Q}}$, σ_j^2, Z_j, $j \in \{c, 1, 2, \ldots, K\}$, Uniform$(0, 1)$ priors for γ_i and γ_i^c, $i \in \{1, 2, \ldots, K\}$, and IG$(3, 0.1)$ priors for h_j^2, $j \in \{c, 1, 2, \ldots, K\}$. In line with theoretical parameter restrictions for nonnegativity and stationarity of the processes, the gamma priors make sure these parameters are nonnegative. The Uniform$(0, 1)$ priors for γ_i and γ_i^c reflect that these parameters denote the probabilities of going into default in the case a country-specific or common credit shock arrives, respectively, and hence should take values between zero and one. In general, the means and variances of all parameters are chosen arbitrarily, but such that the means are small and positive for all parameters and the variances relatively large compared to their means. All-in-all, results are robust against prior specification, since typically the likelihood contribution of the priors is small compared to the likelihood contribution of the data in the acceptance criteria. Furthermore, the posterior standard deviations are also much smaller than the prior standard deviations, indicating that our priors do not impose much information.

3.B.3 *Estimation country-specific components*

In the second substep of the estimation procedure, the common factor part of the model is fixed and the country-specific factors are estimated. In the notation of above, I now denote the fixed common factor part of the model by Ψ,[23] and write the right-hand side of the CDS pricing formula (3.6) as a function F of a vector with the country-specific parameters $\bar{\Theta} = \{h_i^2, \alpha_i^{\mathbb{Q}}, \alpha_i^{\mathbb{Q}}\lambda_{i,\infty}^{\mathbb{Q}}, \sigma_i^2, Z_i, \gamma_i, \gamma_i^c, \alpha_i^{\mathbb{P}}, \alpha_i^{\mathbb{P}}\lambda_{i,\infty}^{\mathbb{P}}\}$, country-specific states

23 More specifically, $\Psi = \{\alpha_c^{\mathbb{P}}, \alpha_c^{\mathbb{P}}\lambda_{c,\infty}^{\mathbb{P}}, \sigma_c^2, \alpha_c^{\mathbb{Q}}, \alpha_c^{\mathbb{Q}}\lambda_{c,\infty}^{\mathbb{Q}}, Z_c, h_c, \{\lambda_{c,t}^{\mathbb{Q}}\}_{t=1}^N, \{\Delta N_{c,t}\}_{t=1}^N\}$.

$X_{i,t} = \{\Delta N_{i,t}, \lambda_{i,t}^{Q}\}$, maturity M, and fixed common component Ψ. That is,

$$CDS_{i,t}(M) = F(\bar{\Theta}, X_{i,t}, M, \Psi). \tag{3.B.10}$$

Apart from this change, the mechanics of the estimation methodology are exactly identical to those of the first and third substeps outlined in this Appendix.

Part II

LABOR INCOME RISK

4 LABOR INCOME RISK AND STOCK RETURNS: THE ROLE OF HORIZON EFFECTS

This chapter is based on:

Eiling, E., De Jong, F., Laeven, R. J. A., Sperna Weiland, R. C. (2019). Labor Income Risk and Stock Returns: The Role of Horizon Effects, Unpublished working paper. University of Amsterdam, Amsterdam, and Tilburg University, Tilburg.

4.1 INTRODUCTION

Labor and stock markets are fundamentally interconnected. Many investors participate in the labor market and a large share of their wealth is typically tied up in risky human capital. Indeed, an estimated 90 to 93 percent of overall wealth is embedded in human capital, making it by far the largest asset class in the economy (Lustig et al., 2013; Palacios, 2015). In addition, for most firms labor is a key input and the human capital of their employees forms an important intangible asset for the firm.[1]

This chapter aims to shed new light on the labor-equity market linkages by examining different time horizons, ranging from one quarter to several years to the very long horizon. There are various economic reasons why linkages between stock markets and labor markets may be different at short, business cycle, and long horizons. The estimated contemporaneous short-term correlation between equity and labor income growth is typically low (e.g., Fama and Schwert, 1977; Cocco et al., 2005). This could be caused by wage stickiness. Wages are usually adjusted infrequently causing labor income data to be smoother than the marginal product of labor. As a result, wage rigidity makes the comovement between wage growth and output (which is perfectly correlated with the

1 The growing literature on labor and finance has developed substantial theoretical and empirical evidence of the interplay between labor markets and asset pricing (among many others, e.g., Danthine and Donaldson, 2002; Santos and Veronesi, 2006; Belo et al., 2014; Dittmar et al., 2016; Donangelo, 2014; Kuehn et al., 2017; Berk and Walden, 2013; Donangelo et al., 2018).

marginal product of labor) lower (e.g., Shimer, 2005; Hall, 2005). Smooth labor income can make it difficult to capture linkages with equity returns based on monthly or quarterly data.

At the same time, we would not expect labor and financial markets to continue to diverge at the very long horizon, as argued by Baxter and Jermann (1997). Indeed, evidence suggests that comovements are more pronounced at longer horizons. Baxter and Jermann (1997) find higher correlations when using annual returns. Storesletten et al. (2004) find that idiosyncratic labor income risk varies countercyclically over the business cycle. Benzoni et al. (2007) assume that labor income and dividends are cointegrated and that labor markets catch up with the aggregate economy in about 20 years. In addition, investors may have different investment horizons due to their career length. Their portfolio decisions are only affected by their human capital during pre-retirement years, when they earn labor income. Different investment horizons may affect the frequency at which systematic risk is priced (Kamara et al., 2016).

Our goal is to empirically test whether the impact of labor income risk on asset pricing displays horizon effects and, if so, to identify the dominant horizon. Using a flexible empirical approach, our analysis includes labor income risk at all horizons, including short, medium, long and very long term. This allows us to let the data speak on which horizon(s) dominate.

As a basis for the empirical tests, we first develop a stylized model in which investors have different career lengths and in which the covariances between labor income growth and stock returns may vary across horizons. In particular, the model shows that optimal portfolio allocation decisions and expected stock returns are explicitly affected by covariances between one-period equity returns and (log) labor income growth rates across different horizons. To decompose aggregate labor income growth into components at different frequencies, we suitably adapt the methodology developed in Ortu et al. (2013, 2015), Bandi and Tamoni (2017), and Boons and Tamoni (2016). This allows us to estimate horizon-specific prices of labor income risk in the cross-section of stock returns.

Our extensive empirical tests reveal that the two- to four-year horizon strongly dominates. Labor income risk at this medium term horizon is an important priced systematic risk factor, while labor income risk at other horizons is not. We start by estimating a multi-factor model that

includes labor income risk at six different horizons, ranging from one quarter, to one, two, four, eight and more than eight years. The exposures to these risk factors peak at the two- to four-year horizon, for which the betas are most statistically significant and economically important. When including labor income risk at each of the six horizons into one model, we find that only the two- to four-year horizon carries a highly significant risk premium, while the other horizons do not. Removing these insignificant horizons does not affect the adjusted R^2 of the model. In fact, the resulting two-factor model with equity market returns and aggregate labor income growth at the two- to four-year horizon can explain a striking 77% of the cross-sectional variation of 25 size book-to-market portfolios. When adding 25 size-investment portfolios, the adjusted R^2 remains high at 71%. This is very similar to the performance of the Fama and French (2015) five-factor model, that has adjusted R^2s of 77% and 76% respectively. By comparison, the standard human capital CAPM with quarterly labor income growth (e.g., Jagannathan and Wang, 1996)[2] can explain very little of the cross-sectional return variation with adjusted R^2s of close to 0% and 7%. Hence, the simple adjustment of measuring labor income risk over a medium term horizon has a striking effect on the model performance.

The estimated price of risk for medium-term labor income risk is negative. This is consistent with our model where agents have "Keeping up with the Joneses" preferences of Abel (1990). Agents care about their wealth relative to their peers. We use country-level aggregate labor income as a measure of the wealth of the peers, similar to Gomez et al. (2009). If these preferences are strong enough, agents are willing to pay more for stocks that have a positive exposure to aggregate labor income growth. Consequently, the price of labor income risk is negative. Gomez et al. (2009) also find negative prices of labor income risk in an international setting and Gomez et al. (2016) confirm this for the U.S. However, these papers do not consider horizon effects, which is our prime focus.

The dominant two- to four-year horizon coincides with the three-year horizon identified by Parker and Julliard (2005) for measuring ultimate consumption risk. When we add the three-year consumption growth factor to our model, we find that ultimate consumption growth does not carry a significant risk premium, while the risk premium estimate for medium term labor income risk remains unaffected. By itself, the ultimate consumption CAPM can explain 30% of the cross-sectional stock

2 See also Mayers, 1972; Campbell, 1996; Palacios-Huerta, 2003.

return variation, which is substantially lower than our model. This suggests that the medium-term labor income risk factor is not a mere proxy for ultimate consumption risk.

Furthermore, we find that the impact of labor income risk on optimal portfolio allocation is also more pronounced at this medium term horizon. Our model gives rise to a simple regression equation that allows us to estimate portfolio adjustments for investors that are exposed to labor income risk across different horizons, varying between one quarter and eight years. The resulting hedging portfolio weights are more significant for investors with horizons of one to four years, both statistically and economically.[3]

An extensive series of robustness tests confirms that our results are robust when estimating the model for the cross-section of 25 size book-to-market and 30 industry portfolios, when using alternative classifications of the horizons, when using univariate rather than multivariate betas in the first-stage regressions, and when using real instead of nominal labor income growth.

The dominant two- to four-year horizon is consistent with various studies that document wage rigidity. Wages can be rigid due to, for instance, infrequent wage setting, labor market frictions or negotiations between workers and the firm. Shimer (2005) and Hall (2005), among others, show that wage rigidity is key for search-and-matching models to match unemployment and vacancy data. Favilukis and Lin (2016b) find that wage rigidity helps to generate both smooth wages and volatile equity returns in a production-based asset pricing model.[4] Notably, the frequency of wage setting coincides with our dominant medium-term horizon. In Favilukis and Lin (2016b), wages are reset every three years.

3 Hedging demand is driven by investors' own labor income risk. Therefore, in this part of the analysis, we take into account that human capital is heterogeneous across investors, in particular across industries of employment (e.g., Katz and Summers, 1989; Neal, 1995). Conventional wisdom suggests that to diversify, one should invest less in stocks from the industry where one works. While we find little evidence of this based on short-term returns, for longer horizons of two to four years, we find significant effects for about one-third of the industries. Specifically, at the four-year horizon, investors working in the services, mining, and finance industries should optimally hold respectively 8.9%, 6.7% and 11.4% less of their own industry stocks and investors from the wholesale industry optimally hold 5.4% more of their own industry stocks as a hedge against their labor income risk.

4 Favilukis and Lin (2016a) show that in the presence of this wage rigidity, wage growth negatively predicts future stock market returns.

Marfe (2018) finds that a measure of labor rigidity (i.e., the labor share) fluctuates countercyclically with a half-life of 3.5 years. Rich and Tracy (2004) empirically show that the median labor contract duration is 35 months for most of their sample. Our study highlights the impact of labor income risk at the two- to four-year horizon for the cross-section of expected stock returns.

This chapter also fits into a growing literature of asset pricing models that allow for lower frequency risks to affect expected returns. Daniel and Marshall (1997) and Parker and Julliard (2005) find that consumption risk at the two- or three-year horizon matters for asset pricing. Bansal and Yaron (2004), among others, show the importance of long-run consumption shocks. Estimates of what constitutes the long run vary. Malloy et al. (2009) consider four to six years, whereas Dew-Becker and Giglio (2016) quantify long run as "centuries". Koijen et al. (2017) and Bandi and Tamoni (2017) find that shocks at the business cycle frequency matter for asset pricing. None of these papers consider labor, and, to the best of our knowledge, we are the first to explore labor income risk across multiple horizons.

The remainder of this chapter is structured as follows. Section 4.2 presents the simple theoretical framework underlying our empirical analysis. Section 4.3 discusses the estimation methodology and provides more details on the data. Section 4.4 presents the empirical results. Section 4.5 discusses a series of robustness tests and Section 4.6 concludes. The Appendices 4.A and 4.B provide details of the model derivation and derive an expression for the standard error corrections in the second-stage cross-sectional regressions, respectively.

4.2 THEORETICAL FRAMEWORK: ASSET PRICING WITH LABOR INCOME RISK AT MULTIPLE HORIZONS

This section develops the theoretical framework that serves as a motivation for examining the relation between expected stock returns and labor income risk across different horizons. The aim of the model is to show that, even in a simple and stylized setting, horizon-specific labor income risk can already play a role in asset pricing. In particular, we consider a CAPM-type setup in which there are multiple cohorts of labor income earning investors with different investment horizons. We show that our framework gives rise to a linear asset pricing model with a standard one-

period equity market risk factor and horizon-specific (log) labor income growth risk factors. Details of the derivations can be found in Appendix 4.A.

We start by considering an individual investor whose evolution of wealth W_t at time $t = 0, 1, \ldots$ is given by

$$W_{t+1} = R_{p,t+1} W_t + L_{t+1}, \tag{4.1}$$

in the spirit of Campbell and Viceira (2002), Chapter 6. Here, $R_{p,t+1} = R_f + \alpha_t'(R_{t+1} - R_f)$ denotes the gross return on the investment portfolio, where R_f is the constant return on a riskless asset, R_t is a vector of equity returns assumed to be jointly lognormal and i.i.d. over time and α_t denotes the proportions of wealth invested in equities at time t, and L_t denotes nontradable labor income at time t assumed to be lognormal. Given the wealth level W_t at time t and a stream of labor income L_{t+i}, $i = 1, \ldots, h$, the investor's wealth at horizon h equals

$$\begin{aligned} W_{t+h} = (R_{p,t+h} \cdots R_{p,t+1}) W_t + (R_{p,t+h} \cdots R_{p,t+2}) L_{t+1} \\ + \ldots + R_{p,t+h} L_{t+h-1} + L_{t+h}, \end{aligned} \tag{4.2}$$

with gross portfolio return in period $t + i$ given by $R_{p,t+i} = R_f + \alpha_{t+i-1}'(R_{t+i} - R_f)$.

We aim to keep the model as flexible as possible in order to let the data speak on whether and how horizon effects play a role. Therefore, we do not a-priori specify the horizon h. Rather, in our analysis we include all horizons, ranging from short, to medium-term, long and very long horizons. Furthermore, we do not specify in advance the sign of the price of labor income risk. Traditionally in the Human Capital CAPM of Mayers (1972), labor income risk carries a positive price of risk. However, several recent papers find a negative price of labor income risk (e.g., Gomez et al., 2009; Gomez et al., 2016; Maio and Min, 2018). To allow for a negative price of risk, we assume that the investor has "Keeping up with the Joneses" (KUJ) preferences.

Following Abel (1990), the investor maximizes the utility of h-period terminal wealth W_{t+h} relative to the aggregate labor income $\bar{L}_{t+h}$. That is, she cares about her wealth relative to the labor income of an average "peer".[5] These relative wealth considerations affect the pricing of

5 The multiplicative habit specification of Abel (1990) has some advantages over an additive specification; see Bilsen et al. (2019) and the references therein for further details.

labor income risk in the cross-section of stock returns. Specifically, investors prefer stocks that positively correlate with aggregate labor income growth, as these help them keep up with their "peers." This results in a negative price of labor income risk, similar to Gomez et al. (2009) and Gomez et al. (2016). On the other hand, following the logic of the Human Capital CAPM of Mayers (1972), the investor's own labor income can also generate a hedging demand for stocks, creating a preference for stocks that are negatively exposed to labor income risk. The traditional hedging demand channel results in a positive price of labor income risk. Our model allows for both mechanisms, where the relative importance of the KUJ effects versus the hedging demand effects will determine the sign of the price of risk. Our empirical analysis helps shed light on this question.

More precisely, the investor aims to maximize over admissible investment strategies the expected utility

$$U(W_{t+h}) = \mathbb{E}_t \left[\left(\frac{W_{t+h}}{\bar{L}_{t+h}^{\psi}} \right)^{1-\gamma} \right] / (1-\gamma), \tag{4.3}$$

with $\gamma > 0$ the relative risk aversion coefficient, and where $0 \leqslant \psi \leqslant 1$ measures the strength of the KUJ effect: for $\psi = 0$, we obtain the standard CRRA function defined over terminal wealth, whereas for $\psi = 1$ we scale wealth by aggregate labor income, both at time $t+h$. Labor income is assumed to be jointly lognormal. Here, $\mathbb{E}_t$ denotes the expectation conditional upon information at time t.[6]

Using a suitable log-linearization of W_{t+h} and a second-order Taylor expansion of the objective function (see Appendix 4.A), we find the optimal portfolio choice rule at time t to be

$$\begin{aligned} \alpha_t = & \frac{W_t + hL_t}{W_t} \frac{1}{\gamma} \mathrm{Var}(r_{t+1})^{-1} \left(\mathbb{E}[r_{t+1}] + \frac{1}{2}\sigma^2 - r_f \right) \\ & + \left(1 - \frac{1}{\gamma}\right) \mathrm{Var}(r_{t+1})^{-1} \left[\psi \frac{W_t + hL_t}{W_t} \mathrm{Cov}_t(r_{t+1}, \bar{l}_{t+h}) \right] \\ & - \left(1 - \frac{1}{\gamma}\right) \mathrm{Var}(r_{t+1})^{-1} \left[\frac{L_t}{W_t} \sum_{i=1}^{h} \mathrm{Cov}_t(r_{t+1}, l_{t+i}) \right], \end{aligned} \tag{4.4}$$

6 Throughout we abstract away from intermediate consumption decisions. This assumption is made to keep the model as simple as possible. The main aim of our theoretical model is to make a case for considering labor income risk across different horizons. Thus, we want to show that horizon effects already arise in the simplest possible setup.

where $r_{t+1} = \log(R_{t+1})$, σ^2 denotes the diagonal of $\mathrm{Var}(r_{t+1})$, and $l_{t+i} = \log(L_{t+i})$ is the log of the investor's labor income at time $t+i$, and $\bar{l}_{t+h} = \log\left(\bar{L}_{t+h}\right)$ is log aggregate labor income at time $t+h$. The first term in (4.4) is the standard speculative demand arising from CRRA portfolio optimization. The second term in (4.4), which depends on the covariance with aggregate labor income risk, arises because of the KUJ utility specification. Due to the "Keeping up with the Joneses" effect, investors prefer stocks that are positively exposed to labor income risk. The third part is the hedging demand that arises from the investor's own labor income risk, which is indirectly exposed to equity risk. To hedge these risks, the investor has to adapt his optimal portfolio holdings and prefers stocks that are negatively exposed to labor income risk. Eqn. (4.4) shows that the portfolio adjustments are driven by the investor's income-to-wealth ratio and the covariances between one-period equity returns and multi-period (log) labor income growth rates.

We now show the implications of the optimal portfolio choice rule (4.4) on equilibrium asset prices. Let there be, at any time t, H cohorts with investment horizon $h = 1, \ldots, H$, initial wealth W_h and current labor income L_h. Assuming stationarity, aggregating asset demands from (4.4) over all cohorts and rewriting (see Appendix 4.A) gives the following asset pricing equation:

$$\begin{aligned}\mathbb{E}[R] = r_f + \gamma \left[\frac{\sum_{h=1}^{H} W_h}{\sum_{h=1}^{H}(W_h + hL_h)}\right] \mathrm{Cov}(r_{t+1}, r_{m,t+1}) \\ + (\gamma - 1) \sum_{h=1}^{H} \left[\frac{\sum_{i=h}^{H} L_i - \psi(W_h + hL_h)}{\sum_{i=1}^{H}(W_i + iL_i)}\right] \mathrm{Cov}\left(r_{t+1}, \bar{l}_{t+h} - \bar{l}_t\right),\end{aligned} \tag{4.5}$$

with $\mathbb{E}[R] = \mathbb{E}[r_{t+1}] + \frac{1}{2}\sigma^2$. Note that aggregate labor income growth at each horizon h is a separate risk factor, whose price of risk depends on the labor income share of all cohorts with investment horizons equal to or longer than h in total wealth. The asset pricing relation can be re-written in beta form, as follows:

$$\mathbb{E}[R] = r_f + \beta_m \lambda_m + \sum_{h=1}^{H} \beta_{l,h} \lambda_{l,h}, \tag{4.6}$$

with $\beta_m = \mathrm{Cov}(r_{t+1}, r_{m,t+1})/\mathrm{Var}(r_{m,t+1})$ and
$\beta_{l,h} = \mathrm{Cov}\left(r_{t+1}, \bar{l}_{t+h} - \bar{l}_t\right)/\mathrm{Var}(\bar{l}_{t+h} - \bar{l}_t)$ the asset return exposures

to the market portfolio return and aggregate labor income growth at horizon h, and where

$$\lambda_m = \gamma \left[\frac{\sum_{h=1}^{H} W_h}{\sum_{h=1}^{H} (W_h + hL_h)} \right] \mathrm{Var}(r_{m,t+1}), \tag{4.7}$$

$$\lambda_{l,h} = (\gamma - 1) \left[\frac{\sum_{i=h}^{H} L_i - \psi(W_h + hL_h)}{\sum_{i=1}^{H} (W_i + iL_i)} \right] \mathrm{Var}(\bar{l}_{t+h} - \bar{l}_t), \tag{4.8}$$

are the prices of market risk and labor income growth risk at horizon h. For $\psi = 0$, i.e., in the absence of the KUJ effect, and $\gamma > 1$, all prices of risk are positive. This is natural, as in that case the assets that correlate strongly with labor income growth are undesirable and require a higher expected return in equilibrium. If the KUJ effect is strong enough, the price of labor income risk may become negative as investors now prefer stocks that are positively exposed to aggregate labor income risk in order to keep up with their "neighbours." Ultimately it is an empirical question whether the hedging or the KUJ effect dominates.

Summarizing, the key ingredients for obtaining horizon effects in our model are: (i) the presence of labor income, (ii) a nonzero covariance between stock returns and labor income growth at a given horizon, and (iii) investors with longer-term investment horizons (i.e., exceeding at least one quarter). Note that the existence of multiple cohorts with heterogeneous investment horizons is not essential for generating horizon effects in our asset pricing equation. Even in the case where we only have one cohort with long-term investment horizon h, the asset pricing equation (4.6) would still contain labor income growth rates up to horizon h. The heterogeneity in investment horizons of the different cohorts mainly serves as a relaxation of the assumption that all investors in the economy have the same horizon. For example, if investors want to accumulate wealth for retirement, their investment horizons depend on how far away they are from their retirement age.

Our model relates to the standard human capital CAPM of Mayers (1972), in which contemporaneous labor income risk is priced. According to that model, the expected excess returns of an asset are given by a linear function of the asset's exposures to equity market risk and human capital risk. The main argument for including human capital is that it is part of the overall wealth portfolio. As the relative value of human capital is unobserved, it is included as a separate factor. In empirical applications, quarterly or monthly labor income growth is often used

as a proxy for human capital returns (see, e.g., Jagannathan and Wang, 1996; Eiling, 2013). In our setting, a similar result would be obtained by assuming that the investment horizon of all investors in the economy is only one period ahead (or, equivalently, by assuming $\lambda_h = 0$ for all $h > 1$) and assuming no KUJ effects (i.e., $\psi = 0$).

4.3 EMPIRICAL METHODOLOGY AND DATA

This section describes the estimation methodology and data. We first show how our model gives rise to a regression specification that allows us to compute horizon-specific hedging demands for individual investors that are exposed to labor income risk. Next, we provide the details of our empirical cross-sectional asset pricing analysis. We show how we can group together labor income growth shocks into components with different degrees of persistence. The advantage of this approach is that it leads to a parsimonious empirical model specification in which we can address a wide range of horizons with relatively few risk factors. We explain how to estimate the horizon-specific labor income risk exposures and market prices of risk in a two-stage cross-sectional regression approach. After that, we provide more details on the data.

4.3.1 *Hedging demand due to labor income risk across horizons*

We show in our theoretical model that horizon-specific labor income risk can affect expected stock returns through a portfolio channel. Therefore, as a preliminary investigation to our asset pricing study, we first analyze how horizon-specific labor income risk affects optimal portfolio allocation.

The portfolio adjustments for labor income risk of an h-period investor are given by the last two terms in Eqn. (4.4). These consist of the "Keeping up with the Joneses" effect and the hedging demand induced by the investor's own labor income that is correlated with equity returns. Our main focus is on the hedging demand, which we estimate unconditionally. More specifically, the hedging demand is denoted by $-q\mathrm{Var}(r_{t+1})^{-1}\mathrm{Cov}(r_{t+1}, \sum_{i=1}^{h}(l_{t+i} - l_t))$, where $q = \left(1 - \frac{1}{\gamma}\right)\frac{L_t}{W_t}$ is a term depending on the investor's relative risk aversion parameter and quarterly wealth-to-income ratio. Rather than focusing on investments

in a riskless bond and a risky stock as in many life-cycle portfolio choice models, we analyze how the composition of the multi-asset equity portfolio should be adjusted for labor income risk. Hence, we consider investments in Z different equity portfolios as potential hedges against labor income risk. We can estimate the hedging weights of an h-period investor by regressing the cumulative sum of her multi-period labor income growth rates on a constant and the excess returns of the Z different equity portfolios, i.e.,

$$\sum_{i=1}^{h} (l_{t+i} - l_t) = \alpha + \sum_{z=1}^{Z} \beta_z r_{z,t+1} + \epsilon_t. \tag{4.9}$$

Here l_{t+i} denotes the investor's own log labor income in period $t+i$, and $r_{z,t+1}$ the quarterly log excess returns of equity portfolio z. Indeed, the regression coefficients are given by $\beta_z = \frac{\mathrm{Cov}\left(r_{z,t+1}, \sum_{i=1}^{h} (l_{t+i} - l_t)\right)}{\mathrm{Var}(r_{z,t+1})}$, which are equivalent to the model-implied hedging demands up to a factor $-q$.

The hedging demand depends on the investor's own labor income risk. Various papers, such as Katz and Summers (1989) and Neal (1995), find that an important source human capital heterogeneity is the industry of employment. Eiling (2013) shows that an investor's industry affiliation is indeed an important factor to take into account in determining optimal portfolio allocation.[7] Therefore, in this part of the analysis, we consider investors whose labor income risk is tied to their industry of employment. To hedge her industry-specific labor income risk, an investor may need to adjust that same industry's weight in her optimal stock portfolio. Other industry weights will possibly also be affected as stocks from other industries can potentially serve as hedges against labor income risk too. We consider nine broad industries spanning aggregate labor income: mining, construction, transport, wholesale trade, retail trade, manufacturing, financial services, services, and the government. We also consider an investor who is only exposed to aggregate labor income risk.

We assume that investors can invest in equity portfolios that mimic average industry equity returns. We create eight of these equity industry portfolios (hence, $Z = 8$) by matching the industries with our labor in-

7 Eiling (2013) also finds that industry-specific human capital helps explain cross-sectional variation in expected stock returns.

come industries based on SIC codes.[8] We do not define equity returns for the government sector as companies in this sector typically do not issue stocks. In this way, we can address the question of whether an investor that is exposed to a certain industry's labor income risk should over- or underinvest in that industry. In a separate test, we consider the equity market return, instead of industry equity returns, as a possible hedge against labor income risk.

The KUJ preferences induce investors to invest more in stocks that covary strongly with aggregate labor income, which proxies for the wealth of their "peers". Hence, these portfolio adjustments do not depend on the investor's own labor income risk. In our model, the KUJ hedging demand of an h-period investor is given by $-\tilde{q}\text{Var}(r_{t+1})^{-1}\text{Cov}_t(r_{t+1}, \bar{l}_{t+h} - \bar{l}_t)$, where $\tilde{q} = \left(1 - \frac{1}{\gamma}\right)\psi\frac{W_t + hL_t}{W_t}$. We can thus estimate these hedging weights unconditionally (up to the factor $\tilde{q}$) by regressing h-month ahead aggregate labor income growth on quarterly (log) equity returns from eight different industries. That is, the regression coefficients of the regression equation

$$\bar{l}_{t+h} - \bar{l}_t = \alpha + \sum_{z=1}^{8} \beta_{k,z} r_{z,t+1} + \epsilon_t \tag{4.10}$$

give the KUJ hedging weights up to the factor $\tilde{q}$. We estimate the KUJ portfolio adjustments in order to see whether this component of the optimal portfolio displays horizon effects.

4.3.2 *Cross-sectional regression methodology*

In the second and main part of our empirical analysis, we test the asset pricing implications of horizon-specific labor income risk. This section explains how we derive a parsimonious empirical specification that allows us to simultaneously consider a wide variety of horizons. Due to multicollinearity issues, we cannot simply include labor income growth

8 The equity portfolio returns are constructed using monthly CRSP data from the period January 1958 until December 2017. For a given industry, a company contributes to the index returns of quarter t if it has a non-zero market capitalization in the last month of quarter $t-1$ and if return data is available for all months in the quarter t. The quarterly return of a company is computed as the compounded monthly returns during that quarter.

at various horizons as factors in our model. Therefore, we first decompose quarterly labor income growth into components that are working on different frequencies, which are by construction (asymptotically) uncorrelated. In our model, only aggregate labor income is priced. Hence, in the asset pricing tests we do not consider industry-specific labor income.

4.3.2.1 *Decomposition of labor income growth across frequencies*

In principle, labor income growth risk at any horizon could be a priced risk factor. However, from an empirical perspective, it would be infeasible to include a separate labor income risk factor for every possible horizon. Therefore, in the empirical specification, we group together labor income growth shocks into components with different degrees of persistence corresponding to six categories: 1-2 quarters, 3-4 quarters, 1-2 years, 2-4 years, 4-8 years and more than eight years. By including these categories we effectively include all horizons.[9]

More specifically, let us denote quarterly aggregate (log) labor income growth as $f_{t+1} = \bar{l}_{t+1} - \bar{l}_t$, with $\bar{l}_t = \log(\bar{L}_t)$ quarterly per worker aggregate log labor income. Then it can be shown that f_{t+1} can be written as the sum of components $f_t^{(j)}$ working on different time scales as follows:

$$f_{t+1} = \sum_{j=1}^{J} f_t^{(j)} + f_t^{(j>J)}, \tag{4.11}$$

where $f_t^{(j)}$ and $f_t^{(j>J)} = \sum_{j>J} f_t^{(j)}$ are the components at time t and scale $j = 1, \ldots, J$, and $j > J$. Ortu et al. (2015) show that such a decomposition holds for any weakly stationary time series. We estimate the components in a procedure similar to Ortu et al. (2013) and Bandi

9 In a robustness test we vary the number of horizons and find that our results remain very similar.

and Tamoni (2017).[10] More precisely, given time series $\{f_t\}_{t\in\mathbb{Z}}$, we first construct moving averages $\pi_t^{(j)}$ of length 2^j as

$$\pi_t^{(j)} = \frac{1}{2^j}\sum_{h=1}^{2^j} f_{t+h}, \tag{4.12}$$

for $j = 0,\dots,J$. Next, we define the component $f_t^{(j)}$ to be the difference between moving averages of length 2^{j-1} and 2^j, i.e.,

$$f_t^{(j)} = \pi_t^{(j-1)} - \pi_t^{(j)}, \tag{4.13}$$

for $j = 1,\dots,J$. As a result, the components $f_t^{(j)}$ can be interpreted as containing those fluctuations with half-life in the interval of $[2^{j-1}, 2^j)$ quarters. We also define $f_t^{(j>J)} = \pi_t^{(J)}$ as containing those fluctuations with half-life exceeding 2^J quarters. In our empirical specification, detailed below, we include a separate risk factor for every scale component $j = 1,2,\dots,J$ and $j > J$. Therefore, for the number of risk factors to remain parsimonious, we need to set the maximum level of persistence. We choose $J = 5$, which means that the scale component $f^{(j>5)}$ groups together shocks with a half-life exceeding $2^5 = 32$ quarters (or 8 years). Table 4.1 provides more details on the mapping between scales j and their corresponding time spans.

The decomposition of the contemporaneous labor income growth factor into different scale components allows us to decompose the covariance between asset i's excess returns and labor income growth across scales:

$$\operatorname{Cov}\left(R_{t+1}^{e,i}, f_{t+1}\right) = \sum_{j=1}^{J}\operatorname{Cov}\left(R_{t+1}^{e,i}, f_t^{(j)}\right) + \operatorname{Cov}\left(R_{t+1}^{e,i}, f_t^{(j>J)}\right). \tag{4.14}$$

The covariance decomposition in (4.14) forms the basis of our empirical analysis on horizon-specific labor income risk. To see this, note that by dividing (4.14) by the variance of f_{t+1}, we obtain a beta decomposition

10 Note that Ortu et al. (2013) and Bandi and Tamoni (2017) define the moving averages $\pi_t^{(j)}$ backwards in time. In our setting, however, it is more intuitive to define them forwards. This implies that the scale factors $f_t^{(j)}$ are random variables whose outcome is realized only at time $t + 2^j$. Likewise, $f_t^{(j>J)}$ is realized at time $t + 2^J$.

of the original contemporaneous beta into a weighted average of scale-specific betas:

$$\begin{aligned}\beta_{i,f} &= \frac{\text{Cov}\left(R^{e,i}_{t+1}, f_{t+1}\right)}{\text{Var}\left(f_{t+1}\right)} \\ &= \sum_{j=1}^{J} w_f^{(j)} \frac{\text{Cov}\left(R^{e,i}_{t+1}, f_t^{(j)}\right)}{\text{Var}\left(f_t^{(j)}\right)} + w_f^{(j>J)} \frac{\text{Cov}\left(R^{e,i}_{t+1}, f_t^{(j>J)}\right)}{\text{Var}\left(f_t^{(j>J)}\right)} \\ &= \sum_{j=1}^{J} w_f^{(j)} \beta_{i,f}^{(j)} + w_f^{(j>J)} \beta_{i,f}^{(j>J)}, \end{aligned} \tag{4.15}$$

where $w_f^{(j)} = \text{Var}(f_t^{(j)})/\text{Var}\left(f_{t+1}\right)$ for all j. The beta decomposition in (4.15) shows that contemporaneous betas, as obtained from typical monthly or quarterly regressions, are largely determined by high-frequency comovements of risk and returns. This is the case, since the variance of high-frequency components is typically much larger than those of lower frequency components, resulting in higher weights $w_f^{(j)}$ for high-frequency comovements. Therefore, it is key to evaluate the scale-specific betas separately in order to accurately identify whether low-frequency risk is priced.

Table 4.1. Mapping between time-scales and time spans. This table provides the mapping between the time-scales j and the corresponding time spans.

Time-scale	Frequency resolution	Interpretation
$j = 1$	1-2 quarters	
$j = 2$	2-4 quarters	
$j = 3$	4-8 quarters	
$j = 4$	8-16 quarters	
$j = 5$	16-32 quarters	
$j = 1:3$	1-8 quarters	High frequency
$j = 4:5$	8-32 quarters	Medium frequency
$j > 5$	$>$ 32 quarters	Low frequency

4.3.2.2 *Horizon-specific risk pricing and cross-sectional regression approach*

The beta decomposition (4.15) allows us to straightforwardly estimate our model with horizon-specific labor income risk factors. Within the standard human capital CAPM, the expected excess return of asset i is given by a linear function of exposures to equity market risk and contemporaneous human capital risk. In empirical applications of this model, it is standard to use aggregate labor income growth as a proxy for human capital returns (see, e.g., Jagannathan and Wang, 1996). This gives rise to the following empirical specification of the traditional human capital CAPM:

$$\mathbb{E}\left[R^{e,i}_{t+1}\right] = \lambda_0 + \beta_{mkt,i}\lambda_{mkt} + \beta_{l,i}\lambda_l + \alpha_i, \tag{4.16}$$

where λ_{mkt} and λ_l are the prices of risk for exposures to equity market and (contemporaneous) labor income risk, respectively, and the $\beta_{mkt,i}$ and $\beta_{l,i}$ the corresponding risk loadings for asset i.

Now it follows directly from (4.15) that we can split the labor income risk factor into scale-specific components, which results in our empirical specification of interest, i.e.,

$$\mathbb{E}\left[R^{e,i}_{t+1}\right] = \lambda_0 + \beta_{mkt,i}\lambda_{mkt} + \sum_{j=1}^{J} \beta^{(j)}_{l,i}\lambda^{(j)}_l + \beta^{(j>J)}_{l,i}\lambda^{(j>J)}_l + \alpha_i. \tag{4.17}$$

Here λ_{mkt} and $\beta_{mkt,i}$ are again the price of equity market risk and asset i's exposure to that risk, respectively. The $\lambda^{(j)}_l$, $j = 1, \ldots, J$, now denote the prices of risk for exposure to labor income growth shocks with fluctuations between 2^{j-1} and 2^j quarters, and $\lambda^{(j>J)}_l$ denotes the price of risk for exposure to labor income shocks with fluctuations beyond 2^J quarters.

The empirical specification in (4.17) can be viewed as a restricted version of the theoretical asset pricing equation (4.6). The general model contains all covariances up to the maximum horizon $H = 2^J$ and is therefore difficult to estimate (we use $J = 5$, which would amount to 32 factors). Our empirical specification captures the essence of the theoretical model: both the empirical and theoretical specifications contain an equity market risk factor and labor income growth risk factors across different horizons.[11]

11 While the empirical specification is consistent with our theoretical asset pricing equation, there is not a one-to-one mapping as the scale components are defined as dif-

To estimate the model, we employ a two-pass cross-sectional regression methodology. In the first stage, we obtain the risk exposures (betas) by running a multivariate ordinary least squares time-series regression of portfolio excess returns on the different risk factors at a quarterly frequency. In particular, we run for each portfolio, i.e., test asset, i the following quarterly multivariate regression:

$$R^{e,i}_{t+1} = \alpha_{0,i} + \beta_{mkt,i} R^{e,mkt}_{t+1} + \sum_{j=1}^{J} \beta^{(j)}_{l,i} f^{(j)}_t + \beta^{(j>J)}_{l,i} f^{(j>J)}_t + \epsilon_{i,t+1}. \quad (4.18)$$

In the second stage, we estimate the market prices of risk by running for each time t the cross-sectional regression specified by (4.17) using the estimated betas from the first stage, i.e.,

$$R^{e,i}_{t+1} = \lambda_{0,t+1} + \hat{\beta}_{mkt,i} \lambda_{mkt,t+1} + \sum_{j=1}^{J} \hat{\beta}^{(j)}_{l,i} \lambda^{(j)}_{l,t+1} + \hat{\beta}^{(j>J)}_{l,i} \lambda^{(j>J)}_{l,t+1} + \eta_{i,t+1}, \quad (4.19)$$

for $t = 0, \ldots, T-1$. The estimated market prices of risk are given by their time-series averages, i.e., $\hat{\lambda}_0 = \frac{1}{T}\sum_{t=1}^{T} \hat{\lambda}_{0,t}$, $\hat{\lambda}_{mkt} = \frac{1}{T}\sum_{t=1}^{T} \hat{\lambda}_{mkt,t}$ and $\hat{\lambda}^{(j)}_l = \frac{1}{T}\sum_{t=1}^{T} \hat{\lambda}^{(j)}_{l,t}$. We derive expressions for the standard errors that account for autocorrelation in the factors and the fact that we use estimated betas in the second-stage regression. Details can be found in Appendix 4.B.[12] Apart from the full model (4.17), we are also interested in several restricted and adapted specifications. In particular, we assess

ferences in moving averages of labor income growth rates. At the same time, Bandi and Tamoni (2017) show that the covariances between one-period equity returns and multi-horizon risk factor returns, as induced by the empirical decomposition (4.11), are consistent with the equivalent covariances in the data.

12 Compared to this chapter, Ortu et al. (2013), Boons and Tamoni (2016) and Bandi and Tamoni (2017) consider a somewhat different setup in which they not only decompose the right-hand side factor of the first-stage regressions (4.18) into different scale components, but also the left-hand side test asset returns. As a consequence, the left-hand side variables of the first-stage regressions are different from the left-hand side variables in the second-stage regressions, and therefore standard Shanken-like corrections to the standard errors cannot be applied to these second-stage regressions. Because of this, these authors resort to using GMM to estimate their models, as this methodology allows them to obtain robust standard errors. It is, however, straightforward to generalize our derivation of the robust standard errors in Appendix 4.B to a setting in which there are different left-hand side variables in the first stage, and therefore the standard two-pass regression methodology also becomes applicable to their setup.

the performance of parsimonious two-factor specifications that include the equity market return and labor income risk for one scale at a time.

4.3.3 *Data and summary statistics*

4.3.3.1 *Labor income data*

The basis of our analysis consists of quarterly (log) labor income growth series, which we define as

$$f_t = \log(L_t) - \log(L_{t-1}), \tag{4.20}$$

where L_t denotes the per worker labor income in quarter t. To obtain multi-horizon human capital returns, we simply sum the quarterly growth rates of the corresponding horizons.

We retrieve quarterly labor income data from the State Quarterly (Q) Table 7, which is published by the Bureau of Economic Analysis. This table provides quarterly wages and salaries for disaggregated industries. We use labor income on nine different industries (based on their SIC codes), which together span aggregate labor income. In particular, the industries that we consider are mining (SIC 1000-1499), construction (SIC 1500-1799), manufacturing (SIC 2000-3999), transportation, communication, and utilities (SIC 4100-4999), wholesale trade (SIC 5000-5199), retail trade (SIC 5200-5299), finance, insurance, and real estate (SIC 6000-6799), services (SIC 7000-8999), and government (SIC 9100-9999).[13] We scale the quarterly wages and salaries for each industry by the average number of workers in each quarter for the corresponding industry using monthly employment data from the Current Employment Statistics survey, published by the Bureau of Labor Statistics. The sample period runs

13 As of 2001, the BEA changed its industry classification system from using SIC codes to NAICS codes. Although it is relatively straightforward to map the NAICS codes to corresponding SIC codes, this change did result in a break in the data. We smooth out the effects of the change in the industry definition as follows: First, we use the labor income data according to the SIC codes to compute quarterly (log) labor income growths until 2001Q4. For 2002Q2 onwards, we use the labor income according to the NAICS codes to compute (log) labor income growths. For 2002Q1, we exploit the fact that labor income based on NAICS codes is available from 1998Q1 onwards and use the 2001Q4 and 2002Q1 NAICS labor income data to compute (log) labor income growth in 2002Q1. To compute longer horizon labor income growth rates, we always start from the quarterly growth rates to ensure we never mix SIC and NAICS labor income levels when calculating growth rates.

from 1958Q1 until 2017Q4, resulting in a total of 240 quarterly observations of per capita labor income. These quarterly series form the basis for constructing labor income growth series components across different scales. Table 4.2 provides summary statistics for the aggregate and industry-specific (log) labor income growth series. For illustrative purposes, we plot the original aggregate labor income growth series and the combined high, medium, and low frequency scale components in Figure 4.1.[14] The figure shows that the scale components become more persistent for lower frequencies, as expected.

Table 4.2. Summary statistics of quarterly labor income growth rates. This table reports summary statistics of quarterly aggregate and industry-specific (log) labor income growth rates. The sample period runs from 1958Q2 until 2017Q4.

	Mean (%)	Std. dev. (%)	Min (%)	Max (%)
f_{agg}	1.06	0.73	-2.82	3.21
f_{min}	1.24	2.44	-8.65	15.70
f_{cnst}	1.05	1.06	-3.05	4.63
f_{man}	1.12	1.15	-5.94	4.73
f_{trnsp}	1.07	1.25	-3.17	7.04
f_{whls}	1.14	1.17	-6.39	4.68
f_{ret}	0.92	0.88	-4.29	4.46
f_{fin}	1.29	2.89	-21.05	12.94
f_{serv}	1.23	1.28	-5.05	8.75
f_{gov}	1.02	0.65	-0.50	3.93

14 To preserve space, I plot only combined high, medium, and low frequency scale components instead of all separate components $j = 1, \ldots, 5, > 5$. In particular, the high-frequency component is defined as the sum of the components at scales $j = 1, 2, 3$ and captures fluctuations in labor income growth with a half-life of between 1 and 8 quarters. Similarly, the medium-frequency component is defined as the sum of the components at scales $j = 4$ and $j = 5$ and captures fluctuations in labor income growth with a half-life of between 8 and 32 quarters.

4.3.3.2 *Other data*

Next to labor income growth rates, our model also contains an equity market risk factor. For this, we use the standard market excess return factor, defined as the value-weighted return of all CRSP stocks listed on the NYSE, AMEX, or NASDAQ minus the one-month Treasury bill rate, which we obtain from Kenneth French's website.[15] We use the monthly series running from April 1958 until December 2017, and convert these to quarterly series by compounding the monthly returns within each quarter.

To test our asset pricing model, we use returns on 25 size and book-to-market portfolios, 25 size and investment portfolios, and 30 industry portfolios. Again, we obtain monthly series for all these portfolios from Kenneth French's website, and we convert these into quarterly excess return series by compounding the monthly excess returns within each quarter. The quarterly return series of the 25 size and book-to-market and the 30 industry portfolios run from 1958Q2 until 2017Q4, whereas the quarterly return series of the size and investment portfolios run from 1963Q3 to 2017Q4. We observe the typical patterns that returns are decreasing in the size dimension (i.e., the 'size-effect'), increasing in the book-to-market dimension (i.e., the 'value premium'), and decreasing in the investment dimension.

We also compare our model to several existing asset pricing models. In particular, we consider the consumption CAPM, the ultimate consumption CAPM of Parker and Julliard (2005) that includes three-year ahead consumption growth, and the Fama and French (1993) three-factor and Fama and French (2015) five-factor models. When considering the (ultimate) consumption CAPM models, we use real (chain-weighted) personal consumption expenditures on nondurable goods per capita, which we obtain from the National Income and Product Accounts (NIPA). The factors for the Fama and French (1993) three-factor and Fama and French (2015) five-factor models are again obtained from Kenneth French's website.

15 http://mba.tuck.dartmouth.edu/pages/faculty/ken.french/data_library.html.

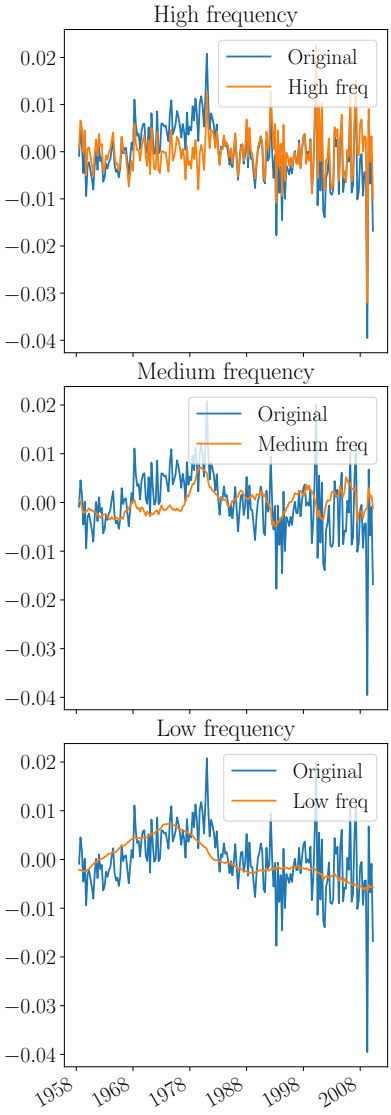


Figure 4.1. Aggregate labor income growth scales at different frequencies. This figure plots quarterly aggregate labor income growth series against its high, medium, and low frequency components. The high frequency component is defined as the sum of the components at scales $j = 1, 2, 3$, and contains fluctuations with a half-life of below 2 year. The medium frequency component is defined as the sum of the components at scales $j = 4, 5$, and contains fluctuations with a half-life of between 2 and 8 years. The low frequency component is defined as the sum of components $j > 5$, and contains fluctuations with a half-life of more than 8 years.

4.4 EMPIRICAL RESULTS

4.4.1 *Horizon-specific hedging demands due to labor income risk*

Based on our model, investors make two types of portfolio adjustments for labor income risk. First, a hedging demand that arises from their own labor income risk which may be correlated with equity returns. Second, a "Keeping up with the Joneses" effect, where investors prefer stocks that deliver high returns when aggregate labor income growth is high, in order to keep up with their "peers". This section explores how these portfolio adjustments for labor income vary across horizons.

We first analyze the hedging demand. In contrast to the asset pricing analysis where only aggregate labor income risk matters, the hedging demand is affected by the investor's own labor income risk. Therefore, in addition to considering an "aggregate" investor (who is only exposed to aggregate labor income risk) we also analyze hedging demand for investors working in different industries. Specifically, we estimate the regression coefficients of expression (4.9). Table 4.3 reports the estimates, multiplied by -1 and expressed in percentage terms. To get the total hedging demand, we need to multiply these coefficients by $q = \left(1 - \frac{1}{\gamma}\right) \frac{L_t}{W_t}$. For a typical US investor, the annual wealth-to-income ratio is approximately six. This means that the $q \approx \left(1 - \frac{1}{10}\right) \frac{1}{24} \approx 0.0375$, assuming a coefficient of relative risk aversion $\gamma = 10$ and multiplying the annual wealth-to-income ratio by four to get the quarterly wealth-to-income ratio. Note that in our specification q is independent of the horizon.

We estimate hedging portfolio weights for investors with horizons of one quarter up to eight years. Table 4.3 shows that the portfolio adjustments for investors with a short term horizon are economically small and mostly insignificant. For instance, we find that an investor who faces aggregate labor income risk and has an investment horizon of one quarter should optimally overweight stocks from the retail industry by a mere 0.08% ($2.12\% \times 0.0375$). However, at longer horizons, the statistical significance and economic magnitudes are substantially larger. We find that at the four-year horizon, this "typical" investor should significantly overweight stocks from the retail industry by 4.91% ($131\% \times 0.0375$) to optimally adjust for her labor income risk.

Intuitively, the longer the horizon, the more the investor is exposed to labor income risk, and the more she should adjust her portfolio. Looking at statistical significance, however, the results indicate that the one- to four-year horizons are most important. At these horizons, there are more individually statistically significant portfolio weights than at the other horizons. Furthermore, as shown in Figure 4.2, the R^2's of the portfolio weight regressions tend to peak at these horizons. All in all, these findings provide suggestive evidence that the medium-term horizons are most important for portfolio allocation decisions.

When analyzing the composition of the hedging portfolios for investors working in different industries, we find several interesting patterns. In general, equities from the services and retail industries serve as good hedges for labor income risk. Most investors should optimally overweight retail stocks and underweight stocks from the services industry to hedge their labor income risk. Furthermore, if human capital and equity returns are positively correlated, investors should invest less in stocks from their industry of employment in order to diversify. While we find little evidence of this based on short term returns, for longer horizons up to four years, we find significant effects for about one-third of the industries. Specifically, at the four-year horizon, investors working in the services, mining, and finance industries should optimally hold respectively 8.9%, 6.7% and 11.4% less of their own industry stocks. At the same time, investors from the wholesale industry should optimally hold 5.4% more of their own industry stocks as a hedge against their labor income risk. This suggests a negative correlation between labor income growth and equity returns within the wholesale industry at the four-year horizon.

The panels of Table 4.3 also report the allocation to the equity market portfolio if that were the only available asset to hedge labor income risk. For the "typical" investor with aggregate labor income risk, the equity market portfolio is not a suitable hedging instrument at any horizon: the coefficient is never statistically significant and the R^2's are very close to zero. This is mostly the case as well for investors facing industry-specific labor income risk, where only a few industry-horizon combinations have a significant portfolio weight and R^2's remain close to zero. This suggests that in order to hedge against labor income risk, one should adjust the composition of the risky stock portfolio, rather than varying the degree of equity market risk exposure of the overall portfolio.

Table 4.4 reports the portfolio adjustments arising from the "Keeping up with the Joneses" preferences. Our goal here is to assess whether this component of the optimal portfolio also displays horizon effects. Indeed, the table shows that similar to the hedging portfolios, the KUJ effect is slightly stronger at the two- to four-year horizon, where the R^2's are somewhat higher.

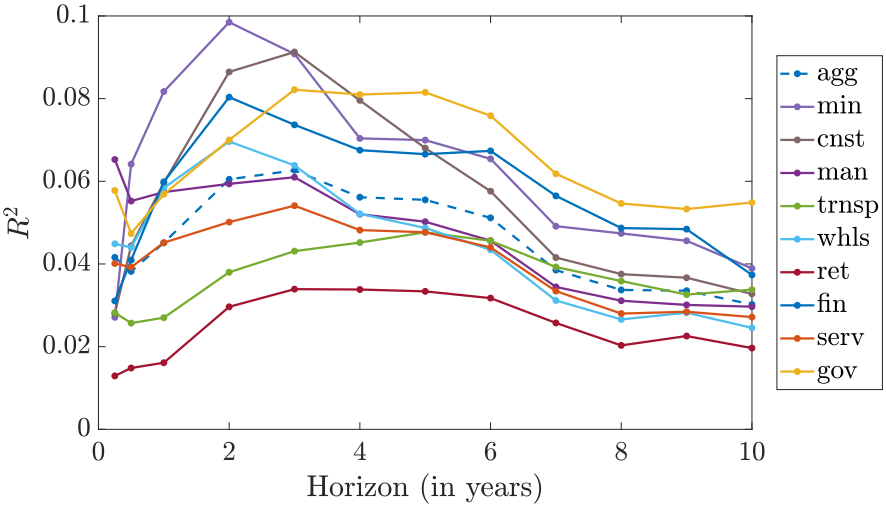


Figure 4.2. Coefficients of determination of portfolio hedging weight regressions across different horizons. This figure plots the R^2s of the portfolio hedging weight regressions of investors that are exposed to either aggregate or industry-specific labor income risk across different horizons.

Table 4.3. Labor income risk induced portfolio hedging weights for different horizons. This table reports the estimated weights (up to a factor q, which depends on the income-to-wealth ratio) for different horizons of equity industry portfolios in the hedging portfolios induced by either aggregate or industry-specific labor income risk from nine different industries. In particular, the table reports the negative regression coefficients (multiplied by 100 to get results in %) of the following regression equation:

$$\sum_{i=1}^{h} \left(l_{k,t+i} - l_{k,t}\right) = \alpha_k + \sum_{z=1}^{8} \beta_z r_{z,t+1} + \epsilon_{k,t}.$$

Here $(l_{k,t+i} - l_{k,t})$ denotes the i-quarter log labor income growth rate in industry k, and $r_{z,t+1}$ the quarterly log returns of industry equity portfolio z. Panels A, B, C, D, E, and F report the results for quarter-, half-, one-, two-, four-, and eight-year horizons, respectively (i.e., $h = 1, 2, 4, 8, 16, 32$ for panels B, C, D, E, and F respectively). For all panels, we also display the hedging weight of the aggregate stock market portfolio. In each regression, we use quarterly data starting from 1958Q2 and ending in 2017Q4 minus $(h-1)$ quarters. The first column gives the weights for an investor that is exposed to aggregate labor income risk. The other columns present hedging portfolio weights for an investor working in one of the nine industries. The last row reports the R^2. $*$, $**$, and $***$ denote significance at the 10%, 5%, and 1% levels. Newey-West standard errors with 1, 1, 3, 7, 15, and 31 lags are used for panels A, B, C, D, E, and F, respectively.

	HC_{agg}	HC_{mining}	HC_{constr}	HC_{manuf}	HC_{transp}	HC_{wholes}	HC_{retail}	$HC_{finance}$	HC_{serv}	HC_{gov}
Panel A: Quarterly horizon, hedging portfolio weights (in %)										
mining	0.33	**-2.85**	-1.26	1.45	-0.10	1.69	-0.38	5.13	-0.90	-1.21**
constr	-0.06	-0.97	**0.31**	-1.08	2.06**	-0.18	-0.10	-3.00	1.59	0.42
manuf	-0.31	0.68	3.47	**-4.41**	0.20	-3.84	-0.63	-4.61	2.05	2.13
transp	1.96	-1.40	0.94	1.73	**1.12**	2.60	0.55	2.05	3.11	2.87***
wholes	0.58	-0.05	1.09	0.28	-0.75	**0.91**	0.35	-1.34	2.56*	0.55
retail	2.12*	-3.09	0.28	5.10***	1.83	3.96**	**0.82**	6.29*	1.15	-1.00

Continuation of Table 4.3

	HC_{agg}	HC_{mining}	HC_{constr}	HC_{manuf}	HC_{transp}	HC_{wholes}	HC_{retail}	$HC_{finance}$	HC_{serv}	HC_{gov}
finance	-1.86	2.56	-0.38	-0.09	-1.79	-1.96	-0.62	**-5.88**	-3.69	-1.31
serv	-1.47	1.02	-2.49*	-1.95	-2.72*	-1.80	0.47	1.64	**-3.70****	-1.01
R^2	4.16%	2.71%	2.80%	6.53%	2.82%	4.49%	1.29%	3.11%	4.01%	5.78%
agg	**0.27**	-3.18	1.20	-0.51	-0.76	-0.01	0.61	-1.33	0.89	0.71
R^2	0.10%	1.22%	0.91%	0.14%	0.27%	0.00%	0.35%	0.15%	0.34%	0.86%
Panel B: Half-year horizon, hedging portfolio weights (in %)										
mining	-0.65	**-12.10*****	-1.68	0.74	-3.02	0.94	-0.03	3.26	-1.19	-2.72**
constr	0.04	1.30	**-0.62**	-1.46	3.75	0.20	-1.36	-2.25	3.10	1.18
manuf	0.08	10.59	6.28	**-4.03**	2.91	-5.05	-1.45	-3.18	-0.22	5.40
transp	4.91*	-2.07	5.73	3.32	**3.95**	4.19	3.03	3.84	8.31**	6.69**
wholes	2.89	-1.04	1.61	2.72	1.10	**4.30**	1.97	0.57	6.41*	2.21
retail	3.33	-9.20	4.23	8.29**	1.12	7.63**	**1.17**	10.19	1.41	-2.23
finance	-4.48	0.71	-3.08	-1.22	-3.99	-5.03	-2.85	**-17.61***	-6.76	-3.17
serv	-4.04	0.46	-6.36*	-6.45*	-6.31**	-5.14	0.91	1.17	**-7.30****	-3.79**
R^2	3.82%	6.41%	4.44%	5.52%	2.57%	4.40%	1.48%	4.09%	3.92%	4.74%
agg	**-0.29**	-8.83*	2.34	-1.66	-2.21	-0.43	0.67	-5.76	0.82	1.43
R^2	0.02%	2.49%	0.69%	0.34%	0.47%	0.02%	0.09%	0.94%	0.07%	0.52%
Panel C: One-year horizon, hedging portfolio weights (in %)										
mining	-4.68	**-32.45*****	-3.84	-2.76	-8.68	-1.32	0.15	3.81	-5.20	-8.37**
constr	0.42	-2.15	**-3.43**	-4.10	5.30	-0.70	-3.59	-0.07	6.96	4.61
manuf	0.57	29.31	11.68	**-1.14**	12.74	-11.94	-3.46	-11.30	-2.82	19.50*
transp	11.72	-3.93	12.99	7.22	**10.66**	7.48	9.22	4.94	16.56	19.70**

Continuation of Table 4.3

	HC_{agg}	HC_{mining}	HC_{constr}	HC_{manuf}	HC_{transp}	HC_{wholes}	HC_{retail}	$HC_{finance}$	HC_{serv}	HC_{gov}
wholes	10.71	2.90	-0.82	9.51	8.84	**14.28**	7.47	13.07	19.35**	6.26
retail	9.21	-0.57	17.19*	17.89**	0.76	21.73***	**2.91**	27.79*	4.84	-3.51
finance	-9.86	-4.95	-1.82	-2.32	-10.03	-10.19	-8.54	**-42.63****	-10.11	-12.04*
serv	-13.53*	-6.20	-15.36*	-19.14**	-18.40**	-14.97*	-0.98	-8.40	**-20.82****	-13.99**
R^2	4.51%	8.17%	5.96%	5.74%	2.70%	5.84%	1.61%	5.99%	4.52%	5.69%
agg	**-1.47**	-15.30	7.11	-4.05	-5.19	-1.19	-0.23	-16.16*	2.40	5.28
R^2	0.06%	1.46%	0.94%	0.30%	0.39%	0.03%	0.00%	1.72%	0.09%	0.83%
Panel D: Two-year horizon, hedging portfolio weights (in %)										
mining	-20.97	**-84.92*****	-18.99	-18.89	-30.20	-13.97	-4.06	9.64	-19.05	-32.68*
constr	0.89	-12.93	**-13.98**	-7.84	12.85	-4.14	-4.95	2.92	14.18	10.18
manuf	24.96	93.40	25.72	**42.78**	68.51	-16.46	4.17	-8.00	9.72	81.52**
transp	32.95	-2.61	32.18	27.04	**34.16**	22.77	31.78	11.28	38.49	64.63**
wholes	36.14	30.39	1.01	25.59	28.53	**47.09**	29.46*	57.68**	61.12**	17.46
retail	35.34*	37.70	71.13***	51.30**	19.13	65.69***	**16.56**	81.78**	26.55	-0.52
finance	-40.23	-38.91	-3.20	-32.22	-45.39	-39.12	-43.82	**-124.81***	-37.19	-45.67
serv	-52.36*	-45.49	-50.70*	-61.89*	-68.48**	-49.19	-17.47	-55.59	**-71.20****	-51.36**
R^2	6.05%	9.85%	8.65%	5.94%	3.80%	6.96%	2.96%	8.04%	5.01%	7.00%
agg	**-5.13**	-30.85	14.68	-3.75	-6.65	-5.88	-2.32	-44.40**	0.36	16.67
R^2	0.07%	0.86%	0.44%	0.03%	0.07%	0.07%	0.02%	2.01%	0.00%	0.74%
Panel E: Four-year horizon, hedging portfolio weights (in %)										
mining	-38.69	**-178.58****	-49.60	-36.57	-92.73	-10.22	13.48	78.56	-29.26	-75.91
constr	-31.52	-83.79	**-74.64**	-39.06	-3.86	-41.14	-28.38	-55.36	-4.48	-10.30

Continuation of Table 4.3

	HC_{agg}	HC_{mining}	HC_{constr}	HC_{manuf}	HC_{transp}	HC_{wholes}	HC_{retail}	$HC_{finance}$	HC_{serv}	HC_{gov}
manuf	121.77	284.55	38.72	**152.00**	264.60	-15.49	42.90	26.32	91.80	289.69*
transp	123.37	91.86	68.71	122.31	**157.07**	108.11	103.31	83.03	130.80	234.24*
wholes	110.32	107.31	36.18	68.17	92.01	**143.90***	89.83	169.80**	185.30**	65.64
retail	131.03*	143.66	232.75***	165.05*	100.50	185.15**	**84.67**	227.86**	122.74	52.74
finance	-136.76	-122.96	-24.61	-112.42	-137.79	-142.88	-143.28	**-305.15***	-146.56	-168.10
serv	-174.24	-165.96	-124.74	-180.68	-235.22	-138.88	-93.47	-184.48	**-236.24***	-201.72*
R^2	5.61%	7.04%	7.96%	5.20%	4.52%	5.21%	3.38%	6.75%	4.82%	8.10%
agg	**8.23**	-9.72	22.65	31.61	26.73	8.34	-0.81	-72.67	8.62	57.34*
R^2	0.02%	0.01%	0.11%	0.18%	0.09%	0.01%	0.00%	0.73%	0.01%	0.74%
Panel F: Eight-year horizon, hedging portfolio weights (in %)										
mining	5.54	**-307.33**	-54.68	-6.95	-222.41	75.68	137.97	291.59	84.45	-36.48
constr	-159.11	-378.72*	**-221.29**	-151.78	-85.33	-136.58	-80.47	-159.11	-130.23	-193.82
manuf	255.93	578.79	27.14	**313.92**	664.33	-60.04	-48.55	19.41	197.52	680.86
transp	293.93	372.43	80.07	252.03	**518.03**	297.51	186.84	286.02	253.95	498.35
wholes	133.87	85.78	7.04	62.16	82.74	**221.91**	164.59	230.52	258.86	31.51
retail	469.99*	512.59	536.16**	539.11*	432.33	513.84*	**346.47***	726.74**	512.13	352.21
finance	-285.26	-276.70	-4.36	-221.51	-237.21	-333.80	-284.82	**-727.50**	-310.02	-327.12
serv	-386.77	-220.57	-161.36	-400.45	-612.74	-288.97	-251.71	-447.31	**-526.95**	-493.99
R^2	3.37%	4.74%	3.75%	3.11%	3.59%	2.66%	2.03%	4.87%	2.80%	5.46%
agg	**51.00**	109.15	62.65	110.41	171.82	66.01	-17.70	-121.46	27.70	114.31
R^2	0.05%	0.11%	0.08%	0.18%	0.29%	0.08%	0.01%	0.21%	0.01%	0.26%

Table 4.4. "Keeping-up-with-the-Joneses" induced portfolio hedging weights for different horizons. This table reports the weights for different horizons of equity industry portfolios in the hedging portfolios induced by investors benchmarking their wealth to aggregate labor income. In particular, the table reports the regression coefficients (multiplied by 100 to get results in %) of the following regression equation:

$$\bar{l}_{t+h} - \bar{l}_t = \alpha + \sum_{z=1}^{8} \beta_z r_{z,t+1} + \epsilon_t.$$

Here $(\bar{l}_{t+h} - \bar{l}_t)$ denotes the h-quarter aggregate log labor income growth rate, and $r_{z,t+1}$ the quarterly log returns of industry equity portfolio z. We report the results for a quarterly horizon of labor income growth rates in the first column (i.e., $h = 1$), as well as the results for half-, one-, two-, four-, and eight-year horizons (i.e., $h = 2, 4, 8, 16, 32$, respectively) in the other columns. The weights are estimated up to a factor $\tilde{q}$, which depends on the income-to-wealth ratio. The last row reports the R^2. $*$, $**$, and $***$ denote significance at the 10%, 5%, and 1% levels. Newey-West standard errors with 1, 1, 3, 7, 15, and 31 lags are used for columns 1, 2, 3, 4, 5, and 6, respectively.

	HC_{agg} (quarter)	HC_{agg} (half-year)	HC_{agg} (one-year)	HC_{agg} (two-year)	HC_{agg} (four-year)	HC_{agg} (eight-year)
mining	-0.33	0.97	2.25	4.61	3.70	-0.41
constr	0.06	-0.12	-0.76	0.86	4.93	7.35
manuf	0.31	-0.37	-0.63	-8.41	-15.53	-12.42
transp	-1.96	-2.94	-4.52	-6.77	-15.88	-13.29
wholes	-0.58	-2.30	-4.33	-7.67	-9.82	-2.00
retail	-2.12*	-1.24	-3.12	-6.75	-15.53*	-22.03
finance	1.86	2.62	3.80	8.21	15.43	11.07
serv	1.47	2.57	5.68**	10.95*	18.15	14.11
R^2	4.16%	3.60%	4.69%	5.61%	5.64%	2.02%

4.4.2 *Asset pricing tests across different horizons*

This section discusses how aggregate labor income risk at different horizons affects the cross-section of expected stock returns. As outlined in Section 4.3, we use a standard two-pass regression methodology, in

which we first estimate risk exposures to horizon-specific labor income risk factors using a multivariate time-series regression, and subsequently use these to estimate market prices of risk in a second-stage cross-sectional regression. Throughout this section, we will focus on the results where we use the standard 25 size and book-to-market portfolios as test assets. Section 4.5 shows that our results are robust to a broader cross-section of test portfolios, including size-investment and industry portfolios.

Table 4.5 reports the estimated first-stage betas and their corresponding Newey-West adjusted t-statistics resulting from multivariate time-series regressions of portfolio excess returns on excess market returns and labor income growth rates across different horizons. The exposures to labor income growth at the two- to four-year horizon (i.e., scale $j = 4$) are substantially more pronounced than at other horizons. First, at the two- to four-year horizon most betas are individually statistically significant, whereas most betas at other horizons are not. Second, a Wald test on the null hypothesis that all betas at a certain horizon are jointly equal to zero and a Wald test on the null hypothesis that all betas at a certain horizon are equal to each other are only rejected for the two- to four-year horizon and the very long horizon (scale $j > 5$). Furthermore, the risk exposures at the two- to four-year horizon are typically negative and show a clear increasing pattern in the size dimension and decreasing pattern in the value dimension. On the other hand, the exposures at other horizons are typically positive and do not show clear cross-sectional dispersion.

Table 4.5. Exposures of 25 size-BM portfolios to labor income risk across different horizons. This table presents the first-stage scale-wise betas with respect to aggregate labor income risk. The betas are estimated by running for each portfolio i the following (quarterly) multivariate time-series regression:

$$R^{e,i}_{t+1} = \alpha_{0,i} + \beta_{mkt,i} R^{e,mkt}_{t+1} + \sum_{j=1}^{J} \beta^{(j)}_{l,i} f^{(j)}_t + \beta^{(j>J)}_{l,i} f^{(j>J)}_t + \epsilon_{i,t+1},$$

where $f^{(j)}_t$ denotes the j-th scale component of aggregate labor income growth and $R^{e,mkt}_{t+1}$ the excess equity market returns. The associated t-statistics (in parentheses) are based on Newey-West adjusted standard errors with 32 lags. $*, **, ***$ indicate significance at the 10%, 5%, and 1% level, respectively. For each scale, the table also reports the p-values of two Wald tests on the joint significance of the corresponding betas. The sample period is 1958Q2 to 2017Q4.

	Betas					t-statistics				
Panel A: Scale j = 1										
	Growth	BM2	BM3	BM4	Value	Growth	BM2	BM3	BM4	Value
Small	1.87	1.94*	2.21**	2.48**	2.26*	(1.10)	(1.69)	(2.18)	(2.35)	(1.81)
SMB2	1.23	0.87	0.54	1.29*	2.05**	(1.56)	(1.02)	(0.82)	(1.85)	(2.42)
SMB3	-0.55	0.92**	1.14	0.20	1.19**	(-0.61)	(1.97)	(1.42)	(0.44)	(1.99)
SMB4	-0.24	0.60	0.35	0.20	1.18	(-0.49)	(1.03)	(0.97)	(0.59)	(1.29)
Big	-0.52	0.64*	0.83	0.83	1.77	(-1.32)	(1.94)	(1.58)	(1.11)	(1.48)
H_0: all betas are zero	(0.1025)									
H_0: all betas are equal	(0.0972)									

Continuation of Table 4.5

	Betas					t-statistics				
Panel B: Scale j = 2										
	Growth	BM2	BM3	BM4	Value	Growth	BM2	BM3	BM4	Value
Small	1.32	1.24	0.42	1.04	1.10	(1.15)	(1.14)	(0.29)	(0.56)	(0.53)
SMB2	0.10	0.34	1.05	1.24	2.34	(0.10)	(0.39)	(0.76)	(0.73)	(1.14)
SMB3	-0.66	0.46	0.44	0.51	0.82	(-0.84)	(0.52)	(0.30)	(0.28)	(0.51)
SMB4	-0.05	-0.14	0.17	0.33	1.64	(-0.07)	(-0.14)	(0.21)	(0.24)	(1.16)
Big	0.21	0.08	0.64	1.35	1.33	(0.19)	(0.11)	(0.95)	(0.71)	(0.69)
H_0: all betas are zero	(0.5460)									
H_0: all betas are equal	(0.8205)									
Panel C: Scale j = 3										
	Growth	BM2	BM3	BM4	Value	Growth	BM2	BM3	BM4	Value
Small	5.27	7.20	2.55	3.05	3.87	(0.85)	(1.44)	(0.87)	(0.92)	(1.51)
SMB2	1.56	1.96	-0.15	1.56	2.42	(0.49)	(0.88)	(-0.06)	(0.72)	(1.03)
SMB3	-0.43	0.97	0.52	1.78	2.35	(-0.23)	(0.57)	(0.25)	(0.82)	(0.78)
SMB4	2.38	-0.70	0.56	1.60	2.38	(0.78)	(-0.54)	(0.42)	(0.85)	(0.97)
Big	-1.06	-1.26	-1.66*	0.53	1.99	(-1.23)	(-0.76)	(-1.64)	(0.19)	(1.10)
H_0: all betas are zero	(0.1442)									
H_0: all betas are equal	(0.1351)									

Continuation of Table 4.5

	Betas					t-statistics				
Panel D: Scale j = 4										
	Growth	BM2	BM3	BM4	Value	Growth	BM2	BM3	BM4	Value
Small	-1.89	-5.52**	-4.57	-6.88**	-7.96**	(-0.63)	(-2.14)	(-1.44)	(-2.03)	(-2.18)
SMB2	-2.16	-4.91*	-6.19**	-6.76**	-7.95**	(-0.97)	(-1.85)	(-2.12)	(-2.42)	(-2.16)
SMB3	-0.96	-4.15**	-7.00***	-7.76**	-7.72***	(-0.50)	(-2.19)	(-3.60)	(-2.56)	(-2.77)
SMB4	-0.54	-4.35**	-3.88	-5.95**	-4.82	(-0.30)	(-2.01)	(-1.41)	(-2.44)	(-1.41)
Big	2.02	-0.06	-3.08**	-0.17	-2.86	(1.07)	(-0.04)	(-1.97)	(-0.06)	(-1.39)
H_0: all betas are zero	(0.0012)									
H_0: all betas are equal	(0.0058)									
Panel E: Scale j = 5										
	Growth	BM2	BM3	BM4	Value	Growth	BM2	BM3	BM4	Value
Small	4.70	5.55*	4.47*	4.61**	4.53*	(1.16)	(1.92)	(1.87)	(2.19)	(1.65)
SMB2	4.56	4.00**	2.18	4.08**	1.93	(1.45)	(2.05)	(1.10)	(2.41)	(1.15)
SMB3	2.57	1.29	2.81	1.35	1.50	(0.86)	(0.53)	(1.64)	(1.05)	(0.59)
SMB4	0.59	1.75	0.87	-1.10	0.32	(0.26)	(1.30)	(0.76)	(-1.22)	(0.25)
Big	-1.98	-0.26	0.33	0.25	-2.63***	(-1.38)	(-0.35)	(0.41)	(0.19)	(-2.84)
H_0: all betas are zero	(0.2574)									
H_0: all betas are equal	(0.4053)									

Continuation of Table 4.5

	Betas					t-statistics				
Panel F: Scale j > 5										
	Growth	BM2	BM3	BM4	Value	Growth	BM2	BM3	BM4	Value
Small	3.85**	1.70	1.28	1.15	1.68	(1.97)	(1.32)	(1.05)	(0.90)	(1.26)
SMB2	1.89	0.59	0.14	0.99	1.57	(1.40)	(0.55)	(0.14)	(1.14)	(1.58)
SMB3	0.68	0.33	0.10	0.32	-0.14	(0.76)	(0.40)	(0.12)	(0.43)	(-0.13)
SMB4	-0.44	-0.93	0.76	0.37	1.06	(-0.64)	(-1.10)	(0.80)	(0.47)	(1.32)
Big	-0.65	-0.49	-0.71	1.60*	-0.62	(-1.18)	(-0.83)	(-1.10)	(1.68)	(-1.33)
H_0: all betas are zero	(0.0000)									
H_0: all betas are equal	(0.0001)									

Table 4.6 reports the estimated market prices of risk at different horizons resulting from the second-stage cross-sectional regressions. We report both Fama and MacBeth (1973) t-statistics, as well as adjusted t-statistics based on our robust standard errors that correct for heteroskedasticity, autocorrelation, and the fact that the betas are estimated in first-stage regressions. All specifications under consideration always contain an intercept term and the equity market risk factor. As the model is estimated using excess returns, the intercept should be statistically insignificant. We note that in most specifications this is indeed the case. At the same time, in all specifications, the price of equity market risk is statistically insignificant.

Our results again point towards a dominant role of labor income risk at the two- to four-year horizon. In the full model specification, in which we include labor income growth rates across all different horizons, the market price of labor income risk at the two- to four-year horizon ($\lambda_l^{(4)}$) is the only one that is statistically significant. Furthermore, when considering two-factor specifications in which we only include the equity market risk factor and one labor income growth factor at one specific horizon, the adjusted R^2 peaks for the two- to four-year horizon. In particular, the adjusted R^2 of the two-factor specification with scale $j = 4$ is 77%, which is even higher than the adjusted R^2 of 76% in the full model specification. By comparison, the adjusted R^2s of the other two-factor specifications with labor income risk at different horizons all range between -4% ($j > 5$) and 27% ($j = 3$). Note that while labor income risk at some of the other horizons carries a significant price of risk when included in a two-factor specification, these are not robust to including the medium-term labor income risk factor. Only the latter remains significantly priced when included along with labor income risk at other horizons.[16]

The cross-sectional results clearly indicate that the two- to four-year horizons is the most relevant. Therefore, our preferred model specification is the two-factor model that includes the equity market risk factor and the labor income growth risk factor at scale $j = 4$. The estimated market price of labor income risk in this two-factor model specification is statistically significant at -0.0020, which means that a (positive) unit

16 While the frequency decomposition results in labor income risk factors that are asymptotically uncorrelated across horizons, the estimated exposures may still be correlated in the second-stage cross-sectional regressions. Hence, the estimated prices of risk can still depend on which other factors are included in the second-stage regression.

Table 4.6. Cross-sectional regressions for 25 size-BM portfolios (multivariate betas). This table reports the second-stage cross-sectional regression results for different model specifications using aggregate labor income growth rates and 25 double-sorted size-BM portfolios as test assets. We report time-series averages of the second-stage market prices of risk (per quarter) with error-in-variable and autocorrelation corrected t-statistics in square brackets (using Newey-West adjustments with 2^j lags). The last column reports the cross-sectional R^2 and adjusted-R^2 (in square brackets). $*, **, ***$ indicate significance at the 10%, 5%, and 1% level, respectively, based on the adjusted t-statistics. The sample period runs from 1958Q2 until 2017Q4.

	λ_0	λ_{mkt}	$\lambda_l^{(1)}$	$\lambda_l^{(2)}$	$\lambda_l^{(3)}$	$\lambda_l^{(4)}$	$\lambda_l^{(5)}$	$\lambda_l^{(j>5)}$	R^2
All	0.0136	-0.0001	0.0004	0.0012	0.0006	-0.0020***	-0.0003	-0.0010	0.83
	[1.20]	[-0.01]	[0.26]	[0.87]	[1.21]	[-2.72]	[-0.39]	[-0.75]	[0.76]
$j = 1$	0.0266	-0.0073	0.0045*						0.29
	[1.54]	[-0.42]	[1.81]						[0.23]
$j = 2$	0.0217	-0.0020		0.0052					0.24
	[0.81]	[-0.08]		[1.04]					[0.18]
$j = 3$	0.0399*	-0.0185			0.0025**				0.33
	[1.75]	[-0.84]			[2.04]				[0.27]
$j = 4$	0.0162	-0.0022				-0.0022**			0.79
	[1.17]	[-0.15]				[-2.16]			[0.77]
$j = 5$	0.0402***	-0.0186					0.0019**		0.21
	[2.96]	[-1.28]					[2.17]		[0.14]
$j > 5$	0.0325***	-0.0094						0.0017	0.04
	[2.92]	[-0.78]						[1.32]	[-0.04]

exposure to this risk factor lowers expected returns with 20 basis points per quarter. As reported in Table 4.6, however, the betas at scale $j = 4$ are typically negative and therefore the risk premium (exposure $\times$ market price of risk) is positive and economically significant. For example, the quarterly risk premium for stocks on small growth firms is $-1.89 \times -0.0020 \approx 40$ basis points, whereas the quarterly risk premium for small value firms is $-7.96 \times -0.0020 \approx 160$ basis points.

In light of our theoretical model, the negative sign of the price of labor income risk is only possible if $\psi > 0$, that is, if agents benchmark their wealth to aggregate labor income, that proxies the labor income of their "peers". When these "Keeping up with the Joneses" preferences are strong enough (and outweigh the hedging demand channel), investors want to invest in securities that are strongly correlated with aggregate labor income. They are therefore willing to pay a higher price for these securities, resulting in a negative labor income risk premium. This is in line with Gomez et al. (2009) and Gomez et al. (2016) who also find strong empirical evidence for such a "Keeping up with the Joneses" benchmark mechanism. However, these papers do not consider horizon effects.[17]

We compare the performance of our model to a number of well-known benchmark asset pricing models. In particular, we consider the static CAPM, the traditional human capital CAPM, which augments the standard CAPM with a (contemporaneous) aggregate labor income growth factor, the consumption CAPM (CCAPM) with quarterly consumption growth, the ultimate consumption CAPM of Parker and Julliard (2005) that includes three-year ahead consumption growth, the Fama and French (1993) three-factor (FF3) and Fama and French (2015) five-factor (FF5) models. Table 4.7 reports the results. For ease of comparison, we also include the results of our preferred specification of of Table 4.6 that includes the equity market return and the two- to four-year labor income

17 Other authors, albeit in different settings, also find empirical evidence of a negative market price of labor income risk. For example, Julliard (2007) shows that a variable capturing expected future labor income growth rates is negatively related to stock market excess returns. He argues that changes in expected future labor income are closely related to time-varying risk premia. In his narrative, high expected future labor income growth rates represent a state of the world where high labor income can be used to finance consumption, thereby decreasing the fear of low stock market returns. Maio and Min (2018) develop a consumption-based asset pricing model that includes labor income growth as a risk factor. Their model implies a negative market price of risk for labor income growth, since higher labor income growth is associated with lower leisure hours as opportunity costs of leisure increase when wages are higher. Other things equal, less leisure time increases the marginal utility of consumption.

risk factor (i.e., scale $j = 4$). We find, consistent with the literature, the cross-sectional fit of the static CAPM to be poor. The intercept is statistically significant and positive, the (adjusted) R^2 is low, and the market price of equity market risk is not statistically significant. The addition of the quarterly aggregate labor income growth factor in the human capital CAPM does not improve the cross-sectional fit. The intercept remains positive and statistically significant, the R^2 hardly improves, and the market price of labor income growth risk is not statistically significant. These results indeed confirm the findings in previous literature that contemporaneous aggregate labor income growth does not seem to play an important role in asset pricing (see, e.g., Eiling, 2013).

Comparing the results of the traditional human capital CAPM with our preferred specification of the scale-specific human capital CAPM clearly highlights the effect of taking into account the 'right' horizon. Indeed, replacing the contemporaneous labor income growth factor of the traditional human capital CAPM with the labor income growth factor at scale $j = 4$ drastically improves the R^2 from 8% to 79%. In fact, this parsimonious two-factor specification seems to capture the cross-section of stock returns even better than the Fama and French (1993) three-factor (FF3) and Fama and French (2015) five-factor (FF5) models, which are particularly designed to perform well on the size-BM portfolios.

The dominant horizon for labor income risk coincides with the three-year horizon of the ultimate consumption CAPM of Parker and Julliard (2005) who use 11-quarter ahead real per capita labor income growth rate as risk factor. Consistent with the findings of Parker and Julliard (2005), we find that this specification greatly outperforms the traditional consumption CAPM with an adjusted R^2 of 30% as opposed to -4%. However, our preferred model with the labor income growth factor at the two- to four-year horizon in turn greatly outperforms the ultimate consumption CAPM with an adjusted R^2 of 77%. Furthermore, when including the (ultimate) consumption CAPM risk factors to our preferred model specification, the market price of two- to four-year labor income risk remains statistically significant, whereas the market prices of the (ultimate) consumption CAPM risk factors become statistically insignificant. This shows that the two- to four-year labor income risk factor is not a mere proxy for ultimate consumption risk. The market price of two- to four-year labor income risk is also robust to including the FF3 and FF5 factors.

Table 4.7. Cross-sectional regressions for 25 size-BM portfolios - comparison with alternative asset pricing models. Panel A of this table evaluates the cross-sectional regression results of six benchmark asset pricing models for quarterly excess returns on 25 size-BM equity portfolios. We consider the standard CAPM, the standard human capital CAPM (HC CAPM), the classic consumption CAPM (CCAPM), the ultimate consumption CAPM of Parker and Julliard (2005), the Fama and French (1993) 3-factor model (FF3), and the Fama and French (2015) 5-factor model (FF5). In Panel B, we report cross-sectional regression results for our preferred specification of the scale-specific human capital CAPM model and several setups that augment this preferred model specification. The cross-sectional regressions are estimated using the Fama and MacBeth (1973) procedure. We report the second-stage cross-sectional regression coefficients and corresponding adjusted t-statistics in square brackets. $*, **, ***$ indicate significance at the 10%, 5%, and 1% level, respectively, based on the adjusted t-statistics. The last column reports the R^2 and adjusted-R^2 (in square brackets).

Panel A: Alternative asset pricing models

	λ_0	λ_{mkt}	λ_l	λ_c	λ_{uc}	λ_{SMB}	λ_{HML}	λ_{RMW}	λ_{CMA}	R^2
CAPM	0.0271***	-0.0036								0.01
	[2.95]	[-0.33]								[-0.03]
HC CAPM	0.0333***	-0.0106	0.0032							0.08
	[2.65]	[-0.81]	[1.49]							[-0.00]
CCAPM	0.0203***			0.0005						0.01
	[2.72]			[0.28]						[-0.04]
Ult. CCAPM	0.0083				0.0347**					0.33
	[0.83]				[2.12]					[0.30]
FF3	0.0284***	-0.0127				0.0066*	0.0134***			0.78
	[2.91]	[-1.13]				[1.69]	[3.23]			[0.65]
FF5	0.0160	-0.0029				0.0083**	0.0127***	0.0090	0.0074	0.82
	[1.06]	[-0.18]				[1.98]	[2.79]	[1.40]	[1.45]	[0.77]

Continuation of Table 4.7

Panel B: Preferred specification augmented with other models

	λ_0	λ_{mkt}	$\lambda_l^{(4)}$	λ_c	λ_{uc}	λ_{SMB}	λ_{HML}	λ_{RMW}	λ_{CMA}	R^2
Preferred	0.0162	-0.0022	-0.0022**							0.79
	[1.17]	[-0.15]	[-2.16]							[0.77]
Preferred	0.0122	0.0025	-0.0025**	0.0034						0.78
+ CCAPM	[0.83]	[0.16]	[-2.07]	[1.27]						[0.75]
Preferred	0.0076	0.0070	-0.0023**		0.0240					0.81
+ ult. CCAPM	[0.42]	[0.36]	[-2.36]		[1.40]					[0.78]
Preferred	0.0185	-0.0043	-0.0012**			0.0077**	0.0125***			0.83
+ FF3	[1.49]	[-0.33]	[-1.98]			[1.97]	[3.02]			[0.79]
Preferred	0.0151	-0.0025	-0.0012*			0.0085**	0.0123***	0.0070	0.0077	0.83
+ FF5	[1.16]	[-0.18]	[-1.84]			[2.02]	[2.72]	[1.20]	[1.41]	[0.78]

4.4.3 *Discussion of the results*

Our goal is to empirically examine whether the relation between labor income risk and expected stock returns displays horizon effects, and if so, to identify the dominant horizon(s). The data paint a clear picture. Labor income risk at the two- to four-year horizon strongly and robustly affects the cross-section of expected stock returns. In sharp contrast, labor income risk at other horizons does not matter.

This dominant horizon is consistent with evidence from existing studies on wage rigidity. Various papers find that for search-and-matching models to generate realistic patterns in unemployment and vacancies, assuming wage rigidity is key (see, e.g., Shimer, 2005; Hall, 2005). In the presence of wage rigidity, wages are smoother than the marginal product of labor, and hence smoother than output. This could result in lower short-term correlations between wage growth and stock returns. Potential drivers of wage rigidity are infrequent wage setting, infrequent negotiations between workers and the firm or other labor market frictions.

At some point, wages are adjusted and reflect again the marginal product of labor. Several studies find that the frequency of wage setting is about every three years, which falls within our dominant horizon. For instance, Marfe (2018) measures labor rigidity as employee compensation over net value added (i.e., the labor share). Without wage rigidity, the labor share would be constant as fluctutuations in net value added are one-to-one reflected in fluctuations in compensation. In contrast, Marfe (2018) finds that the labor share fluctuates countercyclically and has a half-life of 3.5 years. Rich and Tracy (2004) show that for most of their sample of labor contracts, the median duration is 35 months. Furthermore, Favilukis and Lin (2016b) derive a production-based asset pricing model in which wages are reset infrequently. They find that assuming that wages are reset every three years helps generate both smooth wages and volatile stock returns in their model.

When wages are reset to match the marginal product of labor every three years, we would expect higher comovements between wage growth and stock returns at this three-year horizon. This is exactly what comes out of our empirical analysis. Stocks' exposures to labor income risk peak at the two- to four-year horizon. Of course, this medium-term horizon is not the only one that could matter. For instance, longer horizons may matter too, as suggested by Benzoni et al. (2007). They assume a cointe-

grating relationship between wages and dividends where labor markets catch up with the aggregate economy in about 20 years. Rather than taking a stance a-priori, we include all horizons in our empirical analysis and let the data speak on which one(s) dominate. Both the cross-sectional asset pricing tests and the portfolio allocation analysis point towards the same medium-term horizon.

4.5 ROBUSTNESS CHECKS

In this section we discuss a variety of robustness checks, including using different sets of test assets, univariate first-stage betas, real labor income growth and different classifications of the horizons.

4.5.1 *Other test assets*

Next to the standard 25 size-BM portfolios, we also consider different sets of test assets. In particular, we consider a broader cross-section of stock returns and add 25 size-investment and 30 industry portfolios to the 25 size-BM portfolios. Table 4.8 reports the results for the combined 25 size-BM and 25-size investment portfolios, and Table 4.9 reports the results for the 25 size-BM and 30 industry portfolios. As a benchmark, we also present the results of the traditional human capital CAPM, the consumption CAPM, the ultimate consumption CAPM, the Fama and French (1993) three-factor and Fama and French (2015) five-factor models for both sets of test assets in Table 4.10.

In line with our previous results, we again find that the medium frequency scale $j = 4$ strongly dominates. The R^2s of the simple two-factor model that contains the equity market factor and the labor income growth factor at scale $j = 4$ are substantially higher than those of model specifications that include other scale components. In fact, the cross-sectional fit is in both cases much better than those of the traditional human capital CAPM, and (ultimate) consumption CAPM models, and comparable to those of the Fama and French (1993) three-factor and Fama and French (2015) five-factor models, which should, in principle, perform particularly well on the 25 size-BM and 25 size-investment portfolios. The intercepts are not statistically significant for the two-factor model with scale $j = 4$, whereas the intercepts of the Fama and French

(1993) three-factor model are always statistically significantly different from zero, and the intercept of the Fama and French (2015) five-factor model is as well for the 25 size-BM and 30 industry portfolios. Similar to our previous findings, the price of risk for scale $j = 4$ is statistically significantly negative for all test assets and all model specifications. When including the 30 industry portfolios, none of the Fama and French (2015) factors carry significant prices of risk. This suggests that the simple two-factor model outperforms the Fama and French (1993) three-factor and Fama and French (2015) five-factor models for these test assets as well.

Table 4.8. Cross-sectional regressions for 50 combined 25 size-BM and 25 size-investment portfolios (multivariate betas). This table reports the second-stage cross-sectional regression results for different model specifications using aggregate labor income growth rates and a combination of 25 double-sorted size-BM and 25 size-investment portfolios as test assets. We report time-series averages of the second-stage market prices of risk (per quarter) with error-in-variable and autocorrelation corrected t-statistics in square brackets (using Newey-West adjustments with 2^j lags). The last column reports the cross-sectional R^2 and adjusted-R^2 (in square brackets). $*, **, ***$ indicate significance at the 10%, 5%, and 1% level, respectively, based on the adjusted t-statistics. The sample period runs from 1963Q3 until 2017Q4.

	λ_0	λ_{mkt}	$\lambda_l^{(1)}$	$\lambda_l^{(2)}$	$\lambda_l^{(3)}$	$\lambda_l^{(4)}$	$\lambda_l^{(5)}$	$\lambda_l^{(j>5)}$	R^2
All	0.0169	-0.0036	0.0008	0.0006	0.0009	-0.0020**	-0.0006	-0.0006	0.79
	[1.52]	[-0.28]	[0.70]	[0.86]	[1.24]	[-2.34]	[-0.75]	[-0.40]	[0.76]
$j = 1$	0.0260*	-0.0072	0.0043*						0.29
	[1.81]	[-0.47]	[1.81]						[0.26]
$j = 2$	0.0186	0.0009		0.0044					0.25
	[0.95]	[0.05]		[1.51]					[0.21]
$j = 3$	0.0353**	-0.0145			0.0022*				0.26
	[2.30]	[-0.92]			[1.88]				[0.22]
$j = 4$	0.0180	-0.0044				-0.0022*			0.72
	[1.41]	[-0.31]				[-1.94]			[0.71]
$j = 5$	0.0320***	−0.0110					0.0014*		0.12
	[3.08]	[-0.92]					[1.84]		[0.08]
$j > 5$	0.0316***	-0.0100						0.0024	0.06
	[3.06]	[-0.85]						[1.54]	[0.02]

Table 4.9. Cross-sectional regressions for 55 combined 25 size-BM and 30 industry portfolios (multivariate betas). This table reports the second-stage cross-sectional regression results for different model specifications using aggregate labor income growth rates and a combination of 25 double-sorted size-BM and 30 industry portfolios as test assets. We report time-series averages of the second-stage market prices of risk (per quarter) with error-in-variable and autocorrelation corrected t-statistics in square brackets (using Newey-West adjustments with 2^j lags). The last column reports the cross-sectional R^2 and adjusted-R^2 (in square brackets). $*, **, ***$ indicate significance at the 10%, 5%, and 1% level, respectively, based on the adjusted t-statistics. The sample period runs from 1958Q2 until 2017Q4.

	λ_0	λ_{mkt}	$\lambda_l^{(1)}$	$\lambda_l^{(2)}$	$\lambda_l^{(3)}$	$\lambda_l^{(4)}$	$\lambda_l^{(5)}$	$\lambda_l^{(j>5)}$	R^2
All	0.0220**	-0.0045	-0.0015	0.0013	0.0001	-0.0011**	0.0005	0.0005	0.47
	[2.32]	[-0.41]	[-1.24]	[1.35]	[0.37]	[-1.97]	[0.89]	[0.43]	[0.40]
$j = 1$	0.0206***	0.0002	0.0003						0.01
	[2.81]	[0.03]	[0.39]						[-0.03]
$j = 2$	0.0213**	-0.0009		0.0012					0.13
	[2.42]	[-0.09]		[1.40]					[0.09]
$j = 3$	0.0202***	0.0008			0.0001*				0.00
	[2.89]	[0.09]			[0.14]				[-0.03]
$j = 4$	0.0191	-0.0019				-0.0012*			0.36
	[2.11]	[-0.18]				[-1.93]			[0.33]
$j = 5$	0.0223***	-0.0019					0.0007		0.07
	[2.95]	[-0.20]					[1.61]		[0.03]
$j > 5$	0.0202***	0.0008						0.0001	0.00
	[2.83]	[0.08]						[0.08]	[-0.04]

Table 4.10. Cross-sectional regressions for 50 size-BM and size-investment, and 55 size-BM and industry portfolios - comparison with alternative asset pricing models. This table evaluates the cross-sectional regression results of five benchmark asset pricing models for quarterly excess returns on 25 size-BM equity portfolios combined with 25 size-investment portfolios (panel A) and combined with 30 industry portfolios (panel B). We consider the standard human capital CAPM (HC CAPM), the classic consumption CAPM (CCAPM), the ultimate consumption CAPM of Parker and Julliard (2005), the Fama and French (1993) 3-factor model (FF3), and the Fama and French (2015) 5-factor model (FF5). We also add our preferred specification of the scale-specific human capital CAPM model. The cross-sectional regressions are estimated using the Fama and MacBeth (1973) procedure. We report the second-stage cross-sectional regression coefficients and corresponding adjusted t-statistics in square brackets. $*, **, ***$ indicate significance at the 10%, 5%, and 1% level, respectively, based on the adjusted t-statistics. The last column reports the R^2 and adjusted-R^2 (in square brackets).

Panel A: 50 combined 25 size-BM and 25 size-investment portfolios

	λ_0	λ_{mkt}	λ_l	λ_c	λ_{uc}	λ_{SMB}	λ_{HML}	λ_{RMW}	λ_{CMA}	$\lambda_l^{(4)}$	R^2
Preferred spec.	0.0180	-0.0044								-0.0022*	0.72
	[1.41]	[-0.31]								[-1.94]	[0.71]
HC CAPM	0.0309***	-0.0095	0.0038								0.10
	[2.68]	[-0.76]	[1.62]								[0.07]
CCAPM	0.0183**			0.0006							0.01
	[2.35]			[0.32]							[-0.11]
Ultimate CCAPM	0.0034				0.0406*						0.38
	[0.31]				[1.75]						[0.37]
FF3	0.0291***	-0.0140				0.0064	0.0136***				0.76
	[3.35]	[-1.32]				[1.49]	[3.00]				[0.75]
FF5	0.0144	-0.0001				0.0073	0.0121***	0.0058	0.0092***		0.79
	[1.10]	[-0.00]				[1.54]	[2.68]	[1.09]	[2.75]		[0.76]

Continuation of Table 4.10

Panel B: 55 combined 25 size-BM and 30 industry portfolios

	λ_0	λ_{mkt}	λ_l	λ_c	λ_{uc}	λ_{SMB}	λ_{HML}	λ_{RMW}	λ_{CMA}	$\lambda_l^{(4)}$	R^2
Preferred spec.	0.0191	-0.0019								-0.0012*	0.36
	[2.11]	[-0.18]								[-1.93]	[0.33]
HC CAPM	0.0210***	-0.0000	0.0003								0.00
	[2.73]	[-0.00]	[0.21]								[-0.04]
CCAPM	0.0255***			-0.0003							0.00
	[3.90]			[-0.25]							[-0.01]
Ultimate CCAPM	0.0199***				0.0033						0.01
	[3.37]				[0.44]						[-0.01]
FF3	0.0295***	-0.0122				0.0056	0.0081*				0.36
	[3.85]	[-1.29]				[1.41]	[1.91]				[0.33]
FF5	0.0246***	-0.0097				0.0068	0.0072	0.0021	0.0084		0.40
	[3.01]	[-0.92]				[1.60]	[1.54]	[0.40]	[1.33]		[0.34]

4.5.2 *Maximum scale specification*

Throughout our analysis, we defined the maximum scale to be $J = 5$, corresponding to eight years. As a result, we have several scale components ($j = 1, 2, 3, 4, 5$) that capture heterogeneity in labor income risk up to typical business-cycle frequency horizons of eight years, and one residual component ($j > 5$) that groups together long-term labor income risk with horizons beyond eight years. The choice for $J = 5$ was made to strike a balance between allowing for enough flexibility of our analysis within the range of typical business-cycle frequencies on the one hand, and maintaining a tractable empirical specification on the other hand. Therefore, specifying a maximum scale below five would result in losing information on risk at business-cycle frequencies, since more horizons would end up in the residual component. Going beyond scale $j = 5$, on the other hand, further increases the number of factors in the cross-sectional analysis. Furthermore, the construction of factors at scales beyond $j = 5$ requires taking moving averages over 64 (or more) labor income growth rates, resulting in highly persistent series. Nevertheless, in this robustness test we select $J = 4$ and $J = 6$. The results are reported in Table 4.11, panel A ($J = 4$) and panel B ($J = 6$). In both cases, the prices of risk of the $j = 4$ scale component is highly significant, while at other scales it is not. Only the price of risk for very long term labor income risk ($J > 6$) is marginally significant in Panel B. Again, the cross-sectional fit of the parsimonious two-factor model with scale $j = 4$ is remarkable. This confirms that our findings are robust to the specification of the maximum scale.[18]

18 Note that the results of the two-factor preferred specification in this table are slightly different from those in Table 4.6. This is due to a slightly different sample period as a result of a different maximum J.

Table 4.11. Cross-sectional regressions for 25 size-BM portfolios - different maximum scales J. This table reports the second-stage cross-sectional regression results for different model specifications using aggregate labor income growth rates and 25 double-sorted size-BM portfolios as test assets. Panel A reports the results for $J = 4$ and panel B reports the results for $J = 6$. We report time-series averages of the second-stage market prices of risk (per quarter) with error-in-variable and autocorrelation corrected t-statistics in square brackets (using Newey-West adjustments with 2^j lags). The last column reports the cross-sectional R^2 and adjusted-R^2 (in square brackets). $*, **, ***$ indicate significance at the 10%, 5%, and 1% level, respectively, based on the adjusted t-statistics. The sample period runs from 1958Q2 until 2017Q4.

	λ_0	λ_{mkt}	$\lambda_l^{(1)}$	$\lambda_l^{(2)}$	$\lambda_l^{(3)}$	$\lambda_l^{(4)}$	$\lambda_l^{(5)}$	$\lambda_l^{(6)}$	$\lambda_l^{(j>4)}$	$\lambda_l^{(j>6)}$	R^2
Panel A: $J = 4$											
$j = 4$	0.0125	0.0032				-0.0025**					0.81
	[0.93]	[0.22]				[-1.97]					[0.79]
$j > 4$	0.0384***	-0.0146							0.0031*		0.13
	[2.82]	[-1.03]							[1.75]		[0.05]
All	0.0215	-0.0056	0.0021	-0.0018	0.0006	-0.0025**			0.0002		0.84
	[1.11]	[-0.28]	[0.79]	[-1.18]	[0.68]	[-2.09]			[0.07]		[0.79]
Panel B: $J = 6$											
$j = 4$	0.0120	0.0032				-0.0015***					0.79
	[1.03]	[0.25]				[-2.70]					[0.77]
$j > 6$	0.0309***	-0.0058								0.0001	0.02
	[3.01]	[-0.52]								[0.11]	[-0.06]
All	0.0108	0.0051	-0.0009	-0.0001	0.0008	-0.0012***	-0.0000	0.0002		-0.0013*	0.83
	[0.98]	[0.38]	[-1.19]	[-0.07]	[1.58]	[-2.97]	[-0.05]	[0.31]		[-1.80]	[0.75]

4.5.3 *Multivariate versus univariate betas*

All second-stage cross-sectional results discussed so far use betas that are estimated in a single multivariate time-series regression in the first stage. In typical settings, re-estimation of the betas in the first stage is required when considering different model specifications in the second stage, since correlations between included and excluded factors might affect the second-stage results. In our case, however, the components are by construction (asymptotically) uncorrelated across scales, and, therefore, in principle, only the correlations of the scale factors with the equity market factor play a role in the first-stage regressions. Since all specifications we consider include the equity market factor as well, our multivariate regression approach ought to be valid. To test this, we compare our cross-sectional results based on betas obtained by the multivariate first-stage regression with cross-sectional results based on betas obtained by univariate first-stage regressions. Table 4.12 reports the cross-sectional results using univariate betas and the 25 size-BM portfolios as test assets. Indeed, we find all market prices of labor income risk to be of similar magnitudes as those based on multivariate betas as in Table 4.6.

4.5.4 *Nominal versus real labor income growth*

Our cross-sectional results are also robust to using real instead of nominal labor income growth rates. We obtain real labor income by deflating nominal wages by the Personal Consumer Expenditure price deflator as reported by the Bureau of Economic Analysis. Table 4.13 presents the cross-sectional regression results on the 25 size-BM portfolios when considering real aggregate labor income growth across different frequency scales. Again, the cross-sectional fit peaks for the two-factor specification with scale $j = 4$. The estimated value of the price of risk at scale $j = 4$ is also very similar to the estimated value when we take nominal labor income growth (see Table 4.6).

Table 4.12. Cross-sectional regressions for 25 size-BM portfolios (univariate betas). This table reports the second-stage cross-sectional regression results for different model specifications using aggregate labor income growth rates and 25 double-sorted size-BM portfolios as test assets. We report time-series averages of the second-stage market prices of risk (per quarter) with error-in-variable and autocorrelation corrected t-statistics in square brackets (using Newey-West adjustments with 2^j lags). The last column reports the cross-sectional R^2 and adjusted-R^2 (in square brackets). $*, **, ***$ indicate significance at the 10%, 5%, and 1% level, respectively, based on the adjusted t-statistics. The sample period runs from 1958Q2 until 2017Q4.

	λ_0	λ_{mkt}	$\lambda_l^{(1)}$	$\lambda_l^{(2)}$	$\lambda_l^{(3)}$	$\lambda_l^{(4)}$	$\lambda_l^{(5)}$	$\lambda_l^{(j>5)}$	R^2
All	0.0136	0.0041	0.0008	0.0012	0.0011	-0.0023**	0.0002	-0.0014	0.83
	[1.23]	[0.26]	[0.47]	[0.74]	[1.23]	[-2.47]	[0.15]	[-1.11]	[0.76]
$j = 1$	0.0270	-0.0007	0.0043*						0.23
	[1.68]	[-0.04]	[1.65]						[0.16]
$j = 2$	0.0222	-0.0064		0.0053					0.25
	[0.82]	[-0.22]		[1.04]					[0.18]
$j = 3$	0.0349**	-0.0193			0.0015*				0.10
	[2.37]	[-1.21]			[1.83]				[0.02]
$j = 4$	0.0116	0.0142				-0.0025**			0.75
	[0.81]	[0.83]				[-2.01]			[0.73]
$j = 5$	0.0313***	-0.0104					0.0006		0.02
	[3.27]	[-0.98]					[0.99]		[-0.06]
$j > 5$	0.0344***	-0.0093						0.0026*	0.09
	[2.95]	[-0.74]						[1.66]	[0.01]

Table 4.13. Cross-sectional regressions for 25 size-BM portfolios - real aggregate labor income growth (multivariate betas). This table reports the second-stage cross-sectional regression results for different model specifications using real aggregate labor income growth rates and 25 double-sorted size-BM portfolios as test assets. We report time-series averages of the second-stage market prices of risk (per quarter) with error-in-variable and autocorrelation corrected t-statistics in square brackets (using Newey-West adjustments with 2^j lags). The last column reports the cross-sectional R^2 and adjusted-R^2 (in square brackets). $*, **, ***$ indicate significance at the 10%, 5%, and 1% level, respectively, based on the adjusted t-statistics. The sample period runs from 1958Q2 until 2017Q4.

	λ_0	λ_{mkt}	$\lambda_l^{(1)}$	$\lambda_l^{(2)}$	$\lambda_l^{(3)}$	$\lambda_l^{(4)}$	$\lambda_l^{(5)}$	$\lambda_l^{(j>5)}$	R^2
All	0.0179	0.0020	-0.0008	-0.0012	0.0022	-0.0019*	-0.0002	-0.0010	0.53
	[0.71]	[0.07]	[-0.28]	[-0.53]	[1.48]	[-1.81]	[-0.28]	[-0.84]	[0.34]
$j = 1$	0.0276***	-0.0046	0.0009						0.02
	[2.98]	[-0.44]	[0.64]						[-0.07]
$j = 2$	0.0274**	-0.0042		0.0023					0.07
	[2.12]	[-0.32]		[1.37]					[-0.02]
$j = 3$	0.03248	-0.0006			0.0023**				0.25
	[1.25]	[-0.03]			[2.07]				[0.18]
$j = 4$	0.0124	0.0059				-0.0028**			0.41
	[0.77]	[0.33]				[-1.99]			[0.36]
$j = 5$	0.0219*	0.0019					0.0010*		0.06
	[1.70]	[0.13]					[1.65]		[-0.03]
$j > 5$	0.0333**	-0.0133						-0.0026*	0.19
	[2.17]	[-0.79]						[-1.84]	[0.12]

4.6 CONCLUSIONS

Labor and equity markets are interconnected. However, the strength of these connections may vary depending on the horizon. There are various possible economic reasons for these horizon effects, such as career length, wage stickiness or long-term cointegration between wages and dividends. This chapter employs a flexible empirical approach to consider comovements between labor income risk and equity returns at many different horizons. Our goal is to determine whether there are any horizon effects, and if so, to extract the dominant horizon(s) from the data.

Our empirical analysis is guided by a simple theoretical model that we derive, in which investors have different career lengths and the covariance between labor income risk and equity returns may vary across horizons. In this model, expected excess stock returns are explicitly driven by their exposures to labor income risk across different horizons.

When testing this multi-factor model, we include equity market returns and aggregate labor income risk at six different horizons ranging from one quarter up to more than eight years. The result is striking: only labor income risk at the two- to four-year horizon is significantly priced. Even more so, when we focus on the two-factor model that includes labor income risk at this frequency (in addition to equity market returns), we find that the model performs similarly to the Fama and French (2015) five-factor model for the cross-section of size and book-to-market and size-investment portfolios. In sharp contrast, the standard human capital CAPM with quarterly labor income growth barely captures any of the cross-sectional variation in expected returns. Hence, the simple adjustment of measuring labor income risk over a medium-term horizon has a dramatic effect on the model's ability to capture cross-sectional differences in expected stock returns. These findings are in line with a portfolio allocation analysis where we estimate how investors can adjust the composition of their optimal stock portfolios to hedge their labor income risk at different horizons. While portfolio adjustments are small when considering quarterly or annual returns, they are more significant, both statistically and economically, when considering a medium-term horizon.

The price of labor income risk at the medium-term horizon is significantly negative. The negative price of labor income risk is in line with

several recent papers (e.g., Gomez et al., 2009; Gomez et al., 2016; Maio and Min, 2018). One possible explanation is the presence of "Keeping up with the Joneses" preferences, where investors benchmark their terminal wealth to the aggregate labor income of their "peers". This makes stocks that are positively exposed to aggregate labor income risk desirable. We follow Abel (1990) and include these preferences in our model. Our empirical finding that the price of labor income risk is negative suggests that the "Keeping up with the Joneses" effect outweighs the traditional hedging demand channel.

Several papers that examine consumption-based models without labor (e.g., Daniel and Marshall, 1997; Parker and Julliard, 2005) find that these models performs best when using consumption growth at the two- to three-year horizon. This suggests that investors make decisions with a medium-term horizon in mind. Nevertheless, our model strongly outperforms the ultimate consumption CAPM of Parker and Julliard (2005), indicating that the medium term labor income risk factor is not a mere proxy for ultimate consumption risk. Instead, our results are consistent with evidence of wage rigidity where wages are slow to adjust to changes in the marginal product of labor. Various papers suggest that wages are reset every three years (e.g., Rich and Tracy, 2004; Favilukis and Lin, 2016b; Marfe, 2018) which could lead to higher comovements between stock returns and labor income risk at the medium-term horizon. Indeed, we find that stocks' exposures to labor income risk peak at this horizon, both in terms of statistical and economic significance.

4.A MODEL DERIVATIONS

In this Appendix, we provide the details of the derivations in our theoretical model. We first derive a log-linear approximation of multi-period wealth, which is subsequently used in the portfolio optimization problem. When aggregating the resulting optimal portfolio demands over cohorts of investors with heterogenous investment horizons, we finally obtain our asset pricing equation.

4.A.1 *Log-linearization of the wealth dynamics*

For notational convenience, we only discuss the case of one risky asset and hence a scalar portfolio weight α_t. The generalization to multiple risky assets and a vector of portfolio weights is straightforward. Defining the log returns variables $r_f = \log(R_f)$ and $r_{t+1} = \log(R_{t+1})$ we can write the one-period portfolio returns as

$$R_{p,t+1} = R_f[1 + \alpha_t(\exp(r_{t+1} - r_f) - 1)]. \tag{4.A.1}$$

Let $w_t = \log(W_t)$ and $l_{t+i} = \log(L_{t+i})$ denote log wealth and log labor income, respectively. We consider a Taylor expansion of w_{t+h} as a function of $(w_t, l_{t+i}, r_{t+1} - r_f)$.[19] The first derivatives are

$$\frac{\partial w_{t+h}}{\partial w_t} = \frac{(R_{p,t+h} \cdots R_{p,t+1})W_t}{W_{t+h}}, \tag{4.A.2}$$

$$\frac{\partial w_{t+h}}{\partial l_{t+i}} = \frac{(R_{p,t+h} \cdots R_{p,t+i+1})L_{t+i}}{W_{t+h}}, \tag{4.A.3}$$

$$\frac{\partial w_{t+h}}{\partial (r_{t+1} - r_f)} = \frac{\alpha_t R_{t+1}(R_{p,t+h} \cdots R_{p,t+2})W_t}{W_{t+h}}, \tag{4.A.4}$$

and the second derivative with respect to $r_{t+1} - r_f$ is

$$\frac{\partial^2 w_{t+h}}{\partial (r_{t+1} - r_f)^2} = \frac{\alpha_t R_{t+1}(R_{p,t+h} \cdots R_{p,t+2})W_t W_{t+h} - (\alpha_t R_{t+1}(R_{p,t+h} \cdots R_{p,t+2})W_t)^2}{W_{t+h}^2},$$

where we use

$$\frac{\partial R_{p,t+1}}{\partial (r_{t+1} - r_f)} = \frac{\partial^2 R_{p,t+1}}{\partial (r_{t+1} - r_f)^2} = R_f \alpha_t \exp(r_{t+1} - r_f) = \alpha_t R_{t+1}. \tag{4.A.5}$$

19 We do not expand the wealth around $r_{t+i} - r_f$ for $i > 1$ because the associated terms in the Taylor expansion do not depend on α_t.

Evaluating these derivatives in the point $r_{t+1} - r_f = 0$, $R_{p,t+i} = R_f$, and $L_{t+i} = \mathbb{E}_t[L_{t+i}]$ for all i, and defining $\bar{W}_{t+h} = R_f^h W_t + \sum_{j=1}^{h} R_f^{h-j} \mathbb{E}_t[L_{t+j}]$, we find

$$\frac{\partial w_{t+h}}{\partial w_t} = \rho, \tag{4.A.6}$$

$$\frac{\partial w_{t+h}}{\partial l_{t+i}} = \rho_i, \tag{4.A.7}$$

$$\frac{\partial w_{t+h}}{\partial (r_{t+1} - r_f)} = \rho\alpha_t, \tag{4.A.8}$$

$$\frac{\partial^2 w_{t+h}}{\partial (r_{t+1} - r_f)^2} = \rho\alpha_t - (\rho\alpha_t)^2 = \rho\alpha_t(1 - \rho\alpha_t), \tag{4.A.9}$$

with

$$\rho = \frac{R_f^h W_t}{\bar{W}_{t+h}} = \frac{W_t}{W_t + \sum_{j=1}^{h} R_f^{-j} \mathbb{E}_t[L_{t+j}]}, \tag{4.A.10}$$

and

$$\rho_i = \frac{R_f^{h-i} \mathbb{E}_t[L_{t+i}]}{\bar{W}_{t+h}} = \frac{R_f^{-i} \mathbb{E}_t[L_{t+i}]}{W_t + \sum_{j=1}^{h} R_f^{-j} \mathbb{E}_t[L_{t+j}]}. \tag{4.A.11}$$

The Taylor expansion using these derivatives gives the log-linearized wealth at horizon h:

$$w_{t+h} = k(h) + \rho w_t + \rho\alpha_t(r_{t+1} - r_f) + \frac{1}{2}\rho\alpha_t(1 - \rho\alpha_t)\mathrm{Var}(r_{t+1}) + \sum_{i=1}^{h} \rho_i l_{t+i}. \tag{4.A.12}$$

Notice that by definition $\rho + \sum_{i=1}^{h} \rho_i = 1$ so that the log-wealth at $t + h$ can be seen as a weighted average of current log wealth and the present value of the expected stream of labor income up to the horizon h, augmented with the log excess stock return, a convexity effect and a linearization constant $k(h)$.

4.A.2 *Optimal portfolio*

Maximizing (4.3) is now, in a second order approximation, equivalent to maximizing the mean-variance utility function defined over log wealth

$w_{t+h} = \log(W_{t+h})$ and log aggregate labor income $\bar{l}_{t+h} = \log(\bar{L}_{t+h})$,

$$V(w_{t+h} - \psi \bar{l}_{t+h}) = \mathbb{E}_t[w_{t+h} - \psi \bar{l}_{t+h}] + \frac{1}{2}(1-\gamma)\mathrm{Var}_t\left(w_{t+h} - \psi \bar{l}_{t+h}\right). \tag{4.A.13}$$

We only maximize with respect to the short-term portfolio choice variable α_t, since a simple backwards induction argument shows that we can take the portfolio choice rules α_{t+i}, $i = 1, \ldots, h$ as given when considering portfolio choice at time t (see also Campbell, 2018).

Using (4.A.12) and with shorthand notation $\mu = \mathbb{E}_t[r_{t+1}]$ and $\sigma^2 = \mathrm{Var}(r_{t+1})$, the expectation and variance of log wealth minus ψ times labor income growth at horizon h are given by

$$\begin{aligned}\mathbb{E}_t[w_{t+h} - \psi \bar{l}_{t+h}] = k(h) + \rho w_t + \rho \alpha_t(\mu - r_f) + \frac{1}{2}\rho\alpha_t(1-\rho\alpha_t)\sigma^2 \\ + \sum_{i=1}^{h} \rho_i \mathbb{E}_t[l_{t+i}] - \psi \mathbb{E}_t[\bar{l}_{t+h}],\end{aligned} \tag{4.A.14}$$

and

$$\begin{aligned}\mathrm{Var}_t\left(w_{t+h} - \psi \bar{l}_{t+h}\right) = (\rho\alpha_t)^2\sigma^2 + 2\rho\alpha_t \sum_{i=1}^{h} \rho_i \mathrm{Cov}_t(r_{t+1}, l_{t+i}) \\ - 2\rho\alpha_t\psi\mathrm{Cov}_t(r_{t+1}, \bar{l}_{t+h}),\end{aligned} \tag{4.A.15}$$

where we omitted from the expression for the variance all terms that do not depend on α_t. The derivative of the mean-variance utility V defined in (4.A.13) with respect to α_t then is

$$\begin{aligned}\frac{\partial V}{\partial \alpha_t} &= \rho\left(\mu - r_f + \frac{1}{2}\sigma^2\right) - \rho^2\alpha_t\sigma^2 \\ &+ (1-\gamma)\left\{\rho^2\alpha_t\sigma^2 + \rho\sum_{i=1}^{h}\rho_i\mathrm{Cov}_t(r_{t+1}, l_{t+i}) - \rho\psi\mathrm{Cov}_t(r_{t+1}, \bar{l}_{t+h})\right\},\end{aligned}$$

which can be simplified to

$$\begin{aligned}0 = \rho\left(\mu - r_f + \frac{1}{2}\sigma^2\right) - \gamma\rho^2\alpha_t\sigma^2 \\ + (1-\gamma)\rho\left[\sum_{i=1}^{h}\rho_i\mathrm{Cov}_t(r_{t+1}, l_{t+i}) - \psi\mathrm{Cov}_t(r_{t+1}, \bar{l}_{t+h})\right].\end{aligned} \tag{4.A.16}$$

Solving this for α_t gives

$$\alpha_t = \frac{1}{\rho}\frac{\mu - r_f + \frac{1}{2}\sigma^2}{\gamma\sigma^2} - \left(1 - \frac{1}{\gamma}\right)\frac{1}{\rho}\left[\sum_{i=1}^{h}\rho_i \mathrm{Cov}_t(r_{t+1}, l_{t+i}) - \psi \mathrm{Cov}_t(r_{t+1}, \bar{l}_{t+h})\right]/\sigma^2. \tag{4.A.17}$$

Suppose now that labor income in expectation grows at the risk free rate. Then, $R_f^{-i}\mathbb{E}_t[L_{t+i}] = L_t$ and since $\rho_i/\rho = R_f^{-i}\mathbb{E}_t[L_{t+i}]/W_t = L_t/W_t$ we can write

$$\alpha_t = \frac{W_t + hL_t}{W_t}\frac{\mu - r_f + \frac{1}{2}\sigma^2}{\gamma\sigma^2} + \left(1 - \frac{1}{\gamma}\right)\left[\psi\frac{W_t + hL_t}{W_t}\mathrm{Cov}_t(r_{t+1}, \bar{l}_{t+h})\right]/\sigma^2 \tag{4.A.18}$$

$$- \left(1 - \frac{1}{\gamma}\right)\left[\frac{L_t}{W_t}\sum_{i=1}^{h}\mathrm{Cov}_t(r_{t+1}, l_{t+i})\right]/\sigma^2, \tag{4.A.19}$$

where L_t is the current labor income. With multiple assets, we can write the vector of optimal portfolio weights as

$$\alpha_t = \frac{W_t + hL_t}{W_t}\frac{1}{\gamma}\mathrm{Var}(r_{t+1})^{-1}\left(\mu - r_f + \frac{1}{2}\sigma^2\right) + \left(1 - \frac{1}{\gamma}\right)\mathrm{Var}(r_{t+1})^{-1}\left[\psi\frac{W_t + hL_t}{W_t}\mathrm{Cov}_t(r_{t+1}, \bar{l}_{t+h})\right] - \left(1 - \frac{1}{\gamma}\right)\mathrm{Var}(r_{t+1})^{-1}\left[\frac{L_t}{W_t}\sum_{i=1}^{h}\mathrm{Cov}_t(r_{t+1}, l_{t+i})\right], \tag{4.A.20}$$

where σ^2 now denotes the diagonal of $\mathrm{Var}(r_{t+1})$.

4.A.3 *Equilibrium pricing*

We now show the implications of the optimal portfolio choice rule on equilibrium asset prices. Let there be, at any time t, H cohorts with investment horizon $h = 1, \ldots, H$, initial wealth W_h and current labor income L_h. Furthermore, assume that the covariance of the returns with

future individual labor income growth rate is the same as the covariance with future aggregate labor income growth, so that $\mathrm{Cov}_t(r_{t+1}, l_{t+i}) = \mathrm{Cov}_t(r_{t+1}, \bar{l}_{t+i})$ for all i. The optimal portfolio equation is in terms of conditional expectations and variances, given the information at time t. To simplify notation, we now replace all conditional expectations, variances and covariances by their unconditional counterparts, and furthermore define $\gamma_i = \mathrm{Cov}(r_{t+1}, \bar{l}_{t+i} - \bar{l}_t)$, $\mathbb{E}[R] = \mathbb{E}_t[r_{t+1}] + \frac{1}{2}\sigma^2$, and $\Omega = \mathrm{Var}(r_{t+1})$. Then, from Eqn. (4.A.20), the dollar portfolio demand of cohort h is

$$W_h \alpha_h = (W_h + hL_h)\frac{1}{\gamma}\Omega^{-1}(\mathbb{E}[R] - r_f) + \left(1 - \frac{1}{\gamma}\right)\Omega^{-1}\left[\psi(W_h + hL_h)\gamma_h - L_h \sum_{i=1}^{h} \gamma_i\right]. \quad (4.A.21)$$

Adding this up over the cohorts and dividing by aggregate wealth gives the aggregate portfolio weight

$$\alpha_m = \frac{\sum_{h=1}^{H}(W_h + hL_h)}{\sum_{h=1}^{H} W_h}\frac{1}{\gamma}\Omega^{-1}(\mathbb{E}[R] - r_f) + \left(1 - \frac{1}{\gamma}\right)\Omega^{-1}\frac{\sum_{h=1}^{H}\left[\psi(W_h + hL_h)\gamma_h - L_h\sum_{i=1}^{h}\gamma_i\right]}{\sum_{h=1}^{H} W_h}. \quad (4.A.22)$$

Re-writing this with the expected return on the left hand side gives

$$\mathbb{E}[R] - r_f = \gamma\frac{\sum_{h=1}^{H} W_h}{\sum_{h=1}^{H}(W_h + hL_h)}\Omega\alpha_m + (\gamma - 1)\frac{\sum_{h=1}^{H}\left[L_h\sum_{i=1}^{h}\gamma_i - \psi(W_h + hL_h)\gamma_h\right]}{\sum_{h=1}^{H}(W_h + hL_h)}. \quad (4.A.23)$$

Interchanging the summations over h and i in the second term, we find

$$\mathbb{E}[R] - r_f = \gamma\left[\frac{\sum_{h=1}^{H} W_h}{\sum_{i=h}^{H}(W_h + hL_h)}\right]\Omega\alpha_m + (\gamma - 1)\sum_{h=1}^{H}\left[\frac{\sum_{i=h}^{H} L_i - \psi(W_h + hL_h)}{\sum_{i=1}^{H}(W_i + iL_i)}\right]\gamma_h. \quad (4.A.24)$$

Defining the log market portfolio return as $r_{m,t+1} = \alpha_m' r_{t+1}$ and recalling the definitions $\gamma_h = \mathrm{Cov}(r_{t+1}, \bar{l}_{t+h} - \bar{l}_t)$ and $\Omega = \mathrm{Var}(r_{t+1})$ produces Eqn. (4.5) in the main text.

4.B CROSS-SECTIONAL STANDARD ERRORS

In this Appendix, we derive a suitable Shanken-type correction for the second-stage cross-sectional standard errors that accounts for the fact that the betas are estimated in the first-stage time-series regressions and that are robust to heteroskedasticity and autocorrelation in the factors. The exposition follows Cochrane (2005), Chapter 12. We derive the corrected standard errors both for a setting where there is one multivariate first-stage regression, as well as for a setting with multiple univariate first-stage regressions.

4.B.1 *Multivariate betas*

The time-series regressions in the first stage are of the form

$$R_t^e = \alpha + \beta f_t + \epsilon_t, \qquad t = 1, \ldots, T, \tag{4.B.25}$$

where R_t^e is an $N \times 1$ vector of portfolio excess returns, α is an $N \times 1$ vector of intercepts, β is an $N \times K$ coefficient-matrix, f_t a $K \times 1$ vector of factors, which in our case includes the equity market risk factor and the labor income risk factors at scales $j = 1, \ldots, J, > J$, and ϵ_t is an $N \times 1$ vector of zero-mean errors that are allowed to be heteroskedastic and correlated over time. The OLS estimator of the first-stage betas is given by

$$\hat{\beta} = T^{-1} \sum_t \left(R_t^e - \bar{R}^e\right) f_t' \left(\sum_t f_t f_t'\right), \tag{4.B.26}$$

with $\bar{R}^e = T^{-1} \sum_t R_t^e$ an $N \times 1$ vector of average excess returns of the portfolios, and where we assume without loss of generality that $\mathbb{E}[f] = 0$, which can easily be achieved by demeaning the factors.

The second-stage regression is given by (ignoring for a moment an intercept term)

$$\bar{R}^e = \hat{\beta}\lambda + \eta, \tag{4.B.27}$$

with λ an $K \times 1$ vector of market prices of risk, and η an $N \times 1$ vector of zero-mean errors. The OLS estimator of the second-stage regression is given by

$$\hat{\lambda} = (\hat{\beta}'\hat{\beta})^{-1}\hat{\beta}'\bar{R}^e, \tag{4.B.28}$$

and the difference with the true value is

$$\hat{\lambda} - \lambda = (\hat{\beta}'\hat{\beta})^{-1}\hat{\beta}'\eta. \tag{4.B.29}$$

Notice that η can be written as

$$\eta = \bar{R}^e - \beta\lambda - (\hat{\beta} - \beta)\lambda = \bar{R}^e - \mathbb{E}[\bar{R}^e] - (\hat{\beta} - \beta)\lambda. \tag{4.B.30}$$

Hence, the variance of the second-stage regression is

$$\text{Var}(\hat{\lambda}) = (\hat{\beta}'\hat{\beta})^{-1}\hat{\beta}'\text{Var}(\eta)\hat{\beta}(\hat{\beta}'\hat{\beta})^{-1}, \tag{4.B.31}$$

with

$$\text{Var}(\eta) = \text{Var}(\bar{R}^e) + \text{Var}\left((\hat{\beta} - \beta)\lambda\right). \tag{4.B.32}$$

To calculate these variances, we first write

$$\bar{R}^e - \mathbb{E}[\bar{R}^e] = T^{-1}\sum_t u_t, \quad u_t = R_t^e - \mathbb{E}[R^e], \tag{4.B.33}$$

and, hence, $\text{Var}(\bar{R}^e) = T^{-1}\Sigma_u$ with $\Sigma_u = \text{Var}(u_t)$.

The difference between the first-stage estimator and its true value can be expressed as

$$\hat{\beta} - \beta = T^{-1}\sum_t \epsilon_t f_t'\Sigma_f^{-1}, \tag{4.B.34}$$

with $\Sigma_f = T^{-1}\left(\sum_t f_t f_t'\right)$. Hence, the variance of $(\hat{\beta} - \beta)\lambda$ is

$$\text{Var}\left((\hat{\beta} - \beta)\lambda\right) = T^{-2}\text{Var}\left(\sum_t \epsilon_t f_t'\Sigma_f^{-1}\lambda\right). \tag{4.B.35}$$

Using these expressions, the asymptotic variance of $\hat{\lambda}$ is given by

$$\text{Var}(\hat{\lambda}) = T^{-1}(\hat{\beta}'\hat{\beta})^{-1}\hat{\beta}'\left(\Sigma_u + T^{-1}\text{Var}\left(\sum_t \epsilon_t f_t'\Sigma_f^{-1}\lambda\right)\right)\hat{\beta})\hat{\beta}('\hat{\beta})^{-1},$$

(4.B.36)

or

$$\text{Var}(\hat{\lambda}) = T^{-1}(\hat{\beta}'\hat{\beta})^{-1}\hat{\beta}'\left(\Sigma_u + S(\epsilon_t f_t'\Sigma_f^{-1}\lambda\right)\hat{\beta}(\hat{\beta}'\hat{\beta})^{-1}, \quad (4.B.37)$$

with $S(x)$ the long-run covariance matrix of $x_t = \epsilon_t f_t'\Sigma_f^{-1}\lambda$. The matrix Σ_u can be consistently estimated by the sample variance-covariance matrix of the excess returns, and $S(x_t)$ by a Newey-West procedure. Note that these corrected standard errors explicitly account for the errors-in-variables bias and are robust to heteroskedasticity and autocorrelation.

Adding a constant term to the second-stage regression is straightforward. The second-stage regression is then

$$\bar{R}^e = \lambda_0 + \hat{\beta}\lambda + \eta, \quad (4.B.38)$$

which can be succinctly written as

$$\bar{R}^e = Z\gamma + \eta, \quad (4.B.39)$$

with $Z = [\iota\hat{\beta}]$ and $\gamma = (\lambda_0, \lambda)'$. The OLS estimator of the second-stage regression is

$$\hat{\gamma} = (Z'Z)^{-1}Z'\bar{R}^e, \quad (4.B.40)$$

with variance

$$\text{Var}(\hat{\gamma}) = (Z'Z)^{-1}Z'\text{Var}(\eta)Z(Z'Z)^{-1}, \quad (4.B.41)$$

where $\text{Var}(\eta)$ is defined as before.

4.B.2 *Univariate betas*

In the case where we consider univariate betas, we have multiple first-stage regressions of the form:

$$R_t^e = \alpha_j + \beta^{(j)} f_t^{(j)} + \epsilon_t^{(j)}, \qquad \text{for } j = 1, \ldots, J, > J, \qquad t = 1, \ldots, T, \quad (4.B.42)$$

where R_t^e is again the $N \times 1$ vector of portfolio excess returns components, and where $f_t^{(j)}$ is a scalar denoting the labor income growth rate component at scale j.[20] The OLS estimator of the first-stage beta is

$$\hat{\beta}^{(j)} = T^{-1}\sum_t \left(R_t^e - \bar{R}^e\right) f_t^{(j)} / \sigma^2_{f^{(j)}}, \quad (4.B.43)$$

20 The additional regression for the equity market beta can be included in a straightforward manner and for brevity of notation we omit this extra regression in this appendix.

where $\sigma^2_{f^{(j)}} = T^{-1}\sum_t \left(f_t^{(j)}\right)^2$, and where we assume without loss of generality that $\mathbb{E}[f^{(j)}] = 0$, which can easily be achieved by demeaning the factors. We can compactly write the first-stage regressions and the beta estimator as

$$R_t = \alpha + \beta F_t + \epsilon_t, \tag{4.B.44}$$

and

$$\hat{\beta} = T^{-1}\sum_t (R_t - \bar{R})F_t\Sigma_F^{-1}, \tag{4.B.45}$$

with R_t an $N \times K$ matrix ($K = J + 1$) with columns R_t^e, β an $N \times K$ matrix with columns $\beta^{(j)}$, ϵ_t an $N \times K$ matrix with columns $\epsilon_t^{(j)}$, F_t a $J \times J$ diagonal matrix with elements $f_t^{(j)}$, and Σ_F a $J \times J$ diagonal matrix with elements $\sigma^2_{f^{(j)}}$. Hence, the difference between the first-stage estimator and its true value can be written as

$$\hat{\beta} - \beta = T^{-1}\sum_t \epsilon_t F_t \Sigma_F^{-1}. \tag{4.B.46}$$

The second-stage regression is the same as before, $\bar{R}^e = \hat{\beta}\lambda + \eta$. Following the same reasoning as in the multivariate case, the variance of the second-stage estimator for λ is

$$\mathrm{Var}(\hat{\lambda}) = T^{-1}(\hat{\beta}'\hat{\beta})^{-1}\hat{\beta}'\left(\Sigma_u + T^{-1}\mathrm{Var}\left(\sum_t \epsilon_t F_t \Sigma_F^{-1}\lambda\right)\right)\hat{\beta})\hat{\beta}('\hat{\beta})^{-1}, \tag{4.B.47}$$

or

$$\mathrm{Var}(\hat{\lambda}) = T^{-1}(\hat{\beta}'\hat{\beta})^{-1}\hat{\beta}'\left(\Sigma_u + S(\epsilon_t F_t \Sigma_F^{-1}\lambda\right)\hat{\beta}(\hat{\beta}'\hat{\beta})^{-1}, \tag{4.B.48}$$

with $S(x)$ the long-run covariance matrix of $x_t = \epsilon_t F_t \Sigma_F^{-1}\lambda$. The matrix Σ_u can be consistently estimated by the sample variance-covariance matrix of the excess returns, and $S(x_t)$ by a Newey-West procedure.

BIBLIOGRAPHY

Abel, A. B. (1990). Asset prices under habit formation and catching up with the Joneses. *American Economic Review*, 80:38–42.

Acharya, V. V. and Pedersen, L. H. (2005). Asset pricing with liquidity risk. *Journal of Financial Economics*, 77:375–410.

Aït-Sahalia, Y., Cacho-Diaz, J., and Laeven, R. J. A. (2015). Modeling financial contagion using mutually exciting jump processes. *Journal of Financial Economics*, 117:585–606.

Aït-Sahalia, Y., Laeven, R. J. A., and Pelizzon, L. (2014). Mutual excitation in Eurozone sovereign CDS. *Journal of Econometrics*, 183:151–167.

Amihud, Y. and Mendelson, H. (1986). Asset pricing and the bid-ask spread. *Journal of Financial Economics*, 17:223–249.

Ang, A. and Longstaff, F. A. (2013). Systemic sovereign credit risk: Lessons from the U.S. and Europe. *Journal of Monetary Economics*, 60:493–510.

Augustin, P. (2018). The term structure of cds spreads and sovereign credit risk. *Journal of Monetary Economics*, 96:53–76.

Bandi, F. and Tamoni, A. (2017). The horizon of systemic risk: A new beta representation. Unpublished working paper.

Bansal, R. and Yaron, A. (2004). Long run risks: A potential resolution to asset pricing puzzles. *Journal of Finance*, 59:1481–1509.

Bao, J., Pan, J., and Wang, J. (2011). The illiquidity of corporate bonds. *Journal of Finance*, 66:911–946.

Baxter, M. and Jermann, U. (1997). The international diversification puzzle is worse than you think. *American Economic Review*, 87:170–180.

Belo, F., Lin, X., and Bazdresch, S. (2014). Labor hiring, investment, and stock return predictability. *Journal of Political Economy*, 122:129–177.

Benzoni, L., Collin-Dufresne, P., and Goldstein, R. S. (2007). Portfolio choice over the life-cycle when stock and labor markets are cointegrated. *Journal of Finance*, 62:2123–2167.

Berk, J. B. and Walden, J. (2013). Limited capital market participation and human capital risk. *Review of Asset Pricing Studies*, 3:1553–1607.

Berndt, A., Douglas, R., Duffie, D., and Ferguson, M. (2018). Corporate credit risk premia. *Review of Finance*, 22:419–454.

Berndt, A., Douglas, R., Duffie, D., Ferguson, M., and Schranz, D. (2008). Measuring default risk premia from default swap rates and edfs. Unpublished working paper.

Bilsen, S. v., Bovenberg, A. L., and Laeven, R. J. A. (2019). Consumption and portfolio choice under internal multiplicative habit formation. *Journal of Financial and Quantitative Analysis*, Forthcoming.

Black, F. and Cox, K. (1976). Valuing corporate securities: Some effects of bond indenture provisions. *Journal of Finance*, 31:351–367.

Bongaerts, D., De Jong, F., and Driessen, J. (2017). An asset pricing approach to liquidity effects in corporate bond markets. *Review of Financial Studies*, 30:1229–1269.

Boons, M. and Tamoni, A. (2016). Horizon-specific macroeconomic risks and the cross-section of expected returns. Unpublished working paper.

Boswijk, H. P., Laeven, R. J. A., and Lalu, A. (2016). Asset returns with self-exciting jumps: Option pricing and estimation with a continuum of moments. Unpublished working paper. University of Amsterdam, Amsterdam.

Buhler, W. and Trapp, M. (2010). Time-varying credit risk and liquidity premia in bond and CDS markets. Unpublished working paper. University of Mannheim, Mannheim.

Campbell, J. and Taksler, G. (2003). Equity volatility and corporate bond yields. *Journal of Finance*, 58:2321–2349.

Campbell, J. Y. (1996). Understanding risk and returns. *Journal of Political Economy*, 101:298–345.

Campbell, J. Y. (2018). *Financial Decisions and Markets: A Course in Asset Pricing*. Princeton University Press, Princeton.

Campbell, J. Y. and Viceira, L. M. (2002). *Strategic Asset Allocation: Portfolio Choice for Long-Term Investors*. Oxford University Press, New York.

Chen, H., Cui, R., He, Z., and Milbradt, K. (2017). Quantifying liquidity and default risks of corporate bonds over the business cycle. *Review of Financial Studies*, 31:852–897.

Chen, L., Lesmond, D., and Wei, J. (2007). Corporate yield spreads and bond liquidity. *Journal of Finance*, 62:119 – 149.

Chib, S. and Jeliazkov, I. (2001). Marginal likelihood from the Metropolis-Hastings output. *Journal of the American Statistical Association*, 96:270–281.

Cocco, J. F., Gomes, F. J., and Maenhout, P. (2005). Consumption and portfolio choice over the life cycle. *Review of Financial Studies*, 18:491–533.

Cochrane, J. H. (2005). *Asset Pricing*. Princeton University Press, Princeton.

Collin-Dufresne, P., Goldstein, R. S., and J., H. (2010). Is credit event risk priced? modeling contagion via the updating of beliefs. NBER Working Paper.

Cremers, M., Driessen, J., and Maenhout, P. (2008). Explaining the level of credit spreads: Option-implied jump risk premia in a firm value model. *Review of Financial Studies*, 21:2209–2242.

Daniel, K. and Marshall, D. (1997). Equity-premium and risk-free rate puzzles at long horizons. *Macroeconomic Dynamics*, 1:452–484.

Danthine, J. P. and Donaldson, J. (2002). Labor relations and asset returns. *Review of Economic Studies*, 69:41–64.

De Jong, F. and Driessen, J. (2012). Liquidity risk premia in corporate bond markets. *Quarterly Journal of Finance*, 2.

De Jong, F. and Schotman, P. C. (2010). Price discovery in fragmented markets. *Journal of Financial Econometrics*, 8:1–28.

Dew-Becker, I. and Giglio, S. (2016). Asset pricing in the frequency domain: Theory and empirics. *Review of Financial Studies*, 29:2029–2068.

Dick-Nielsen, J. (2014). How to clean enhanced trace data. Supplement to "Liquidity biases in TRACE" (2009).

Dick-Nielsen, J., Feldhütter, P., and Lando, D. (2012). Corporate bond liquidity before and after the onset of the subprime crisis. *Journal of Financial Economics*, 103:471–492.

Dittmar, R. F., Palomino, F., and Yang, W. (2016). Leisure preferences, long-run risks, and human capital returns. *Review of Asset Pricing Studies*, 6:88–134.

Donangelo, A. (2014). Labor mobility: Implications for asset pricing. *Journal of Finance*, 69:1321–1346.

Donangelo, A., Eiling, E., and Palacios, M. (2010). Aggregate asset-pricing implications of human capital mobility in general equilibrium. Unpublished working paper. University of Texas, Texas, University of Amsterdam, Amsterdam, and Vanderbilt University, Nashville.

Donangelo, A., Gourio, F., Kehrig, M., and Palacios, M. (2018). The cross-section of labor leverage and equity returns. *Journal of Financial Economics*, Forthcoming.

Driessen, J. (2005). Is default event risk priced in corporate bonds? *Review of Financial Studies*, 18:165–195.

Duffee, G. R. (1999). Estimating the price of default risk. *Review of Financial Studies*, 12:197–226.

Duffie, D., Pan, J., and Singleton, K. J. (2000). Transform analysis and asset pricing for affine jump-diffusions. *Econometrica*, 68:1343–1376.

Duffie, D., Pedersen, L. H., and Singleton, K. J. (2003). Modeling sovereign yield spreads: A case study of Russian debt. *Journal of Finance*, 58:119–160.

Duffie, D. and Singleton, K. J. (1999). Modeling term structures of defaultable bonds. *Review of Financial Studies*, 12:687–720.

Duffie, D. and Singleton, K. J. (2003). *Credit Risk: Pricing, Measurement and Management*. Princeton University Press, Princeton.

Edward, A. K., Harris, L. E., and Piwowar, M. S. (2007). Corporate bond market transaction costs and transparency. *Journal of Finance*, 62:1421–1451.

Eiling, E. (2013). Industry-specific human capital, idiosyncratic risk, and the cross-section of expected stock returns. *Journal of Finance*, 68:43–84.

Eiling, E., De Jong, F., Laeven, R. J. A., and Sperna Weiland, R. C. (2019). Labor income risk and stock returns: The role of horizon effects. Unpublished working paper. University of Amsterdam, Amsterdam, and Tilburg University, Tilburg.

Elton, E. J., Gruber, M. J., Agrawal, D., and Mann, C. (2001). Explaining the rate spread on corporate bonds. *Journal of Finance*, 56:247–277.

Eraker, B. (2004). Do stock prices and volatility jump? Reconciling evidence from spot and option prices. *Journal of Finance*, 59:1367–1403.

Eraker, B., Johannes, M., and Polson, N. (2003). The impact of jumps in volatility and returns. *Journal of Finance*, 58:1269–1300.

Ericsson, J. and Renault, O. (2006). Liquidity and credit risk. *Journal of Finance*, 61:2219–2250.

Errais, E., Giesecke, K., and Goldberg, L. R. (2010). Affine point processes and portfolio credit risk. *SIAM Journal of Financial Mathematics*, 1:642–665.

Fama, E. and French, K. (1993). Common risk factors in the returns on stocks and bonds. *Journal of Financial Economics*, 33:3–56.

Fama, E. and French, K. (2015). A five-factor asset pricing model. *Journal of Financial Economics*, 116:1–22.

Fama, E. and MacBeth, J. (1973). Risk, return, and equilibrium: Empirical tests. *Journal of Political Economy*, 81:607–636.

Fama, E. and Schwert, G. (1977). Human capital and capital market equilibrium. *Journal of Financial Economics*, 4:115–146.

Favilukis, J. and Lin, X. (2016a). Does wage rigidity make firms riskier? Evidence from long-horizon return predictability. *Journal of Monetary Economics*, 78:80–95.

Favilukis, J. and Lin, X. (2016b). Wage rigidity: A quantitative solution to several asset pricing puzzles. *Review of Financial Studies*, 29:148–192.

Feller, W. (1951). Two singular diffusion problems. *Annals of Mathematics*, 54:173–182.

Friewald, N., Jankowitsch, R., and Subrahmanyam, M. G. (2012). Illiquidity or credit deterioration: A study of liquidity in the US corporate bond markets during financial crises. *Journal of Financial Economics*, 105:18–36.

Glosten, L. and Milgrom, J. (1985). Bid, ask, and transaction prices in a specialist market with heterogeneously informed traders. *Journal of Financial Economics*, 14:71–100.

Gomez, J.-P., Priestley, R., and Zapatero, F. (2009). Implications of keeping up with the Joneses behavior for the equilibrium cross-section of stock returns: International evidence. *Journal of Finance*, 64:2703–2737.

Gomez, J.-P., Priestley, R., and Zapatero, F. (2016). Labor income, relative wealth concerns, and the cross-section of stock returns. *Journal of Financial and Quantitative Analysis*, 51:1111–1133.

Hall, R. E. (2005). Employment fluctuations with equilibrium wage stickiness. *American Economic Review*, 95:50–65.

Hawkes, A. (1971a). Spectra of some self-exciting and mutually exciting point processes. *Biometrika*, 58:83–90.

Hawkes, A. G. (1971b). Point spectra of some mututally exciting point processes. *Journal of the Royal Statistical Society*, B 33:438 – 443.

Hawkes, A. G. (1971c). Spectra of some self-exciting and mutually exciting point processes. *Biometrika*, 58:83 – 90.

Hawkes, A. G. and Oakes, D. (1974). A cluster representation of a self-exciting process. *Journal of Applied Probability*, 11:493 – 503.

He, Z. and Milbradt, K. (2014). Endogenous liquidity and defaultable bonds. *Econometrica*, 82:1443–1508.

He, Z. and Xiong, W. (2012). Rollover risk and credit risk. *Journal of Finance*, 67:391–430.

Ho, T. and Stoll, H. (1981). Optimal dealer pricing under transactions and return uncertainty. *Journal of Financial Economics*, 9:47–73.

Houweling, P., Mentink, A., and Vorst, T. (2005). Comparing possible proxies of corporate bond liquidity. *Journal of Banking and Finance*, 29:1331–1358.

Huang, J.-Z. and Huang, M. (2012). How much of the corporate-yield spread is due to credit risk? *Review of Asset Pricing Studies*, 2:153–202.

Jagannathan, R. and Wang, Z. (1996). The conditional CAPM and the cross-section of expected returns. *Journal of Finance*, 51:3–53.

Janosi, T., Jarrow, R. A., and Yildirim, Y. (2002). Estimating expected losses and liquidity discounts implicit in debt prices. *Journal of Risk*, 5:1–38.

Jarrow, R. A., Lando, D., and Yu, F. (2005). Default risk and diversification: Theory and empirical implications. *Mathematical Finance*, 15:1–26.

Jarrow, R. A. and Turnbull, S. M. (1995). Pricing derivatives on financial securities subject to credit risk. *Journal of Finance*, 50:53–85.

Johannes, M., Kumar, R., and Polson, N. (1999). State dependent jump models: How do US equity indices jump? Unpublished working paper. University of Chicago, Chicago.

Johannes, M. and Polson, N. (2009). MCMC methods for financial econometrics. In Aït-Sahalia, Y. and Hansen, L., editors, *Handbook of Financial Econometrics*, pages 1–72. Elsevier.

Julliard, C. (2007). Labor income risk and asset returns. Unpublished working paper.

Kamara, A., Korajczyk, R., Lou, X., and Sadka, R. (2016). Horizon pricing. *Journal of Financial and Quantitative Analysis*, 51:1769–1793.

Katz, L. and Summers, L. (1989). Industry rents: Evidence and implications. *Brookings Papers on Economic Activity. Microeconomics*, page 209–290.

Koijen, R., Lustig, H., and Van Nieuwerburgh, S. (2017). The cross-section and time series of stock and bond returns. *Journal of Monetary Economics*, 88:50–69.

Kuehn, L. A., Simutin, M., and Wang, J. J. (2017). A labor capital asset pricing model. *Journal of Finance*, 72:2131–2178.

Lando, D. (1998). On cox processes and credit risky securities. *Review of Derivatives Research*, 2:99–120.

Leland, H. (1994). Corporate debt value, bond convenants, and optimal capital structure. *Journal of Finance*, 49:1213–1252.

Leland, H. E. and Toft, K. B. (1996). Optimal capital structure, endogenous bankruptcy, and the term structure of credit spreads. *Journal of Finance*, 51:987–1019.

Lin, H., Wang, J., and Wu, C. (2011). Liquidity risk and expected corporate bond returns. *Journal of Financial Economics*, 99:628–650.

Liu, J., Longstaff, F. A., and Mandell, R. (2006). The market price of risk in interest rate swaps: The roles of default and liquidity risks. *Journal of Business*, 79:2337–2360.

Longstaff, F. A., Mithal, S., and Neis, E. (2005). Corporate yield spreads: Default risk or liquidity? New evidence from the credit default swap market. *Journal of Finance*, 60:2213–2253.

Longstaff, F. A., Pan, J., Pedersen, L., and Singleton, K. J. (2011). How sovereign is sovereign credit risk? *American Economic Journal: Macroeconomics*, 3:75–103.

Longstaff, F. A. and Schwartz, E. S. (1995). A simple approach to valuing risky fixed and floating rate debt. *Journal of Finance*, 50:789–819.

Lustig, H., Van Nieuwerburgh, S., and Verdelhan, A. (2013). The wealth-consumption ratio. *The Review of Asset Pricing Studies*, 3:38–94.

Maio, P. and Min, B.-K. (2018). Wage growth and equity risk premia. Unpublished working paper.

Malloy, C. J., Moskowitz, T. J., and Vissing-Jørgensen, A. (2009). Long-run stockholder consumption risk and asset returns. *Journal of Finance*, 64:2427–2479.

Marfe, R. (2018). Labor rigidity and the dynamics of the value premium. Unpublished working paper.

Mayers, D. (1972). Nonmarketable assets and capital market equilibrium under uncertainty. In Jensen, M. C., editor, *Studies in the Theory of Capital Markets*, pages 223–248. Praeger, New York, NY.

Merton, R. C. (1974). On the pricing of corporate debt: The risk structure of interest rates. *Journal of Finance*, 29:449–470.

Monfort, A., Pegoraro, F., Renne, J.-P., and Roussellet, G. (2018). Affine modeling of credit risk, pricing of credit events and contagion. Unpublished working paper.

Monfort, A. and Renne, J.-P. (2014). Decomposing Euro-area sovereign spreads: Credit and liquidity risks. *Review of Finance*, 18:2103–2181.

Nagler, F. (2019). Yield spreads and the corporate bond rollover channel. *Review of Finance*, Forthcoming.

Nashikkar, A., Subrahmanyam, M., and Mahanti, S. (2011). Liquidity and arbitrage in the market for credit risk. *Journal of Financial and Quantitative Analysis*, 46:627–656.

Neal, D. (1995). Industry-specific human capital: Evidence from displaced workers. *Journal of Labor Economics*, 13:653–677.

Ogata, Y. and Akaike, H. (1982). On linear intensity models for mixed doubly stochastic Poisson and self-exciting point processes. *Journal of the Royal Statistical Society*, 44:102 – 107.

Ortu, F., Tamoni, A., and Tebaldi, C. (2013). Long-run risk and the persistence of consumption shocks. *Review of Financial Studies*, 26:2876–2915.

Ortu, F., Tamoni, A., and Tebaldi, C. (2015). A persistence-based Wold-type decomposition for stationary time series. Unpublished working paper.

Palacios, M. (2015). Human capital as an asset class: Implications from a general equilibrium model. *Review of Financial Studies*, 28:978–1023.

Palacios-Huerta, I. (2003). The robustness of the conditional CAPM with human capital. *Journal of Financial Econometrics*, 1:272–289.

Pan, J. and Singleton, K. J. (2008). Default and recovery implicit in the term structure of sovereign cds spreads. *Journal of Finance*, 63:2345–2384.

Parker, J. and Julliard, C. (2005). Consumption risk and the cross section of expected returns. *Journal of Political Economy*, 113:185–222.

Remolona, E., Scatigna, M., and Wu, E. (2008). The dynamic pricing of sovereign risk in emerging markets: Fundamentals and risk aversion. *Journal of Fixed Income*, 17:57–71.

Reynaud-Bouret, P. and Schbath, S. (2010). Adaptive estimation for Hawkes processes: Application to genome analysis. *Annals of Statistics*, 38:2781 – 2822.

Rich, R. and Tracy, J. (2004). Uncertainty and labor contract durations. *Review of Economics and Statistics*, 86:270–287.

Rigobon, R. and Fernandez-Arias, E. (2000). Financial contagion in emerging markets. In *Wanted: World Financial Stability*, pages 33–47. Inter-American Development Bank.

Saita, L. (2005). The puzzling price of corporate default risk. Unpublished working paper. Stanford University, Stanford.

Santos, T. and Veronesi, P. (2006). Labor income and predictable stock returns. *Review of Financial Studies*, 19:1–44.

Shimer, R. (2005). The cyclical behavior of equilibrium unemployment and vacancies. *American Economic Review*, 95:25–49.

Sperna Weiland, R., Laeven, R. J. A., and De Jong, F. (2019). Feedback between credit and liquidity risk in the US corporate bond market. Unpublished working paper. University of Amsterdam, Amsterdam, and Tilburg University, Tilburg.

Sperna Weiland, R. C. (2018). Credit risk premia in sovereign credit default swaps. Unpublished working paper. University of Amsterdam, Amsterdam.

Storesletten, K., Telmer, C., and Yaron, A. (2004). Cyclical dynamics of idiosyncratic labor market risk. *Journal of Political Economy*, 112:695–717.

Yu, F. (2002). Modeling expected return on defaultable bonds. *Journal of Fixed Income*, 12:69–81.

Zinna, G. (2013). Sovereign default risk premia: Evidence from the default swap market. *Journal of Empirical Finance*, 21:15–35.

SUMMARY

ESSAYS ON MACRO-FINANCIAL RISKS

This dissertation consists of several studies on macro-financial risks. In particular, Part I, consisting of Chapters 2 and 3, discusses issues related to credit risk. Part II, consisting of Chapter 4, focuses on the relationship between labor income risk and stock returns.

Chapter 2

In Chapter 2, we investigate the interplay between credit and liquidity risk in the US corporate bond market. We develop a novel reduced-form framework that allows us to explicitly quantify the interactions between credit and liquidity risk and their impact on bond prices, yield spreads, and investment risk. In particular, we propose to use mutually exciting processes to construct a dynamic feedback mechanism between the two risk types. The cross-excitation between credit risk build-up and liquidity dry-up allows the model to accommodate a positive credit-liquidity feedback loop in which credit and liquidity shocks tend to cluster in a potentially asymmetric fashion. We develop a Bayesian estimation procedure and use US bond transaction data to estimate the model.

We find strong evidence for asymmetric feedback between credit and liquidity risk that is more pronounced during the most turbulent times and for bonds with lower credit ratings. For example, our yield-spread decomposition reveals that the impact of liquidity shocks on credit risk is mostly negligible. The effects of credit shocks on liquidity, on the other hand, are much larger and economically important. The credit-induced liquidity component contributes for 0.50 (AAA/AA) up to 0.73 (B and lower) percentage points to average 2007-2009 yield spreads and for up to 2.05 (B and lower) percentage points in the most distressed period. In a case study on Ford Motor Company, we find that the credit-induced liquidity component accounts for over 60% in relative terms of yield spreads during the peak of the crisis.

Chapter 3

In Chapter 3, I consider the pricing of risk premia related to sovereign credit risk. In particular, I consider 'distress risk premia', defined as the compensation that investors demand for being exposed to unexpected variations in credit risk, as well as risk premia related to (unpredictable) default events themselves. In order to estimate these risk premia, I develop a new model for the term-structure of sovereign credit risk in which sovereign defaults can be triggered by shocks in either a common or country-specific factor. By modeling both these factors as self-exciting processes, the model is able to capture several features observed in the data, namely the high degree of commonality in sovereign credit risk, the clustering of credit shocks over time and across countries, and jump-like increases in sovereign CDS spreads. In the empirical analysis, I use CDS term structure data on 28 geographically dispersed countries with credit ratings ranging from A to B, and historical default probabilities per rating class as reported by S&P.

The estimated model allows for a decomposition of CDS spreads along two dimensions. First, CDS spreads can be decomposed into country-specific and systemic risk components. Second, CDS spreads can be decomposed into expected default risk, distress risk premia, and default event risk premia components. In relative terms, I find that for all rating classes on average approximately 65% of five-year CDS spreads can be attributed to country-specific risk and 35% to systemic risk. The decomposition into risk premia components shows that the default event risk premium component is substantial and seems to matter (relatively) most for short-term CDS spreads of countries with lower credit ratings. Distress risk premia, on the other hand, are priced more heavily in long-term CDS spread of countries with higher credit ratings. Combining the two decomposition dimensions reveals that differences in CDS decompositions across rating classes are mainly caused by differences in sovereign-specific risk rather than differences in exposure to systemic risk.

Chapter 4

In Chapter 4, we study the asset pricing implications of labor income risk, thereby focusing in particular on possible horizon effects that might

play a role. To test for horizon effects, we use a flexible empirical framework that allows us to include labor income risk at multiple horizons simultaneously. We find a clear distinctive role for the two- to four-year horizon. Not only does labor income risk at this medium horizon carry a highly significant price of risk, whereas at other horizons it does not, also the ability to explain the cross-sectional differences of stock returns peaks at this horizon. For example, a simple two-factor model that includes the contemporaneous equity market return factor and labor income risk factor at the medium term horizon can explain a striking 71% of the cross-sectional variation in 25 size book-to-market and 25 size-investment portfolios. By contrast, the standard human capital CAPM with quarterly labor income growth can only explain 7% of the cross-sectional variation in these portfolios. This means that simply changing the horizon over which labor income risk is measured has a dramatic impact on the model performance. Furthermore, we document similar horizon effects in optimal portfolio allocation, where labor income risk generates significant adjustments to the composition of the optimal risky equity portfolio at the medium term horizon. Our results are consistent with wage stickiness, where wages are reset every two to four years.

SAMENVATTING (SUMMARY IN DUTCH)

ESSAYS OVER MACRO-FINANCIËLE RISICO'S

Dit proefschrift bestaat uit verschillende studies over macro-financiële risico's. In het bijzonder behandelt deel I, bestaande uit Hoofdstukken 2 en 3, verschillende onderwerpen met betrekking tot kredietrisico. Deel II, bestaande uit Hoofdstuk 4, focust zich op de relatie tussen arbeidsinkomensrisico en de rendementen op aandelen.

Hoofdstuk 2

In Hoofdstuk 2 onderzoeken wij het samenspel tussen krediet- en liquiditeitsrisico's in de bedrijfsobligatiemarkt in de Verenigde Staten. We ontwikkelen een nieuw raamwerk dat ons in staat stelt om de interacties tussen krediet- en liquiditeitsrisico's te kwantificeren en na te gaan wat de impact van die interacties is op obligatieprijzen, rendementsverschillen en investeringsrisico's. In het bijzonder stellen wij voor om zogenaamde 'elkaar-versterkende processen' te gebruiken om zo een dynamisch terugkoppelingsmechanisme tussen de twee risicotypes te construeren. Het elkaar-versterkende effect tussen de opbouw van kredietrisico enerzijds en het opdrogen van marktliquiditeit anderzijds stelt het model in staat om een positieve krediet-liquiditeitsterugkoppeling te faciliteren waarin krediet- en liquiditeitschokken de neiging hebben om op een mogelijk asymmetrische manier te clusteren. We ontwikkelen een Bayesiaanse schattingsprocedure en gebruiken transactiedata van bedrijfsobligaties in de VS om het model te schatten.

We vinden sterk bewijs voor asymmetrische terugkoppelingeffecten tussen krediet- en liquiditeitsrisico die typisch sterker zijn tijdens de meest turbulente periodes en voor obligaties met lagere kredietscores. Ter illustratie, onze decompositie van rendementsverschillen onhult dat de impact van liquiditeitsschokken op kredietrisico over het algmeen te verwaarlozen valt. De effecten van kredietschokken op liquiditeit zijn echter veel groter en economisch significant. Het zogenaamde kredietveroorzaakte liquiditeitsdeel voegt voor 0.50 (AAA/AA) tot 0.73 (B en

lager) percentagepunten toe aan de gemiddelde rendementsverschillen tijdens de periode 2007-2009, maar dit loopt op tot 2.05 (B en lager) percentagepunten in de meest stressvolle periode. In een casus over Ford Motor Company vinden we, in relatieve termen, dat tijdens het hoogtepunt van de financiële crisis het krediet-veroorzaakte liquiditeitsdeel meer dan 60% van de rendementsverschillen voor zijn rekening neemt.

Hoofdstuk 3

In Hoofdstuk 3 beschouw ik het inprijzen van risicopremies gerelateerd aan het kredietrisico van landen. In het bijzonder beschouw ik zowel 'distress-risicopremies', gedefinieerd als de compensatie die investeerders vragen omdat ze blootgesteld zijn aan onverwachte variaties in kredietrisico, als risicopremies gerelateerd aan de daadwerkelijke (onverwachte) default-gebeurtenissen. Om deze risicopremies te schatten, ontwikkel ik een nieuw model voor de termijnstructuur van het kredietrisico van landen, waarin defaults kunnen worden veroorzaakt door schokken in een gemeenschappelijke factor of door schokken in een landspecifieke factor. Door deze factoren als zichzelf-versterkende processen te modelleren, is het model in staat om verschillende kenmerken van de data te accommoderen, namelijk 1) de hoge graad van gemeenschappelijkheid van kredietrisico van verschillende landen, 2) het clusteren van kredietschokken over tijd en tussen landen, en 3) sprongachtige toenames in de CDS premies van landen. In de empirische analyse gebruik ik zowel CDS termijnstructuurdata van 28 geografisch verspreide landen met kredietboordelingen variërend van A tot B, als historische defaultkansen per kredietbeoordelingsklasse zoals gerapporteerd door S&P.

Het geschatte model faciliteert een decompositie van CDS premies in twee dimensies. Allereerst kunnen CDS premies onderverdeeld worden in landspecifieke en systemische risicocomponenten. Ten tweede kunnen CDS premies onderverdeeld worden in verwachte defaultrisicocomponenten, distress-risicopremies, en default-gebeurtenisrisicopremies. In relatieve termen vind ik dat voor alle kredietbeoordelingsklassen gemiddeld ongeveer 65% van de vijfjaars CDS premies toegeschreven kan worden aan landspecifieke risico's en 35% aan systemische risico's. De decompositie in risicopremiecomponenten laat zien dat de default-gebeurtenisrisicopremie substantieel is en (in relatieve termen) voornamelijk

uit lijkt te maken voor de korte termijn en voor landen met lage kredietbeoordelingen. Distress-risicopremies, daarentegen, zijn zwaarder ingeprijsd in lange termijn CDS premies van landen met hogere kredietbeoordelingen. Een combinatie van de twee decomposities onthult dat de verschillen in de decomposities van landen met verschillende kredietboordelingen voornamelijk worden verzoorzaakt door landspecifieke risico's en niet door verschillen in de blootstellingen aan systemische risico's.

Hoofdstuk 4

In Hoofdstuk 4 bestuderen we de implicaties van arbeidsinkomensrisico op activawaardering, waarbij we specifiek focussen op de rol van mogelijke horizoneffecten. Om te testen of er eventuele horizoneffecten zijn, gebruiken we een flexibel empirisch raamwerk dat ons in staat stelt om het arbeidsinkomensrisico voor verschillende horizonnen gelijktijdig mee te nemen. We vinden een duidelijk afgebakende rol voor de twee- tot vierjaars horizon. Arbeidsinkomensrisico op deze middellange horizon heeft niet alleen een erg significante risicoprijs, terwijl dat niet het geval is voor andere horizonnen, maar ook het vermogen om de cross-sectionele verschillen tussen aandelenrendementen te verklaren piekt op deze horizon. Bijvoorbeeld, een simpel twee-factor model bestaande uit een gelijktijdige aandelenmarktrendementsfactor en de arbeidsinkomensgroeivoet op de middellange horizon kan een verbluffende 71% van de cross-sectionele variatie binnen 25 portfolios gesorteerd op grootte en boek-ten-opzichte-van-marktwaarde en 25 portfolios gesorteerd op grootte en investeringsintensiteit verklaren. Ter vergelijking, het standaard menselijk kapitaal CAPM met een kwartaalijkse arbeidsinkomensgroeivoet kan slechts 7% van de cross-sectionele variatie binnen die portfolios verklaren. Dit betekent dus dat simpelweg het veranderen van de horizon waarover arbeidsinkomensrisico wordt gemeten een dramatisch effect heeft op de werking van het model. Verder documenteren we vergelijkbare horizoneffecten voor de optimale portefeuillesamenstelling, waar arbeidsinkomensrisico op de middellange horizon significante aanpassingen binnen de optimale aandelenportefeuille genereert. Onze resultaten zijn consistent met de stroperigheid van lonen ten gevolge van het opnieuw vaststellen van lonen elke twee tot vier jaar.

LIST OF CO-AUTHORS

Rob C. Sperna Weiland conducted the research for Chapter 2, which is based on Sperna Weiland et al. (2019), in collaboration with Prof. dr. Roger J.A. Laeven and Prof. dr. Frank C.J.M. de Jong. Sperna Weiland made a very significant independent contribution to this research, but greatly appreciates and acknowledges the input and work of his co-authors. In particular, Laeven and Sperna Weiland developed the idea underlying the model. Laeven suggested the use of an MCMC estimation framework, which was further developed by Sperna Weiland. De Jong suggested the construction of the measurement equation. Sperna Weiland conducted all the elements of the analysis. Sperna Weiland and Laeven wrote the paper.

Rob C. Sperna Weiland conducted the research for Chapter 3, which is based on the single-authored paper Sperna Weiland (2018). Sperna Weiland greatly appreciates and acknowledges the input of his advisors Prof. dr. Roger J.A. Laeven, Prof. dr. Frank C.J.M. de Jong, and Prof. dr. Peter J.C. Spreij.

Rob C. Sperna Weiland conducted the research for Chapter 4, which is based on Eiling et al. (2019), in collaboration with Dr. Esther Eiling, Prof. dr. Frank C.J.M. de Jong and Prof. dr. Roger J.A. Laeven. Sperna Weiland made a very significant independent contribution to this research, but greatly appreciates and acknowledges the input and work of his co-authors. In particular, Eiling developed the initial idea of analyzing horizon effects in labor and finance. De Jong developed the model. Laeven suggested the KUJ preference structure. Sperna Weiland conducted all the elements of the analysis. Sperna Weiland, Eiling, and Laeven wrote the paper.

TINBERGEN INSTITUTE RESEARCH SERIES

The Tinbergen Institute is the Institute for Economic Research, which was founded in 1987 by the Faculties of Economics and Econometrics of the Erasmus University Rotterdam, University of Amsterdam and VU University Amsterdam. The Institute is named after the late Professor Jan Tinbergen, Dutch Nobel Prize laureate in economics in 1969. The Tinbergen Institute is located in Amsterdam and Rotterdam. The following books recently appeared in the Tinbergen Institute Research Series:

694. P. GAL, *Essays on the role of frictions for firms, sectors and the macroeconomy*
695. Z. FAN, *Essays on International Portfolio Choice and Asset Pricing under Financial Contagion*
696. H. ZHANG, *Dealing with Health and Health Care System Challenges in China: Assessing Health Determinants and Health Care Reforms*
697. M. VAN LENT, *Essays on Intrinsic Motivation of Students and Workers*
698. R.W. POLDERMANS, *Accuracy of Method of Moments Based Inference*
699. J.E. LUSTENHOUWER, *Monetary and Fiscal Policy under Bounded Rationality and Heterogeneous Expectations*
700. W. HUANG, *Trading and Clearing in Modern Times*
701. N. DE GROOT, *Evaluating Labor Market Policy in the Netherlands*
702. R.E.F. VAN MAURIK, *The Economics of Pension Reforms*
703. I. AYDOGAN, *Decisions from Experience and from Description: Beliefs and Probability Weighting*
704. T.B. CHILD, *Political Economy of Development, Conflict, and Business Networks*
705. O. HERLEM, *Three Stories on Influence*
706. J.D. ZHENG, *Social Identity and Social Preferences: An Empirical Exploration*

707. B.A. LOERAKKER, *On the Role of Bonding, Emotional Leadership, and Partner Choice in Games of Cooperation and Conflict*

708. L. ZIEGLER, *Social Networks, Marital Sorting and Job Matching. Three Essays in Labor Economics*

709. M.O. HOYER, *Social Preferences and Emotions in Repeated Interactions*

710. N. GHEBRIHIWET, *Multinational Firms, Technology Transfer, and FDI Policy*

711. H. FANG, *Multivariate Density Forecast Evaluation and Nonparametric Granger Causality Testing*

712. Y. KANTOR, *Urban Form and the Labor Market*

713. R.M. TEULINGS, *Untangling Gravity*

714. K.J.VAN WILGENBURG, *Beliefs, Preferences and Health Insurance Behavior*

715. L. SWART, *Less Now or More Later? Essays on the Measurement of Time Preferences in Economic Experiments*

716. D. NIBBERING, *The Gains from Dimensionality*

717. V. HOORNWEG, *A Tradeoff in Econometrics*

718. S. KUCINSKAS, *Essays in Financial Economics*

719. O. FURTUNA, *Fiscal Austerity and Risk Sharing in Advanced Economies*

720. E. JAKUCIONYTE, *The Macroeconomic Consequences of Carry Trade Gone Wrong and Borrower Protection*

721. M. LI, *Essays on Time Series Models with Unobserved Components and Their Applications*

722. N. CIURILĂ, *Risk Sharing Properties and Labor Supply Disincentives of Pay-As-You-Go Pension Systems*

723. N.M. BOSCH, *Empirical Studies on Tax Incentives and Labour Market Behaviour*

724. S.D. JAGAU, *Listen to the Sirens: Understanding Psychological Mechanisms with Theory and Experimental Tests*

725. S. ALBRECHT, *Empirical Studies in Labour and Migration Economics*

726. Y. ZHU, *On the Effects of CEO Compensation*

727. S. XIA, *Essays on Markets for CEOs and Financial Analysts*

728. I. SAKALAUSKAITE, *Essays on Malpractice in Finance*

729. M.M. GARDBERG, *Financial Integration and Global Imbalances*

730. U. THŰMMEL, *Of Machines and Men: Optimal Redistributive Policies under Technological Change*

731. B.J.L. KEIJSERS, *Essays in Applied Time Series Analysis*

732. G. CIMINELLI, *Essays on Macroeconomic Policies after the Crisis*

733. Z.M. LI, *Econometric Analysis of High-frequency Market Microstructure*

734. C.M. OOSTERVEEN, *Education Design Matters*

735. S.C. BARENDSE, *In and Outside the Tails: Making and Evaluating Forecasts*

736. S. SÓVÁGÓ, *Where to Go Next? Essays on the Economics of School Choice*

737. M. HENNEQUIN, *Expectations and Bubbles in Asset Market Experiments*

738. M.W. ADLER, *The Economics of Roads: Congestion, Public Transit and Accident Management*

739. R.J. DÖTTLING, *Essays in Financial Economics*

740. E.S. ZWIERS, *About Family and Fate: Childhood Circumstances and Human Capital Formation*

741. Y.M. KUTLUAY, *The Value of (Avoiding) Malaria*

742. A. BOROWSKA, *Methods for Accurate and Efficient Bayesian Analysis of Time Series*

743. B. HU, *The Amazon Business Model, the Platform Economy and Executive Compensation: Three Essays in Search Theory*

COLOPHON

This document was typeset using `classicthesis` (version 4.2) developed by André Miede.