

HEC MONTRÉAL

École affiliée à l'Université de Montréal

Four Essays on Applications of Filtering Methods in Finance

par
Jean-François Bégin

Thèse présentée en vue de l'obtention du grade de Ph. D. en administration
(option Méthodes Quantitatives)

Novembre 2016

© Jean-François Bégin, 2016

HEC MONTRÉAL

École affiliée à l'Université de Montréal

Cette thèse intitulée :

Four Essays on Applications of Filtering Methods in Finance

Présentée par :

Jean-François Bégin

a été évaluée par un jury composé des personnes suivantes :

Laurent Charlin
HEC Montréal
Président-rapporteur

Geneviève Gauthier
HEC Montréal
Directrice de recherche

Debbie Dupuis
HEC Montréal
Membre du jury

Yacine Aït-Sahalia
Princeton University
Examineur externe

Bernard Gauthier
HEC Montréal
Représentant du directeur de HEC Montréal

Résumé

Cette thèse se concentre sur quelques applications des méthodes de filtrage en finance. Elle est divisée en quatre essais.

Dans le premier essai, les primes de swaps de défaillance sont étudiées avant, pendant et après la crise financière à l'aide d'un modèle de risque de crédit flexible. Il comprend un régime statistique propre à chaque firme qui capte les changements dans la volatilité du ratio d'endettement. La relation négative entre les taux de recouvrement et les probabilités de défaut est modélisée à l'aide d'un recouvrement endogène qui dépend de la santé financière de la firme. Une méthode de filtrage est adaptée pour le présent problème d'estimation: nous généralisons l'algorithme de détection-estimation de Tugnait (1982) afin de tenir en compte la non-linéarité des équations de la représentation état-espace. Le modèle est utilisé pour montrer l'importance de la liquidité spécifique à une émission obligataire dans la différence entre les rendements des obligations et des swaps de défaillance.

Dans le second essai, le risque de crédit est analysé dans le secteur financier. Ceci nécessite un modèle complexe capturant les différents déterminants du risque de crédit (et plus particulièrement, des crises financières). Le modèle utilisé généralise celui proposé dans le premier essai en permettant aux différentes firmes d'être corrélées. Une méthode d'estimation en deux étapes basée sur des méthodes de filtrage est proposée. Nos résultats empiriques démontrent une augmentation de la corrélation durant la dernière crise financière. De plus, une mesure du risque systémique est calculée et ce risque est analysé dans deux sous-secteurs: les banques et les assureurs.

Dans un autre ordre d'idée, la littérature récente propose des conclusions empiriques contradictoires sur la relation entre le risque idiosyncratique et les rendements d'actions. Le troisième essai investigue les primes de risque sur actions. Un modèle discret à volatilité stochastique avec sauts est construit; les composantes gaussienne et de sauts de chaque firme dépendent de composantes systématiques. Le modèle est estimé

sur 260 firmes en utilisant les rendements et les prix d'options du marché, ainsi que les rendements et les prix d'options pour chacune de ces firmes. Une estimation en deux étapes basée sur des méthodes de filtrage est accomplie. Tel qu'espéré, le risque systématique est un déterminant clé expliquant 60% des primes de risque en moyenne, alors que le risque idiosyncratique explique plus de 40%. Nous montrons aussi que ce risque idiosyncratique provient exclusivement de la composante de sauts.

Dans le quatrième essai, une nouvelle mesure basée sur des prix d'options à haute fréquence est développée. Analogue aux mesures de variance réalisée pour les rendements, les séries chronologiques de prix d'options peuvent être utilisées pour calculer la variance réalisée d'une option (*ORV*). Un filtre particulière inspiré de l'algorithme d'échantillonnage avec rééchantillonnage par importance (Gordon et al., 1993) est spécifiquement construit pour tenir compte des différentes sources d'information. Selon une étude de simulation, les *ORV* sont non-redondantes puisqu'elles contiennent de l'information sur la taille et le moment des sauts dans la variance. Dans une étude empirique, les données du S&P 500 sont employées pour estimer un modèle diffusif avec sauts.

Mots clés: filtrage; filtres déterministes; filtres particuliers; risque de crédit; risque systémique; secteur financier; primes de risques; risque idiosyncratique; risque systématique; évaluation d'options; données à haute fréquence.

Méthodes de recherche: économétrie; modélisation mathématique; recherche quantitative; analyse multivariée

Abstract

This thesis focuses on applications of filtering methods in finance. It is divided in four essays.

In the first essay, credit default swap (CDS) premiums are investigated before, during and after the last financial crisis with a flexible credit risk model. It includes a firm-specific statistical regime that accommodates for changes in the leverage uncertainty. The negative relationship between default probabilities and recovery rates is modelled with an endogenous recovery rate that depends on the firm's financial health. A filtering method tailored for the issue at hand is implemented. Using the term structure of CDS premiums for 225 companies, a firm-by-firm estimation is performed. Based on regression tests and yield-to-maturity spreads implied by our credit risk model, we find that bond-specific liquidity is an important driver of the bond-CDS basis and its contribution is even more significant during the past crisis.

In the second essay, firm-specific credit risk is analyzed in the financial services sector. This requires an elaborate model that captures the main determinants of credit risk (and more specifically, financial crises). It extends the model proposed in the first essay by allowing for linkages between co-movements of firm leverages. To estimate the model, we develop a two-stage filtering procedure. We find evidence of larger correlations between firm leverage co-movements during the high-volatility regime which suggests the existence of greater interconnectedness during the last crisis. Finally, systemic risk measures of two subsectors, namely insurance and banking, are computed and compared.

In another vein, the recent literature provides conflicting empirical evidence on the relationship between idiosyncratic risk and equity returns. In the third essay, equity risk premiums are investigated. A discrete-time jump-diffusion model is developed in which a firm's Gaussian and jump innovations have systematic components. The model is estimated on 260 firms, based on market returns and option prices, as well

as returns and option prices for each of these firms. A two-step estimation based on stochastic filtering techniques is implemented. First, we find that systematic risk factors explain close to 60% of the risk premium on average, while idiosyncratic factors explain more than 40%. Second, we show that the contribution of idiosyncratic risk to the equity risk premium arises exclusively from the jump risk component. Tail risk, interpreted as jumps, thus plays a central role in the pricing of idiosyncratic risk.

In the fourth essay, the information content of high frequency option prices is studied. A measure of option price variation, the option realized variance, is proposed as an observable source that summarizes information from intraday option prices. A particle filter inspired from the sequential importance resampling method (Gordon et al., 1993) is specifically constructed to account for the various model inputs. We conduct an extensive simulation study of this new variable and document empirically its incremental information content. Our results show that the information contained in these variances improves the inference of latent variables such as the instantaneous variance and jumps. In our empirical study, we employ data from the S&P 500 index to document several properties of the option realized variance.

Keywords: filtering; deterministic filters; particle filters; credit risk; systemic risk; financial sector; risk premiums; idiosyncratic risk; systematic risk; option valuation; high frequency data.

Research methods: econometrics; mathematical modeling; quantitative research; multivariate analysis.

Contents

Résumé	v
Abstract	vii
Contents	ix
List of Figures	xv
List of Tables	xix
Abbreviations	xxi
Acknowledgements	xxiii
1 Introduction	1
2 Firm-Specific Credit Risk Modelling in the Presence of Statistical Regimes and Noisy Prices	5
Abstract	5
2.1 Introduction and Literature Review	6
2.2 Credit Risk Model	10
2.3 Estimation Method and Pricing Technique	12
2.3.1 Estimation Technique	12
2.3.1.1 Unscented Kalman Filter-Based Method	13
2.3.2 Credit-Sensitive Derivative Pricing	15
2.4 Estimation Results	16
2.4.1 Data	16
2.4.2 Estimated Parameters	18
2.4.3 In- and Out-of-Sample Performances	19
2.4.3.1 In-Sample Performance	21
2.4.3.2 Out-of-Sample: Forecasting 2013	23
2.4.4 Most Likely Statistical Regimes Through Time	24
2.5 Term Structure Comparison: Recovery Rate Versus Regime-Switching	26
2.6 Bond-CDS Basis	28
2.6.1 Measuring the Basis	29

2.6.2	Data	30
2.6.3	Bond-Specific Liquidity as a Driver of the Basis	32
2.7	Concluding Remarks	35
2.A	Derivative Pricing	36
2.A.1	Trinomial Lattice Approach	36
2.A.2	Coupon Bond Prices	38
2.A.3	Credit Default Swap Premiums	38
2.B	Simulation Study for the DEA-UKF Method	39
2.C	Liquidity: Additional Material	41
	References	44
3	Credit and Systemic Risks in the Financial Services Sector: Evidence from the 2008 Global Crisis	51
	Abstract	51
3.1	Introduction and Review of the Literature	52
3.2	Multivariate Credit Risk Model	55
3.2.1	Markov-Switching Dynamics	56
3.2.2	Default Intensity	57
3.3	Data and Assumptions	58
3.4	Firm-Specific Credit Risk	61
3.4.1	Default Probabilities	64
3.4.2	Recovery Risk	67
3.5	Dependence	68
3.6	Systemic Risk	74
3.6.1	Systemic Risk Measures	75
3.6.2	Granger Causality Tests	78
3.7	Concluding Remarks	80
3.A	Endogenous Correlation Coefficients	82
3.B	Calculation of Systemic Risk Measures	84
3.C	Systemic Risk Measures: Additional Results	85
3.C.1	Linear Granger Causality Tests	85
3.C.2	Nonlinear Granger Causality Tests	86
3.D	Advantages of the Regime-Switching Framework over the “One-Regime” Equivalent	87
3.D.1	Parameter Estimates	87
3.D.2	Fit	87
3.D.3	Default Probabilities	90
3.D.4	Recovery Rates	91
3.D.5	Dependence	91
3.D.6	Systemic Risk Measures	95
	References	95
4	Idiosyncratic Jump Risk Matters: Evidence from Equity Returns and Options	101

Abstract	101
4.1 Introduction	102
4.2 Model	105
4.2.1 Stock Returns	106
4.2.2 Pricing Kernel and Risk Premiums	108
4.2.3 Option Prices	110
4.3 Joint Estimation Using Returns and Option Prices	111
4.3.1 Data	112
4.3.2 Joint Estimation	115
4.4 Empirical Results	117
4.4.1 Market	117
4.4.2 Idiosyncratic <i>Jump</i> Risk Matters	122
4.4.2.1 Averages Across Industries	127
4.4.3 Commonality in Idiosyncratic Jump Risk	129
4.4.4 On the Importance of Accounting for Equity-Specific Jumps	130
4.4.5 Realized Premium based Portfolios of Stocks	133
4.4.5.1 Characteristics of the Quintile Portfolios and Double-Sort Portfolios	136
4.5 Conclusion	137
4.A Innovations' Cumulant Generating Functions	140
4.A.1 Continuous Component	140
4.A.2 Jump Component	140
4.B Innovations' Risk Neutral Cumulant Generating Functions	141
4.C Risk Premiums	142
4.D Risk Neutral Conditional Variances and Jump Intensities	143
4.E Moment Generating Function of Risk-Neutral Excess Returns	145
4.F Particle Filter	146
4.G Stock Fundamentals	148
References	149
5 Extracting Latent States with High Frequency Option Prices	155
Abstract	155
5.1 Introduction	156
5.2 Framework	159
5.2.1 Model	160
5.2.2 Links Between Theoretical and Empirical Quantities	161
5.2.2.1 Index Prices as a First Source of Information	162
5.2.2.2 Option-Based Information	164
5.3 Filtering and Estimation	167
5.3.1 Filtering Algorithm	168
5.3.2 Estimation	169
5.4 Simulation-Based Results	171
5.4.1 Filter's Performance as a Function of Information	172

5.4.2	Information Embedded in <i>ORV</i>	176
5.4.3	Option Quadratic Variation Increments Approximation	179
5.5	Exploring Option Realized Variances Empirically	179
5.5.1	Data	181
5.5.2	Understanding Option Realized Variances	183
5.5.3	Predicting Index Return Variations	189
5.6	Option Pricing Implications	192
5.6.1	Parameter Estimates	192
5.6.2	Informativeness of Data Sources	194
5.6.3	In- and Out-of-Sample Assessment	195
5.6.3.1	In-Sample	195
5.6.3.2	Out-of-Sample	199
5.7	Concluding Remarks	201
5.A	Pricing	201
5.A.1	Radon-Nikodym Derivative	201
5.A.2	Model Under the Risk-Neutral Measure \mathbb{Q}	202
5.A.3	Pricing	204
5.B	Estimation	205
5.B.1	Drift Term	205
5.B.2	Filtering Procedure	206
5.B.2.1	Intraday Simulation	206
5.B.2.2	Time Aggregation	208
	References	209
6	Concluding Remarks	215
A	Appendices of <i>Idiosyncratic Jump Risk Matters: Evidence from Equity Returns and Options</i>	217
A.1	Normal-Inverse Gaussian Distribution	217
A.1.1	Interpretation of the Jump Intensity Parameter	218
A.1.2	Returns' Conditional Moments	218
A.1.3	Conditional Variance and Jump Intensity Variable Moments	220
A.2	Detailed Proofs of Section 4.B's Results	222
A.3	Detailed Proofs of Section 4.C's Results	223
A.3.1	Market Drift Under \mathbb{P}	223
A.3.2	Stock Drift Under \mathbb{P}	224
A.4	Calculation Associated with Section 4.D	225
A.5	Detailed Proofs of Section 4.E's Results	227
A.5.1	Moment Generating Function	227
A.5.2	European Call Option Price	230
B	Appendices of <i>Extracting Latent States with High Frequency Option Prices</i>	233
B.1	Model Construction	233
B.1.1	General Properties of Moment Generating Functions	233

B.1.2	Exponential Martingales	234
B.1.2.1	Proofs	234
B.1.3	Risk-Neutral Innovations	236
B.1.3.1	Proofs	237
B.1.4	Log-Equity Price Drift	240
B.2	Option Pricing	247
B.2.1	Moment Generating Function of Log-Equity Price Variation .	247
B.2.2	Quadratic Variation of Option Prices	248
B.2.2.1	Derivative Calculation for ΔOQV	249
B.3	Simulation-Based Results: Additional Results	251
B.3.1	Filtering Tests	251
B.3.2	Option Data: How Much Is Enough?	251
B.4	Descriptive Statistics of Realized Variations	257
Bibliography		258

List of Figures

2.1	Evolution of the mean five-year CDS premiums and of the difference between the average ten-year CDS premiums and the average one-year CDS premiums.	17
2.2	SSE/SST: Evolution of the in-sample performance for IG and HY. . .	22
2.3	Proportion of firms in the high-volatility statistical regime.	25
2.4	Credit spreads in basis points as a function of maturities, credit ratings and eras.	27
2.5	Time series of the average bond-CDS basis and the average composite λ measure for IG and HY.	32
2.6	Histogram of the regressions' R -squared values.	35
2.7	The branching to default at a typical node in the tree when the number of regimes is set to two.	37
3.1	Evolution of the average CDS level in basis points and of the average CDS slope in basis points for both subsectors.	60
3.2	Time series of the proportion of firms in the high-volatility regime across both insurance and banking subsectors.	64
3.3	One-year default probabilities computed using the credit risk model and the default count approach and filtered regimes for AIG and LEH.	65
3.4	Average one-year default probabilities computed using the credit risk model and the default count approach.	68
3.5	Time series of the average one-year expected recovery rate for insurance and banking companies.	69
3.6	Heat maps of $\rho_{1,1}$, $\rho_{1,2}$, $\rho_{2,1}$ and $\rho_{2,2}$ for the 35 firms.	71
3.7	Histogram of $\rho_{1,1}$, $\rho_{1,2}$, $\rho_{2,1}$ and $\rho_{2,2}$ for three categories: correlations across insurance firms' leverages (Insurance/Insurance), correlations across banking companies' leverages (Banking/Banking), and correlations between insurance and banking firms' leverages (Insurance/Banking).	73
3.8	Time series of median leverage correlation coefficients for three groups: correlations across insurance firms (Insurance/Insurance), correlations across banking companies (Banking/Banking), and correlations between insurance and banking firms (Insurance/Banking).	74
3.9	Time series of the systemic risk measure contribution and the relative systemic risk measure contribution using the full model.	78

3.10	Sample autocorrelation functions for squared differenced systemic risk contributions and squared GARCH noise terms.	80
3.11	Average one-year default probabilities computed using the regime-switching credit risk model and its “one-regime” equivalent.	90
3.12	Average one-year expected recovery rates using the regime-switching credit risk model and its “one-regime” equivalent.	92
3.13	Heat map of ρ for the 35 firms using “one-regime” equivalent model of Boudreault et al. (2014) along with the correlation coefficients if both firms are in regime 1 and 2, respectively.	92
3.14	Time series of median leverage correlation coefficients for three groups: correlations across insurance firms (Insurance/Insurance), correlations across banking companies (Banking/Banking), and correlations between insurance and banking firms (Insurance/Banking).	93
3.15	Histogram of ρ for three categories: correlations across insurance firms’ leverages (Insurance/Insurance), correlations across banking companies’ leverages (Banking/Banking), and correlations between insurance and banking firms’ leverages (Insurance/Banking).	94
3.16	Time series of the systemic risk measure contribution and the relative systemic risk measure contribution using both regime-switching and “one-regime” equivalent models.	94
4.1	S&P 500 and ATM implied volatilities.	116
4.2	Filtered innovations and variances for the market model.	120
4.3	Annualized normal and jump risk premiums for the market model.	122
4.4	Decomposition of the equity risk premium by firm.	126
4.5	Time series decomposition of the average equity risk premium by industry.	128
4.6	Decomposition of the equity risk premium by firm.	132
5.1	Example of instantaneous variance densities using different data sources.	173
5.2	Option realized variance against option quadratic variance increments for a call-equivalent delta of 0.50 and a maturity of 30 business days, on days without (top panels) and without jumps (bottom panels).	180
5.3	Realized volatility, square root of the bipower variation, ATM option implied volatility for options with maturity closest to 30 business days and ATM option realized volatility (square root of <i>ORV</i>) for options with maturity closest to 30 business days.	184
5.4	Option realized variances and fitted surface as a function of both the option Black and Scholes (1973) call-equivalent delta and business days to maturity.	185
5.5	Option realized variance surface characteristics and realized measures.	186
5.6	Principal components of the option realized variance surface.	188
5.7	Filtered log-equity price jumps (left panels) and variance jumps (right panels).	196

B.1	Filtered instantaneous variance using various data source.	252
B.2	Filtered quadratic variation using various data source.	253
B.3	Filtered integrated variance using various data source.	254
B.4	Filtered log-equity jumps using various data source.	255
B.5	Filtered variance jumps using various data source.	256

List of Tables

2.1	Descriptive statistics on the distribution of firm-specific parameters and noise terms across the portfolio of firms of the CDX indices. . . .	20
2.2	SSE/SST: In- and out-of-sample performance for IG and HY.	23
2.3	Descriptive statistics on the bond-CDS basis and the bond-specific liquidity measure λ	31
2.4	Bond-CDS basis regression results.	34
2.5	Simulation studies on the DEA-UKF methodology.	40
2.6	Descriptive statistics for the liquidity proxies.	41
3.1	Insurance and banking firms.	59
3.2	Descriptive statistics on the distribution of firm-specific parameters and noise terms.	62
3.3	Descriptive statistics of one-year default probability estimates across periods.	67
3.4	Average systemic risk measures on three different periods (i.e. pre-crisis, crisis and post-crisis) using different modelling assumptions. . .	77
3.5	Linear and nonlinear Granger causality tests after GARCH filtering. .	81
3.6	Linear and nonlinear Granger causality tests before GARCH filtering.	86
3.7	First-stage parameter estimates: “one-regime” equivalent of Boudreault et al. (2013).	88
3.8	Relative root mean square errors in percentage: comparison between the regime-switching model and the “one-regime” equivalent of Boudreault et al. (2013).	89
3.9	Relative root mean square errors in percentage: comparison between the regime-switching model and the “one-regime” equivalent of Boudreault et al. (2013), continued.	90
4.1	Description of the SPX index option data (1996–2015).	113
4.2	Description of firms option data (1996–2015).	114
4.3	Index parameters estimated using returns and option data.	118
4.4	Valuation errors on the options used for the estimation.	121
4.5	Firm parameters estimated using returns and option data.	124
4.6	Correlation between the parameters of stocks.	125
4.7	Firm parameters estimated using returns and option data for each of the eight largest GICS industries.	129

4.8	Firm parameters estimated using returns and option data: Neglecting idiosyncratic jumps.	134
4.9	Excess returns of portfolios based on idiosyncratic jump risk premiums.	138
4.10	Description of quintile portfolios.	139
4.11	Double sort: Excess returns of portfolios based on idiosyncratic jump risk premiums.	139
5.1	Model parameters used in the simulation study.	171
5.2	RMSE for instantaneous variance, integrated variance, quadratic variance, log-equity jump, and instantaneous variance jump across 100 simulated paths.	174
5.3	Average jump times in percentages.	176
5.4	Information content of <i>ORV</i> when jumps are absent.	177
5.5	Information content of <i>ORV</i> when jumps are present.	178
5.6	RMSE, relative RMSE and regression <i>R</i> -squared for ΔOQV approximation across 500 days.	181
5.7	Correlation matrix for option realized variance characteristics and realized measures.	187
5.8	Option realized volatility characteristics, realized measures, and principal component regressions.	189
5.9	Correlation matrix of realized values used in the predictive regressions.	190
5.10	Predictive regressions.	191
5.11	S&P 500 parameter estimates.	193
5.12	Posterior standard deviation of filtered values with and without <i>ORV</i>	195
5.13	In-sample performance (2004–2012).	198
5.14	One-day ahead out-of-sample performances, in terms of RIVRMSE and RORVRMSE (2013).	200
B.1	RMSE for various quantities across 100 simulated paths when using only a limited number of options.	257
B.2	Descriptive statistics of realized variance, bipower variation and option realized variance	257

Abbreviations

ATM	At-The-Money
BV	Bipower Variation
CAPM	Capital Asset Pricing Model
CDS	Credit Default Swap
CRSP	Center for Research in Security Prices
DEA	Detection Estimation Algorithm
DGP	Data Generating Process
DP	Diks and Panchenko
ERP	Equity Risk Premium
FISD	Mergent Fixed Income Securities Database
GARCH	Generalized Autoregressive Conditional Heteroskedasticity
GICS	Global Industry Classification Standard
HY	High Yield
IG	Investment Grade
IRC	Imputed Round-trip Cost
ITM	In-The-Money
IV	Implied Volatility
IVRMSE	Implied Volatility Root Mean Square Error
MLE	Maximum Likelihood Estimation
MCMC	Markov-Chain Monte Carlo
NAIC	National Association of Insurance Commissioners
NIG	Normal-Inverse Gaussian
OIS	Overnight Indexed Swap

OPRA	Option Price Reporting Authority
ORV	Option Realized Variance
OTM	Out-of-The-Money
PD	Probabilities of Default
RIVRMSE	Relative Implied Volatility Root Mean Square Error
RMSE	Root Mean Square Error
RND	Radon-Nikodym Derivative
RV	Realized Variance
SIC	Standard Industrial Classification
SIR	Sequential Importance Resampling
SSE	Sum of Squared Errors
SST	Sum of Squared Total
TARP	Troubled Asset Relief Program
TRACE	Trade Reporting And Compliance Engine
UKF	Unscented Kalman Filter
VaR	Value at Risk
YTM	Yield To Maturity

Acknowledgements

First and foremost, I would like to express my sincere gratitude to my advisor Professor Geneviève Gauthier. Without her constant feedback and guidance, this PhD would not have been achievable. Throughout my thesis-writing period, she provided sound advice, good company, and lots of great ideas. I would have been lost without her.

I also wish to thank members of the examining committee, Professors Laurent Charlin, Debbie Dupuis and Yacine Aït-Sahalia, for their insightful comments.

My sincere thanks also go to Professors Diego Amaya, Mathieu Boudreault and Christian Dorion who helped and gave valuable research advice. I am thankful for all the useful discussions and brainstorming sessions we had during my PhD.

I would like to thank the many people who have taught me mathematics, statistics and finance: my undergraduate teachers (especially Professor Manuel Morales and Doctor Charles Dugas), and my graduate teachers (especially Professors Bruno Rémillard, Georges Dionne and Chantal Labbé).

A very special thank you to Professors Mylène Bédard and Patrice Gaillardetz for providing career guidance many times during my time at Université de Montréal and HEC Montréal.

I am indebted to my schoolmates for all the emotional support, entertainment, and caring they provided.

I would also like to thank my parents and my sisters for supporting me throughout the writing of this thesis, and life in general. Last but not least, I wish to express my sincere gratitude to Sabrina who has been by my side throughout this PhD, living every single minute of it. You gave me strength when I needed some. Thank you.

Chapter 1

Introduction

Financial modelling often prescribes fundamental nonlinear relationships between model variables. A number of these variables are not directly observed, but are rather inferred through a mathematical model from other variables that are observed, or directly measured. Finding “best estimates” for the unobserved variables often relies on filters. Filtering methods are well-suited for the inference of the model variables’ posterior probability distribution as they are efficient and intuitive.

Even though this technology was historically used in engineering applications¹, nowadays a large number of financial practitioners and academics see in these methods a proper way of dealing with unobservable quantities (or so-called latent variables). Indeed, a great deal of latent variables are present in financial modelling: default intensity in credit risk models, instantaneous volatility in asset models, short rate in interest rate models, regime in regime-switching models and convenience yield in commodity futures models to name a few. Creal’s (2012) survey reviews and introduces some of the most standard methods to deal with latent variables.

State-space representations are convenient means for studying dynamic systems. This representation consists of two equations: the observation (or measurement) equation and the transition equation.

Definition 1.1. (State-Space Representation).

Let \mathbf{x}_t be the latent variables at time t and \mathbf{y}_t the observations at time t . The observation and transition equations are respectively given by

$$\mathbf{y}_t = m_t(\mathbf{x}_t, \boldsymbol{\varepsilon}_t) \tag{1.1}$$

¹Guidance, navigation and control of vehicles, especially spacecraft and aircraft, signal processing, robotic motion planning and control and trajectory optimization to name a few.

$$\mathbf{x}_t = s_t(\mathbf{x}_{t-1}, \boldsymbol{\eta}_t) \quad (1.2)$$

where the functions m_t and s_t are possibly nonlinear, but known. The sequences $\{\boldsymbol{\varepsilon}_t\}_{t \in \mathbb{N}}$ and $\{\boldsymbol{\eta}_t\}_{t \in \mathbb{N}}$ are two sequences of independent random variables. Generally, we can write in closed-form both observation and transition densities denoted by $f(\mathbf{y}_t | \mathbf{x}_t; \boldsymbol{\Theta})$ and $f(\mathbf{x}_t | \mathbf{x}_{t-1}; \boldsymbol{\Theta})$, respectively. The two densities also typically depend upon a vector of unknown parameters $\boldsymbol{\Theta}$.

The filtering problem consists in estimating the latent states $\mathbf{x}_{0:t} = \{\mathbf{x}_0, \dots, \mathbf{x}_t\}$ in dynamical systems when partial observations are made, and random perturbations are present. Uncertainty on the state variable is formulated as a joint conditional probability distribution $f(\mathbf{x}_{0:T} | \mathbf{y}_{1:T}; \boldsymbol{\Theta})$ known as the joint smoothing distribution. The latter is defined as

$$f(\mathbf{x}_{0:T} | \mathbf{y}_{1:T}; \boldsymbol{\Theta}) = \frac{f(\mathbf{x}_{0:T}, \mathbf{y}_{1:T}; \boldsymbol{\Theta})}{f(\mathbf{y}_{1:T}; \boldsymbol{\Theta})}. \quad (1.3)$$

Filtering methods give us practical ways to compute such posterior distributions.

Most filtering methods are recursive, as the posterior distribution of $\mathbf{x}_{0:t-1}$ is explicitly used to find the posterior distribution of $\mathbf{x}_{0:t}$. Indeed, it is easy to show that

$$\begin{aligned} f(\mathbf{x}_{0:t} | \mathbf{y}_{1:t}; \boldsymbol{\Theta}) &= \frac{f(\mathbf{y}_t | \mathbf{x}_{0:t}, \mathbf{y}_{1:t-1}; \boldsymbol{\Theta}) f(\mathbf{x}_{0:t} | \mathbf{y}_{1:t-1}; \boldsymbol{\Theta})}{f(\mathbf{y}_t | \mathbf{y}_{1:t-1}; \boldsymbol{\Theta})} \\ &= \frac{f(\mathbf{y}_t | \mathbf{x}_{0:t}, \mathbf{y}_{1:t-1}; \boldsymbol{\Theta}) f(\mathbf{x}_t | \mathbf{x}_{0:t-1}, \mathbf{y}_{1:t-1}; \boldsymbol{\Theta})}{f(\mathbf{y}_t | \mathbf{y}_{1:t-1}; \boldsymbol{\Theta})} f(\mathbf{x}_{0:t-1} | \mathbf{y}_{1:t-1}; \boldsymbol{\Theta}) \\ &= \frac{f(\mathbf{y}_t | \mathbf{x}_t; \boldsymbol{\Theta}) f(\mathbf{x}_t | \mathbf{x}_{0:t-1}; \boldsymbol{\Theta})}{f(\mathbf{y}_t | \mathbf{y}_{1:t-1}; \boldsymbol{\Theta})} f(\mathbf{x}_{0:t-1} | \mathbf{y}_{1:t-1}; \boldsymbol{\Theta}). \end{aligned}$$

Therefore, we can find an approximation of the time t posterior density $f(\mathbf{x}_{0:t} | \mathbf{y}_{1:t}; \boldsymbol{\Theta})$ from $f(\mathbf{x}_{0:t-1} | \mathbf{y}_{1:t-1}; \boldsymbol{\Theta})$.

In some applications, deterministic filters can be applied. For instance, the Kalman (1960) filter is a recursive algorithm that builds on the linearity and the Gaussianity of the state-space representation. If m_t and s_t are nonlinear functions, we could use deterministic extensions of the Kalman filter: the extended Kalman filter (EKF), the unscented Kalman filter (UKF), etc. However, a more robust and general approach uses Monte Carlo simulation: the so-called particle filters or the sequential Monte Carlo (SMC) methods. Particle filtering methodology uses a genetic type mutation-selection sampling approach, with a set of particles to represent the posterior distribution of some stochastic process given some partial observations.

The state-space representation typically depends upon a vector of unknown parameters Θ that need to be estimated from the observed data $\mathbf{y}_{1:T} = \{\mathbf{y}_1, \dots, \mathbf{y}_T\}$. Even though Θ is unknown, it is common in the literature to run filtering algorithms assuming a fixed value of the parameter vector. Thus, as a by-product of the filter and given a value of Θ , we are able to compute a likelihood function. Therefore, this output can be used to estimate Θ by maximizing the likelihood function. Simply put, we change the parameter vector recursively until we reach $\hat{\Theta}$, the maximum likelihood estimator of Θ (i.e. the parameter vector associated with the highest likelihood estimation).

The conditional density of the observations at time t given the past observations is written as an integral:

$$f(\mathbf{y}_t | \mathbf{y}_{1:t-1}; \Theta) = \int f(\mathbf{y}_t | \mathbf{x}_t, \mathbf{y}_{1:t-1}; \Theta) f(\mathbf{x}_t | \mathbf{y}_{1:t-1}; \Theta) d\mathbf{x}_t \quad (1.4)$$

where \mathbf{x}_t are the state variables at time t and \mathbf{y}_t are the observations at time t . For most models, such integrals cannot be solved analytically. Fortunately, most filtering methods can compute these integrals numerically as a by-product of the filter, allowing us to obtain an estimation of the likelihood function.

In this thesis, I rely heavily on such filter-based likelihood estimations to find $\hat{\Theta}$. I propose four different applications of filters in finance: two in the credit risk literature, one in the option pricing literature, and a final contribution in the high frequency literature. The purpose of this introduction is not to detail each chapter; each essay has its own introduction explaining the relevant research questions and the different methods used to answer them. Rather, this introduction gives an overview of the essays of this thesis.

The first essay is about credit risk. A firm-by-firm estimation of a regime-switching hybrid credit risk model is achieved using a generalization of Tugnait's (1982) detection-estimation algorithm (DEA). The DEA is designed to filter a state and a regime simultaneously in the context of linear state-space representations. In this essay, the DEA is generalized to account for nonlinearities by running the UKF of Julier and Uhlmann (1997) instead of the classic Kalman (1960) filter. Credit default swaps are used as inputs. Some interesting results regarding the bond-CDS basis are given. For instance, liquidity was an important driver of the basis during the past crisis.

In the second essay, a multivariate generalization of the credit risk model proposed in the first essay is implemented. Using weekly credit default swap premiums for 35 financial firms, we analyze the credit risk of each of these companies and their

statistical linkages, placing special emphasis on the 2005–2012 period. Moreover, we study the systemic risk affecting both the banking and insurance subsectors.

The third essay focuses on equity risk premiums. A discrete-time stock price model is estimated. The latter allows for firm-specific Gaussian and jump components to be related to systematic components. A two-step estimation is performed. First, the index parameters are inferred using a particle filter similar to the one of Gordon et al. (1993). Then, each firm's parameters are estimated using a similar stochastic filter.

In the fourth essay, a novel measure based on high frequency option prices is proposed. Analogous to realized variance measures for returns, high frequency option time series permits the calculation of the option realized variance (*ORV*). A continuous-time jump-diffusion model is estimated using a particle filter. The latter is inspired by Gordon et al.'s (1993) sequential importance resampling scheme. A simulation-based study shows that out-of-the-money *ORVs* contain new information about jump sizes and times. Estimation using real data from the S&P 500 index is done.

Chapter 2

Firm-Specific Credit Risk Modelling in the Presence of Statistical Regimes and Noisy Prices

Abstract^{*}

Security prices are important inputs for estimating credit risk models. Yet, to obtain an accurate firm-specific credit risk assessment, one needs a reliable model and a methodology that filters all elements unrelated to the firm's fundamentals from observed market prices.

We first introduce a flexible hybrid credit risk model defined in a Markov-switching environment. It captures firm-specific changes in the leverage uncertainty during crises as well as the negative relationship between creditworthiness and recovery rates. Second, estimation is performed using maximum likelihood by accounting for latent regimes and unobserved noise included in security prices.

Using CDS premiums for 225 firms of both CDX North American IG and HY indices, we present two different empirical applications. The effects of stochastic recovery and the presence of regimes on theoretical credit spread curves are investigated. We also apply the model to corporate bond credit spreads to assess the importance of bond-specific liquidity in the bond-CDS basis.

Keywords: risk analysis; credit risk; maximum likelihood estimation; regime-switching; filtering.

^{*}Joint work with Mathieu Boudreault and Geneviève Gauthier. Boudreault is affiliated with UQAM and Gauthier with HEC Montréal.

2.1 Introduction and Literature Review

The Global Financial Crisis of 2008 is considered by various economists as the worst crisis since the Great Depression. Worldwide, one hundred and one corporate issuers rated by Moody's defaulted on a total of \$281.2 billion of debt in 2008 (Moody's, 2009). For the sake of comparison, only 18 issuers defaulted in 2007 on a total of \$6.7 billion of debt. Also, the year 2008 included the largest defaulter in history: Lehman Brothers.

Indeed, defaults are rare events and the lack of direct observations brings an additional challenge to estimating credit risk. To overcome this challenge, a common practice is to use credit ratings as a proxy for firm-specific credit risk. However, in the aftermath of the crisis, credit ratings were criticized and many stakeholders started advocating for market-based valuation of credit risk.

Although market prices are forward-looking measures that are updated frequently by market participants, observed prices can diverge from their theoretical values as risks unrelated to their fundamentals could be priced (e.g. illiquidity, asymmetric information, price discreteness). These additional risks are hereafter referred to as market noise. Filtering techniques are powerful tools for extracting fundamental values from noisy security prices.

In the literature of credit risk, some recent estimation procedures allow for noisy observations. Indeed, Duan and Fulop (2009) propose an estimation technique for Merton's (1974) model that allows for trading noise in observed equity prices. The maximum likelihood estimation (MLE) is carried out by a particle filter. It is argued that ignoring noise could non-trivially inflate one's estimate of the asset volatility. Huang and Yu (2010) introduce an alternative estimation technique based on Markov-chain Monte Carlo (MCMC) methods for the same credit risk framework. According to the authors, noises are positively correlated with firms' values. Finally, in a hybrid credit risk framework, Boudreault et al. (2013) use the unscented Kalman filter (UKF) to extract the latent creditworthiness from credit default swap (CDS) premiums.¹

Firm-specific estimation of credit risk in the previously cited papers is simplified because these frameworks are based on single-factor models. As the last financial crisis showed, credit spreads and other solvency indicators exhibit structural changes.

¹More recently, Kwon and Lee (2015) use MCMC to estimate the Black and Cox (1976) model, whereas Guarin et al. (2014) propose a filtering technique to extract the time-varying default risk from CDS premiums.

Regime-switching dynamics are able to capture behaviour changes similar to those associated with a crisis by allowing for dynamic levels of uncertainty. Such regime-switching approaches are also very intuitive in the context of crises; in its most simplistic form, one could have a regime associated with “good” times and another with “bad” times. However, firm-specific MLE in the presence of regimes and noisy market prices is complicated because both regimes and noises are not directly observed.

This paper contributes to the current credit risk literature by proposing a Markov-switching model and a consistent estimation technique for the assessment of credit risk. We design a flexible model that is able to capture the desired empirical facts. Our hybrid credit risk model includes a firm-specific statistical regime variable that accommodates for changes in the leverage uncertainty. The endogenous random recovery rates are negatively related to default probabilities as they both depend on the firm’s financial health. Finally, the framework accounts for the presence of noise as market prices potentially include measurement errors.

Continuous-path structural models have failed to appropriately represent short-term credit spreads, mainly because the default is a predictable stopping time. To solve this issue, structural models with incomplete information were proposed; the presence of a surprise element that adds randomness to the default trigger is an important feature of these models. For instance, Duffie and Lando (2001) suppose that bond investors cannot observe the issuer’s assets directly and receive periodic and imperfect accounting reports instead. Jarrow and Protter (2004) show that the incomplete knowledge of the firm’s assets and liabilities leads to an inaccessible default time.

Other solutions combine ideas from both structural and reduced-form approaches to obtain hybrid models: firms’ liabilities and assets are modelled as stochastic processes and the default time is given by the first jump of a Cox process for which the intensity depends on the firm’s fundamentals. In this spirit, Madan and Unal (2000) use a two-factor hazard rate model which links the intensity to the value of the firm; this leads to a non-predictable default time. Bakshi et al. (2006a) use Duffie and Singleton’s (1999) reduced-form framework with factors modelled as Vasicek (1977) processes; one of these processes is the firm’s leverage. We use a similar approach: the default intensity of a firm is a convex function of its leverage.

As documented by various authors, CDS premium dynamics and credit spreads change during financial crises. It thus makes sense to include this risk in our framework,

especially as our sample period includes the last crisis. The dynamics of CDS premiums are investigated by Huang and Hu (2012) during the 2008 crisis by applying a smooth-transition autoregressive model. Maalaoui Chun et al. (2014) study regimes by applying a regime-detection technique that distinguishes between level and volatility regimes in credit spreads and show that most breakpoints occur around economic downturns, thus linking the statistical regimes to financial crises. Furthermore, Alexander and Kaeck (2008) show that CDS premiums display pronounced regime-specific behaviour.

The negative correlation between the firm's leverage and the recovery rate is also of paramount importance: as the firm's financial health becomes precarious (i.e. the firm's leverage rises), its probability of default increases and the recovery rate decreases, respectively. According to Altman et al. (2005) and Acharya et al. (2007), a negative correlation between probabilities of default and recovery rates exists, and both variables seem to be driven by the same factor. Other contributions show the importance of the negative relationship between probabilities of default and recovery rates. Using BBB-rated corporate bonds, regression analyses and the information on all companies that have defaulted between 1981 and 1999, Bakshi et al. (2006b) find that, on average, a 4% increase in the risk-neutral hazard rate is associated with a 1% decline in risk-neutral recovery rates. Further, an econometric model developed by Bruche and González-Aguado (2010) tries to assess by how much one underestimates credit risk if the negative relationship between probabilities of default and recovery rates is ignored. Using a simple structural model, Bade et al. (2011) find that default and recovery are highly negatively correlated; the recovery is modelled as a stochastic quantity that depends on observable risk factors and a systematic random variable.

The contributions of this paper are manifold. First of all, as stated above, we propose a hybrid credit risk model that includes a firm-specific statistical regime, among other realistic features.

Second, the model is estimated by maximum likelihood using a filtering approach that we design for the issue at hand: two latent variables (i.e. leverage and regime) shall be filtered simultaneously in addition to the model's parameters, which need to be estimated for each firm. The proposed filter has numerous advantages. First, it is possible to use multiple data sources to help disentangle some effects that cannot be separated using a single product.² Second, using time series of derivatives captures the dynamics under both objective and risk-neutral probability measures simultaneously.

²For instance, liquidity issues in some tenors, recovery rate uncertainty, default risk, etc.

Finally, its numerical implementation is fast and efficient because it is deterministic, as opposed to particle filters and MCMC approaches. Since the estimation procedure is based on security prices, we also propose a numerical scheme based on trinomial lattices that allow an efficient pricing in the framework.

The model is then estimated with CDS premiums for each of the 225 firms of both CDX North American IG and HY indices. The estimated results are used in two empirical illustrations.

First, the effects of stochastic recovery and the presence of regimes on theoretical credit spread curves are investigated before, during and after the financial crisis. We find that recovery uncertainty and its negative relationship with default probability have a major impact on mid- and long-term credit spreads. In addition, the presence of regimes modifies the short-term shape of the average credit spread curves during the financial turmoil, especially for highly rated firms.

Second, hedging credit risk using CDSs is at the mercy of basis risk as different risk factors could impact bonds and CDS premiums. For instance, factors such as bond-specific liquidity could have an effect on bond spreads making the hedging strategy less effective. Therefore, using credit risk inferred from CDS premiums, we analyze the bond-CDS basis during the 2008 financial crisis. Based on regression tests and yield-to-maturity spreads implied by our credit risk model, we find that bond-specific liquidity is an important driver of the bond-CDS basis and its contribution is even more significant during the past crisis.

The rest of this paper is organized as follows. The joint default and loss model is presented in Section 2.2. Section 2.3 explains the model estimation approach and the pricing methodology. The estimation results are discussed in Section 2.4. In Section 2.5, the impact on spreads of endogenous random recovery and regime-switching risks is analyzed. The bond-CDS basis is investigated during the last crisis in Section 2.6. Finally, Section 2.7 concludes.

2.2 Credit Risk Model

In this paper, the shell approach of Boudreault et al. (2013) is extended to capture the time-varying nature of the leverage volatility.³ The shell is essentially an intensity process that depends on the firm's leverage, which is modelled by Markov-switching dynamics in this study. This flexible approach allows for an endogenous recovery rate that is both stochastic and negatively correlated with the firm's probability of default. Even though the extension looks straightforward, the estimation and CDS pricing procedures under the new Markov-switching generalization are cumbersome and require the implementation of sophisticated numerical methods.

Let L_t and A_t be the time t market value of the firm's liabilities and assets respectively. The leverage ratio is defined as the quotient of these two values.⁴ Because the asset volatility is not constant in time, the model should be flexible enough to capture the changes in the state variable dynamics. Hence, the market value of the firm's log-leverage is characterized by the following regime-switching dynamics:

$$\log\left(\frac{L_t}{A_t}\right) \equiv x_t = x_{t-1} + \left(\mu_{s_t}^{\mathbb{P}} - \frac{1}{2}\sigma_{s_t}^2\right)\Delta + \sigma_{s_t}\sqrt{\Delta}\epsilon_t^{\mathbb{P}}, \quad (2.1)$$

where s_t is the regime prevailing during the time interval $((t-1)\Delta, t\Delta]$, Δ is the time step between two consecutive observations, $\{\epsilon_t^{\mathbb{P}}\}_{t=1}^{\infty}$ is a sequence of independent standardized Gaussian random variables under the statistical probability measure \mathbb{P} , and $\mu_1^{\mathbb{P}}, \dots, \mu_K^{\mathbb{P}}, \sigma_1, \dots, \sigma_K$ are parameters that drive the leverage dynamics under the K possible regimes. The information structure is captured with the filtration $\{\mathcal{G}_t\}_{t=0}^{\infty}$ generated by the noise series $\{\epsilon_t^{\mathbb{P}}\}_{t=1}^{\infty}$ and the regimes $\{s_t\}_{t=0}^{\infty}$.

Because the leverage ratio is one of the potential drivers of default, it is incorporated in an intensity process $\{H_t\}_{t=0}^{\infty}$ that characterizes the potential default:

$$H_t = \beta + \left(\frac{L_t/A_t}{\theta}\right)^{\alpha} \quad (2.2)$$

where α, β and θ are positive constants. The default intensity increases with the leverage ratio, making the default more likely to happen. The intensity used in this study generalizes assumptions made by others. For instance, Bakshi et al. (2006a) use an

³Volatility is important in credit risk models. For instance, Zhang et al. (2009) show that volatility risk predicts 48% of the variation in the CDS spreads using a Merton-type structural model and a calibration approach.

⁴The leverage ratio L_t/A_t is not constrained to lie within the unit interval since L_t is the liabilities value and not the debt value. The liabilities value L_t could thus be larger than A_t .

intensity that depends linearly on the firm's leverage. In Doshi et al. (2013), the intensity is modelled as a quadratic function of the leverage.⁵ As usual in intensity-based models, the default time arises as soon as the intensity accumulation reaches a random level determined by an exponentially distributed random variable E_1 of mean 1 independent of $\{\mathcal{G}_t\}_{t=0}^\infty$:

$$\tau = \inf \left\{ t \in \{1, 2, \dots\} : \sum_{u=0}^{t-1} H_u \Delta > E_1 \right\}. \quad (2.3)$$

The recovery rate is determined by the firm's assets and liabilities at default. Following the intuition of structural models, it is characterized by the amount recovered by the creditors divided by the amount available upon default. At default, the creditors receive the amount the firm owe them (i.e. L_τ) or the asset value, minus legal and restructuring fees (i.e. $(1 - \kappa)A_\tau$), depending on how much each quantity is worth. Thus, the recovery rate at default time is proxied by

$$R_\tau = \frac{\min((1 - \kappa)A_\tau; L_\tau)}{L_\tau} = \min\left((1 - \kappa)\frac{A_\tau}{L_\tau}; 1\right) = \min((1 - \kappa)e^{-x_\tau}; 1), \quad (2.4)$$

where κ represents the legal and restructuring fees, expressed as a proportion of the asset value at default time. Consequently, an endogenous random recovery rate negatively correlated to default probabilities is implied by Equation (2.4).

The unobserved regime is modelled as a time-homogenous Markov chain with transition probabilities

$$p_{ij} = \mathbb{P}(s_t = j | s_{t-1} = i), \quad i, j \in \{1, 2, \dots, K\}. \quad (2.5)$$

The market model is incomplete, implying that there is an infinite number of pricing measures. Among these measures, we restrict the choices to those preserving the model structure:

$$x_t = x_{t-1} + \left(\mu_{s_t}^{\mathbb{Q}} - \frac{1}{2} \sigma_{s_t}^2 \right) \Delta + \sigma_{s_t} \sqrt{\Delta} \epsilon_t^{\mathbb{Q}} \quad (2.6)$$

and

$$q_{ij} = \mathbb{Q}(s_t = j | s_{t-1} = i), \quad i, j \in \{1, 2, \dots, K\} \quad (2.7)$$

where $\{\epsilon_t^{\mathbb{Q}}\}_{t=1}^\infty$ is a sequence of independent standardized Gaussian random variables under \mathbb{Q} .

⁵The quadratic intensity case is a specific case of Equation (2.2) where $\alpha = 2$.

2.3 Estimation Method and Pricing Technique

The model introduced in Section 2.2 is estimated on a firm-by-firm basis using MLE. The estimation is far from being straightforward since each firm's leverage ratios and statistical regimes are unobservable and should be inferred from observed quantities. Also, noisy observations lay additional stress upon the estimation procedure. To solve the problem at hand, a well-chosen combination of numerical methods and deterministic filters is employed.

First, a filtering method that simultaneously infers both latent quantities is needed. Theoretically, it would be possible to compute the filtered leverages given a regime path; however, this approach suffers from the curse of dimensionality. We extend Tugnait (1982) by allowing for non-linearities in the state-space representation, while considering only the most likely regime paths at each time step.

Second, the filter uses complex derivative prices (i.e. CDS) as measurements. Hence, a fast-pricing scheme must be implemented to ensure the feasibility of the estimation step.

2.3.1 Estimation Technique

A state-space representation is commonly used to define a model for which the state is a Markov process and observed quantities (i.e. CDS premiums) are related to the state variable. The log-leverage dynamics is given by Equation (2.1):

$$x_t = x_{t-1} + \left(\mu_{s_t}^{\mathbb{P}} - \frac{1}{2} \sigma_{s_t}^2 \right) \Delta + \sigma_{s_t} \sqrt{\Delta} \epsilon_t^{\mathbb{P}}.$$

The measurement equation becomes

$$y_t^{(i)} = \log \left(g^{(i)}(x_t, s_t) \right) + v_t^{(i)}, \quad (2.8)$$

with $y_t^{(i)}$ the observed i -year credit default swap log-premium, $v_t^{(i)}$ being a Gaussian random variable (having a standard deviation $\delta^{(i)}$) independent across maturities and $g^{(i)}$ the theoretical premium of an i -year credit default swap (in basis points). To reduce the number of parameters to be estimated, independence in the pricing error between each CDS has been assumed. Note that the CDS pricing formula (i.e. function $g^{(i)}$) is a nonlinear function of the log-leverage x_t . To maintain positive premiums, a logarithmic transformation is applied.

2.3.1.1 Unscented Kalman Filter-Based Method

If leverage time series were observable, the regimes could be painlessly filtered (for a review of classic methods, see Elliott et al., 1995). However, it is not the case and filtering regimes based on a latent time series is not straightforward.

Tugnait (1982) considers the problem of state estimation and system structure detection for discrete stochastic systems with parameters that may change among a finite set of values. His detection-estimation algorithm (DEA) is designed to filter a state and a regime simultaneously in the context of linear state-space representations. In this paper, the DEA is generalized to account for nonlinearities by running the UKF of Julier and Uhlmann (1997) instead of the classic Kalman (1960) filter.⁶ According to Christoffersen et al. (2014), the UKF may prove to be a good approach for a variety of problems in fixed income pricing. Moreover, for some applications in finance, it seems to significantly outperform the extended Kalman filter and behaves well when compared to the much more computationally intensive particle filter. To the best of our knowledge, this study is the first using the DEA coupled with the UKF.

The rationale behind the method is quite intuitive: if one could observe the regime path, the UKF could be applied directly. However, since it is not the case, the filtered value of the log-leverage would require the computation of a sum over each possible regime path. The number of terms in the summation would increase exponentially making its calculation futile. Since the regime process is Markovian and tends to forget its origin, a good approximation considers the most likely regime paths, which is Tugnait's (1982) idea. Precisely, we collapse from K^M to K^{M-1} regime paths by keeping the paths that yield the most likely sequences in terms of *a posteriori* probabilities. This would lead to minimum probability of error according to van Trees (1968).

This method is similar to the one proposed by Kim (1994). The main difference between the DEA and Kim (1994) is the collapsing scheme. The latter proposes to take averages across the different regime paths to reduce the number of sequences instead of keeping the most likely ones. Another related method would be the one of Fearnhead and Clifford (2003). They propose an optimal procedure in which all paths with probability greater than some cutoff are kept and the others are resampled (as in a conventional particle filter). However, since Fearnhead and Clifford (2003) rely on

⁶The parameters of the UKF technique have been assumed to be $\kappa_{\text{UKF}} = 2$, $\alpha_{\text{UKF}} = 0.1$, and $\beta_{\text{UKF}} = 0$.

particle filtering, the estimation using this stochastic filter would be more cumbersome.⁷

Suppose that at stage $t - 1$, K^{M-1} regime paths are available. At stage t and given the K^{M-1} paths, all possible extensions of these regime sequences shall be considered. For instance, all the possible extensions $s_{0:t}^j \equiv \{s_0^j, s_1^j, \dots, s_t^j\}$ of the a^{th} path $s_{0:t-1}^a$ are given by

$$s_{0:t}^j = \{s_{0:t-1}^a, s\}, \quad s \in \{1, 2, \dots, K\}.$$

Doing this for each of the K^{M-1} paths shall create K^M new regime paths. Then, using the UKF, it is possible to compute the filtered value of x_t based on the first t observations and the regime path $s_{0:t}^j$:

$$x_{t|t}^j \equiv \mathbb{E}^{\mathbb{P}}(x_t | \mathbf{y}_{1:t}, s_{0:t}^j)$$

where $\mathbf{y}_{1:t} \equiv \{\mathbf{y}_1, \mathbf{y}_2, \dots, \mathbf{y}_t\}$ and \mathbf{y}_t is a size n vector containing log-premiums $y_t^{(i)}$. Moreover, the posterior probability of observing this specific regime path based on the first t observations is

$$\mathbb{P}(s_{0:t}^j | \mathbf{y}_{1:t}) = \frac{f(\mathbf{y}_t | s_{0:t}^j, \mathbf{y}_{1:t-1}) \mathbb{P}(s_t^j | s_{0:t-1}^j, \mathbf{y}_{1:t-1}) \mathbb{P}(s_{0:t-1}^j | \mathbf{y}_{1:t-1})}{\sum_{k=1}^{K^M} f(\mathbf{y}_t | s_{0:t}^k, \mathbf{y}_{1:t-1}) \mathbb{P}(s_t^k | s_{0:t-1}^k, \mathbf{y}_{1:t-1}) \mathbb{P}(s_{0:t-1}^k | \mathbf{y}_{1:t-1})}$$

where $\mathbb{P}(s_t^j | s_{0:t-1}^j, \mathbf{y}_{1:t-1})$ can be computed using p_{ij} of Equation (2.5) and $s_{0:t}^k$ is the k^{th} extended path. In addition,

$$f(\mathbf{y}_t | s_{0:t}^j, \mathbf{y}_{1:t-1}) = \exp\left(-\frac{1}{2} \log(\det \mathbf{V}_t^j) - \frac{1}{2} (\mathbf{e}_t^j)^\top (\mathbf{V}_t^j)^{-1} (\mathbf{e}_t^j) - \frac{n}{2} \log(2\pi)\right),$$

$\mathbf{e}_t^j = \mathbf{y}_t - \mathbf{y}_{t|t-1}^j$ is the difference between the observations and the forecasted value of the observations at stage t , and $\mathbf{V}_t^j = \mathbb{E}[(\mathbf{e}_t^j)(\mathbf{e}_t^j)^\top | s_{0:t}^j, \mathbf{y}_{1:t-1}]$. This leads to the following state estimate:

$$x_{t|t} = \sum_{k=1}^{K^M} x_{t|t}^k \mathbb{P}(s_{0:t}^k | \mathbf{y}_{1:t}).$$

Finally, we collapse the K^M regime paths at time t to obtain K^{M-1} new regime paths according to Tugnait's (1982) idea. Simply put, we select the K^{M-1} regime paths yielding the highest posterior probabilities $\mathbb{P}(s_{0:t}^j | \mathbf{y}_{1:t})$ across the K^M remaining paths.

⁷Indeed, a change in the parameters can cause the filter to select a different set of particles inducing discontinuity in the likelihood function. The discontinuity is created from the resampling stage within the update state of the particle filter. For more details, see Creal (2012).

The quasi-likelihood function is given as a by-product of running the unscented Kalman filter:

$$\prod_{t=1}^T \left(\sum_{k=1}^{K^M} f(\mathbf{y}_t | s_{0:t}^k, \mathbf{y}_{1:t-1}) \mathbb{P}(s_t^k | s_{0:t-1}^k, \mathbf{y}_{1:t-1}) \mathbb{P}(s_{0:t-1}^k | \mathbf{y}_{1:t-1}) \right).$$

In our application, quasi-likelihood means that the first two moments of the posterior distribution are matched and a posterior Gaussian distribution has been assumed.

A simulation study (see Section 2.B) concludes that the DEA-UKF method yields virtually no differences between estimated and true parameter values on average. Moreover, the filtered values recovered by the filter are very close to their true values, with no differences between the two.

2.3.2 Credit-Sensitive Derivative Pricing

The derivative securities priced within this model take into consideration default, recovery and regime risks. However, stochastic recovery rates and regime-switching dynamics prevent us from using closed-form solutions for CDS premiums and corporate coupon bond prices. Even though Monte Carlo methods could be applied, lattice-based methods are more suited for the problem at hand since these yield accurate prices in a timely fashion.

The use of lattice methods in Markov-switching environment dates back to Bollen (1998) who proposes a pentanomial lattice. However, more recently, Yuen and Yang (2010) use a different idea: they adopt identical trinomial lattices across every regime, but change the risk-neutral weights as the regime state shifts. This would allow the trinomial trees to recombine across the different regimes. It works for many kinds of options, but credit-sensitive derivatives could not be priced directly using this methodology. Yuen and Yang's (2010) method is thus extended to price credit-sensitive securities (i.e. bond and CDS). An additional branch is added at each node to account for the potential default of the firm. This idea is similar to the one proposed by Schönbucher (2002). Details about the pricing scheme are provided in Section 2.A.

2.4 Estimation Results

Even though the question of whether CDS premiums include liquidity effects or counterparty risk is controversial (see Arora et al., 2012, Brigo et al., 2011, Bühler and Trapp, 2009b, Qiu and Yu, 2012, Tang and Yan, 2007), a large number of authors still use them to capture credit risk. Longstaff et al. (2005) argue that CDSs are of a contractual nature that affords relative ease of transacting large notional amounts compared to the corporate bond market. Moreover, an investor can liquidate a position by entering into a new swap in the opposite direction instead of selling his current position. Therefore, liquidity is less relevant given the ability to replicate swap cash flows using another CDS. Moreover, Mahanti et al. (2008), Dionne and Maalaoui Chun (2013) and Guarin et al. (2014) use CDS premiums as pure measures of credit risk.

We follow this literature and use CDS premiums since they are good proxies for credit risk. Moreover, the filtering approach adopted in this paper allows for potential lack of liquidity in specific tenors to be absorbed by the noise terms. Thus, the selected methodology would help to reduce the impact of illiquidity in our dataset, if ever there is any. Also, to minimize the impact of such risk, the study also focuses on firms that are part of the widely followed CDX indices.

2.4.1 Data

The investigation is performed on 225 firms of the CDX North American Investment Grade (IG) and High Yield (HY) indices provided by the Markit Group on September 20, 2013.⁸ The indices span multiple credit ratings and sectors. The weekly term structure of CDS premiums from January 5, 2005, to December 25, 2013, is also provided by Markit for a maximum of 469 weeks. CDS premiums up to the end of 2012 are used in the estimation; the last year of data (i.e. 2013) is kept for an out-of-sample analysis. We use Wednesday CDS premiums.⁹ Prices for maturities of one, two, three, five, seven and ten years are available for most firms (125 IG and 100 HY).¹⁰ This yields a grand total of 487,796 observations in our final sample of CDS

⁸CDS.NA.IG.21.V1 and CDS.NA.HY.21.V1 respectively.

⁹We focus on Wednesday CDSs because it is the least likely day to be a holiday and it is least likely to be affected by weekend effects. For more details on the advantages of using Wednesday data, see Dumas et al. (1998).

¹⁰However, 15 firms (4 IG and 11 HY) were removed from the sample since there were not enough observations for the estimation procedure (i.e. less than 100 observations per maturity).




Figure 2.1: Evolution of the mean five-year CDS premiums and of the difference between the average ten-year CDS premiums and the average one-year CDS premiums.

The CDS premiums were taken from both CDX.NA.IG.21.V1 and CDS.NA.HY.21.V1 portfolios, between January 2005 and December 2012. The grey surface corresponds to the financial crisis (July 2007 to March 2009).

premiums. The ‘No Restructuring’ clause is selected to capture only the credit risk of the firm.

To illustrate how CDS premiums move over time, the first panel of Figure 2.1 shows the evolution of the mean five-year CDS premium taken across firms for both IG and HY portfolios. Premiums were more or less stable during the pre-crisis era. They increased during the financial crisis: the mean five-year premiums jumped to a high of 321 basis points (bps) for IG firms and 1,422 bps for HY firms. In the post-crisis era, the premiums decreased, but did not reach their pre-crisis levels.

The second panel of Figure 2.1 shows the evolution of the average slope proxied by the difference between the average ten-year CDS premiums and the average one-year CDS premiums taken across firms for both IG and HY portfolios. The slope of the CDS premiums’ term structure became negative during the crisis for IG and HY firms on average. According to Figure 2.1, both the level and the slope of the CDS premiums change during the financial crisis. The regime-switching component of the proposed framework is required to capture these important changes in behaviour.

The initial state value x_0 has been set to the book value of the log-leverage using

Compustat.¹¹ Also, to price CDS, we follow standard industry practice as well as Carr and Wu (2010) and we use the interest rate curve defined by the LIBOR and swap rates. One-, two-, three- and six-month LIBOR are selected as well as one-, two-, three-, four-, five-, seven- and ten-year swap rates.¹²

Saunders and Allen (2010) decompose the recent financial crisis in three periods. The first period corresponds to the credit crisis in the mortgage market (June 2006 to June 2007), the second one covers the period of the liquidity crisis (July 2007 to August 2008) and the third period covers the default crisis period (September 2008 to March 2009). This study focuses on the second and the third periods; thus, the financial crisis started in July 2007 and finished in March 2009 throughout this paper. Notice that these deterministic sample's subperiods are not related to the firm-specific regimes anyhow.

2.4.2 Estimated Parameters

The leverage model used hereafter is a simplified version of Equation (2.1); we only consider two regimes. Considering more regimes is theoretically feasible; however, the number of parameters increase drastically and the numerical optimization of the firm-specific likelihood function becomes unmanageable. We also consider the same drift parameter across both regimes ($\mu^P = \mu_1^P = \mu_2^P$ and $\mu^Q = \mu_1^Q = \mu_2^Q$). Indeed, the drift parameter estimators of the latent variable is rather inaccurate and create numerical instability due to the short span of the time series used.¹³ Besides this caveat, the regime-switching framework still captures the variation of volatility across the different regimes (i.e. uncertainty).

The parameters are estimated using a quasi-MLE on a firm-by-firm basis. Overall, the parameters to be estimated for each company are $\phi = \{\mu^P, \mu^Q, \sigma_1, \sigma_2, p_{12}, q_{12}, p_{21}, q_{21},$

¹¹Initial leverages are approximated using their book values. More precisely, total liabilities divided by total assets is taken. Both quantities are acquired by Compustat, which is available from Wharton Research Data Services (WRDS). The fourth quarter of 2004's accounting data is selected to compute the proxy since Q4 of 2004 predates January 2005 (the beginning of our sampling period). In the database, the total liabilities are identified by LTQ and the total assets by ATQ. In addition, the firm's ticker symbol is matched to that of Markit's CDS premiums through the reference entity's name to ensure that the right information for each firm is used. For one firm, no data is available; this firm is thus removed from the sample.

¹²These rates are provided by the Federal Reserve of St. Louis website via FRED (Federal Reserve Economic Data). The series IDs are USD1MTD156N, USD2MTD156N, USD3MTD156N, USD6MTD156N, DSWP1, DSWP2, DSWP3, DSWP4, DSWP5, DSWP7 and DSWP10.

¹³Even in a 'one-regime' framework where the log-leverage is assumed to be observed, the precision of the drift parameter estimate is proportional to the square root of the sampling period length. Hence, long-time series are required to pin down these parameters.

$\alpha, \beta, \theta, \kappa, \delta^{(1)}, \delta^{(2)}, \delta^{(3)}, \delta^{(5)}, \delta^{(7)}, \delta^{(10)}$. Thus, a single set of parameters is used for each firm to explain the default risk and loss given default; this contrasts with calibration techniques where the credit default swap term structure is fitted at every available period.

Table 2.1 shows descriptive statistics on firm-specific estimated parameters. All the firms have a positive intensity process since β is always positive. The parameter α is always greater than 1.67, confirming the convex relationship between the intensity process and the leverage ratio. The volatility parameters σ_1 and σ_2 are very different one from another, on average. The low-volatility regimes normally correspond with the end of the pre-crisis era and the post-crisis period; its average value is around 11%. The dispersion across the different σ_1 is rather small (i.e. 4%). For most of the firms, the high-volatility regimes correspond to the financial crisis. The volatility is roughly 36% on average. The dispersion of the second regime volatility parameter is higher (i.e. 8%), which means that this parameter is quite dependent on the firm's condition during turmoil. The regimes are persistent since both \mathbb{P} - and \mathbb{Q} -transition probability matrices are concentrated over the main diagonal.

Noise term standard errors for mid- and long-term tenors are small on average: for instance, 13.2% for two-year maturity and 3.4% for five-year. It is the highest for the one-year CDS on average. Two reasons can explain this result: elements not necessarily related to the entity's true default and recovery risks, as well as fitting error due to model misspecification. Credit default swaps with a maturity of five years are the ones with the smallest noise standard errors on average. This makes sense empirically: five-year is known to be the most liquid tenor for CDS.

2.4.3 In- and Out-of-Sample Performances

The regime-switching hybrid default model (RS) is compared to three other models: the “one-regime” (1R) equivalent of our model, a regime-switching structural version (SA) of the proposed framework, and a regime-switching reduced-form model (RFA).

The so-called 1R model is the one introduced in Boudreault et al. (2013); the main difference between this model and ours is that we use regime-switching dynamics to model the firm's leverage.

Table 2.1: Descriptive statistics on the distribution of firm-specific parameters and noise terms across the portfolio of firms of the CDX indices.

Panel A: Descriptive statistics on leverage dynamics parameters (under \mathbb{P}) and fees κ .							
	$\mu^{\mathbb{P}}$	p_{12}	p_{21}	σ_1	σ_2	κ	Obs.
Mean	0.0169	0.0128	0.0190	0.1147	0.3572	0.6099	210
SD	0.0485	0.0161	0.0223	0.0376	0.0806	0.1728	
10%	-0.0087	0.0028	0.0048	0.0738	0.2526	0.4159	
25%	-0.0012	0.0038	0.0075	0.0934	0.3253	0.5000	
50%	0.0000	0.0083	0.0110	0.1111	0.3541	0.5819	
75%	0.0144	0.0149	0.0203	0.1351	0.3773	0.7126	
90%	0.0709	0.0274	0.0459	0.1614	0.4656	0.8474	
IG	0.0084	0.0132	0.0194	0.1089	0.3738	0.5825	121
HY	0.0285	0.0123	0.0184	0.1226	0.3346	0.6471	89
Panel B: Descriptive statistics on leverage dynamics parameters (under \mathbb{Q}) and intensity.							
	$\mu^{\mathbb{Q}}$	q_{12}	q_{21}	α	θ	β	Obs.
Mean	0.0165	0.0061	0.0397	11.4807	1.4385	0.0088	210
SD	0.0302	0.0064	0.0287	7.9199	0.3090	0.0225	
10%	-0.0050	0.0016	0.0081	6.0294	1.0724	0.0002	
25%	-0.0002	0.0028	0.0136	7.5511	1.3113	0.0007	
50%	0.0104	0.0046	0.0382	9.9483	1.4823	0.0029	
75%	0.0258	0.0074	0.0579	12.2609	1.5215	0.0083	
90%	0.0447	0.0123	0.0760	17.1065	1.7277	0.0170	
IG	0.0034	0.0060	0.0547	12.3016	1.3643	0.0066	121
HY	0.0342	0.0062	0.0193	10.3646	1.5394	0.0117	89
Panel C: Descriptive statistics on error terms.							
	$\delta^{(1)}$	$\delta^{(2)}$	$\delta^{(3)}$	$\delta^{(5)}$	$\delta^{(7)}$	$\delta^{(10)}$	Obs.
Mean	0.2503	0.1318	0.0810	0.0339	0.0470	0.0762	210
SD	0.0647	0.0362	0.0297	0.0263	0.0400	0.0460	
10%	0.1749	0.0904	0.0420	0.0002	0.0003	0.0282	
25%	0.2049	0.1063	0.0665	0.0138	0.0095	0.0366	
50%	0.2446	0.1300	0.0817	0.0353	0.0457	0.0690	
75%	0.2972	0.1559	0.1023	0.0452	0.0661	0.1001	
90%	0.3442	0.1763	0.1152	0.0706	0.1061	0.1425	
IG	0.2310	0.1253	0.0768	0.0335	0.0657	0.0964	121
HY	0.2765	0.1407	0.0867	0.0345	0.0217	0.0487	89

For each of the 210 firms, the parameters of the model are estimated using weekly CDS premiums of maturities one, two, three, five, seven and ten years, using the DEA-UKF filtering technique. The mean, standard deviation (SD) and quantiles are computed across firms. The last two rows compute the mean across firms of CDX.NA.IG.21.V1 and CDX.NA.HY.21.V1 portfolios. The δ s represent the standard deviation of the noise terms present in the observation equation of the filter. IG refers to Investment Grade and HY means High Yield.

Letting $\beta = 0$ and $\alpha \rightarrow \infty$, the hybrid approach leads to a pure structural framework in which the log-leverage ratio is modelled by a regime-switching discretized version of an arithmetic Brownian motion. The structural model is nested in the hybrid framework; thus, its overall in-sample performance is expected to be weaker than the full model. However, it is possible that local performances dominate those of the full model.

The pure reduced-form version of the model (i.e. $\theta \rightarrow \infty$) is not the one used in this

paper since this model would have constant intensity and this is obviously too restrictive. Instead, the log-intensity process is modelled by a regime-switching discretized version of an arithmetic Brownian motion under the \mathbb{Q} measure:

$$\log(H_t) = \begin{cases} \log(H_{t-1}) + \left(\mu_H^{\mathbb{Q}} - \frac{1}{2}\sigma_{H,1}^2\right)\Delta + \sigma_{H,1}\sqrt{\Delta}\epsilon_t^{\mathbb{Q}} & \text{if } s_t = 1 \\ \log(H_{t-1}) + \left(\mu_H^{\mathbb{Q}} - \frac{1}{2}\sigma_{H,2}^2\right)\Delta + \sigma_{H,2}\sqrt{\Delta}\epsilon_t^{\mathbb{Q}} & \text{if } s_t = 2 \end{cases}$$

where $\mu_H^{\mathbb{Q}}$ is the drift parameter and $\sigma_{H,1}$ and $\sigma_{H,2}$ are volatility parameters for regimes 1 and 2, respectively. Moreover, $\{\epsilon_t^{\mathbb{Q}}\}_{t=1}^{\infty}$ is a sequence of independent standardized Gaussian random variables under \mathbb{Q} . Since endogenous recoveries no longer make sense in the context of this model, we opt for a constant exogenous recovery rate $R_t = R \in [0, 1]$. This new parameter is estimated among all other parameters in the filtering procedure.

The models' performances are compared using the sum of squared errors:

$$\text{SSE} = \sum_{t=t_1}^{t_2} \sum_{i=1}^{N_t} (m_{t,i} - o_{t,i})^2 \quad (2.9)$$

where $m_{t,i}$ is the theoretical i -year CDS premium at time t , $o_{t,i}$ is the observed i -year CDS premium at time t , and N_t is the number of CDSs considered at time t . The parameter estimates are obtained using the entire sample and are kept fixed at any point in time. To standardize these fitting performances, the SSE is divided by the total variation of the observed CDS premiums:

$$\text{SST} = \sum_{t=t_1}^{t_2} \sum_{i=1}^{N_t} (o_{t,i} - \bar{o})^2 \text{ where } \bar{o} = \frac{1}{\sum_{t=t_1}^{t_2} N_t} \sum_{t=t_1}^{t_2} \sum_{i=1}^{N_t} o_{t,i}. \quad (2.10)$$

Consequently, the performance measure is the ratio SSE/SST . Better models have ratios closer to zero.¹⁴

2.4.3.1 In-Sample Performance

The in-sample study is performed over 417 weeks, starting in January 2005 and ending in December 2012.

¹⁴We also apply a loose criterion to eliminate outliers: if the absolute difference between the model premium and the observed premium is more than 5 times the observed premium, it is considered an outlier: $|m_{t,i} - o_{t,i}| > 5o_{t,i}$. For the in-sample analysis, seven observations out of 489,796 observations are removed. Thus, we remove 0.0014% of our initial CDS sample.

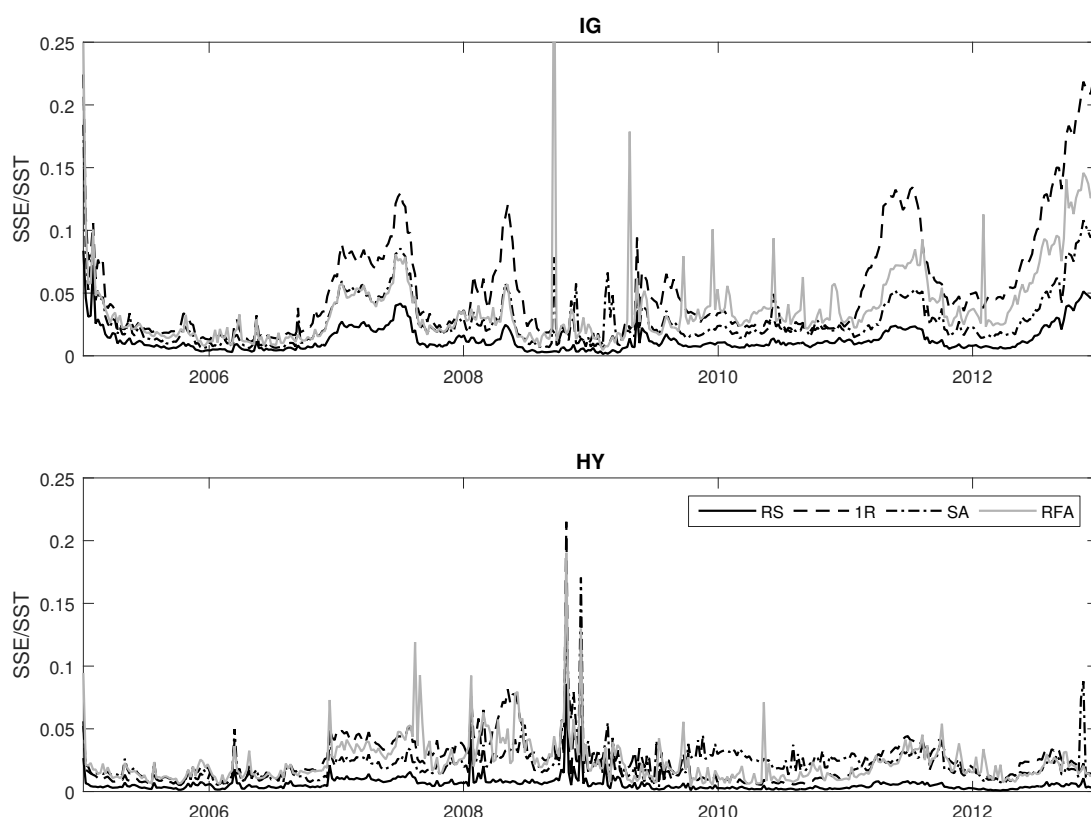


Figure 2.2: SSE/SST: Evolution of the in-sample performance for IG and HY.

These figures show the evolution of the SSE/SST ratio calculated by week. For each firm, a single set of parameters is estimated using CDS premiums with maturities 1, 2, 3, 5, 7, and 10 years between January 2005 and December 2012. IG stands for Investment Grade and HY for High Yield. RS stands for regime-switching hybrid model, 1R means ‘one-regime’ equivalent of our model, SA is the regime-switching structural model and RFA refers to the regime-switching reduced-form model. Notice that on September 17, 2008, the three benchmark models seem to yield large pricing errors on IG firms. This is mainly caused by errors on AIG’s credit default swap premiums (on September 15, 2008, AIG’s credit rating was downgraded from AA- to A-; the downgrade had an important impact on CDS premiums).

Figure 2.2 shows the evolution of the SSE/SST ratio calculated by week. The measure used in these figures combines the six tenors. The in-sample performance of the full model is very good with an average SSE/SST of 3.33%. It is better when compared to the other models: indeed, the performance measure of the full model is 1.83 times lower than the one of the reduced-form approach, 2.05 than the structural model, and 3.45 than the ‘one-regime’ equivalent. During the financial crisis of 2008, the SSE/SST ratio spikes when we consider the three benchmark models. However, the full model does well during these turbulent times: there are no apparent spikes in SSE/SST for RS during the crisis era.

Panels A, B and C of Table 2.2 shows the ratio calculated by maturity and by period. The performance of the regime-switching hybrid model remains good, even during the

crisis era. Systematically, the full model outperforms the other ones which means that each risk involved in the full model is important to capture credit risk. For IG firms, the performance is sound, even for one-year CDS premiums.

Table 2.2: SSE/SST: in- and out-of-sample performance for IG and HY.

Panel A: Pre-crisis (in-sample)								
	IG				HY			
	RS	1R	SA	RFA	RS	1R	SA	RFA
1	0.089	0.176	0.171	0.162	0.023	0.065	0.081	0.051
2	0.023	0.053	0.050	0.045	0.014	0.030	0.029	0.026
3	0.007	0.019	0.022	0.016	0.005	0.011	0.010	0.013
5	0.003	0.012	0.012	0.008	0.002	0.004	0.008	0.006
7	0.001	0.008	0.010	0.008	0.001	0.006	0.014	0.007
10	0.004	0.018	0.021	0.017	0.004	0.014	0.029	0.017
Panel B: Crisis (in-sample)								
	IG				HY			
	RS	1R	SA	RFA	RS	1R	SA	RFA
1	0.046	0.091	0.098	0.103	0.142	0.527	0.263	0.233
2	0.021	0.045	0.037	0.043	0.084	0.274	0.189	0.162
3	0.009	0.020	0.017	0.022	0.022	0.067	0.050	0.043
5	0.001	0.007	0.004	0.017	0.002	0.007	0.007	0.006
7	0.005	0.022	0.010	0.026	0.001	0.004	0.005	0.007
10	0.013	0.051	0.027	0.051	0.005	0.014	0.015	0.017
Panel C: Post-crisis (in-sample)								
	IG				HY			
	RS	1R	SA	RFA	RS	1R	SA	RFA
1	0.055	0.150	0.197	0.118	0.089	0.176	0.171	0.162
2	0.024	0.064	0.076	0.051	0.023	0.053	0.050	0.045
3	0.014	0.033	0.041	0.030	0.007	0.019	0.022	0.016
5	0.004	0.010	0.015	0.011	0.003	0.012	0.012	0.008
7	0.002	0.009	0.015	0.012	0.001	0.008	0.010	0.008
10	0.011	0.030	0.038	0.034	0.004	0.018	0.021	0.017
Panel D: 2013 (out-of-sample)								
	IG				HY			
	RS	1R	SA	RFA	RS	1R	SA	RFA
1	2.017	5.156	1.974	2.183	0.139	0.189	0.129	0.156
2	0.679	0.949	0.885	0.722	0.080	0.125	0.100	0.099
3	0.261	0.259	0.349	0.289	0.045	0.066	0.075	0.051
5	0.032	0.084	0.028	0.021	0.038	0.046	0.040	0.046
7	0.072	0.203	0.061	0.054	0.041	0.062	0.043	0.053
10	0.100	0.288	0.092	0.110	0.050	0.081	0.060	0.070

This table shows the SSE/SST ratio calculated by maturity and period (i.e. pre-crisis, crisis, post-crisis, 2013). For each firm, a single set of parameters is estimated using CDS premiums with maturities of 1, 2, 3, 5, 7 and 10 years between January 2005 and December 2012. IG stands for Investment Grade and HY for High Yield. RS stands for regime-switching hybrid model, 1R means ‘one-regime’ equivalent of our model, SA is the regime-switching structural model and RFA refers to the regime-switching reduced-form model.

2.4.3.2 Out-of-Sample: Forecasting 2013

In our out-of-sample study, the parameters are estimated using CDS premiums from January 2005 to December 2012. Then, CDS premiums observed in 2013 are used to

evaluate the out-of-sample fit of the model. The out-of-sample measures are calculated for every CDS premium observed in 2013.¹⁵

Panel D of Table 2.2 shows that our new model produces adequate one-week-ahead forecasts. The full model's curves are lower than the other ones for most maturities and risk classes, which is good. For IG, SSE/SST ratios are higher for short maturities. The main reason for this behaviour is that 1-year SST is much smaller than the other SSTs, even though the 1-year SSE is somewhat smaller than those computed for other maturities. Therefore, it is natural to observe high SSE/SSTs for 1-year credit default swaps in 2013.

2.4.4 Most Likely Statistical Regimes Through Time

The approach of Viterbi (1967) is adapted to the context of a hidden regime and a latent variable to extract the most probable regime path for each firm. The regime path that maximizes the likelihood function given the estimated parameters is constructed recursively. At this point, we find relevant to stress that the filtered statistical regimes depend on firm-specific information and shall not only be related to financial or economic cycles. Therefore, one should not misinterpret the notion of filtered regime: these are computed for each firm and are solely related to the firm's financial health. By taking the average proportion of firms in the high-volatility regime, we hope to find a systematic trend across the firms, although it is possible that some companies were only sparsely affected by the crisis or that some were in precarious positions during pre- and post-crisis eras.

Figure 2.3 shows the proportion of firms considered in the high-volatility regime for each week and for each risk class. In our context, being in the second (high-volatility) regime is synonymous with more uncertainty in the firm's leverage. The fast transition from the first (low-volatility) regime to the second (high-volatility) regime during 2007 and 2008 is obviously very natural due to the financial conditions at that time. At the beginning of the sample, many HY firms are in the high-volatility regime, meaning that the filtered leverage is more uncertain for these firms. Note that uncertainty is typically related to the slope of the CDS term structure; for instance, in a study based on sovereign CDS, Augustin (2012) argues that slopes tend to be positive in good times and invert during economic distress. We compare the slopes of the term structure of

¹⁵As before, observations satisfying the rejection criterion are removed from the analysis. Overall, 41 out of 67,373 observations are discarded since they are considered to be outliers by at least one of the four models.

CDS premiums (i.e. ten-year minus one-year CDS premiums) during the pre-crisis era and find that they tend to be smaller when firms are in the second regime. Therefore, the filtered statistical regimes are consistent with this observation. Moreover, the beginning of our sample corresponds to a period of general widening in yield spreads in debt markets due to the downgrade of General Motors and Ford's ratings (Saldías, 2013).




Figure 2.3: Proportion of firms in the high-volatility statistical regime.

Each week, the proportion of firms considered in the second regime is computed for the 121 IG and the 89 HY ones. The grey surface corresponds to the financial crisis (July 2007 to March 2009).

For HY firms, the moment when the proportion of firms in the second regime increased corresponds to the beginning of the financial crisis (i.e. July 2007). During the first six months of the crisis, the proportion of firms in the high-volatility regime goes from 8% to almost 60%. For IG firms, the transition happens a little later and is consistent with the beginning of the NBER economic recession in the United States (i.e. December 2007). At the end of the financial crisis (or the recession), both IG and HY firms remain in the high-volatility regime for several weeks. These results are in line with the recent literature. Indeed, using a different approach, Maalaoui Chun et al. (2014) detect some persistence in the volatility regimes of credit spreads. They are also consistent with those of Garzarelli (2009) and Mueller (2008). These authors document that credit spread levels increase before the onset of the NBER recession and somehow persist until long after the recession is over.

2.5 Term Structure Comparison: Recovery Rate Versus Regime-Switching

This section assesses the relative importance of recovery rate and regime-switching risks in credit spread curves. Essentially, we compare the difference between risky and riskless zero-coupon yields for various model specifications of the model.

The time t value of risky zero-coupon bond (see Section 2.A.2) is given by $V(t, T, \widehat{x}_t, \widehat{s}_t; R_\tau, 0)$ where \widehat{x}_t is the time t filtered log-leverage and \widehat{s}_t is the most probable regime at time t .¹⁶

To compare the relative importance of recovery-rate risk and regime-switching risk, four different formulations of the model are estimated. The first one (RS-ENDR) is described in Section 2.2. The second model (RS-EXOR) is a variation of the full model: instead of using endogenous recovery rates, a constant exogenous recovery rate $R_t = R \in [0, 1]$ is incorporated into the set of parameters to be estimated. The third model is the “one-regime” equivalent with endogenous recovery rates (1R-ENDR) presented in Boudreault et al. (2013). The last model (1R-EXOR) is a ‘one-regime’ equivalent of the full model with a constant exogenous recovery rate to be estimated.

In this section, the results are broken down by credit rating. These ratings are attributed by Standard and Poor’s between January 31, 2005, and December 31, 2012, and are available from Compustat in WRDS. The rating used is identified as ‘S&P Domestic Long Term Issuer Credit Rating’ (SPLTICRM) in the database.¹⁷

Figure 2.4 shows average credit spreads across credit ratings and models.¹⁸ There has been an important rise in average credit spreads during the financial turmoil that affects mainly the short and mid terms. According to any of the tested models, the two-year average credit spread of A-rated firms is about 10 times larger during the crisis period than before. For riskier firms rated BBB, BB and B, the average credit spread is about 4 times larger. When considering the five-year credit spread, it is 4 times larger for

¹⁶Normally, we would need to take a weighted average of the different prices, given the regime at t ; the weights correspond to $\mathbb{Q}(s_t = s | \mathbf{y}_{1:t})$, the risk-neutral probability of being in regime s at time t conditional on $\mathbf{y}_{1:t}$. However, these probabilities cannot be computed readily, and this is why we use the most probable regime at time t as a proxy. We could also use the probability under the physical measure as a proxy; however, this shall lead to the same estimate since the probability of being in a regime is most of the time in the neighbourhood of 1 or 0.

¹⁷The firm’s ticker symbol is matched to the data in Compustat. For two firms, no rating is available, leaving a sample of 208 firms whenever results are presented by credit rating.

¹⁸Curves obtained from AAA- and AA-rated firms are not shown in the figure since the number of firms with such rating is too small to be representative.

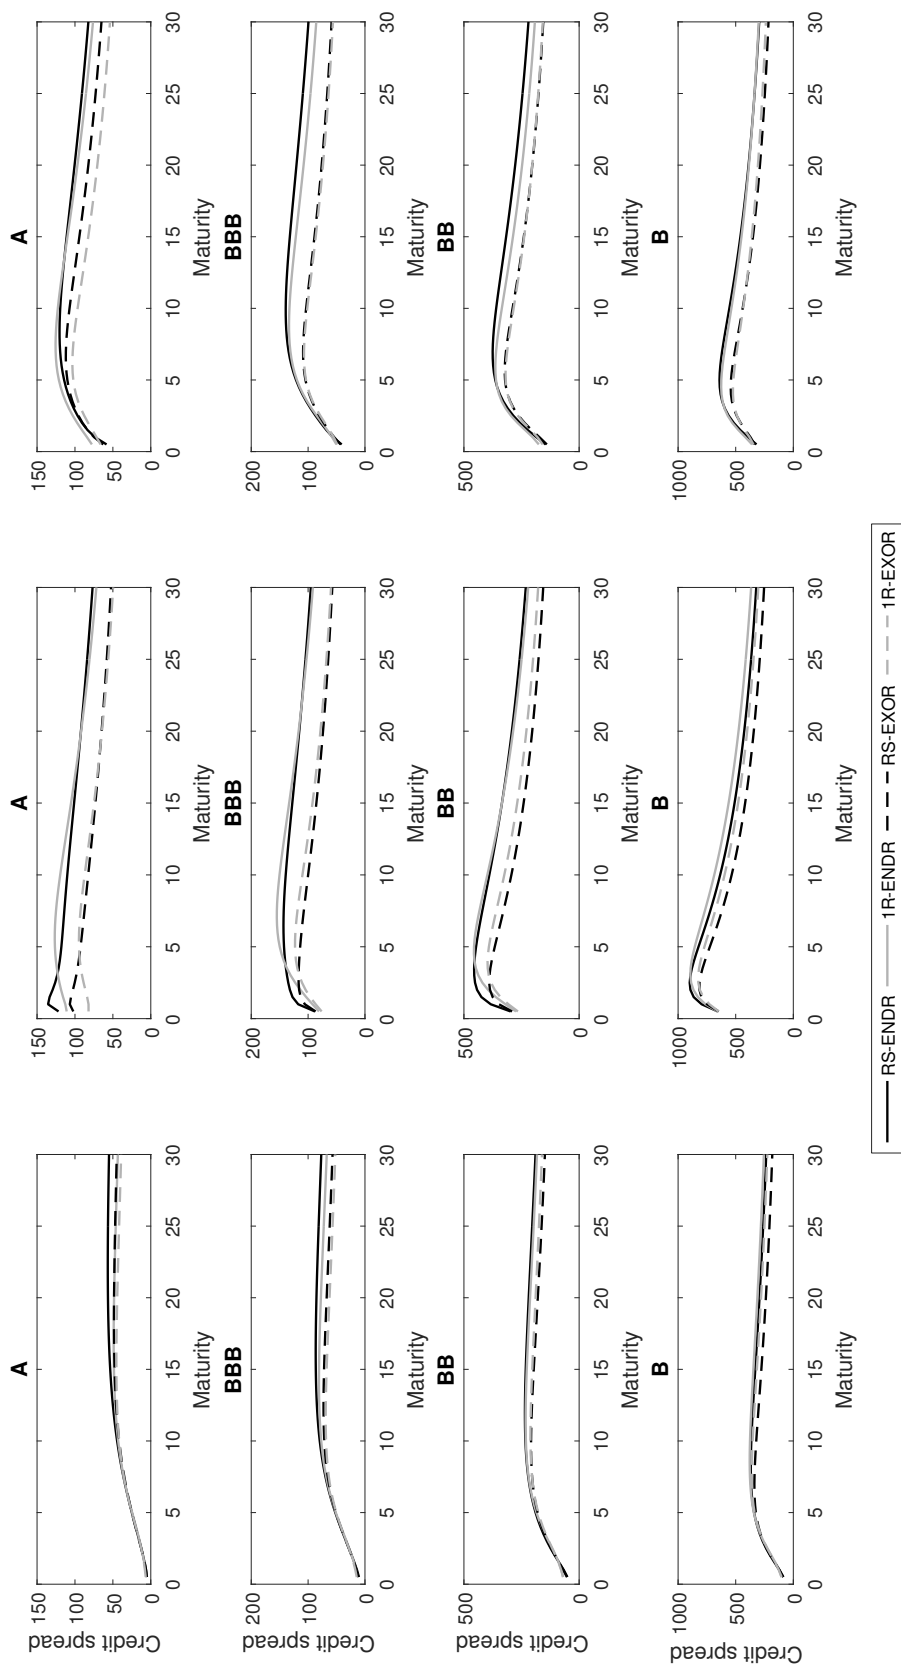


Figure 2.4: Credit spreads in basis points as a function of maturities, credit ratings and eras.

For each of the 208 firms, the parameters of the model have been estimated using weekly CDS and the DEA-UKF technique. The data are available for January 2005 to December 2012. The term structures are evaluated each week. Then, we average the spreads across the time periods (i.e. pre-crisis, crisis, post-crisis era). Each graph shows average term structures for four different models. Two versions of our regime-switching model are used: RS-ENDR and RS-EXOR. Moreover, two versions of the one-regime equivalent are used: 1R-ENDR and 1R-EXOR.

the A-rated firms and about 2.5 times larger for the credit ratings BBB, BB and B. There is a reduction of the credit spread in the post-crisis period, yet it never reaches the pre-crisis levels.

Recovery uncertainty and its negative relationship with default probability have a major impact on mid- and long-term credit spreads. Indeed, in a regime-switching environment, the five-year credit spread increases by 11% (from 23 to 34%) with the endogenous recovery assumption when compared to an equivalent model with a constant recovery rate during the crisis. On a longer horizon, the effect is even more significant, ranging between 12% and 69% depending on the credit rating or the period considered. The dependence between default probabilities and recovery rates have been documented empirically (i.e. lower-rated firms have lower recovery rates), but most modelling approaches neglect the accrued risk associated with this negative relationship. Thus, assuming constant recovery rates seriously impacts credit spread curves, especially over the long run.

The presence of regimes has a second-order effect on average credit spread curves and is mainly attributed to how these averages are constructed. The parameters of the ‘one-regime’ model capture the behaviour during good and bad times. Unlike the previous approach, the two-regime model allows for a distinct set of parameters for each state. However, the curves presented in Figure 2.4 are constructed by averaging all weekly curves on a given period and correspond to a mix of various firms and regime weights. Even then, the presence of regimes affects the long-term credit spread of highly rated firms in the pre- and post-crisis periods. During the financial turmoil, the presence of regimes modifies the short-term shape of the average credit spread curves, especially for highly rated firms.

2.6 Bond-CDS Basis

Many investors, such as insurers and pension funds, were considerably affected by the last crisis as their liabilities mainly comprise long-term commitments funded by long-term fixed income instruments like corporate bonds. According to a NAIC¹⁹ special report on the insurance industry’s use of derivatives, credit risk is hedged primarily with credit default swaps (NAIC, 2015). Generally, one would buy some single-name

¹⁹National Association of Insurance Commissioners.

CDSs that hedge against default of a reference entity.²⁰ However, this scheme is subject to basis risk as distinct risk constituents could impact bond and CDS spreads.

The bond-CDS basis measures the extent to which these two spreads differ from one another. Broadly speaking, the basis is computed as the difference between measures of the CDS spread and the bond spread, both with the same maturity dates. Therefore, a deeper understanding of the bond-CDS basis starts with a better understanding of credit risk and how the latter is involved in the pricing of both bonds and CDSs. As noted in Section 2.4, the proposed model outperforms classic benchmarks and yields small errors on CDS premiums. In this section, we use the proposed model (along with the estimated parameters) to adequately capture CDS credit risk and compute accurate estimates of credit spreads.

In the literature, most researchers focused on market-wide and firm-specific factors. Fontana (2010) finds that the basis is time-varying and negatively correlated with the LIBOR–OIS²¹ spread and OIS–Treasury bill spread, which are proxies for the increased funding costs and flight-to-liquidity, respectively. Bühler and Trapp (2009a) explore the impact of bond-specific liquidity measures by using an indirect proxy. As a matter of fact, they use the bond yield volatility of a specific portfolio as a proxy for the portfolio’s liquidity.

Yet, direct bond-specific drivers have not been investigated much. Bai and Collin-Dufresne (2011) use the bond yield bid-ask spread, bond liquidity beta and bond liquidity market beta as measures of the bond-specific liquidity. They find that the negative basis is mainly explained by bond trading liquidity risk during the crisis. In this spirit, we propose to evaluate the importance of bond-specific liquidity in the basis. However, unlike Bai and Collin-Dufresne (2011), we use different bond-specific liquidity and bond-CDS basis measures.

2.6.1 Measuring the Basis

As stated in Elizalde et al. (2009), the three most common bond spreads to compute the bond-CDS basis are the Z-spread, the par asset swap spread (ASW) and the par equivalent CDS spread (PECS). The first two have features that make them difficult to compare properly to CDS premiums. As a matter of fact, they do not take into account

²⁰Fabozzi and Mann (2012) state that credit derivatives are mainly used to transfer and hedge credit risk.

²¹Overnight indexed swap.

expected recovery rates and the term structure of the probabilities of default, while PECS does.

These three measures take the CDS premiums as-is and subtract a computed bond spread, given some assumptions. However, to benefit from our credit risk model, we need to compute the basis the other way around. Actually, using the model introduced above, bond-specific theoretical YTM spreads can readily be obtained.²² Since the model is estimated exclusively on CDS premiums, this would give us a bond-equivalent of the risk embedded in the firm's CDSs, while using the exact bond maturity date. To obtain a measure of the bond-CDS basis, we subtract the bond observed YTM spread from the model's bond-equivalent YTM spread:

$$\text{Bond-CDS basis} = \underbrace{\text{Bond-equivalent YTM spread}}_{\text{From the model}} - \underbrace{\text{Bond YTM spread}}_{\text{From the market}} \quad (2.11)$$

This measure of the bond-CDS spread obviously accounts for recovery rates and the term structure of the default probabilities as does our credit risk model.

2.6.2 Data

We obtain bond information from Bloomberg and from Mergent Fixed Income Securities Database (FISD). The same 225 firms are considered. In order to find the right firms that match our CDS premiums, the tickers are manually matched. The selected bonds are senior, non-callable, non-putable bullet bonds with fixed coupon rates. Then, the trading data are acquired by the Trade Reporting And Compliance Engine (TRACE). We keep issues with at least 100 trades during the period considered (i.e. from 2005 to 2012).

Dick-Nielsen's (2009) algorithm is used to filter out the errors in TRACE. Omitting this step might result in high liquidity biases: if TRACE data are not cleaned up before use, the number of transactions will be too high.²³ Our final bond subsample contains 1,046 issues from 97 firms (66 IG and 31 HY), for a total of 2,264,566 observations.

²²The difference between a yield-to-maturity of the corporate bond and the linearly interpolated maturity-matched risk-free rate calculated on the same day.

²³The filter is divided into three steps. First, true duplicates are deleted (i.e. intra-day trades with the same unique message sequence number). Then, reversals (a trade cancellation for a trade report that was originally submitted to TRACE on a previous date) are also deleted. Finally, same-day corrections are deleted. These are identified using the report's trade status.

To obtain the spreads (in percent) from our bond sample, LIBOR and swap rates are used. We use one-, two-, three- and six-month LIBOR as well as one-, two-, three-, four-, five-, seven-, ten- and 30-year swap rates. We linearly interpolate the different rates to obtain the corresponding rate.²⁴ These rates are provided by the Federal Reserve of St. Louis website via FRED. Dick-Nielsen et al. (2012) also use swap rates to proxy risk-free rates.

Bond YTM's are computed on clean prices from which the corresponding maturity-matched risk-free rate is removed. Then, the daily spreads are averaged over each month. Note that we winsorize the 0.5% higher and lower values of observed YTM spreads.

Bond-equivalent YTM spreads are obtained using the regime-switching hybrid credit risk model. One spread is computed for every week (since the time step of our estimation method was $\Delta = 1/52$) and the monthly theoretical spread is the average of the weekly ones. Like CDS premiums, bond prices are numerically computed using the trinomial lattice (see Subsection 2.3.2).

Table 2.3: Descriptive statistics on the bond-CDS basis and the bond-specific liquidity measure λ .

	Basis			λ		
	All	IG	HY	All	IG	HY
Mean	-0.7483	-0.6689	-1.1839	-0.0178	-0.0162	-0.0261
SD	1.6004	1.2754	2.7325	3.0399	3.0138	3.1736
5%	-2.8854	-2.4519	-6.1766	-3.1579	-3.1709	-3.0787
10%	-1.9839	-1.8242	-3.6912	-2.7747	-2.7887	-2.7214
25%	-1.1719	-1.1283	-1.6904	-1.9154	-1.9080	-1.9506
50%	-0.6284	-0.6265	-0.6419	-0.7667	-0.7449	-0.8687
75%	-0.1686	-0.1807	-0.1103	0.9018	0.9198	0.7994
90%	0.4202	0.3993	0.5945	3.4023	3.4081	3.3875
95%	1.3282	1.2464	1.7583	5.7441	5.6507	6.1544

This table shows statistics for the bond-CDS basis (in percent) and the bond-specific liquidity measure λ . The two quantities are calculated monthly for each bond from January 2005 to December 2012. The mean, standard deviation (SD) and quantiles are computed across firms. IG refers to Investment Grade and HY means High Yield.

The leftmost columns of Table 2.3 show the descriptive statistics for the bond-CDS basis as computed by Equation (2.11). The average basis is negative for both IG and HY bonds. The standard deviation is greater for HY bonds than for IG. Top panels of Figure 2.5 exhibits the average monthly bond-CDS basis for both IG and HY bonds. The basis is negative for IG firms throughout the sample; there is a decrease in the basis during the 2005–2012 period, although this decline is somewhat small. For HY

²⁴The rate that matches the product's maturity.

bonds, the basis is also negative for most months. There is a strong decrease in the basis during the crisis, on average. At the end of the crisis, the bond-CDS basis tends to increase to reach pre-crisis levels in 2011.

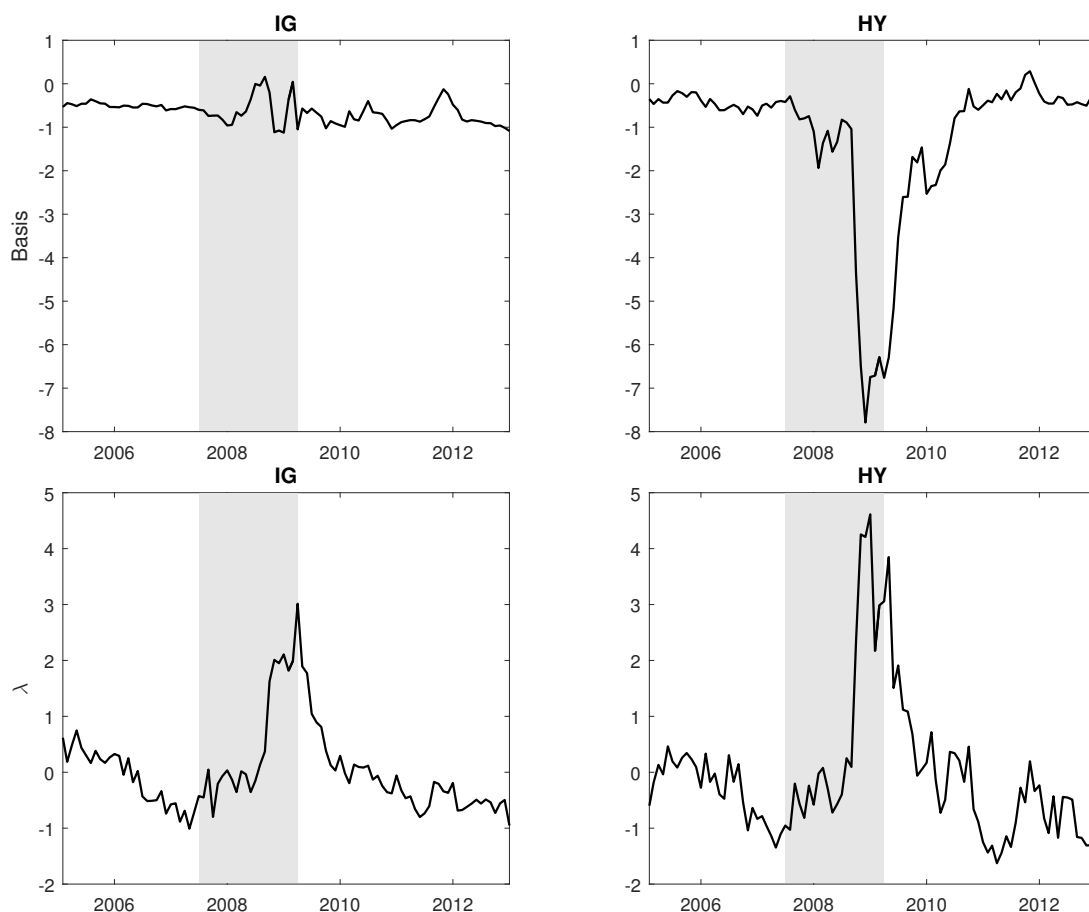


Figure 2.5: Time series of the average bond-CDS basis (in percent) and the average composite λ measure for IG and HY.

For each month, the average over each variable is computed across our dataset of 1,046 bond issues (paying no attention to the maturity of the bonds). To proxy the bond-specific liquidity, we use Dick-Nielsen et al. (2012)'s λ measure in the study. The latter is a construction made of four different liquidity proxies: the Amihud (2002) measure, the Amihud risk, the imputed round-trip cost (IRC) and the IRC risk. The liquidity proxy is computed on a monthly basis from January 2005 to December 2012.

2.6.3 Bond-Specific Liquidity as a Driver of the Basis

Liquidity is usually vaguely defined as the degree to which an asset or security can be bought or sold in the market quickly without affecting the asset's price (Tang and Yan, 2007). However, it has multiple facets and cannot be defined efficiently by a single statistic. There are three oft-cited dimensions of the liquidity risk. The first one is tightness: if the bid-ask spread is small, it is assumed that the market is liquid. The

second dimension is depth: it is related to the amount of security that can be traded without affecting the price. Finally, the third dimension is called resiliency: a market is liquid in this specific dimension if price recovers quickly after a demand or supply shock.

Many recent papers focus on liquidity issues in the bond market. Using a reduced-form approach, Longstaff et al. (2005) use credit default swap premiums to measure the size of the default component in corporate spreads. The authors find that the majority of the spread is due to default risk and that the non-default component is time varying and strongly related to measures of bond-specific illiquidity. Dick-Nielsen et al. (2012) analyze corporate bond spreads during 2005–2009 using a new robust illiquidity measure based on four different proxies. They show that the spread contribution from illiquidity increases dramatically with the onset of the subprime crisis. This effect is slow and persistent for investment-grade bonds while it is stronger but shorter-lived for speculative-grade bonds.

To proxy the bond-specific liquidity, we use Dick-Nielsen et al.'s (2012) λ measure in the study. The latter is a construction made of four different liquidity proxies: the Amihud (2002) measure, the Amihud risk, the imputed round-trip cost (IRC) and the IRC risk. To be precise, each proxy is normalized and then summed to create the new λ measure. Note that the liquidity proxy is computed on a monthly basis.

To be certain that this measure is robust, seven different liquidity proxies used in Dick-Nielsen et al. (2012) are computed and a principal component analysis is applied on these variables.²⁵ The first component explains 77% of the variation and is positively correlated with Dick-Nielsen et al.'s (2012) λ measure (i.e. 55%). Therefore, it seems fair to use this variable as a liquidity proxy.

The rightmost columns of Table 2.3 show the descriptive statistics of the liquidity proxy λ . The composite λ measure has a zero mean by construction; however, its distribution seems highly asymmetrical. The 5th percentile is around -3.18 and the 95th percentile is at 5.74. Note that the same asymmetric behaviour is true for the four constituents of the λ measure. Bottom panels of Figure 2.5 shows the average value of the liquidity proxy through time for both IG and HY bonds. For both risk classes, the average λ increases during the crisis, although the increase for HY is much more severe. The increase for HY bonds is consistent with the decrease of the basis on average: as the average basis decreases, the bond-specific liquidity measure increases.

²⁵The Amihud (2002) measure, the Amihud risk, the imputed round-trip cost (IRC), the IRC risk, the Roll (1984) measure, the turnover rate of a bond and the proportion of zero trading days. For more details on these variables, see Dick-Nielsen et al. (2012) and Section 2.C.

For IG, the relationship is less clear when we aggregated issue-specific bases and proxies.

To assess the importance of bond-specific liquidity, a dummy variable regression using monthly observations is run. In the latter, we regress the liquidity proxy on the bond-CDS basis by issue:

$$(\text{Bond-CDS basis})_{it} = \gamma_0^{(i)} \mathbb{I}_{\{\text{Pre-crisis}\}}(t) + \gamma_1^{(i)} \mathbb{I}_{\{\text{Pre-crisis}\}}(t) \lambda_{it} + \gamma_2^{(i)} \mathbb{I}_{\{\text{Crisis}\}}(t) + \gamma_3^{(i)} \mathbb{I}_{\{\text{Crisis}\}}(t) \lambda_{it} \\ + \gamma_4^{(i)} \mathbb{I}_{\{\text{Post-crisis}\}}(t) + \gamma_5^{(i)} \mathbb{I}_{\{\text{Post-crisis}\}}(t) \lambda_{it} + \epsilon_{it}, \quad (2.12)$$

where $\gamma_j^{(i)}$ are the issue-specific regressors. We only consider issues for which there are at least three data points in each era.

We report the average and the standard deviation of the regression estimates as well as the average R -squared in Table 2.4. For the three periods, the average coefficients are negative and statistically significant at a significance level of 5%. The liquidity proxy average coefficient is -0.0080 in the pre-crisis era, when liquidity was less important in general. Then, in the crisis era, the average liquidity proxy coefficient jumped to -0.2369 , which is 30 times higher than the pre-crisis average coefficient. In the post-crisis period, the coefficient increases, but does not reach pre-crisis levels meaning that liquidity is an important determinant of the bond-CDS spread. The R -squared values of these regressions are around 39% on average. Therefore, the unique liquidity factor accounts for a little less than half of the variations, on average.

Table 2.4: **Bond-CDS basis regression results.**

	Pre-crisis		Crisis		Post-crisis	
	Intercept	Liquidity	Intercept	Liquidity	Intercept	Liquidity
Average coefficient	-0.6228	-0.0080	-0.9389	-0.2369	-1.3437	-0.0788
SD	0.0281	0.0040	0.0972	0.1087	0.1621	0.0281
Average R^2	0.3942					
Number of regressions	268					

Using our dataset of bond issues, the regressions

$$(\text{Bond-CDS basis})_{it} = \gamma_0^{(i)} \mathbb{I}_{\{\text{Pre-crisis}\}}(t) + \gamma_1^{(i)} \mathbb{I}_{\{\text{Pre-crisis}\}}(t) \lambda_{it} + \gamma_2^{(i)} \mathbb{I}_{\{\text{Crisis}\}}(t) + \gamma_3^{(i)} \mathbb{I}_{\{\text{Crisis}\}}(t) \lambda_{it} \\ + \gamma_4^{(i)} \mathbb{I}_{\{\text{Post-crisis}\}}(t) + \gamma_5^{(i)} \mathbb{I}_{\{\text{Post-crisis}\}}(t) \lambda_{it} + \epsilon_{it},$$

are estimated for each bond issue. The conclusions of the statistical test $H_0 : \bar{\gamma}_j^{(i)} = 0$ against $H_1 : \bar{\gamma}_j^{(i)} \neq 0$, $j = 0, 1, \dots, 5$, are reported. Estimates in bold are significant at a confidence level of 95%. On average, 60 data points are used in each regression (on a maximum of 84 months). The averages and SD are based on 268 regressions. These regressions needed at least 3 data points in each era to be considered.




Figure 2.6: **Histogram of the regressions' R -squared values.**

For each issue, we run the regression of Equation (2.12). The R -squared are then computed from these regressions. The basis and liquidity proxy are computed on a monthly basis from January 2005 to December 2012.

A histogram of the regressions' R -squared values is given in Figure 2.6. About 60% of the regressions have an R -squared of more than 30% while using only the liquidity factor as a regressor (along with some dummies). Indeed, this is consistent with the idea that liquidity is a major determinant of the bond-CDS basis.

To conclude this section, our regression-based tests are rather clear: bond-specific liquidity is an important driver of the bond-CDS basis and bond hedgers should account for this additional risk when hedging their positions. Every coefficient is statistically and economically significant, meaning that bond-specific liquidity was (and probably still is today) an important determinant of the basis. This result is consistent with Bai and Collin-Dufresne (2011), as their liquidity proxies are highly significant during the crisis.

2.7 Concluding Remarks

Regime-switching dynamics are able to capture behaviour changes similar to those associated with a crisis. However, firm-specific MLE in the presence of regimes and noisy prices is problematic as both regimes and noises are not directly observed.

This paper contributes to the present credit risk literature by proposing a Markov-switching model and a filter-based estimation technique. A flexible credit risk model is designed to capture empirical evidences observed during the last decade. A regime-switching variable is included to accommodate behaviour changes during the financial

crisis. A negative dependence between the endogenous recovery rates and the firm's default probability addresses the Altman et al. (2005) empirical findings.

A firm-by-firm estimation procedure based on a filtering method and the principle of maximum likelihood deals with latent variables and noise. Tugnait (1982) is extended to allow for nonlinearities in the state-space representation and a fast lattice-based pricing scheme is implemented to ensure the feasibility of the estimation step.

The in-sample performance of the model reveals that it is flexible enough to adjust to various firms and financial cycles. An out-of-sample study concludes that the model is reliable and outperforms other benchmarks.

The negative dependence between the default probabilities and the recovery rates greatly impacts mid- and long-term credit spreads. Indeed, in a regime-switching environment, the five-year credit spread increases with the endogenous recovery assumption when compared to a model with a constant recovery rate. Also, during the financial turmoil, the short-term shape of the average credit spread curves is changed by the presence of regimes.

Based on model-implied yield-to-maturity spreads and regression tests, bond-specific liquidity appears to be a significant driver of the bond-CDS basis. For instance, during the crisis era, the average liquidity proxy coefficient reached -0.2369, which is 30 times higher than the pre-crisis average coefficient. This conclusion is consistent with our prior belief that bond hedgers must account for additional basis risk when hedging their positions in fixed income markets. The exact way bond-specific liquidity would change the hedging strategy is a question by itself and is left for future research.

2.A Derivative Pricing

2.A.1 Trinomial Lattice Approach

Yuen and Yang (2010) propose a trinomial lattice approach for Markov-switching dynamics. A “up-across-down” branching structure is chosen with $x_u = xe^{\sigma\sqrt{\Delta}}$, $x_m = x$, and $x_d = xe^{-\sigma\sqrt{\Delta}}$ where x is the actual value of the log-leverage process at a typical node. Moreover, when the number of regimes is K , the value of σ is given by $\sigma = \max_{1 \leq i \leq K} \sigma_i + (\sqrt{1.5} - 1)\bar{\sigma}$ where $\bar{\sigma}$ is the arithmetic mean of σ_i . This suggestion is based on the values used in the binomial and trinomial trees in the literature.




Figure 2.7: The branching to default at a typical node in the tree when the number of regimes is set to two.

The weights for the i th regime are

$$\pi_m^i = 1 - \frac{1}{\lambda_i^2}, \quad \pi_u^i = \frac{e^{\mu_i^Q \Delta} - e^{-\sigma \sqrt{\Delta}} - (1 - 1/\lambda_i^2)(1 - e^{-\sigma \sqrt{\Delta}})}{e^{\sigma \sqrt{\Delta}} - e^{-\sigma \sqrt{\Delta}}},$$

$$\pi_d^i = \frac{e^{\sigma \sqrt{\Delta}} - e^{\mu_i^Q \Delta} - (1 - 1/\lambda_i^2)(e^{\sigma \sqrt{\Delta}} - 1)}{e^{\sigma \sqrt{\Delta}} - e^{-\sigma \sqrt{\Delta}}},$$

where $\lambda_i = \frac{\sigma}{\sigma_i}$. These weights are different for each regime. Schönbucher's (2002) lattice that deals with credit-sensitive instruments is adapted in the trinomial lattice approach for Markov-switching dynamics by adding an additional branch at each node. According to Equations (2.2) and (2.3), the default probability is

$$p = 1 - \exp\left(-\left(\beta + \left(\frac{e^x}{\theta}\right)^\alpha\right)\Delta\right),$$

if the log-leverage value at this typical node is x .

Figure 2.7 shows the different branches to be considered to use this numerical scheme when $K = 2$. Note that even if the figure contains two different trees, these trees represent the same lattice: only the weights change across the different regimes.

2.A.2 Coupon Bond Prices

To price coupon bonds on a single firm, recovery specifications are crucial. Again, the endogenous recovery of Subsection 2.2 is utilized. Bond investors will recover a fraction R_τ of an equivalent Treasury bond at default time τ . Given a coupon rate of c , a maturity T , an initial log-leverage of x_t , an initial regime of s_t , and a face value of 1, the coupon bond price is given by

$$\begin{aligned}
 & V(t, T, x_t, s_t; R_\tau, c) \\
 &= \mathbb{E}^\mathbb{Q} \left[\frac{c}{2} \sum_{t_i^*} P(t, t_i^*) \mathbb{I}_{\{t \leq t_i^* < \tau\}} + P(t, T) \mathbb{I}_{\{\tau > T\}} + P(t, \tau) R_\tau P(\tau, T) \mathbb{I}_{\{t < \tau \leq T\}} \middle| \mathcal{F}_t \right] \\
 &= \mathbb{E}^\mathbb{Q} \left[\frac{c}{2} \sum_{t_i^*} P(t, t_i^*) \exp \left(- \sum_{t \leq u < t_i^*} H_u \Delta \right) \middle| \mathcal{G}_t \right] \mathbb{I}_{\{\tau > t\}} + \mathbb{E}^\mathbb{Q} \left[P(t, T) \exp \left(- \sum_{t \leq u < T} H_u \Delta \right) \middle| \mathcal{G}_t \right] \mathbb{I}_{\{\tau > T\}} \\
 &+ \mathbb{E}^\mathbb{Q} \left[\sum_{t \leq k < T} R_k P(t, T) \exp \left(- \sum_{t \leq u < k} H_u \Delta \right) (1 - \exp(-H_k \Delta)) \middle| \mathcal{G}_t \right] \mathbb{I}_{\{\tau > t\}} \quad (2.13)
 \end{aligned}$$

where $P(t, T)$ is the time t value of a risk-free zero-coupon bond maturing at T and t_i^* are coupon payment dates. Note that a zero-coupon bond can be computed by replacing the coupon rate c by zero in Equation (2.13).

2.A.3 Credit Default Swap Premiums

A CDS is a credit derivative that compensates the buyer in the event of a default (or other credit events). In the most basic type of CDS, the protection seller provides a payment of par-minus-recovery (settled in cash) on default. This shall cover the loss incurred by a typical bondholder. In exchange, the protection buyer pays a periodic premium that ceases if a default occurs. Normally, these premiums are paid four times per year. The premium of such a contract is fixed by setting the expected present value of losses equal to the expected present value of the premiums.

The endogenous recovery rate defined in Section 2.2 introduces a negative relation between recovery and default risks. Given that the CDS matures at time T and that $P(t, T)$ is the time t value of a risk-free zero-coupon bond maturing at T , the protection leg (expected present value of the losses) is

$$\mathbb{E}^\mathbb{Q} [P(t, \tau)(1 - R_\tau) \mathbb{I}_{\{t < \tau \leq T\}} | \mathcal{F}_t] =$$

$$\mathbb{E}^{\mathbb{Q}} \left[\sum_{t \leq k < T} (1 - R_k) P(t, k) \exp \left(- \sum_{t \leq u < k} H_u \Delta \right) (1 - \exp(-H_k \Delta)) \middle| \mathcal{G}_t \right] \mathbb{I}_{\{\tau > t\}}, \quad (2.14)$$

where the filtration $\{\mathcal{F}_t\}_{t=1}^{\infty}$ is defined as $\mathcal{F}_t = \sigma(\mathcal{G}_t, \mathcal{H}_t)$ and $\mathcal{H}_t = \sigma(\mathbb{I}_{\{\tau \leq s\}} : s \leq t)$.

To simplify the presentation, assume that a premium of 1 is paid. In this case, the premium leg (expected present value of the premiums) is given by

$$\mathbb{E}^{\mathbb{Q}} \left[\sum_{t_i^*} P(t, t_i^*) \mathbb{I}_{\{t \leq t_i^* < \tau\}} \middle| \mathcal{F}_t \right] = \mathbb{E}^{\mathbb{Q}} \left[\sum_{t_i^*} \delta_i^* P(t, t_i^*) \exp \left(- \sum_{t \leq u < t_i^*} H_u \Delta \right) \middle| \mathcal{G}_t \right] \mathbb{I}_{\{\tau > t\}} \quad (2.15)$$

where t_i^* are premium payment dates and $\delta_i^* = t_i^* - t_{i-1}^*$. The periodic premium is the ratio of (2.14) over (2.15).

The numerical scheme explained in Section 2.A.1 shall come in handy to compute credit default swap premiums in the proposed framework. To price credit default swaps, a time step of 1 week (i.e. $\Delta = 1/52$) is used in the tree.

2.B Simulation Study for the DEA-UKF Method

The efficiency of the method is assessed by two simulation studies.²⁶ The first Monte Carlo experiment compares the filtered leverages and regimes to their true values. To this end, 500 paths of 500 leverage observations are simulated. Then, we compute the logarithm of the CDS premiums for six tenors and add Gaussian random noises as stated in Equation (2.8).²⁷ To each path, we apply the DEA-UKF and compare the extracted filtered states to their true values, which are known because the data are simulated.

Panel A of Table 2.5 exhibits the summary statistics for the true and filtered states. Relative to the true states, the DEA-UKF filtered values of leverages are close to the true state values. Even for higher moment behaviour, the filtered states are virtually identical to the true ones. The root mean square error (RMSE) on filtered leverages (with respect to their true values) is 0.0034. With respect to the most probable regime, the DEA-UKF is capturing the right one 98.3% of the time.

²⁶In this paper, M is set to 5. This choice appears to be a good compromise between accuracy and efficiency. Moreover, the results seem to be robust to different choices of M greater than 4.

²⁷As in our dataset, we use 1-, 2-, 3-, 5-, 7- and 10-year CDS premiums.

Table 2.5: Simulation studies on the DEA-UKF methodology.

Panel A: Summary statistics for the true and filtered states.						
	True			DEA-UKF		
Filtered leverage						
Mean	0.6560			0.6561		
SD	0.2092			0.2091		
Skewness	0.5443			0.5448		
Kurtosis	3.3402			3.3430		
Panel B: Comparison of true and estimated parameters.						
	True	Mean	SD	t	CR	t (CR)
$\mu^{\mathbb{P}}$	0.0100	0.0098	0.0025	-2.0340	0.9700	2.0520
$\mu^{\mathbb{Q}}$	0.0100	0.0101	0.0004	3.1893	0.9480	-0.2052
σ_1	0.0800	0.0802	0.0048	0.9666	0.9440	-0.6156
$p_{1,2}$	0.0500	0.0493	0.0131	-1.1558	0.9520	0.2052
$q_{1,2}$	0.0500	0.0499	0.0021	-1.1614	0.9340	-1.6416
σ_2	0.3000	0.3000	0.0033	-0.1028	0.9420	-0.8208
$p_{2,1}$	0.0500	0.0495	0.0140	-0.7612	0.9580	0.8208
$q_{1,2}$	0.0500	0.0497	0.0033	-2.1076	0.9160	-3.4883
α	10.0000	9.9955	0.3133	-0.3184	0.9520	0.2052
θ	1.2000	1.2010	0.0090	2.4801	0.9400	-1.0260
β	0.0100	0.0100	0.0010	-0.6704	0.9540	0.4104
κ	0.5000	0.5000	0.0002	0.9253	0.9660	1.6416
$\delta^{(1)}$	0.1500	0.1496	0.0052	-1.9237	0.9480	-0.2052
$\delta^{(2)}$	0.1000	0.0998	0.0035	-1.0133	0.9340	-1.6416
$\delta^{(3)}$	0.0500	0.0501	0.0022	0.6300	0.9440	-0.6156
$\delta^{(5)}$	0.0300	0.0299	0.0017	-1.7315	0.9500	0.0000
$\delta^{(7)}$	0.0500	0.0498	0.0018	-2.0463	0.9540	0.4104
$\delta^{(10)}$	0.1000	0.0999	0.0031	-0.3991	0.9480	-0.2052

For each path, leverage paths were generated using Equation (2.1). The experiment consisted of 500 paths of 500 weekly observations. Mean corresponds to the mean estimated parameter. SD stands for standard deviation of the parameter sample. The column labeled t represents the t -statistics of the following hypothesis testing: $H_0 : \bar{\phi} - \phi = 0$, $H_1 : \bar{\phi} - \phi \neq 0$. The columns labeled CR exhibit the coverage ratio of the 95% confidence interval of each individual parameters. The last column labeled t (CR) shows the t -statistics of the following hypothesis testing: $H_0 : CR = 0.95$, $H_1 : CR \neq 0.95$.

In a second test, we investigate potential parameter estimation issues. The parameters are estimated with the DEA-UKF leading to a sample of 500 estimates.

Panel B of Table 2.5 reports true values, averages and standard deviations of parameter estimates for the DEA-UKF estimation method. The filter yields virtually no differences between estimated and true parameter values on average. The difference for each parameter is also inside the 99% margin of error constructed from the sample standard deviation for most parameters, except for μ^Q .

Table 2.5 also exhibits the coverage rate (CR) of the 95% confidence interval for each parameter and their t -statistics. The CR corresponds to the proportion of 95% confidence intervals containing the true parameter value. Confidence intervals for the DEA-UKF method are in line with theory (i.e. close to 95%). Also, most CR are not statistically different of 0.95 at a significance level of 5%.

Finally, a multivariate confidence region is constructed by inverting the likelihood ratio statistic. Let $U(\mathbf{y}) = \{\hat{\phi} : 2L(\hat{\phi}; \mathbf{y}) - 2L(\phi_0; \mathbf{y}) < F_{\text{crit}}\}$, where L is the log-likelihood

Table 2.6: Descriptive statistics for the liquidity proxies.

Panel A: Summary statistics for liquidity proxies and the theoretical yield spread.									
	λ	Amihud	A. risk	IRC	IRC risk	Roll	Turnover	Zero	Spread
Mean	0.011	0.483	0.713	0.007	0.005	0.015	0.032	0.712	1.785
SD	3.250	1.756	2.362	0.007	0.006	0.018	0.101	0.283	2.660
1%	-3.560	0.000	0.000	0.000	0.000	0.001	0.000	0.043	0.080
5%	-3.158	0.002	0.000	0.000	0.000	0.003	0.000	0.091	0.134
25%	-1.915	0.043	0.039	0.003	0.001	0.007	0.001	0.545	0.423
50%	-0.767	0.158	0.225	0.005	0.004	0.012	0.012	0.818	1.066
75%	0.902	0.403	0.627	0.009	0.007	0.019	0.035	0.952	1.917
95%	5.744	1.733	2.812	0.018	0.014	0.037	0.111	1.000	6.403
99%	13.010	5.566	7.785	0.031	0.026	0.069	0.299	1.000	12.361
Panel B: Correlation matrix for liquidity proxies and the theoretical yield spread.									
	λ	Amihud	A. risk	IRC	IRC risk	Roll	Turnover	Zero	Spread
λ	1.000								
Amihud	0.317	1.000							
A. risk	0.539	0.425	1.000						
IRC	0.418	0.094	0.118	1.000					
IRC risk	0.528	0.032	0.196	0.589	1.000				
Roll	0.275	0.204	0.206	0.391	0.357	1.000			
Turnover	-0.022	-0.058	-0.026	0.007	0.008	-0.026	1.000		
Bond zero	0.013	0.169	0.016	-0.024	-0.168	0.027	-0.132	1.000	
Spread	0.161	0.059	0.154	0.291	0.281	0.300	0.057	-0.207	1.000

This table shows statistics for corporate bond liquidity proxies. The proxies are calculated monthly for each bond from January 2005 to December 2012. Panel A shows quantiles for the proxies. Panel B reports correlation among the proxies. Amihud means Amihud measure, A. risk stands for Amihud risk, IRC corresponds to IRC measure, Roll means Roll measure, Zero is for bond zero-trading days, and Spread stands for theoretical YTM spread in percentage.

function, $\hat{\phi}$ is an estimator, \mathbf{y} denotes the data and F_{crit} is the critical value given by a chi-square distribution with 18 degrees of freedom (i.e. for a 95% confidence region, $F_{\text{crit}} = 28.86$). Notice that for the MLE, U should contain the true parameter vector 95% of the time asymptotically, under usual regularity conditions. According to our tests, the coverage rate is 94.4% for the filter. Therefore, the results obtained using the DEA-UKF are consistent with the theory.

2.C Liquidity: Additional Material

In the spirit of Dick-Nielsen et al. (2012), eight different liquidity measures are presented: Amihud (2002) measure, Amihud risk, imputed roundtrip cost (IRC), IRC risk, Roll's (1984) measure, turnover rate of a bond, proportion of zero trading days, and Dick-Nielsen et al.'s (2012) λ measure. Each measure is computed on a monthly basis; moreover, these measures are calculated for each bond considered in the subsample.

Amihud (2002) constructs an illiquidity measure based on a model proposed by Kyle (1985). The measure ascertains the price impact of a trade per unit traded. It is defined as the daily average of absolute returns r_j divided by the trade size Q_j (in million of dollars) of consecutive transactions:

$$\text{Amihud}_t = \frac{1}{N_t} \sum_{j=1}^{N_t} \frac{|r_j|}{Q_j}$$

where N_t is the number of returns on day t , $r_j = (P_j - P_{j-1})/P_{j-1}$, and P_j is the price j^{th} trade of the day. The monthly Amihud measure is defined by taking the median of the daily measures.

Amihud risk is simply the standard deviation of the daily Amihud measure over one month.

Feldhütter (2012) proposes an alternative measure of transaction costs based on the so-called Imputed Roundtrip Trades (IRT). Often, a corporate bond is traded two or three times within a very short period of time; this is likely to occur because a dealer matches a buyer and a seller and collects the bid-ask spread as a fee. If two or three trades in a given bond with the same trade size take place on the same day, and there are no other trades with the same size on that day, it is assumed that the transactions are part of an IRT. The imputed roundtrip cost is then defined as

$$\text{IRC} = \frac{P_{\max} - P_{\min}}{P_{\max}}$$

where P_{\max} is the largest price in the IRT and P_{\min} is the smallest one. A daily estimate of roundtrip costs is the average of roundtrip costs on that day. In addition, the monthly IRC is simply the average of daily roundtrip costs.

IRC risk is the standard deviation of the daily IRC measure over one month.

Roll (1984) suggests that, under certain assumptions, the percentage bid-ask spread equals two times the square root of minus the covariance between two consecutive returns:

$$\text{Roll}_t = 2 \sqrt{-\text{cov}(r_t, r_{t-1})}$$

where t is the time period for which the measure is calculated. If the covariance is negative, this measure is not well-defined; the negative values are therefore discarded. A daily Roll measure is calculated on days with at least one transaction using a rolling window of 28 days. The monthly Roll measure is computed by taking the median

within the month. The intuition behind Roll's (1984) measure is that the bond price bounces back and forth between the bid and the ask prices and higher bid-ask spread shall lead to a higher negative correlation.

The monthly turnover rate of bonds is defined by the ratio of the total trading volume during the month divided by the amount outstanding:

$$\text{Turnover}_t = \frac{\text{Total trading volume}_t}{\text{Amount outstanding}}$$

where t is the considered month.

The proportion of zero-trading days is the percentage of days during a month where the bond did not trade.

Finally, Dick-Nielsen et al.'s (2012) λ measure is a composite factor that loads evenly on Amihud measure, Amihud risk, IRC measure, and IRC risk. Precisely, each measure L_t^j is normalized (i.e. $\tilde{L}_t^j = (L_t^j - \mu^j)/\sigma^j$) and then added up to obtain

$$\lambda_t = \sum_{j=1}^4 \tilde{L}_t^j.$$

This is also done on a monthly basis.

Panel A of Table 2.6 shows a summary of the distribution of the considered liquidity proxies. The composite λ measure has a zero mean by construction; however, its distribution seems highly asymmetrical. The first percentile is around -3.6 and the 99th percentile is at 13.0. The same asymmetric behaviour is true for the four constituents of the λ measure.

The median Amihud measure is 0.158. This means that a trade of 300,000 dollars implies a median move in the bond price of 4.7%. This value is somewhat larger than the one obtained by Dick-Nielsen et al. (2012): they found roughly 0.13%. However, they consider different firms and a different period; they also remove small trades. Han and Zhou (2008) find that a trade of 300,000 dollars moves the price by 10.2% on average. The median roundtrip cost in percentage of the price is 0.5% according to the IRC measure. The roundtrip costs for the 5% most liquid bonds is 0.03%, which is coherent with the results of Dick-Nielsen et al. (2012). The average number of bond zero-trading days is 71.2%. This is consistent with popular thinking: corporate bond market is an illiquid market. Moreover, Dick-Nielsen et al. (2012) find similar figures for the bond zero-trading days proxy. The average monthly turnover rate is 3.2%,

meaning that for the average bond in our sample it takes between 2 and 3 years to turn over once.

In Panel B of Table 2.6, a correlation matrix for the different liquidity proxies and the theoretical YTM spread is presented. Some of the liquidity proxies do not seem to be highly correlated with one another. Nonetheless, some others are mildly correlated. Note that the λ measure is a construction based on the Amihud measure, the Amihud risk, the IRC measure, and the IRC risk, therefore it is natural that the λ measure is moderately correlated with the other four. The theoretical YTM spread is somehow related with some of our measures: for instance, the correlation between the Roll measure and the YTM spread is 30%. The liquidity measure is construction made from bond prices; thus, it is very intuitive to have moderate correlation between the liquidity measures and the yield spread.

References

- Acharya, V. V., S. T. Bharath, and A. Srinivasan. 2007. Does industry-wide distress affect defaulted firms? Evidence from creditor recoveries. *Journal of Financial Economics* 85:787–821.
- Alexander, C., and A. Kaeck. 2008. Regime dependent determinants of credit default swap spreads. *Journal of Banking & Finance* 32:1008–1021.
- Altman, E. I., B. Brady, A. Resti, and A. Sironi. 2005. The link between default and recovery rates: Theory, empirical evidence, and implications. *Journal of Business* 78:2203–2228.
- Amihud, Y. 2002. Illiquidity and stock returns: Cross-section and time-series effects. *Journal of Financial Markets* 5:31–56.
- Arora, N., P. Gandhi, and F. A. Longstaff. 2012. Counterparty Credit Risk and the Credit Default Swap Market. *Journal of Financial Economics* 103:280–293.
- Augustin, P. 2012. The term structure of CDS spreads and sovereign credit risk. *Working paper*.
- Bade, B., D. Rösch, and H. Scheule. 2011. Default and recovery risk dependencies in a simple credit risk model. *European Financial Management* 17:120–144.
- Bai, J., and P. Collin-Dufresne. 2011. The determinants of the CDS-bond basis during the financial crisis of 2007–2009. *Working paper*.

- Bakshi, G., D. Madan, and F. X. Zhang. 2006a. Investigating the role of systematic and firm-specific factors in default risk: Lessons from empirically evaluating credit risk models. *Journal of Business* 79:1955–1987.
- Bakshi, G., D. Madan, and F. X. Zhang. 2006b. Understanding the role of recovery in default risk models: Empirical comparisons and implied recovery rates. *Working paper*.
- Black, F., and J. C. Cox. 1976. Valuing corporate securities: Some effects of bond indenture provisions. *Journal of Finance* 31:351–367.
- Bollen, N. P. 1998. Valuing options in regime-switching models. *Journal of Derivatives* 6:38–49.
- Boudreault, M., G. Gauthier, and T. Thomassin. 2013. Recovery rate risk and credit spreads in a hybrid credit risk model. *Journal of Credit Risk* 9:3–39.
- Brigo, D., M. Predescu, A. Capponi, T. R. Bielecki, and F. Patras. 2011. Liquidity modeling for credit default swaps: An overview. In *Credit risk frontiers: Subprime crisis, pricing and hedging, CVA, MBS, ratings, and liquidity*, pp. 585–617. John Wiley and Sons.
- Bruche, M., and C. González-Aguado. 2010. Recovery rates, default probabilities, and the credit cycle. *Journal of Banking & Finance* 34:754–764.
- Bühler, W., and M. Trapp. 2009a. Explaining the bond-CDS basis: The role of credit risk and liquidity. *Working paper* pp. 1–35.
- Bühler, W., and M. Trapp. 2009b. Time-varying credit risk and liquidity premia in bond and CDS markets. *Working paper*.
- Carr, P., and L. Wu. 2010. Stock options and credit default swaps: A joint framework for valuation and estimation. *Journal of Financial Econometrics* 8:409–449.
- Christoffersen, P., C. Dorion, K. Jacobs, and L. Karoui. 2014. Nonlinear Kalman filtering in affine term structure models. *Management Science* 60:2248–2268.
- Creal, D. 2012. A survey of sequential Monte Carlo methods for economics and finance. *Econometric Reviews* 31:245–296.
- Dick-Nielsen, J. 2009. Liquidity biases in TRACE. *Journal of Fixed Income* 19:43–55.

- Dick-Nielsen, J., P. Feldhütter, and D. Lando. 2012. Corporate bond liquidity before and after the onset of the subprime crisis. *Journal of Financial Economics* 103:471–492.
- Dionne, G., and O. Maalaoui Chun. 2013. Default and liquidity regimes in the bond market during the 2002–2012 period. *Canadian Journal of Economics* 46:1160–1195.
- Doshi, H., J. Ericsson, K. Jacobs, and S. M. Turnbull. 2013. Pricing credit default swaps with observable covariates. *Review of Financial Studies* 26:2049–2094.
- Duan, J.-C., and A. Fulop. 2009. Estimating the structural credit risk model when equity prices are contaminated by trading noises. *Journal of Econometrics* 150:288–296.
- Duffie, D., and D. Lando. 2001. Term structures of credit spreads with incomplete accounting information. *Econometrica* 69:633–664.
- Duffie, D., and K. J. Singleton. 1999. Modeling term structures of defaultable bonds. *Review of Financial Studies* 12:687–720.
- Dumas, B., J. Fleming, and R. E. Whaley. 1998. Implied volatility functions: Empirical tests. *Journal of Finance* 53:2059–2106.
- Elizalde, A., S. Doctor, and Y. Saltuk. 2009. Bond-CDS basis handbook. Tech. rep., JP Morgan European Credit Derivatives Research.
- Elliott, R. J., L. Aggoun, and J. B. Moore. 1995. *Hidden Markov models: Estimation and control*. Springer Science & Business Media.
- Fabozzi, F. J., and S. V. Mann. 2012. *The handbook of fixed income securities*. McGraw Hill Professional.
- Fearnhead, P., and P. Clifford. 2003. On-line inference for hidden Markov models via particle filters. *Journal of the Royal Statistical Society* 65:887–899.
- Feldhütter, P. 2012. The same bond at different prices: Identifying search frictions and selling pressures. *Review of Financial Studies* 25:1155–1206.
- Fontana, A. 2010. The negative CDS-bond basis and convergence trading during the 2007/09 financial crisis. *Working paper*.
- Garzarelli, F. 2009. The 2007-09 credit crisis and its aftermath. *Working paper*.

- Guarin, A., X. Liu, and W. L. Ng. 2014. Recovering default risk from CDS spreads with a nonlinear filter. *Journal of Economic Dynamics and Control* 38:87–104.
- Han, S., and H. Zhou. 2008. Effects of liquidity on the nondefault component of corporate yield spreads: evidence from intraday transactions data. *Working paper*.
- Huang, A. Y., and W.-C. Hu. 2012. Regime switching dynamics in credit default swaps: Evidence from smooth transition autoregressive model. *Physica A: Statistical Mechanics and its Applications* 391:1497–1508.
- Huang, S. J., and J. Yu. 2010. Bayesian analysis of structural credit risk models with microstructure noises. *Journal of Economic Dynamics and Control* 34:2259–2272.
- Jarrow, R., and P. Protter. 2004. Structural versus reduced form models: a new information based perspective. *Journal of Investment Management* 2:1–10.
- Julier, S., and J. Uhlmann. 1997. A new extension of the Kalman Filter to nonlinear systems. In *SPIE proceedings series*, pp. 182–193. Society of Photo-Optical Instrumentation Engineers.
- Kalman, R. E. 1960. A new approach to linear filtering and prediction problems. *Journal of Basic Engineering* 82:35–45.
- Kim, C.-J. 1994. Dynamic linear models with Markov-switching. *Journal of Econometrics* 60:1–22.
- Kwon, T. Y., and Y. Lee. 2015. Estimating structural credit risk models when market prices are contaminated with noise. *Applied Stochastic Models in Business and Industry* 32:18–32.
- Kyle, A. S. 1985. Continuous auctions and insider trading. *Econometrica* 53:1315–1336.
- Longstaff, F. A., S. Mithal, and E. Neis. 2005. Corporate yield spreads: default risk or liquidity? New evidence from the credit default swap market. *Journal of Finance* 60:2213–2253.
- Maalaoui Chun, O., G. Dionne, and P. Francois. 2014. Detecting regime shifts in credit spreads. *Journal of Financial and Quantitative Analysis* 49:1339–1364.
- Madan, D., and H. Unal. 2000. A two-factor hazard rate model for pricing risky debt and the term structure of credit spreads. *Journal of Financial and Quantitative Analysis* 35:43–66.

- Mahanti, S., A. Nashikkar, M. Subrahmanyam, G. Chacko, and G. Mallik. 2008. Latent liquidity: a new measure of liquidity, with an application to corporate bonds. *Journal of Financial Economics* 88:272–298.
- Merton, R. C. 1974. On the pricing of corporate debt: the risk structure of interest rates. *Journal of Finance* 29:449–470.
- Moody's. 2009. Corporate Default and Recovery Rates, 1920-2008. Tech. rep., Moody's. URL <http://v2.moodys.com/cust/content/content.ashx?source=StaticContent/Free%20Pages/Credit%20Policy%20Research/documents/current/20074000000578875.pdf>.
- Mueller, P. 2008. Credit spreads and real activity. *Working paper*.
- NAIC. 2015. Update on the insurance industry's use of derivatives and exposure trends. Tech. rep., National Association of Insurance Commissioners. URL http://www.naic.org/capital_markets_archive/150807.htm.
- Qiu, J., and F. Yu. 2012. Endogenous liquidity in credit derivatives. *Journal of Financial Economics* 103:611–631.
- Roll, R. 1984. A simple implicit measure of the effective bid-ask spread in an efficient market. *Journal of Finance* 39:1127–1139.
- Saldías, M. 2013. Systemic risk analysis using forward-looking distance-to-default series. *Journal of Financial Stability* 9:498–517.
- Saunders, A., and L. Allen. 2010. *Credit risk management in and out of the financial crisis: New approaches to value at risk and other paradigms*. John Wiley & Sons.
- Schönbucher, P. J. 2002. A tree implementation of a credit spread model for credit derivatives. *Journal of Computational Finance* 6:1–38.
- Tang, D. Y., and H. Yan. 2007. Liquidity and credit default swap spreads. *Working paper*.
- Tugnait, J. K. 1982. Detection and estimation for abruptly changing systems. *Automatica* 18:607–615.
- van Trees, H. L. 1968. *Detection, estimation and linear modulation theory*. John Wiley & Sons.
- Vasicek, O. 1977. An equilibrium characterization of the term structure. *Journal of Financial Economics* 5:177–188.

- Viterbi, A. J. 1967. Error bounds for convolutional codes and an asymptotically optimum decoding algorithm. *IEEE Transactions on Information Theory* 13:260–269.
- Yuen, F. L., and H. Yang. 2010. Option pricing with regime switching by trinomial tree method. *Journal of Computational and Applied Mathematics* 233:1821–1833.
- Zhang, B. Y., H. Zhou, and H. Zhu. 2009. Explaining credit default swap spreads with the equity volatility and jump risks of individual firms. *Review of Financial Studies* 22:5099–5131.

Chapter 3

Credit and Systemic Risks in the Financial Services Sector: Evidence from the 2008 Global Crisis

Abstract^{*}

We develop a portfolio credit risk model that includes firm-specific Markov-switching regimes as well as individual stochastic and endogenous recovery rates. Using weekly credit default swap premiums for 35 financial firms, we analyze the credit risk of each of these companies and their statistical linkages, putting emphasis on the 2005–2012 period. Moreover, we study the systemic risk affecting both the banking and insurance subsectors.

Keywords: credit risk; systemic risk; financial sector; insurance; banking; default probability; correlation; unscented Kalman filter (UKF).

^{*}Joint work with Mathieu Boudreault, Delia Alexandra Doljanu and Geneviève Gauthier. Boudreault is affiliated with UQAM, Doljanu with National Bank of Canada, and Gauthier with HEC Montréal.

3.1 Introduction and Review of the Literature

The financial crisis of 2008 shed light on how the interconnectedness of large financial institutions can seriously affect their solvency. Yet, even if the financial crisis is behind us, we still need to comprehend its aftereffects. In this spirit, we assess the evolution of major determinants of financial crises during the 2005–2012 period, namely the roles of leverage, losses and linkages.

This paper offers new insights into the role of credit and systemic risks affecting both insurance and banking subsectors. Among others, we investigate the changes in correlation through time and the contribution of insurance and banking firms in the risk of collapse of an entire financial system. To this end, we construct a multivariate credit risk model that accounts for firm-specific financial health. The framework embeds oft-cited stylized facts such as leverage volatility (modelled via statistical regimes), recovery rates negatively related to default probabilities, and pairwise regime-dependent correlations. We also propose a consistent and reliable method to estimate within the multivariate framework.

Various credit risk models have been proposed in the literature. They have been historically divided into two categories: structural and reduced-form models.¹ Even though the reduced-form approach provides a better fit to market data than the structural approach does, it lacks the economic and financial intuition of the structural framework. To overcome the limitations of both traditional approaches while retaining the main strengths of each, hybrid credit risk models have emerged in the literature.² In this paper, we adopt a credit risk framework that belongs to this last class of models, linking the default intensity to the capital structure of the firm through its leverage ratio. More precisely, to model the leverages, losses and linkages adequately, a regime-switching extension of the multivariate hybrid credit risk model of Boudreault et al. (2014) is proposed: it allows for firm-specific statistical regimes that accommodate for changes in the leverage volatility and an endogenous stochastic recovery rate that is negatively related to the default probabilities, and therefore impacts the loss distribution. Regime-switching dynamics are required to capture the various changes in behaviour through

¹Structural models link the credit events to the firm's economic fundamentals by assuming that default occurs when the firm's value falls below some boundary. Reduced-form models consider the surprise element of the default trigger exogenously given through a default intensity process.

²For instance, Duffie and Lando (2001), Çetin et al. (2004), Giesecke and Goldberg (2003) and Giesecke (2006) use incomplete information models in a way that firm assets and the default barrier are not observable by investors. Another segment of the literature focuses on modelling the default time as the first jump of a Cox process for which the intensity depends on the firm's fundamentals (e.g. Madan and Unal, 2000).

time, and more particularly during crises.

Generally, studies of individual firms' solvency have mostly focused on balance sheet information (Allen et al., 2002), credit ratings (Gupton et al., 2007), or distance to default (Bharath and Shumway, 2008). The financial services sector is no exception to the rule. Indeed, Harrington (2009) employs, among other things, balance sheet information to assess the role of AIG and the insurance subsector in the recent crisis. Milne (2014) uses the distance to default to investigate the solvency of European banks, concluding that the distance to default measure performs poorly as a market-based signal for bank risk. In our study, we employ weekly single-name credit default swap (CDS) premiums of 35 major financial institutions over 2005–2012. The use of market data are worthwhile: CDS premiums contain forward-looking information and are frequently updated by market participants as the information becomes available. Accordingly, they are more appropriate to detect sudden changes in solvency or occurrence of crises.³ In particular, we find that AIG's 1-year default probability (PD) spikes to 42% on September 10, 2008, a week before its near-default. On average, the banking subsector's 1-year PD increases from 0.5% to 4.6% during the crisis, while the insurance subsector's PD increases from 0.4% to 4.2%.

Although numerous single-firm approaches exist for measuring credit risk, financial institutions are intertwined and, therefore, credit risk assessment of the financial services sector requires an examination of the interconnectedness of its institutions. There are several ways to look at the interconnectedness of companies: correlation in the firm's assets or default intensity through copulas or common factors (e.g. Frey and McNeil, 2003, Hull et al., 2010, Li, 2000, Meine et al., 2016), exposure to other common risks such as jumps (e.g. Duffie and Gârleanu, 2001) or other contagion mechanisms (e.g. Davis and Lo, 2001) such as network approaches (e.g. Billio et al., 2012, Markose et al., 2012, Nier et al., 2007). This study models dependence through pairwise regime-dependent correlations of leverage co-movements. We link the regimes to the firm-specific correlation coefficients as one of our main goals is to capture the increase in pairwise correlation during the last financial crisis. Our empirical results show that linkage varies over time. We find evidence of larger correlations between firm leverage co-movements during the high-volatility regime which suggests the existence of greater interconnectedness during the last crisis. Moreover, the regime-dependent linkage structure varies across subsectors.

³Moreover, these are superior to ratings-based methods because rating revisions tend to lag behind the market and default probabilities based on the latter depend on aggregated default counts (i.e. not firm-specific).

Since the financial crisis, many multivariate credit risk frameworks have been used to investigate systemic risk in the financial sector. Notably, Huang et al. (2009) and Huang et al. (2012) construct a systemic risk measure inferred from CDS spreads and equity price co-movements.⁴ Using a network approach, principal component analysis and Granger-causality networks, Billio et al. (2012) quantify the interdependence among four groups of financial institutions during the recent crisis. Their empirical results suggest that the banking and insurance subsectors are more important sources of interconnectedness than other financial institutions. Another contribution in that field is the systemic expected shortfall proposed by Acharya et al. (2010) that measures the expected loss to each institution conditional on the undercapitalization of the entire financial system. Other measures of systemic risk applied to financial institutions have been proposed by Adrian and Brunnermeier (2009) and Saldías (2013).⁵

With respect to systemic risk in the insurance subsector, Weiß and Mühlnickel (2014) use the Systemic Risk Index measure developed by Acharya et al. (2012) and find that the contribution of insurers to systemic risk is only determined by the insurer's size, whereas Cummins and Weiss (2014) show that non-core activities of U.S. insurers may pose systemic risk.⁶ Finally, closer to our systemic risk study, Chen et al. (2014) discuss systemic risk in the insurance and banking subsectors using Huang et al.'s (2009) measure along with CDS premiums and high-frequency equity returns. They find a unidirectional causal effect from banks to insurers when accounting for heteroskedasticity.

In our study, as the proposed model captures firm-specific credit risk and dependence across the firms, it serves as a building block to construct a systemic risk measure inspired from Acharya et al. (2010). We find increases in systemic risk contributions for both insurance and banking subsectors during the crisis period. In line with Chen et al. (2014) and Billio et al. (2012), we also detect a unidirectional causal effect from banks to insurers when accounting for heteroskedasticity. Therefore, even if our methodology differs and our data extends over the aftermath of the crisis, our results suggest that the direction of the causal relationship is robust.

⁴Huang et al. (2009) propose the use of the so-called "distress insurance premium." This theoretical price of insurance against distressed losses is calculated as the risk-neutral expectation of portfolio credit losses that equal or exceed a minimum share of the sector's total liabilities.

⁵Adrian and Brunnermeier (2009) introduce the concept of CoVaR that measures the value at risk (VaR) of the financial system conditional on the distress of a specific firm. Saldías (2013) develops a forward-looking measure based on the gap between portfolio and average distance to default series to monitor systemic risk in Europe.

⁶Core activities refer to insurance underwriting, reserving, claims settlement and reinsurance. Non-core activities are associated with banking activities engaged in by some insurers.

The contributions of this paper are threefold. First, firm-specific credit risk is modelled in the financial services sector by the means of a multivariate credit risk model which captures the main determinants of credit risk (and more specifically, financial crises). Second, we provide a consistent estimation method for the multivariate model. Estimation of the model's parameters is a crucial step to adequately measure both credit and systemic risks. Indeed, as defaults are rare events, a lack of direct observations brings an extra challenge when firm-specific credit risk needs to be estimated.⁷ Finally, new insights into the financial services sector's credit and systemic risks are provided, especially regarding sector-wide linkages and systemic risk in the insurance and banking subsectors during the last financial crisis.

The remainder of this paper is organized as follows: Section 3.2 explains the multivariate credit risk model used. In Section 3.3, the CDS dataset is described. The firm-specific credit risk results are discussed in Section 3.4. Section 3.5 shows the results regarding the linkages between firms. Section 3.6 provides an assessment of the systemic risk in both the insurance and banking subsectors. Finally, Section 3.7 concludes.

3.2 Multivariate Credit Risk Model

As discussed in the introduction, to adequately capture credit and systemic risks requires the incorporation of some desired features, namely the “L”s of financial crises: leverage, losses and linkages. In this spirit, the proposed multivariate Markov-switching model combines the regime-switching univariate framework of Bégin et al. (2016) and the portfolio hybrid default risk approach of Boudreault et al. (2014).

We use regime-switching dynamics in this study as crises are the type of events that occur suddenly and cannot be well captured by a highly persistent autoregressive framework such as DCC-type processes (see Engle, 2002, for more details on the dynamic conditional correlation framework).

⁷Numerous studies construct proxies for default probabilities, recovery rates and other models' inputs based on aggregated information across ratings, balance sheet data and equity returns. More recently, a number of authors have implemented filtering approaches to deal with the latent nature of some models' variables and the presence of noise in the market data. For instance, see Duan and Fulop (2009), Huang and Yu (2010) and Boudreault et al. (2013).

3.2.1 Markov-Switching Dynamics

As a starting point to the model, the time t market value of the i^{th} firm's assets and the present value of the i^{th} firm's liabilities are denoted by $A_t^{(i)}$ and $L_t^{(i)}$ respectively. To capture changes in the asset and liabilities dynamics, a regime-switching variable is incorporated. This would allow for the flexibility needed to model periods of turmoil. Hence, $s_t^{(i)}$ is the hidden state of the regime prevailing at time t . As emphasized by the notation, the regime dynamics are firm-specific.

The leverage ratio $X_t^{(i)} = L_t^{(i)}/A_t^{(i)}$ follows a first-order two-state Markov-switching process such as

$$\log(X_t^{(i)}) = \log(X_{t-1}^{(i)}) + \left(\mu^{(i)} - \frac{1}{2}(\sigma_{s_t^{(i)}}^{(i)})^2\right)\Delta t + \sigma_{s_t^{(i)}}^{(i)} \sqrt{\Delta t} \varepsilon_t^{(i)}, \quad i \in \{1, 2, \dots, N\} \quad (3.1)$$

where Δt represents the time between two consecutive observations, and $\{\varepsilon_t^{(i)}\}_{t=1}^{\infty}$ is a standardized Gaussian noise series. The drift $\mu^{(i)}$ as well as regime diffusions $\sigma_1^{(i)}$ and $\sigma_2^{(i)}$ are firm-specific parameters to be estimated. Note that $\{s_t^{(i)} : i = 1, 2, \dots, N\}$ are independent first-order Markov chains. If $p_{11}^{(i)}$ denotes $\mathbb{P}(s_t^{(i)} = 1 \mid s_{t-1}^{(i)} = 1)$ and $p_{22}^{(i)}$ denotes $\mathbb{P}(s_t^{(i)} = 2 \mid s_{t-1}^{(i)} = 2)$, the regime state $s_t^{(i)}$ has the following transition matrix:

$$\mathbf{P}^{(i)} = \begin{bmatrix} p_{11}^{(i)} & 1 - p_{11}^{(i)} \\ 1 - p_{22}^{(i)} & p_{22}^{(i)} \end{bmatrix}. \quad (3.2)$$

When it comes to a portfolio approach, one must consider the interrelation among firms that can lead to clusters of defaults and may significantly impact the future value distribution of the portfolio. To this end, the model captures the firms' interconnections through the correlation between noise terms of log-leverage described in Equation (3.1), i.e.

$$\rho_{s_t^{(i)}}^{(i,j)} = \text{Corr}^{\mathbb{P}}(\varepsilon_t^{(i)}, \varepsilon_t^{(j)}) \quad (3.3)$$

with $\mathbf{s}_t \in \{s_t^{(i)}, s_t^{(j)}\}$, or $\mathbf{s}_t \in \{(1, 1), (1, 2), (2, 1), (2, 2)\}$. Thus, four correlation values have to be estimated for each pair of firms depending on their specific regimes, i.e. $\rho_{s_t}^{(i,j)} = (\rho_{1,1}^{(i,j)}, \rho_{1,2}^{(i,j)}, \rho_{2,1}^{(i,j)}, \rho_{2,2}^{(i,j)})$.

Depending on the modelling objective, the log-leverage dynamics evolve either under risk-neutral pricing measure \mathbb{Q} , or under physical measure \mathbb{P} for risk management purposes. The market model is incomplete, implying that there are an infinite number of pricing measures. Among these measures, we restrict the choices to those preserving

the model structure by having different $\mu^{(i)}$, $p_{11}^{(i)}$ and $p_{22}^{(i)}$ under both measures \mathbb{P} and \mathbb{Q} , meaning that the regime risk is priced.

3.2.2 Default Intensity

The multivariate Markov-switching model is based on a hybrid default risk framework that combines features of both structural and reduced-form approaches. The model also features an endogenous stochastic recovery rate that depends on the firm's default probability.

More precisely, the model first relies on the assumption that default is driven by an intensity process H_t that depends on the leverage ratio X_t such that

$$H_t^{(i)} = \beta^{(i)} + \left(\frac{X_t^{(i)}}{\theta^{(i)}} \right)^{\alpha^{(i)}} \quad (3.4)$$

where $\alpha^{(i)} > 0$, $\beta^{(i)} > 0$ and $\theta^{(i)} > 0$ are firm-specific constants to be estimated. Furthermore, the intensity process allows the default time to be defined as a reduced-form default trigger, that is, the first jump of a Cox process:

$$\tau^{(i)} \equiv \inf \left\{ t \in \{1, 2, \dots\} : \sum_{u=0}^{t-1} H_u^{(i)} \Delta t > E_1^{(i)} \right\} \quad (3.5)$$

where $E_1^{(i)}$ is an independent exponential random variable with mean 1. Since $\alpha^{(i)}$, $\beta^{(i)}$ and $\theta^{(i)}$ are positive constants, the likelihood of default tends to increase with the leverage ratio. Notice that parameters α and θ gauge the sensitivity of the firm's survival against its leverage ratio. The convexity of the default intensity is guided by α , while the critical leverage threshold is defined by θ . The parameter β captures a portion of the default drivers, and ensures that H_t is a positive function when $\beta > 0$. With all other variables being the same, the larger the β , the greater the intensity and default probability.

This framework allows for an endogenous recovery rate that depends on the capital structure of the firm at the time of default. Considering liquidation and legal fees as a fraction $\kappa^{(i)}$ of the market asset value at default, the debtholders receive the smallest amount between the value of liabilities and what remains from the liquidation of assets: $\min((1 - \kappa^{(i)})A_\tau^{(i)}; L_\tau^{(i)})$. Given the leverage dynamics, the random behaviour of

the recovery rate at the time of default is

$$R_{\tau}^{(i)} = \min \left(\left(1 - \kappa^{(i)} \right) \frac{1}{X_{\tau}^{(i)}}; 1 \right). \quad (3.6)$$

The endogenous recovery rate distribution is consistent with the empirical literature, as it is a decreasing function of the leverage ratio, meaning that default probability is negatively correlated with the recovery rate at the moment of default.⁸ The stochastic behaviour of the recovery rate as well as regime-switching dynamics imply that CDS premiums cannot be calculated in closed form. Therefore, a numerical method based on a trinomial lattice approach is used. Details on the method used to price CDS are available in Chapter 2 of this thesis.

3.3 Data and Assumptions

Since the late 1990s, the credit risk market has substantially grown and the CDS has become a new instrument for investors to manage and measure their risk. Considering that the CDS premium is directly linked to the credit quality of the bond issuer, it is expected to reflect an adequate measure of credit risk. In the recent literature, many authors challenge this argument (see Bielecki et al., 2011, Friewald et al., 2012, among others). However, empirical studies suggest that credit risk is one of the most important risks involved in the CDS spread and therefore, provides a good proxy for studying a firm's credit risk.⁹ In this study, CDS premiums are used as inputs in a filtering procedure to estimate the Markov-switching hybrid credit risk model.

CDS premiums are provided by Markit for tenors of 1, 2, 3, 5, 7, and 10 years. As our model embeds two latent variables (the leverage and the hidden regime), more information is needed to adequately infer both these unobservable quantities. Therefore, we will use all the available tenors, even though most authors only use 5-year premiums.

We select the companies listed under the “Financial” classification in the database. Further selection is performed by keeping only insurance and banking firms with at least two years of data; this step is accomplished using each firm's Standard Industrial Classification (SIC) main code. This study is thus based on 35 financial sector firms.

⁸For instance, see Altman et al. (2005).

⁹Ericsson et al. (2009), Tang and Yan (2007) and Longstaff et al. (2005) show that a significant portion of CDS spreads can be directly attributed to credit risk.

Table 3.1: **Insurance and banking firms.**

Insurance firms	Banking firms
ACE Limited (ACE) Allstate Corporation (ALL) American International Group, Inc (AIG) Aon Corporation (AOC) Berkshire Hathaway, Inc (BRK) Chubb Corporation (CB) Genworth Financial, Inc (GNWTH) Hartford Financial Services Group (HIG) Liberty Mutual Insurance Company (LIBMUT) Lincoln National Corporation (LNC) Loews Corporation (LTR) Marsh & McLennan Companies, Inc (MMC) MetLife, Inc (MET) Prudential Financial, Inc (PRU) Safeco Corporation (SAFC) XL Capital Limited (XL)	American Express Company (AXP) Bear Stearns Companies, Inc (BSC) Bank of America Corporation (BACORP) Capital One Financial Corporation (COF) Charles Schwab Corporation (SCH) Citigroup, Inc (C) Deutsche Bank AG (DB) Federal Home Loan Mortgage Corporation (FHLMC) Federal National Mortgage Association (FNMA) Goldman Sachs Group, Inc (GS) JPMorgan Chase & Co (JPM) Lehman Bros Holdings, Inc (LEH) Merrill Lynch & Co, Inc (MER) Morgan Stanley (MWD) SunTrust Banks, Inc (STI) US Bancorp (USB) WA Mut, Inc (WM) Wachovia Corporation (WB) Wells Fargo & Co (WFC)

The weekly term structure of CDS data starts on January 5, 2005, and ends on December 26, 2012, for a maximum of 417 observations. Premiums correspond to Wednesday data as it is the least likely day to be a holiday and is also less likely to be affected by weekend effects.¹⁰ The CDS's tier is chosen as senior and refers to the level of debt in the capital structure of the reference entities. Furthermore, the selected restructuring clause is XR, meaning that all restructuring events are excluded as trigger events.

Throughout the paper, firms are divided into two categories: insurance firms and banking companies.¹¹ Table 3.1 lists these various companies, including 16 insurance companies and 19 banking firms. The majority of these institutions are large publicly traded companies.

Figure 3.1 exhibits the weekly average 5-year CDS premiums for both subsectors, and the weekly average CDS term structure slope where, for a given firm, the slope is proxied by the difference between the 10-year and 1-year CDS premiums. Among the firms, AIG, Lincoln National and Washington Mutual have the largest average premiums, reaching maximum values of 3,336.2 basis points (bps), 2,695.5 bps and 5,207.8 bps for 5-year tenors respectively. During the sample period, the market considered AIG, Lincoln National and Washington Mutual the riskiest firms. This is consistent with the near-collapse of AIG, Lincoln National's stock drop and the failure of Washington Mutual, which was the largest commercial bank failure in American history. Conversely, Fannie Mae and Freddie Mac, for which CDSs data were considered up

¹⁰For more details on the advantages of using Wednesday data, see Dumas et al. (1998).

¹¹The range of SIC codes for insurance firms is between 6300 and 6499. The banking subsector's SIC code ranges from 6000 to 6299.

to September 2008, have the narrowest average premiums. Although the CDS holders triggered the default clauses for both entities, the debt was implicitly guaranteed by the U.S. government which mitigated the risk associated with these firms in the CDS market.

In addition to CDS data, the model requires other inputs such as the risk-free interest rate and the firms' initial leverages. The risk-free interest rate is assumed to be constant over time at 1.75%.¹² The leverages as of January 2005 are approximated from the total liabilities divided by the total assets of each firm in the sample.¹³




Figure 3.1: Evolution of the average CDS level in basis points and of the average CDS slope in basis points for both subsectors.

The CDS premiums were taken from Markit for the 16 insurance firms and 19 banking companies selected, between January 2005 and December 2012. The grey surface corresponds to the financial crisis (July 2007 to March 2009). The CDS level is proxied by the weekly average of 5-year CDS premiums. The CDS slope is proxied by the average difference between the 10-year CDS premiums and the 1-year CDS premiums.

Saunders and Allen (2010) break the recent financial crisis into three periods. The first period corresponds to the credit crisis in the mortgage market (June 2006 to June 2007), the second one covers the period of the liquidity crisis (July 2007 to August 2008), and the third the default crisis period (September 2008 to March 2009). This study focuses on the second and the third periods; thus, the financial crisis started in July 2007 and finished in March 2009 throughout this paper.

¹²This value represents the average rate of the daily 1-month and 3-month Treasury constant maturity series obtained from the Federal Reserve Bank of St. Louis (via FRED).

¹³More specifically, the firms' financial information is extracted from the Wharton Research Data Services (WRDS) Compustat database. In the database, the total liabilities are identified by LTQ and the total assets by ATQ.

3.4 Firm-Specific Credit Risk

Since leverage ratios and Markov regimes are unobservable variables, a filtering procedure is needed. We infer the latent variables from observable CDS premiums. However, estimating all firms simultaneously is not numerically feasible. The estimation is thus broken down into two stages. First, the firm-specific parameters are estimated. The second stage then focuses on the interrelation between firms while keeping the firm-specific parameters fixed. This approach is similar to the Inference Function for Margin (IFM) estimator proposed by Joe (2014). Also, an unreported Monte Carlo study shows that the two-stage approach produces unbiased estimators for all parameters.

If leverage time series were observable, the regimes could easily be filtered (for a review of classic methods, see Elliott et al., 1995). However, this is not the case and filtering regimes based on a latent time series is not straightforward. An extension of Tugnait's (1982) detection-estimation algorithm (DEA) is designed to filter both unobserved variables simultaneously.¹⁴ In addition to being adequate from a statistical point of view, this filter provides firm-specific model parameters based on maximum likelihood estimators. For more information on the method, refer to Bégin et al. (2016).

The set of Markov-switching parameters to be estimated for each firm in the first stage is

$$\phi_1 = (\mu^{\mathbb{P}}, \mu^{\mathbb{Q}}, \sigma_1, \sigma_2, p_{11}^{\mathbb{P}}, p_{22}^{\mathbb{P}}, p_{11}^{\mathbb{Q}}, p_{22}^{\mathbb{Q}}, \alpha, \beta, \theta, \kappa, \delta^{(1)}, \delta^{(2)}, \delta^{(3)}, \delta^{(5)}, \delta^{(7)}, \delta^{(10)})$$

where $\delta^{(1)}$, $\delta^{(2)}$, $\delta^{(3)}$, $\delta^{(5)}$, $\delta^{(7)}$, and $\delta^{(10)}$ are standard errors of the noise terms for tenors of 1, 2, 3, 5, 7 and 10 years, respectively. The filter-based methodology allows us to recover both real probability \mathbb{P} and risk-neutral \mathbb{Q} parameters. Descriptive statistics of the model parameters are presented in Table 3.2.¹⁵

The regime variable is an important feature of our model.¹⁶ Empirical results show

¹⁴To account for nonlinearities in the state-space representation, the unscented Kalman filter (UKF) of Julier and Uhlmann (1997) is applied instead of the classic Kalman (1960) filter.

¹⁵Note that we consider the same drift parameter across both regimes in our model. Indeed, the drift parameter estimators of the latent variable are rather inaccurate and create numerical instability due to the short span of the time series used. Even in a “one-regime” framework where the log-leverage is assumed to be observed, the precision of the drift parameter estimate is proportional to the square root of the sampling period length.

¹⁶We illustrate the advantages of the regime-switching model over the “one-regime” equivalent (i.e. the model proposed by Boudreault et al., 2014) in the Section 3.D. Overall, the regime-switching model

Table 3.2: Descriptive statistics on the distribution of firm-specific parameters and noise terms.

	$\mu^Q(\%)$	$\mu^P(\%)$	σ_1	σ_2	$p_{11}^Q(\%)$	$p_{11}^P(\%)$	$p_{22}^Q(\%)$	$p_{22}^P(\%)$	κ
Average	-0.074	0.008	0.070	0.347	99.570	97.675	93.987	96.617	0.565
SD	0.169	0.261	0.015	0.015	0.194	1.699	1.663	3.222	0.065
Minimum	-0.905	-0.242	0.032	0.306	98.932	90.481	89.098	87.055	0.443
10%	-0.163	-0.136	0.054	0.328	99.338	96.146	92.328	92.082	0.475
25%	-0.073	-0.085	0.058	0.339	99.462	96.765	92.936	95.036	0.517
50%	-0.039	-0.035	0.069	0.352	99.617	98.105	93.861	97.675	0.554
75%	-0.008	0.011	0.080	0.359	99.685	98.633	95.160	99.011	0.629
90%	0.027	0.071	0.088	0.360	99.795	99.207	96.178	99.548	0.646
Maximum	0.086	1.385	0.099	0.362	99.885	99.994	96.466	99.999	0.669
Insurance									
Average	-0.027	-0.065	0.075	0.352	99.537	97.864	93.759	96.157	0.579
SD	0.058	0.082	0.014	0.010	0.150	0.992	1.823	2.483	0.065
Banking									
Average	-0.114	0.069	0.065	0.343	99.598	97.516	94.179	97.004	0.553
SD	0.219	0.338	0.014	0.017	0.224	2.138	1.539	3.758	0.065
	α	θ	$\beta(\%)$	$\delta^{(1)}$	$\delta^{(2)}$	$\delta^{(3)}$	$\delta^{(5)}$	$\delta^{(7)}$	$\delta^{(10)}$
Average	10.724	1.349	0.088	0.244	0.142	0.086	0.052	0.035	0.063
SD	2.543	0.103	0.212	0.059	0.037	0.025	0.018	0.023	0.021
Minimum	7.260	1.166	0.000	0.143	0.062	0.036	0.009	0.006	0.034
10%	8.370	1.211	0.000	0.180	0.102	0.056	0.034	0.008	0.037
25%	8.780	1.269	0.000	0.198	0.120	0.072	0.041	0.017	0.046
50%	10.191	1.346	0.022	0.237	0.140	0.083	0.050	0.037	0.062
75%	11.864	1.411	0.092	0.284	0.161	0.100	0.059	0.044	0.071
90%	15.341	1.507	0.176	0.330	0.202	0.120	0.072	0.065	0.099
Maximum	17.648	1.574	1.238	0.349	0.215	0.145	0.104	0.103	0.120
Insurance									
Average	10.148	1.365	0.057	0.251	0.136	0.084	0.043	0.032	0.059
SD	1.482	0.109	0.073	0.061	0.038	0.027	0.012	0.016	0.015
Banking									
Average	11.209	1.336	0.114	0.239	0.147	0.088	0.059	0.038	0.066
SD	3.138	0.099	0.281	0.058	0.036	0.023	0.019	0.028	0.026

For each of the 35 firms, the parameters of the model are estimated using weekly CDS premiums with maturities 1, 2, 3, 5, 7 and 10 years, using the DEA-UKF filtering technique. The mean, standard deviation (SD) and quantiles are computed across firms. The last four rows compute the mean and SD across insurance and banking sectors. The δ s represent the standard deviation of the noise terms present in the filter's observation equation.

strong persistence for both low- and high-volatility regimes. Indeed, transition probabilities p_{11}^P and p_{22}^P are greater than 87% for all firms, with the majority exceeding 97%. In particular, Fannie Mae, Freddie Mac and Merrill Lynch transition probabilities p_{22}^P reach virtually 100%, suggesting permanent regime changes during the crisis. This is because CDS data are truncated at the effective acquisition date, which corresponds to the high-volatility regime. Both the insurance and banking subsectors tend to have similar transition probabilities on average.

The average uncertainty parameters related to the first and second regimes (σ_1 and σ_2) are about 7% and 35%, respectively, implying a large difference between the two

is statistically superior to its “one-regime” equivalent. For instance, the relative root mean square error of the “one-regime” model is 29% higher than the one of the regime-switching framework.

regimes. The univariate step procedure also allows the firm-specific constants α , β and θ , which define the intensity process of Equation (3.4), to be estimated. All firms have positive values for each constant. The estimated α has minimum and maximum values of 7.3 and 17.7, respectively, implying that the intensity process is strongly convex with the leverage ratio. The convexity of the relationship is higher on average for banking firms when compared with insurance companies. Finally, the critical leverage value θ lies between 1.17 and 1.57, which is realistic given that a portion of the default risk is captured by parameter β and the leverage ratio affects the default intensity in a nonlinear fashion.

Table 3.2 also shows the descriptive statistics of parameter κ , which is related to liquidation and legal fees. The estimated value across firms ranges between 44% and 67%, and represents a fraction of the market asset value at default.

Standard errors of the trading noise are relatively low for tenors of longer than two years with an average value lying between 3.5% and 8.6%. However, short tenors have higher variations that may be related to lower trading frequency of 1- and 2-year CDS contracts. One can also mention the very high-volatility period during which the analysis is performed, implying higher standard errors than a stable period would generate. The average standard errors are comparable across both subsectors.

Figure 3.2 depicts the proportion of firms in the high-volatility (turbulent) regime across both insurance and banking subsectors. This proportion raises rapidly at the onset of the crisis for banking firms: it goes from 21% to 84% in the first six months of the crisis, with a sizable increase in the week following the credit crunch (A). The transition for insurance companies happens later in early 2008: the proportion of firms in the high-volatility regime is virtually 100% from March 2008 to September 2008. For both subsectors, there is some persistence in the proportion during the post-crisis era. This observation is consistent with volatility regime persistence noted in Maalaoui Chun et al. (2014), Garzarelli (2009) and Mueller (2008). Interestingly, the banking subsector's proportion of firms in the high-volatility regime increases during the European debt crisis (from 2009 to 2012). For the same years, the insurance subsector's proportion remains at zero. Finally, even though the filtered statistical regimes depend only on firm-specific information, they suggest a rather important link with the crisis, on average.




Figure 3.2: **Time series of the proportion of firms in the high-volatility regime across both insurance and banking subsectors.**

Based on the firm-specific parameters, the most probable regimes are extracted and aggregated across the two sectors. The different letters correspond to major events during the crisis: (A) The credit crunch begins in earnest (August 1, 2007). (B) The Federal Reserve Board approves the financing arrangement between JPM and BSC (March 14, 2008). (C) LEH files for Chapter 11 bankruptcy protection. MER is taken over by the BACORP. AIG almost defaulted the next day (September 15, 2008). (D) Three large U.S. life insurance companies seek TARP funding: LNC, HIG and GNWTH (November 17, 2008). (E) The U.S. Treasury Department, Federal Reserve, and FDIC announce a package of guarantees, liquidity access, and capital for BACORP (January 16, 2009).

3.4.1 Default Probabilities

The evolution of PDs estimated by the model is investigated (hereafter PD_{model}). This quantity is related to the first “L” of financial crises: leverage. The credit risk framework links default probabilities to firms’ leverage ratios through the intensity process described in Equation (3.4). Since the firm’s leverage is not directly observable from market data, CDS premiums are used to infer the model’s latent variables (i.e. hidden regimes and leverages). Therefore, the model estimates a forward-looking measure of the firm-specific default probability.¹⁷

Throughout this subsection, we compare the model’s estimates with PD computed using a default count approach (PD_{DC}). The latter are based on historical data rather than current market conditions. Default counts are aggregated over time by rating categories across the banking, finance and insurance industries from January 2002 to December 2012 in transition matrices, which can be compounded for multiple periods to produce n -year default probabilities.¹⁸ Finally, Moody’s ratings, extracted for the

¹⁷The model estimates are computed using the trinomial lattice approach and the estimated parameters under the physical measure \mathbb{P} .

¹⁸A generator estimation approach with a window length of three years ex ante data is used. See Dionne et al. (2010) for more details.

35 firms on a monthly basis from January 2005 to December 2012, allow us to readily obtain the PD_{DC} . Default probabilities computed using this approach are not firm-specific, but depend on aggregated information across firms with the same rating.

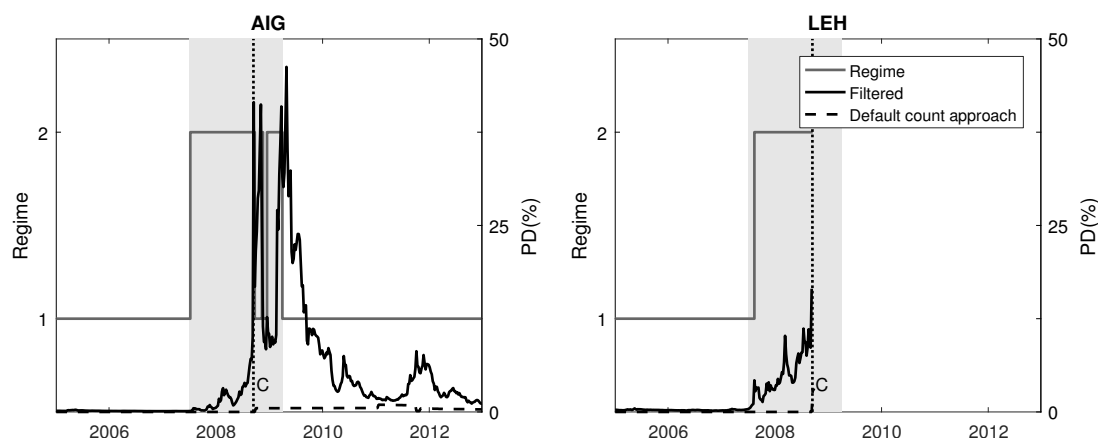


Figure 3.3: One-year default probabilities computed using the credit risk model and the default count approach and filtered regimes for AIG and LEH.

This figure shows the time series of 1-year default probabilities over the period of time 2005–2012. The time series of the model are inferred from CDS premiums market data. The default count approach time series are obtained from monthly transition matrices for banking, finance and insurance industries using the generator estimation approach with window length of three years ex ante default data. The letter (C) corresponds to a major event during the crisis: LEH files for Chapter 11 bankruptcy protection. MER is taken over by the BACORP. AIG almost defaulted the day after (September 15, 2008).

Below, we examine the time-varying behaviour of PD_{model} by focusing our analysis on a few important firm-specific events of the last crisis. The results for AIG and LEH are given in Figure 3.3 and firm-specific averages across the three periods are shown in Table 3.3. Let us first take the case of AIG, which almost defaulted on September 16, 2008. One month prior to that date, the model derives 1-year PD of approximately 6% for AIG. Then, the estimate reaches 10% on September 10, followed by a spike of 42% one week later. A similar behaviour is observed for Lehman Brothers prior to its collapse on September 15, 2008. Indeed, high levels are reached four months prior to the bankruptcy event (9% for the 1-year PD_{model}), followed by a jump in the probability of default of approximately 11% on September 10, 2008. When it comes to acquired firms such as Bear Stearns, Merrill Lynch, Wachovia and Washington Mutual, the same characteristic jump pattern is displayed close to major events preceded by relatively large PDs. Moreover, one can observe higher estimates when examining PD measures of firms that have been acquired during the crisis and distressed firms in comparison to the others in the sample.

In opposition, PD_{DC} suffers from two main caveats. First, as it is computed using a rolling window and past data, it tends to lag behind the market, which explains why the

probabilities are smaller during the crisis than in its aftermath (see Table 3.3). Second, it does not reflect firm-specific default probabilities, given that it is based on sector-wide aggregated data. This is another reason why many firms have the same PD_{DC} in Table 3: firms tend to have the same credit rating during that period. Contrarily, the CDS-implied default probabilities are forward looking and strongly reacts at the onset of the financial crisis and during the European crisis of 2012.

Figure 3.4 exhibits persistence of high PDs in the aftermath of the Great Financial Crisis. It also shows PD_{model} and PD_{DC} for 1-year horizons, averaged across both subsectors and various events that happened during the sampling period. Before the crisis, the average PD_{model} of both banking and insurance subsectors are at the same level. The average PD_{model} of the banking sector starts to rise just before the onset of the crisis and jumps at the credit crunch (Event A in Figure 3.4). The average levels of insurance subsector PD_{model} have been less affected at the beginning of the crisis, but strongly react halfway through, reaching levels higher than those of the banking subsector. Indeed, between September 3, 2008 and October 8, 2008, the insurance subsector's average PD_{model} increased by 6.5%, while the banking sector's average only increased by 2.7%. During that month (C in Figure 3.4), Lehman Brothers went bankrupt, Bank of America bought Merrill Lynch and the Federal Reserve Board authorized the Federal Reserve Bank of New York to lend up to \$85 billion to AIG; these events could explain this increase to some extent. Also, other insurance firms such as Genworth, Lincoln National, Hartford Financial Services and XL Capital have a large PD_{model} in the second half of the crisis period.¹⁹ Note that these firms were rather compromised at the end of the crisis, bringing the average to higher levels. On November 17, 2008 (D in Figure 3.4), three large U.S. life insurance companies seek TARP funding; accordingly, the insurance subsector's PDs decrease for a couple of months. There is modest persistence in the aftermath of the crisis as PD_{model} slowly reverts back to a level still above that measured at the beginning of the sampling period. Then, by mid-2009, both subsectors had similar 1-year PD_{model} ; however, the banking subsector's probabilities slightly increase during the European debt crisis.

On average, the insurance subsector 1-year levels of PD_{model} are lower than those of the banking subsector: the difference between the two subsectors is 0.17% in the pre-crisis era, 0.40% during the crisis, and 0.34% afterwards. Even though some insurance firms are quite exposed during the crisis (e.g. companies selling bond insurance and CDS), some are less affected by the turmoil (e.g. property and casualty insurers).

¹⁹ Average increases of 32%, 14%, 8% and 10%, respectively, from the first to the second half of the crisis period for 1-year PDs.

Table 3.3: Descriptive statistics of one-year default probability estimates across periods.

Firm	Pre-crisis		Crisis		Post-crisis	
	Filtered	Default count	Filtered	Default count	Filtered	Default count
ACE	0.570	0.000	1.529	0.144	1.708	0.353
ALL	0.249	0.000	1.502	0.144	1.069	0.353
AXP	0.307	0.000	4.807	0.144	2.382	0.353
AIG	0.139	0.000	7.458	0.144	6.498	0.548
AOC	0.499	0.004	0.905	0.269	1.474	0.792
BSC	0.668	0.000	6.486	0.001	-	-
BRK	0.394	0.000	3.674	0.000	3.812	0.051
BACORP	0.280	0.000	2.158	0.059	4.889	0.474
COF	1.035	0.003	8.860	0.144	3.593	0.790
SCH	0.318	0.000	1.315	0.144	0.972	0.353
CB	0.211	0.000	1.199	0.144	0.944	0.353
C	0.226	0.000	3.257	0.080	4.582	0.401
DB	1.233	0.000	1.494	0.000	-	-
FHLMC	0.121	0.000	1.434	0.000	-	-
FNMA	0.085	0.000	0.614	0.000	-	-
GNWTH	0.542	0.000	17.912	0.239	11.438	0.792
GS	0.297	0.000	2.662	0.086	3.045	0.353
HIG	0.139	0.000	3.974	0.179	3.822	0.792
JPM	0.550	0.000	3.807	0.009	3.821	0.053
LEH	0.257	0.000	4.918	0.135	-	-
LIBMUT	0.329	0.004	2.092	0.269	2.102	0.792
LNC	0.205	0.000	6.237	0.149	4.328	0.792
LTR	0.438	0.003	1.283	0.144	1.676	0.353
MMC	0.771	0.004	1.537	0.269	1.667	0.792
MER	0.695	0.000	6.407	0.139	6.567	0.474
MET	0.283	0.000	3.207	0.144	3.050	0.353
MWD	0.934	0.000	9.841	0.144	7.985	0.401
PRU	0.174	0.000	5.322	0.149	3.321	0.792
SAFC	0.183	0.004	1.057	0.269	2.386	0.792
STI	0.197	0.000	2.122	0.144	3.127	0.783
USB	0.201	0.000	1.450	0.009	1.574	0.052
WM	1.753	0.000	17.147	0.975	-	-
WB	0.199	0.000	3.020	0.084	-	-
WFC	0.284	0.000	0.921	0.009	1.038	0.353
XL	0.459	0.000	7.696	0.268	4.917	1.040
Insurance	0.348	0.001	4.161	0.183	3.294	0.586
Banking	0.514	0.000	4.561	0.151	3.631	0.403

This table shows descriptive statistics for 1-year default probabilities for each firm, across the different periods. Model time series are inferred from CDS premium market data. The default count approach time series are obtained from monthly transition matrices for the banking, finance and insurance industries using the generator estimation approach with a window length of three years ex ante default data. For some firms, no default probabilities are available during the post-crisis era. These are firms that either defaulted or were acquired prior to the end of the crisis. All values are reported as a percentage.

3.4.2 Recovery Risk

The second “L” of financial crises is losses. It is modelled implicitly in the credit risk framework through the endogenous recovery rate of Equation (3.6) and depends on the firm’s financial health. Thus, the recovery rate changes over time and from one firm to another.

Figure 3.5 exhibits the average 1-year expected recovery rate for both subsectors. In general, the expected recovery rate is lower for insurance firms over a 1-year time horizon. The average 1-year expected rate for insurance and banking subsectors is 43.6% and 49.3%, respectively. Across both subsectors, there is a decrease in the




Figure 3.4: Average one-year default probabilities computed using the credit risk model and the default count approach.

This figure shows the time series of 1-year average default probabilities across the portfolio over the 2005–2012 period. Model time series are inferred from CDS premium market data. The default count approach time series are obtained from monthly transition matrices for the banking, finance and insurance industries using the generator estimation approach with a window length of three years ex ante default data. The different letters correspond to major events during the crisis: (A) The credit crunch begins in earnest (August 1, 2007). (B) The Federal Reserve Board approves the financing arrangement between JPM and BSC (March 14, 2008). (C) LEH files for Chapter 11 bankruptcy protection. MER is taken over by the BACORP. AIG almost defaults the day after (September 15, 2008). (D) Three large U.S. life insurance companies seek TARP funding: LNC, HIG and GNWTH (November 17, 2008). (E) The U.S. Treasury Department, Federal Reserve, and FDIC announce a package of guarantees, liquidity access, and capital for BACORP (January 16, 2009).

recovery rate over time: from 54.6% in the pre-crisis era to 41.9% during the crisis, on average.

Lastly, note that the average recovery rate calculated in this study is consistent with those of Altman et al. (2005) and Vazza and Gunter (2012). Indeed, Altman et al. (2005) find an average recovery at default of 53% and 35% for senior secured and unsecured bonds, respectively. Also, in Vazza and Gunter (2012), senior secured and unsecured bonds have an average discounted recovery rate of 56.4 and 42.9%, respectively, during the 1987–2012 period.

3.5 Dependence

Through the regime-dependent leverage correlation, we account explicitly for potential linkage between the various firms investigated. This dimension is important in modelling financial crises.



Figure 3.5: Time series of the average one-year expected recovery rate for insurance and banking companies.

Based on filtered regimes, we compute the average one-year expected recovery rate each week for each firm from Equation (3.6). We then take the sample average across both sectors.

At this point, we consider it relevant to stress that the linkage between the firms in our framework has two dimensions. In a direct manner, the correlation induces links between the firm's leverages, and ergo, their default probabilities. This would increase the likelihood of default clusters for positively correlated firms in periods of turmoil. Also, in an indirect way, the potential losses are also correlated, as they depend on the firms' financial health. Therefore, troubled firms that are highly linked (i.e. large positive correlation) would have recovery rates that decrease at the same time.

Below, we discuss some of the results obtained from the multivariate extension of the univariate Markov-switching framework. As a starting point to the multivariate step, suppose that we have N firms across the portfolio and correlations are recovered from leverage ratios of all possible pairs of firms (i, j) , with $1 \leq i, j \leq N$. Thus, the number of estimated values is $N(N - 1)/2$ for each regime state leading to $2N(N - 1)$ total values. The set of parameters for the bivariate estimation stage is

$$\phi_2 = (\rho_{1,1}^{(i,j)}, \rho_{1,2}^{(i,j)}, \rho_{2,1}^{(i,j)}, \rho_{2,2}^{(i,j)})$$

for each pair of firms. Since the leverage ratio time series are inferred from the set of CDS premiums by the DEA-UKF methodology, recovering a correlation from smoothed leverage data would result in underestimated coefficients. Therefore, dependence among firms must be captured endogenously or prior to the filtering process. Details on the estimation of endogenous correlation coefficients are presented in Section 3.A.

At this moment, we feel the need to stress that the estimated correlation coefficients might be larger than the levels typically seen in credit risk models. Three reasons explain these differences: the rather challenging sampling period, the fact that we use CDS premiums instead of equity returns to estimate the coefficients and an estimation technique that accounts for the presence of noise in market prices.²⁰

The heat maps of Figure 3.6 summarize $\rho_{1,1}$, $\rho_{1,2}$, $\rho_{2,1}$ and $\rho_{2,2}$ for each pair of firms.²¹ The results highlight positive pairwise correlations when both firms are in the same regime, with some minor exceptions for Charles Schwab and Deutsche Bank (i.e. seven coefficients out of 1,190 coefficients estimated are negative). In the stable regime (top left panel), the top left 16×16 correlations suggest a higher degree of interconnectedness in the insurance subsector. Regarding the banking firms, the bottom right 19×19 coefficients display more heterogeneity. Also, Freddie Mac and Fannie Mae strongly move together, but are not significantly connected to the rest of the subsector.

Results also display a higher degree of leverage interdependence when the regime switches from stable to volatile regimes for both entities: accordingly, the bottom-right panel of Figure 3.6 (volatile regime) is much darker than the top-left one (stable regime).

Correlation coefficients are lower for pairs of firms that are not in the same regime (top-right and bottom-left panels of Figure 3.6). As one of the two firms is more uncertain than the other, it makes sense for this pair to be less interconnected.

For the rest of this section, we break the sample down into three categories: correlations between the leverages of two insurance firms (Insurance/Insurance), correlations between the leverages of two banking firms (Banking/Banking) and correlations between the leverages of one insurance firm and one banking company (Insurance/Banking).

²⁰In a one-regime framework, Boudreault et al. (2015) compare correlation coefficients obtained using filtering schemes to the ones based on equity returns. The authors conclude that equity return correlations could be different from the ones inferred from filtering methods while using CDS premiums as inputs.

²¹In some specific cases, CDS premiums dynamics as well as univariate parameter estimates are such that the joint probability of pairwise firms to be in different regimes is too low. Consequently, the estimation procedure is unable to adequately recover the coefficients leading to missing values in the correlation matrix (white squares in the heat maps). We were unable to adequately recover 42 coefficients out of 2,380, which is less than 2% of all the correlation coefficients. As an example, one can see that during the period of time Fannie Mae and Freddie Mac are in the first regime, few institutions belong to the second one. Conversely, when both government-sponsored entities are in the high-volatility regime, the model is able to estimate correlations for almost all pairwise combinations.




Figure 3.6: Heat maps of $\rho_{1,1}$, $\rho_{1,2}$, $\rho_{2,1}$ and $\rho_{2,2}$ for the 35 firms.

This figure shows $\rho_{1,1}$, $\rho_{1,2}$, $\rho_{2,1}$ and $\rho_{2,2}$. The first 16 rows and columns correspond to the insurance subsector and last 19 ones to banks. Values in white are correlation coefficients that could not be estimated (i.e. not enough data).

Figure 3.7 shows the histogram of $\rho_{1,1}$, $\rho_{1,2}$, $\rho_{2,1}$ and $\rho_{2,2}$ for the three categories: Insurance/Insurance, Banking/Banking and Insurance/Banking. For the three categories, there is an increase in the average correlation when both firms move from the stable to the volatile regimes. For correlation coefficients between two insurance firms, the average goes from 60% for both firms being in the first regime to 80% in the second one, for an increase of about 20%. For banking firms' correlation coefficients, the average increase is about 26%, from 54% to 79%. As shown in Section 3.4, this regime is associated with the last financial crisis for most firms. The correlation between the

leverages of one insurance firm and one banking company is lower in general, with averages of 45% and 74% for the stable and volatile regimes, respectively.

Interestingly, if one firm is in the first regime and the other is in the second, correlations seem lower than if both firms are in the same regime. Therefore, if firms are in different regimes, it means that one is going worse than the other and their comovements should be less related (i.e. firms should be less interconnected).

Roughly speaking, firms become much more interconnected in the high-volatility regime. Also, the general shape of the empirical distribution of correlation coefficients also changes considerably from one regime to the other.

For insurance firms (i.e. Insurance/Insurance), the stable regime correlation coefficients are distributed around its average and the empirical distribution is unimodal. For the turbulent regime, the distribution becomes left-skewed and its mode shifts to the right, meaning that the majority of insurance firms are highly correlated.

For banking firms (i.e. Banking/Banking), the low-volatility regime empirical distribution displays bimodality. This could be explained by two clusters of banking companies. The top-left panel of Figure 3.6 shows that the first-regime correlations is much more heterogeneous for financial institutions: some banks are largely correlated, while others exhibit lower levels of dependence. This dependence behaviour could be the very consequence of a bank's primary activities. To verify this conjecture, we further divide the banking subsector into two categories: correlations between two investment banks and correlations between two commercial banks. We find that the average correlation is about 44% when both firms are commercial banks and 57% when both firms are investment banks. These averages are consistent with the clustering mentioned above. In calm times, commercial banks are less interconnected whereas investment banks are more correlated. This result is a consequence of the riskier nature of activities carried out by investment banks, when compared to commercial banks. During the high-volatility regime, even firms that have low correlation in the stable regime are now highly interconnected. The empirical distribution of turbulent regime correlation coefficients is unimodal and left-skewed (Figure 3.7).

For correlation coefficients between the leverages of insurance and banking companies (i.e. Insurance/Banking), the first regime distribution displays lower correlation than the two other categories. However, during the turbulent regime, correlation increases and the distribution is also left-skewed. Interestingly, this would mean that firms less interconnected in the stable regime could be highly correlated in the turbulent one.

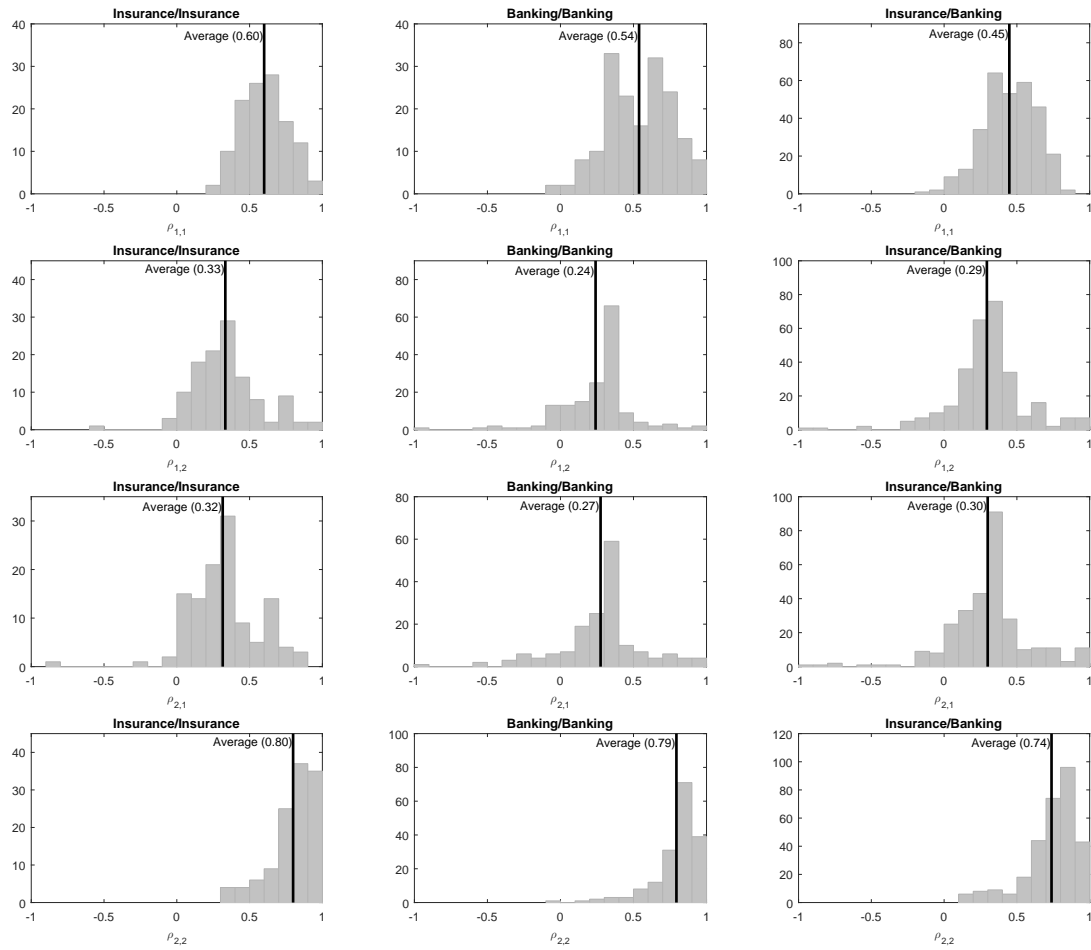


Figure 3.7: Histogram of $\rho_{1,1}$, $\rho_{1,2}$, $\rho_{2,1}$ and $\rho_{2,2}$ for three categories: correlations across insurance firms' leverages (Insurance/Insurance), correlations across banking companies' leverages (Banking/Banking), and correlations between insurance and banking firms' leverages (Insurance/Banking).

These figures show the empirical distribution of the $\rho_{1,1}$, $\rho_{1,2}$, $\rho_{2,1}$ and $\rho_{2,2}$ for three groups. The horizontal bar represents the sample mean.

Figure 3.8 exhibits the time series of median leverage correlation coefficients across firms and for the three categories of subsectors.²² As expected, the median correlation increases during the crisis, and decreases afterwards. Over 2005–2012, the banking subsector's correlations are larger, with a median about 3% higher than the insurance subsector correlations on average. The Insurance/Insurance and Banking/Banking curves are similar in the pre-crisis era. However, at the onset of the crisis, the Banking/Banking median correlation increases rapidly as the Insurance/Insurance curve remains somewhat similar. Then, at the beginning of 2008, the Insurance/Insurance

²²We take the median since it is less influenced by extreme values. The average would produce a similar pattern, however the series would be more volatile.




Figure 3.8: Time series of median leverage correlation coefficients for three groups: correlations across insurance firms (Insurance/Insurance), correlations across banking companies (Banking/Banking), and correlations between insurance and banking firms (Insurance/Banking).

Based on filtered regimes, we compute the median correlation coefficients across the three groups.

median correlation coefficient starts rising as the crisis becomes much more systemic. Note that there is some persistence in the post-crisis era for the three time series.

In summary, empirical results show that firms' leverages are more correlated during the high-volatility regime, suggesting more dependence within both subsectors during the last crisis. These results present major implications for risk management practices since the increased dependence could lead to important consequences in credit-sensitive portfolios. They would also have a major impact on systemic risk measures.

3.6 Systemic Risk

The model of Section 3.2 coupled with estimated parameters of Sections 3.4 and 3.5 is able to adequately grasp the systemic risk embedded in the financial services firms, as it was carefully constructed to capture three of Billio et al.'s (2012) "L"s: leverage, loss and linkage. In this study, the systemic risk measure is defined as the expected value of a loss over a period of three months given that it is higher than the 99th percentile of the loss distribution.²³ It is related to the systemic risk measure of Acharya et al. (2010), which is a function of the marginal expected shortfall. Our measure is also related to the one adopted by Huang et al. (2009) and Chen et al. (2014); however, in their respective studies, the conditional expectation threshold is determined by a

²³This systemic risk measure is analogous to the concept of expected shortfall or conditional tail expectation.

fraction of the total liabilities. Moreover, their prices of insurance against distressed losses are computed under the risk-neutral measure. In this study, we use the physical probabilities, as we focus on real-life expected losses. Similarly to Huang et al. (2009) and Chen et al. (2014), our measure is forward-looking as it is based on CDS data and does not require a large sample of firms.²⁴

We further divide our systemic risk measure into two components: the contribution of insurance and banking firms, respectively. These components correspond to the notion of marginal expected shortfall and could be described as the subsector's losses when the whole financial service sector is doing poorly.

To this end, we construct a theoretical debt portfolio that includes the total liabilities of each financial firm. The value of the total liabilities (i.e. LTQ) is extracted from the Compustat database; it is available for each quarter, in millions. Therefore, we use linear interpolation to obtain the total liabilities value for each week. The sum of all firm's liabilities is about $\$ 1.9 \times 10^7$ millions, on average.

As the framework is rather complicated, we rely on a Monte Carlo procedure to calculate the measures of systemic risk (*SR*). Section 3.B provides the steps to compute the systemic risk measure along with the systemic risk contributions for each subsector which represent the subsector's respective expected losses given that the whole sector's losses are higher than the 99th percentile. Intuitively, it informs us about each subsector's systemic risk importance.

The relative contributions *RSR* are scaled versions of the nominal price measures. Practically, we simply divide the contributions by the sum of the total liabilities for each respective subsector. This allows us to compare the relative systemic risk contributions of each subsector readily as they share the same scale.

3.6.1 Systemic Risk Measures

Before commenting on each subsector's systemic risk contributions, we assess the importance of the loss and correlation assumptions made in the framework of Section

²⁴In fact, the evolution of our measure is very similar to the price of financial distress (DIP) of Huang et al. (2009) even if both measures are based on somewhat different quantities. Our measure is also highly correlated with the DIP; the correlation between *RSR* and the DIP at 5% is 79.5% and 91.2% for insurance and banking subsectors, respectively.

3.2. To do so, the systemic risk measure SR is computed using different modelling assumptions such as independence between firms' leverage ratios and constant recovery rates.²⁵

For each modelling assumption stated above, Table 3.4 exhibits the average systemic risk contribution across different eras. Before the crisis, the expected loss, given that the 99% VaR has been reached, is estimated at \$ 23,846 billion which corresponds to 0.9% of the total liabilities across the 35 firms. During the crisis, the conditional expected loss rose to \$ 113,959 billion of dollars, or 3.9% of the total liabilities.

During the pre-crisis period, using an endogenous recovery rate increases the measures by 6% on average.²⁶ Using the regime-dependent correlations have only minor impacts on the systemic risk contributions during this period: when the sector is healthy as a whole, correlation should not have large repercussions on loss distributions as defaults rarely occur.

During the crisis era, the regime-dependent correlations and the endogenous recovery rate both have a major impact on the systemic risk measure: on average, they increase the systemic risk contributions by factors of 14% and 11%, respectively. The dependence assumption has important consequences on the insurance subsector, increasing the risk contributions by 23% on average. For the same subsector, the endogenous recovery assumption increases the measures by 3%. For the banking subsector, endogenous recovery has a major impact on RSR with an average rise of 18%. Regime-dependent correlation increases the measure by a factor of 6% on average.

During the post-crisis period, the correlation assumption has the most significant effect on the insurance subsector's measures, with increases of 7% on average. For the banking subsector, the endogenous recovery assumption is the most significant one: the contributions increase by a factor of 9% on average.

In summary, it is now clear that linkages and losses are essential in explaining the rise of systemic risk during the last financial crisis. These two financial crises' "L"s have important ramifications on the measures for both subsectors. For instance, independence completely underestimates the tail risk in general, and especially in periods of turmoil. Thus, regime-dependent correlation is important in explaining systemic risk during periods of crisis.

²⁵In practice, we use the average recovery rate instead of the endogenous recovery rate of Equation (3.6), which varies with the leverage ratio. It removes the negative correlation between default probabilities and recovery rates.

²⁶This is found by taking $\frac{1}{4}(0.902/0.842 - 1 + 0.902/0.843 - 1 + 0.906/0.862 - 1 + 0.905/0.862 - 1) = 0.0605$.

Table 3.4: Average systemic risk measures on three different periods (i.e. pre-crisis, crisis and post-crisis) using different modelling assumptions.

Panel A: Insurance subsector						
	Pre-crisis		Crisis		Post-crisis	
	Nominal	Unit (%)	Nominal	Unit (%)	Nominal	Unit (%)
Regime correlation, endogenous recovery	23,846.2	0.902	113,959.0	3.976	64,550.8	2.301
Regime correlation, exogenous recovery	22,256.4	0.842	107,442.7	3.737	63,126.4	2.250
Independence, endogenous recovery	23,849.9	0.902	91,074.6	3.150	59,771.0	2.126
Independence, exogenous recovery	22,273.8	0.843	90,886.8	3.137	59,308.5	2.111
Panel B: Banking subsector						
	Pre-crisis		Crisis		Post-crisis	
	Nominal	Unit (%)	Nominal	Unit (%)	Nominal	Unit (%)
Regime correlation, endogenous recovery	101,571.4	0.906	782,154.6	6.166	961,985.4	9.789
Regime correlation, exogenous recovery	96,387.0	0.862	659,127.7	5.188	880,730.0	8.955
Independence, endogenous recovery	101,468.7	0.905	735,895.8	5.790	952,787.0	9.694
Independence, exogenous recovery	96,373.3	0.862	631,118.2	4.964	875,675.5	8.903

The theoretical debt portfolio includes the total liabilities of each financial firm. The value of the total liabilities (i.e. LTQ) is extracted from the Compustat database. The systemic risk measures are computed using Monte Carlo methods and 5×10^5 paths over a span of three months. Each systemic risk measure is computed for four different scenarios: regime-dependent correlation with endogenous recovery rates (full model), regime-dependent correlation with endogenous recovery rates, independence assumption with endogenous recovery rates, and independence assumption with exogenous recovery rates. Systemic risk measures in nominal units are given in millions.

Now focusing on the contributions given by the model described in Section 3.2, we display the time series of nominal and unit price contributions in Figure 3.9. The top panel represents the systemic risk contributions in nominal terms. The banking subsector contribution always lies above the insurance subsector's time series: not surprisingly, the banking subsector's total liabilities are larger than those of the insurance subsector, implying larger marginal expected shortfalls.

The unit price contributions are given in the bottom panel of Figure 3.9. The two subsectors' contributions are similar during the pre-crisis era. However, at the beginning of the crisis, the banking contribution rises quickly, capturing the increase in systemic risk for this subsector. Halfway through the crisis, the insurance contribution jumps from 1% to almost 20%: this rise is consistent with AIG's near-default and the increased credit risk in Lincoln National, XL and Genworth. In the post-crisis era, the banking systemic risk remains high and increases from 9% to 15% during the European debt crisis. The insurance subsector's contribution slowly decreases to reach pre-crisis levels at the end of the sample.

The sample correlation between the relative systemic risk contribution of insurers and banks, RSR_t^{Ins} and RSR_t^{Bnk} , respectively, is about 30% for the 2005–2012 period, which is rather low. When we consider only the crisis era, this sample correlation escalates

to 69%, implying large comovements in both subsectors' contributions. For the remainder of the paper, we focus on the unit price contributions.




Figure 3.9: Time series of the systemic risk measure contribution and the relative systemic risk measure contribution using the full model.

The systemic risk measures are computed using Monte Carlo methods and 5×10^5 paths over a span of three months. The theoretical debt portfolio that includes the total liabilities of each financial firm.

3.6.2 Granger Causality Tests

At this point, it would be interesting to look at causality: as in Chen et al. (2014), we would like to test whether a subsector's contribution could be used to forecast the other's systemic risk. As a starting point, linear Granger (1969) causality tests are employed. The latter involves F -tests to determine whether lagged data on a variable Y provides any statistically significant information on another variable X in the presence of lagged values of X . In this spirit, the null hypothesis of this statistical test should read: Y does not Granger-cause X .

Even though this test is very popular in the empirical literature, the linear Granger causality test does not capture nonlinear and higher-order causal relationships. To grasp these nonlinear effects, we also use nonlinear Granger causality tests. A general version of the nonlinear causality tests was first developed by Baek and Brock (1992) and then modified by Hiemstra and Jones (1994). However, Diks and Panchenko

(2006) show that the Hiemstra and Jones's (1994) statistical test could overreject the null hypothesis given that the rejection probabilities may tend to one as the sample size increases. They also propose a new nonparametric test for nonlinear Granger causality that avoids the over-rejection issue. Therefore, in this paper, we use Diks and Panchenko's (2006) (hereinafter DP) statistic to test causality in the nonlinear case.

Granger causality tests require stationary time series; however, the subsector systemic risk measures are both non-stationary. By visual inspection of Figure 3.9, it is explicit that these series are not stationary.²⁷ Therefore, we difference both series. ADF tests are done on the differenced time series and the null hypothesis is rejected for both series this time.²⁸

Also, as noted in Hiemstra and Jones (1994), heteroskedasticity could lead to a substantial bias. By visually inspecting the autocorrelation functions of squared differenced contributions (top panels of Figure 3.10), we conclude that there is conditional heteroskedasticity in both time series. We follow Chen et al. (2014) and deal with it by using a generalized autoregressive conditional heteroskedasticity (GARCH) model. For insurance and banking subsectors' contributions, we estimate a GARCH(1,1) model and extract the Gaussian noise processes. To assess if there is any residual heteroskedasticity, we plot the autocorrelation functions of the squared noise terms (bottom panels of Figure 3.10). It seems that the GARCH(1,1) model sufficiently accounts for the conditional heteroskedasticity in the original time series.

Using the noise process for both subsectors' contributions (i.e. post-GARCH filtering), we run the linear and nonlinear Granger causality tests. Table 3.5 shows the various results for both causality tests and for both subsectors.²⁹

In terms of linear Granger tests, the systemic risk of banking firms causes the systemic risk of insurance companies. For the opposite relationship, we cannot reject the null hypothesis: we cannot conclude that the systemic risk of insurers Granger-causes the systemic risk of banks.

For the nonlinear case, the banking subsector's systemic risk only Granger-causes the insurance subsector's systemic risk when the lag length is equal to one (at a confidence

²⁷Indeed, augmented Dickey-Fuller (ADF) tests fail to reject the null hypothesis for both systemic risk contributions.

²⁸For the insurance subsector: ADF-statistic of -26.45 and a p -value below 0.1%. For the banking subsector: ADF-statistic of -24.24 and a p -value below 0.1%.

²⁹Section 3.C provides results for linear and nonlinear Granger tests based on differenced time series (i.e. before GARCH filtering).

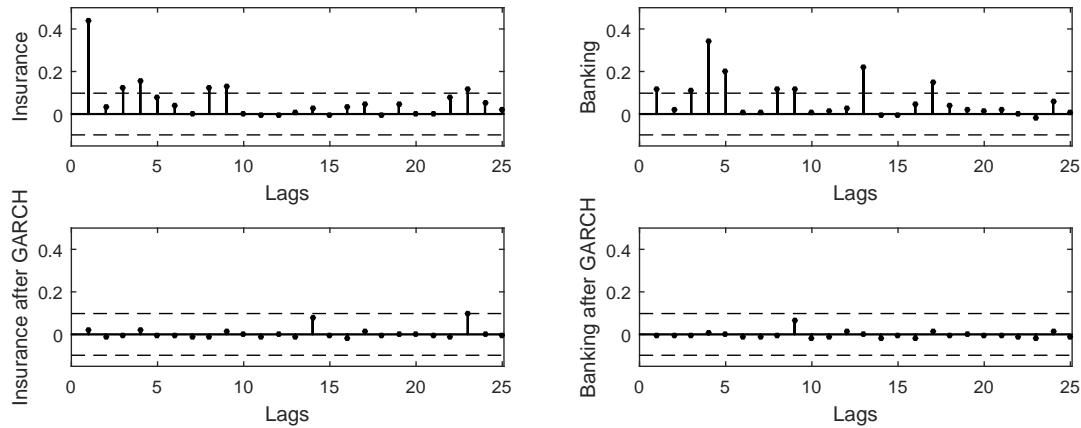


Figure 3.10: **Sample autocorrelation functions for squared differenced systemic risk contributions and squared GARCH noise terms.**

The systemic risk measure contributions are computed using Monte Carlo methods and 5×10^5 paths over a span of three months. The theoretical debt portfolio that includes the total liabilities of each financial. The GARCH noise terms are computed by fitting a GARCH(1,1) model to the differenced systemic risk contributions and by extracting the Gaussian noise terms.

level of 95%). However, the insurer's systemic risk does not Granger-cause the bank's systemic risk for any lag length.

These results are in line with Chen et al. (2014): there is a unidirectional causal effect from banks to insurers when accounting for heteroskedasticity while the opposite relationship (from insurers to banks) is not statistically significant. Therefore, even if our methodology differs and our data extend over the aftermath of the crisis, our results suggest that the direction of the causal relationship is robust. The systemic risk results are also consistent with Billio et al. (2012) who find that banks tend to have a much more important role in the transmission of shocks.³⁰

3.7 Concluding Remarks

Unlike conventional empirical studies of credit risk, this paper focuses on the financial services sector. To adequately model three financial crises "L"s, a Markov-switching extension of the hybrid credit risk model of Boudreault et al. (2014) is proposed. The latter allows for firm-specific statistical regimes that accommodate for changes in the leverage uncertainty, an endogenous stochastic recovery rate that is negatively related

³⁰As a robustness test, we use a GJR-GARCH instead of a GARCH(1,1) model. We find that the unidirectional relationship from banks to insurers in terms of Granger-causality is robust to this change of the conditional variance dynamics.

Table 3.5: Linear and nonlinear Granger causality tests after GARCH filtering.

Panel A: Linear Granger causality.							
X = Insurance, Y = Banking				X = Banking, Y = Insurance			
L_X	L_Y	F	p -value	L_X	L_Y	F	p -value
1	1	4.991	0.026	1	1	1.464	0.227
Panel B: Nonlinear Granger causality using Diks and Panchenko (2006).							
X = Insurance, Y = Banking				X = Banking, Y = Insurance			
L_X	L_Y	DP	p -value	L_X	L_Y	DP	p -value
1	1	1.745	0.040	1	1	0.789	0.215
2	2	1.418	0.078	2	2	1.382	0.084
3	3	0.892	0.186	3	3	0.239	0.406
4	4	0.826	0.204	4	4	0.649	0.258
5	5	0.866	0.193	5	5	0.106	0.458
6	6	0.783	0.217	6	6	0.408	0.342
7	7	0.828	0.204	7	7	0.638	0.262
8	8	0.636	0.262	8	8	0.545	0.293

This table provides the various statistics and p -values associated with Granger causality tests after GARCH filtering. Broadly speaking, the null hypothesis is $H_0 : Y$ does not Granger-cause X . Therefore, the leftmost columns show whether the systemic risk of banks does Granger-cause systemic risk of insurers. Moreover, the rightmost columns show whether the systemic risk of insurers does Granger-cause systemic risk of banks. L_X and L_Y are the number of lags of X and Y , respectively. For the linear case, they are determined using the Bayesian information criterion. Values in bold denote a significance level of 5%. We use a lead length of 1 and a bandwidth of 0.5 in Diks and Panchenko (2006). The GARCH noise terms are computed by fitting a GARCH(1,1) model to the differenced systemic risk contributions and by extracting the Gaussian noise terms. F stands for F -statistic and DP for DP-statistic.

to the default probabilities, and pairwise correlations of leverages' co-movements. The firm-by-firm estimation of the model is based on the entire term structure of single-name CDS premiums of 35 major financial institutions and uses a two-stage filtering technique.

The model provides a framework that reacts quickly to new information and is well adapted to measure firm-specific credit risk, even during financial turmoil. We find that the banking subsector's default probabilities are higher during the first half of the crisis era. Halfway through the crisis period, the insurance subsector's average PDs are tremendously affected.

Our results indicate an increase in correlation during the high-volatility regime in comparison with the stable regime for 33 out of the 35 firms within the portfolio. It suggests the existence of a strong linkage among many financial institutions under study during the last crisis.

Finally, the empirical study presented in this paper finds supportive evidence of increased systemic risk within the financial services sector during the last global crisis. There is a unidirectional causal effect from banks to insurers when accounting for

heteroskedasticity. However, the opposite relationship (from insurers to banks) is not statistically significant.

Possible extensions of the framework would account for a fourth financial crises' "L": liquidity. This would obviously require the incorporation of an efficient liquidity proxy that captures this dimension. The inclusion of this aspect is not trivial. We leave this question for further research.

3.A Endogenous Correlation Coefficients

To obtain endogenous correlation estimates, correlation coefficients are introduced into the covariance matrix of the augmented state vector on which the unscented transformation is performed. As a starting point, one can write the second order moment of the augmented state vector as

$$\mathbf{P}_{t|t}^a = \begin{bmatrix} \mathbf{P}_{t|t} & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{\Sigma}_{s_t} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{R} \end{bmatrix}_{2(D+2) \times 2(D+2)}$$

where $[\mathbf{P}_{t|t}]_{2 \times 2}$ is the covariance matrix of predicted state variables $(\hat{x}_{t|t}^{(i)}, \hat{x}_{t|t}^{(j)})$ updated at each time step, $[\mathbf{\Sigma}_{s_t}]_{2 \times 2}$ is the covariance matrix of leverage noise terms associated with regimes $\mathbf{s}_t = (s_t^{(i)}, s_t^{(j)})$, and $[\mathbf{R}]_{2D \times 2D}$ is the trading noise variance matrix. Furthermore, dimension D refers to the number of CDS maturities available for each firm. More precisely, covariance and variance matrices can be expressed as

$$\mathbf{\Sigma}_{s_t} = \begin{bmatrix} (\sigma_{s_t}^{(i)})^2 & \sigma_{s_t}^{(i)} \sigma_{s_t}^{(j)} \rho_{s_t}^{(i,j)} \\ \sigma_{s_t}^{(i)} \sigma_{s_t}^{(j)} \rho_{s_t}^{(i,j)} & (\sigma_{s_t}^{(j)})^2 \end{bmatrix} \times \Delta t \text{ and } \mathbf{R} = \text{diag}(\delta^2).$$

Note that $\text{diag}(\delta^2)$ is the operator that creates a square matrix with diagonal elements corresponding to δ^2 , and $\delta = [\delta^{(i,1)}, \delta^{(i,2)}, \delta^{(i,3)}, \delta^{(i,5)}, \delta^{(i,7)}, \delta^{(i,10)}, \delta^{(j,1)}, \delta^{(j,2)}, \delta^{(j,3)}, \delta^{(j,5)}, \delta^{(j,7)}, \delta^{(j,10)}]$ is the vector of the noise terms' standard deviation. By maximizing the joint bivariate log-likelihood function, one obtains the correlation coefficient estimates.

According to Hamilton (1994) and considering M parallel UKF in the bivariate framework, the (quasi-) log-likelihood function based on observations $\mathbf{y}_t = (\mathbf{y}_t^{(i)}, \mathbf{y}_t^{(j)})$ up to

time step T for all possible paths M is computed by

$$\sum_{t=1}^T \sum_{l=1}^M \log(f(\mathbf{y}_t | Y_{t-1}; \phi_2))$$

where $\phi_2 = \{\rho_{(1,1)}^{(i,j)}, \rho_{(1,2)}^{(i,j)}, \rho_{(2,1)}^{(i,j)}, \rho_{(2,2)}^{(i,j)}\}$, and the conditional likelihood $f(\mathbf{y}_t | Y_{t-1}; \phi_2)$ given $Y_{t-1} = \{\mathbf{y}_1^{(i)}, \dots, \mathbf{y}_{t-1}^{(i)}, \mathbf{y}_1^{(j)}, \dots, \mathbf{y}_{t-1}^{(j)}\}$ is the probability density function of a $2D$ -variate Gaussian distribution valued at $(\mathbf{y}_t^{(i)}, \mathbf{y}_t^{(j)})$ with mean and covariance obtained from the filtering procedure. More specifically, the mean $(\hat{\mathbf{y}}_{t|t-1}^{(i)}, \hat{\mathbf{y}}_{t|t-1}^{(j)})$ is a $(1 \times 2D)$ vector obtained from $\mathbb{E}[(\mathbf{y}_t^{(i)}, \mathbf{y}_t^{(j)}) | Y_{t-1}]$, and the covariance matrix of dimension $(2D \times 2D)$ is $\mathbf{P}_{yy} = \text{Cov}[(\mathbf{y}_t^{(i)}, \mathbf{y}_t^{(j)}), (\mathbf{y}_t^{(i)}, \mathbf{y}_t^{(j)}) | Y_{t-1}]$.³¹ By using Bayes' rule, one can express the conditional likelihood function as

$$f(\mathbf{y}_t | Y_{t-1}; \phi_2) = \frac{\sum_{\mathbf{s}_t} f(\mathbf{y}_t, \mathbf{s}_t, Y_{t-1}; \phi_2)}{f(Y_{t-1}; \phi_2)} = \sum_{\mathbf{s}_t} f(\mathbf{y}_t | \mathbf{s}_t, Y_{t-1}; \phi_2) \times f(\mathbf{s}_t | Y_{t-1}; \phi_2).$$

The conditional likelihood of $\mathbf{y}_t = (\mathbf{y}_t^{(i)}, \mathbf{y}_t^{(j)})$ is computed analytically using the $2D$ -variate Gaussian density function. From the Markov property, the likelihood function given $\mathbf{y}_{t-1} = (\mathbf{y}_{t-1}^{(i)}, \mathbf{y}_{t-1}^{(j)})$ and the actual regimes $\mathbf{s}_t = (s_t^{(i)}, s_t^{(j)})$ of firms i and j can be expressed as

$$f(\mathbf{y}_t | \mathbf{s}_t, Y_{t-1}; \phi_2) = f(\mathbf{y}_t | \mathbf{s}_t, \mathbf{y}_{t-1}; \phi_2) = \frac{1}{(2\pi)^D |\mathbf{P}_{yy}|^{1/2}} \exp\left(-\frac{1}{2} e_{\mathbf{s}_t}^\top \mathbf{P}_{yy}^{-1} e_{\mathbf{s}_t}\right)$$

where $e_{\mathbf{s}_t}$ is the error between observations and their forecast values. Second, the conditional likelihood of $\mathbf{s}_t = (s_t^{(i)}, s_t^{(j)})$ given Y_{t-1} is obtained recursively. Let $\eta_t^\top = f(\mathbf{y}_t | \mathbf{s}_t, Y_{t-1}; \phi_2)$ and $\xi_{t|t-1} = f(\mathbf{s}_t | Y_{t-1}; \phi_2)$ be two vectors of size 4×1 . Then, one can use the following recursion equations

$$\xi_{t+1|t} = \mathbf{P}^{(i,j)\top} \xi_{t|t} \text{ and } \xi_{t|t} = \frac{\eta_t(\times) \xi_{t|t-1}}{\eta_t^\top \xi_{t|t-1}}$$

where (\times) refers to the element-by-element multiplication and $\mathbf{P}^{(i,j)}$ is the following transition matrix

$$\mathbf{P}^{(i,j)} = \begin{bmatrix} p_{11}^{(i)} p_{11}^{(j)} & p_{11}^{(i)} p_{12}^{(j)} & p_{12}^{(i)} p_{11}^{(j)} & p_{12}^{(i)} p_{12}^{(j)} \\ p_{11}^{(i)} p_{21}^{(j)} & p_{11}^{(i)} p_{22}^{(j)} & p_{12}^{(i)} p_{21}^{(j)} & p_{12}^{(i)} p_{22}^{(j)} \\ p_{21}^{(i)} p_{11}^{(j)} & p_{21}^{(i)} p_{12}^{(j)} & p_{22}^{(i)} p_{11}^{(j)} & p_{22}^{(i)} p_{12}^{(j)} \\ p_{21}^{(i)} p_{21}^{(j)} & p_{21}^{(i)} p_{22}^{(j)} & p_{22}^{(i)} p_{21}^{(j)} & p_{22}^{(i)} p_{22}^{(j)} \end{bmatrix}.$$

³¹The two moments are computed as a by-product of the UKF methodology.

An estimate of ϕ_2 is obtained by maximizing the log-likelihood function:

$$\hat{\phi}_2 = \operatorname{argmax} \left\{ \sum_{t=1}^T \sum_{l=1}^M \sum_{s_t} \ln(\xi_{t|t-1}) - D \ln(2\pi) - \frac{1}{2} \ln |\mathbf{P}_{yy}| - \frac{1}{2} e_{s_t}^\top \mathbf{P}_{yy}^{-1} e_{s_t} \right\}.$$

3.B Calculation of Systemic Risk Measures

Algorithm 1 (Calculation of systemic risk measures).

1. Generate 500,000 log-leverage paths of three months (i.e. 13 weeks), along with default indicators and losses given default.

- (a) For each firm i , generate the time $t + u$ log-leverage such that

$$\log(X_{t+u}^{(i)}) = \log(X_{t+u-1}^{(i)}) + \left(\mu^{(i)} - \frac{1}{2} \left(\sigma_{s_{t+u}}^{(i)} \right)^2 \right) \Delta t + \sum_{j=1}^N R_{s_{t+u}}^{(i,j)} \sqrt{\Delta t} \varepsilon_{t+u}^{(j)}$$

where $R_{s_{t+u}}^{(i,j)}$ is the $(i, j)^{\text{th}}$ entry of the (lower triangular) Cholesky decomposition of the regime-dependent covariance matrix and $\varepsilon_{t+u}^{(j)}$ are standardized Gaussian random variables.³²

- (b) Determine if the remaining firms default. This step is performed using the model's PD over the next week:

$$\text{PD}_{t+u-1, t+u}^{(i)} = 1 - \exp \left[-\Delta t \left(\beta^{(i)} + \left(\frac{X_{t+u-1}^{(i)}}{\theta^{(i)}} \right)^{\alpha^{(i)}} \right) \right].$$

Firm i defaults if $U_{t+u}^{(i)} \leq \text{PD}_{t+u-1, t+u}^{(i)}$ where $U_{t+u}^{(i)}$ is a uniformly distributed random variable $U_{t+u}^{(i)}$ on $[0, 1]$.

- (c) For each firm i , compute the loss given default

$$\text{LGD}_{t+u}^{(i)} = \begin{cases} TL_{t+u}^{(i)} \left(1 - \min \left[\frac{1}{X_{t+u}^{(i)}} (1 - \kappa^{(i)}); 1 \right] \right) & \text{if the firm defaults} \\ 0 & \text{otherwise} \end{cases}$$

where $TL_{t+u}^{(i)}$ is the i^{th} firm total liabilities at time $t + u$. The i^{th} firm is removed from the set of active companies if it defaults.

³²Since the correlation coefficients are estimated in a pairwise manner, it is possible that the full correlation matrix is not positive-definite. Following the literature, we find the closest correlation matrices in the Frobenius norm. In this paper, Qi and Sun's (2006) method is applied. The nearest symmetric correlation matrix is the closest to the estimated correlation matrix in the sense of the Frobenius norm. Qi and Sun's (2006) method is highly efficient and converges readily.

2. Aggregate each firm's losses and compute the total losses across the firms:

$$L_{t,t+13}^{(i)} = \sum_{u=1}^{13} \text{LGD}_{t+u}^{(i)} \text{ and } L_{t,t+13} = \sum_{i=1}^N L_{t,t+13}^{(i)}.$$

3. Compute the systemic risk measure (SR) for the financial service sector as the sample average across the 500,000 log-leverage paths of

$$L_{t,t+13} \mathbb{I}(L_{t,t+13} > \text{VaR}_{0.99}(L_{t,t+13})),$$

where $\text{VaR}_{0.99}(L_t)$ represents the 99th percentile of the total losses distribution and $\mathbb{I}(\cdot)$ the indicator function.

4. Finally, calculate the systemic risk contribution (so-called nominal price, in millions) for each subsector (i.e. the marginal expected shortfall). This is the sample average across the 500,000 log-leverage paths of

$$\left(\sum_{i \in \text{SS}} L_{t,t+13}^{(i)} \right) \mathbb{I}(L_{t,t+13} > \text{VaR}_{0.99}(L_{t,t+13})),$$

where $\text{SS} \in \{\text{Ins}, \text{Bnk}\}$. Also, compute the relative systemic risk contribution RSR_t^{SS} (so-called unit price, as a percentage) for each subsector by dividing each systemic risk contribution by the subsector's total liabilities.

3.C Systemic Risk Measures: Additional Results

3.C.1 Linear Granger Causality Tests

First, let us focus on whether the systemic risk of banking firms Granger-causes the systemic risk of insurance companies. The results of this statistical test is given in Panel A of Table 3.6. As usual, we search for the optimal number of lags based on the Bayesian information criterion (BIC).³³ The statistical test reports whether the coefficients of the lagged RSR^{Bnk} are jointly significantly different from zero. The F -statistic has a value of 6.199 and the null hypothesis is rejected at a level of 1%, meaning that

³³For the restricted model (i.e. the one using only insurance subsector contributions' lagged values), we find one lag for insurers. Then, using the unrestricted model (i.e. the one that includes banking subsector contributions' in lagged values), we detect five lags for banks.

Table 3.6: **Linear and nonlinear Granger causality tests before GARCH filtering.**

Panel A: Linear Granger causality.							
X = Insurance, Y = Banking				X = Banking, Y = Insurance			
L_X	L_Y	F	p -value	L_X	L_Y	F	p -value
1	5	6.199	0.000	5	3	5.013	0.002
Panel B: Nonlinear Granger causality using Diks and Panchenko (2006).							
X = Insurance, Y = Banking				X = Banking, Y = Insurance			
L_X	L_Y	DP	p -value	L_X	L_Y	DP	p -value
1	1	3.369	0.000	1	1	1.981	0.024
2	2	2.785	0.003	2	2	1.559	0.060
3	3	2.528	0.006	3	3	1.076	0.141
4	4	1.957	0.025	4	4	0.251	0.401
5	5	1.888	0.030	5	5	0.251	0.401
6	6	1.617	0.053	6	6	-0.309	0.621
7	7	1.589	0.056	7	7	-0.216	0.585
8	8	1.596	0.055	8	8	-0.677	0.751

This table provides the various statistics and p -values associated with Granger causality tests. Broadly speaking, the null hypothesis is $H_0 : Y$ does not Granger-cause X . Therefore, the leftmost columns show whether the systemic risk of banks does Granger-cause systemic risk of insurers. Moreover, the rightmost columns show whether the systemic risk of insurers does Granger-cause systemic risk of banks. L_X and L_Y are the number of lags of X and Y respectively. For the linear case, they are determined using the Bayesian information criterion. Values in bold denote a significance level of 5%. We use a lead length of 1 and a bandwidth of 0.5 in Diks and Panchenko (2006). The results are robust to other choices. F stands for F -statistic and DP for DP-statistic.

the banking subsector systemic risk Granger-causes the insurance subsector systemic risk.

For the opposite relationship (i.e. systemic risk of insurers Granger-causes the systemic risk of banks), the F -statistic is 5.013 and the null hypothesis is again rejected with a p -value lower than 1%.³⁴ This statistical conclusion would imply that the insurance subsector systemic risk Granger-causes that of the banking subsector.

In summary, our results show a compelling interconnectedness between both subsectors. This conclusion is similar to the findings of Chen et al. (2014).

3.C.2 Nonlinear Granger Causality Tests

The Diks and Panchenko (2006)'s Granger causality test requires the user to select some values such as the lead length, lag lengths L_X and L_Y , and bandwidth. Unfortunately, there is no method to define their optimal values. Following Hiemstra and Jones (1994) and Diks and Panchenko (2006), we set the lead length at 1 and $L_X = L_Y$, using a maximum of eight common lags. The bandwidth is set to 0.5.

³⁴We find five lags for banks and three for insurers using the same procedure as above.

Panel B of Table 3.6 shows the results of the nonlinear Granger causality tests. If we focus first on whether the systemic risk of banking firms Granger-causes the systemic risk of insurance companies, we find that we reject the null hypothesis for lag lengths below six at a confidence level of 5%. For the opposite relationship, we reject the null hypothesis for a lag length of one only.

Again, this would imply interconnectedness between both subsectors; however, connections from the banking subsector to the insurance subsector are somewhat stronger using the nonlinear statistical tests since we reject more often (i.e. for more lags).

3.D Advantages of the Regime-Switching Framework over the “One-Regime” Equivalent

In this section, we investigate the benefits of using a regime-switching framework in terms of fit, default probabilities, recovery rates and systemic risk measures.

3.D.1 Parameter Estimates

Table 3.7 exhibits the firm-specific parameters for each of the 35 companies considered in this study using “one-regime” equivalent model of Boudreault et al. (2013). A comparison with Table 2 of the main paper shows that the noise terms’ standard deviations are larger for short and long maturities, meaning that the overall fitting of the CDS term structure deteriorates when the regime-switching feature is removed. Parameter σ of the “one-regime” equivalent model is always between σ_1 and σ_2 .

3.D.2 Fit

We now consider the difference between the model and the observed CDS premiums.

Tables 3.8 and 3.9 show the relative root mean square errors (RRMSEs) in percentage for each tenor and each firm considered. In general, the regime-switching model yields smaller RRMSEs for most tenors. Overall, the RRMSE of the “one-regime” equivalent is 29% higher than the one of RS. The regime-switching model is doing better in terms of fitting the whole term structure of CDS premiums. For instance, for 5-year

Table 3.7: First-stage parameter estimates: “one-regime” equivalent of Boudreault et al. (2013).

Firm	$\mu^Q(\%)$	$\mu^P(\%)$	σ	κ	α	θ	$\beta(\%)$	$\delta^{(1)}$	$\delta^{(2)}$	$\delta^{(3)}$	$\delta^{(5)}$	$\delta^{(7)}$	$\delta^{(10)}$
ACE	-0.117	0.165	0.070	0.073	12.512	1.459	0.090	0.241	0.132	0.039	0.085	0.121	0.143
ALL	-1.080	-0.170	0.097	0.393	10.694	1.341	0.063	0.379	0.240	0.129	0.008	0.058	0.088
AXP	-1.955	-1.868	0.108	0.492	12.623	1.663	0.040	0.401	0.221	0.123	0.010	0.056	0.099
AIG	-1.016	0.030	0.137	0.407	11.980	1.325	1.509	0.338	0.155	0.077	0.108	0.148	0.199
AOC	0.046	-3.338	0.059	0.084	14.459	1.340	0.004	0.302	0.189	0.103	0.002	0.059	0.099
BSC	-3.593	-2.452	0.245	0.627	4.364	1.342	0.873	0.299	0.162	0.057	0.101	0.115	0.227
BRK	-0.815	-0.506	0.106	0.397	9.864	1.496	0.082	0.252	0.124	0.070	0.066	0.090	0.116
BACORP	-0.875	-0.735	0.102	0.408	9.166	1.339	0.129	0.336	0.161	0.065	0.078	0.106	0.137
COF	-0.460	0.565	0.074	0.201	15.017	1.268	0.082	0.391	0.220	0.127	0.006	0.052	0.104
SCH	-0.962	0.425	0.128	0.395	6.112	1.460	1.470	0.260	0.141	0.082	0.099	0.104	0.098
CB	-1.389	-1.042	0.155	0.436	6.340	1.562	0.550	0.376	0.216	0.151	0.034	0.065	0.078
C	-0.963	-0.376	0.111	0.448	9.443	1.287	0.186	0.275	0.118	0.038	0.081	0.106	0.130
DB	-1.087	-0.310	0.104	0.401	6.622	1.357	0.077	0.186	0.102	0.045	0.063	0.043	0.054
FHLMC	-0.669	0.190	0.090	0.409	8.593	1.319	0.012	0.267	0.122	0.075	0.079	0.113	0.187
FNMA	-1.226	-0.128	0.130	0.453	5.976	1.570	0.069	0.306	0.158	0.089	0.077	0.083	0.137
GE-GNWT	0.028	0.724	0.110	0.454	7.769	1.367	0.022	0.255	0.093	0.025	0.114	0.159	0.197
GS	-1.042	-0.164	0.137	0.384	7.006	1.388	2.855	0.313	0.170	0.057	0.089	0.123	0.148
HIG	-0.543	0.096	0.127	0.509	8.001	1.343	0.180	0.305	0.129	0.063	0.070	0.093	0.130
JPM	-1.315	-1.546	0.146	0.474	8.009	1.259	0.746	0.453	0.268	0.142	0.015	0.071	0.113
LEH	-3.663	-3.566	0.230	0.593	4.807	1.431	0.039	0.298	0.116	0.024	0.098	0.114	0.196
LIBMUT	-0.491	0.625	0.164	0.533	5.096	1.403	0.256	0.285	0.175	0.088	0.032	0.064	0.078
LNC	-1.058	-9.548	0.067	0.976	18.057	1.735	0.000	0.281	0.127	0.059	0.052	0.066	0.103
LTR	-0.695	-1.532	0.108	0.433	8.377	1.279	0.008	0.306	0.198	0.115	0.011	0.067	0.122
MMC	0.538	-0.025	0.159	0.585	3.746	1.350	0.000	0.243	0.096	0.009	0.089	0.144	0.176
MER	-1.327	-0.800	0.148	0.510	5.396	1.312	0.014	0.443	0.256	0.136	0.002	0.054	0.111
MET	-0.972	-0.638	0.158	0.360	12.557	1.488	2.953	0.420	0.194	0.125	0.002	0.053	0.089
MWD	-1.124	-0.064	0.150	0.566	5.156	1.170	0.728	0.422	0.256	0.135	0.006	0.061	0.100
PRU	-0.766	0.058	0.126	0.507	7.593	1.320	0.142	0.415	0.235	0.124	0.027	0.062	0.109
SAFC	0.306	0.804	0.096	0.333	6.206	1.469	0.077	0.229	0.111	0.027	0.107	0.118	0.192
STI	-0.909	-0.201	0.092	0.605	11.040	1.339	0.035	0.277	0.222	0.114	0.076	0.111	0.089
USB	-0.721	0.359	0.105	0.481	8.064	1.235	0.013	0.244	0.218	0.129	0.091	0.105	0.147
WM	-1.328	-0.170	0.088	0.193	18.037	1.468	0.214	0.267	0.114	0.064	0.091	0.108	0.135
WB	-1.203	-0.239	0.105	0.377	8.856	1.369	0.022	0.186	0.106	0.049	0.033	0.056	0.091
WFC	-0.364	-1.240	0.041	0.489	25.445	0.832	0.000	0.271	0.130	0.001	0.128	0.157	0.184
XL	-0.180	0.806	0.063	0.089	15.971	1.328	0.031	0.298	0.159	0.084	0.020	0.041	0.085
Insurance	-0.513	-0.330	0.113	9.951	1.413	0.373	0.411	0.308	0.161	0.080	0.052	0.088	0.125
Banking	-1.317	0.470	0.123	9.460	1.337	0.400	0.448	0.310	0.172	0.082	0.064	0.091	0.131
All firms	-0.949	0.104	0.118	9.684	1.372	0.388	0.431	0.309	0.167	0.081	0.059	0.090	0.128

The table shows parameters estimates obtained from CDS data from January 2005 to December 2012 by applying filtering techniques and quasi-likelihood maximization. The following parameter are reported: the drifts μ under both measures \mathbb{P} and \mathbb{Q} , the diffusion σ , the constants α , β and θ that define the intensity process, the liquidation and legal fees parameter κ , and finally the standard errors of the noise terms for tenors of 1, 2, 3, 5, 7 and 10 years.

Table 3.8: **Relative root mean square errors in percentage: comparison between the regime-switching model and the “one-regime” equivalent of Boudreault et al. (2013).**

		1-year	2-year	3-year	5-year	7-year	10-year	All
ACE	1R	22.81	11.63	2.39	8.73	12.43	14.11	13.48
	RS	19.48	14.69	9.37	4.46	1.00	3.62	10.93
ALL	1R	38.37	23.44	13.17	1.03	5.89	8.87	19.62
	RS	26.47	21.82	14.23	3.60	0.63	4.11	15.33
AXP	1R	35.92	20.67	12.82	0.35	5.68	9.46	18.27
	RS	18.58	10.50	6.15	5.67	6.15	10.41	10.58
AIG	1R	39.29	15.50	5.74	9.89	15.31	20.67	20.72
	RS	33.18	14.30	7.26	3.17	3.24	6.47	15.39
AOC	1R	31.18	19.30	11.02	0.16	6.46	9.74	16.34
	RS	18.06	10.70	5.80	3.21	3.78	6.85	9.54
BSC	1R	23.67	13.15	5.98	12.60	14.80	19.98	16.05
	RS	19.44	13.50	7.76	4.58	0.96	3.89	10.47
BRK	1R	26.94	11.73	5.85	5.71	8.85	11.20	13.75
	RS	24.57	13.60	7.81	5.15	3.15	6.68	12.45
BACORP	1R	41.37	16.73	5.27	6.26	10.07	13.11	19.71
	RS	44.76	22.48	11.28	5.19	1.29	3.85	21.13
COF	1R	32.33	20.31	12.55	2.00	5.65	10.31	17.11
	RS	25.59	17.49	11.84	5.52	0.64	4.38	13.85
SCH	1R	29.61	15.99	7.74	7.74	8.23	9.34	15.32
	RS	23.95	15.19	6.34	7.14	5.45	7.75	12.81
CB	1R	42.10	22.95	13.61	1.74	5.95	8.32	20.78
	RS	21.74	13.40	7.69	4.46	2.25	4.43	11.22
C	1R	29.97	11.92	2.53	7.36	10.54	12.74	15.13
	RS	39.14	22.04	11.81	4.44	1.41	3.68	19.11
DB	1R	17.64	7.60	3.76	5.71	3.38	5.08	8.69
	RS	19.00	8.47	4.73	6.77	3.99	7.14	9.73
FHLMC	1R	30.66	11.40	5.22	6.59	10.90	18.56	16.35
	RS	22.66	14.69	10.38	7.11	2.41	5.91	12.44
FNMA	1R	33.69	15.34	7.54	6.46	8.37	13.95	17.00
	RS	21.14	17.33	10.27	4.49	2.42	6.28	12.37
GNWTH	1R	31.54	9.39	1.62	11.46	16.89	22.22	18.24
	RS	57.09	26.23	15.77	4.63	1.18	5.31	26.61
GS	1R	33.53	13.72	4.67	7.98	12.41	15.24	17.24
	RS	24.95	13.63	7.81	2.86	3.08	5.90	12.40
HIG	1R	33.21	12.68	5.36	5.92	8.98	13.15	16.23
	RS	20.16	11.55	8.33	3.71	3.54	6.76	10.66
JPM	1R	52.32	28.35	15.51	0.92	7.19	10.83	25.66
	RS	26.32	21.05	11.94	5.40	1.39	5.31	14.93
LEH	1R	25.78	11.97	3.52	11.22	13.32	19.14	15.76
	RS	18.53	9.92	5.35	2.32	3.58	5.73	9.32
LIBMUT	1R	33.39	16.55	9.73	1.76	5.42	8.80	16.30
	RS	18.91	10.72	6.71	2.96	2.37	5.65	9.69
LNC	1R	31.54	12.30	5.30	4.72	6.41	9.95	14.92
	RS	16.54	8.03	3.14	4.86	3.87	6.06	8.40
LTR	1R	29.52	19.49	10.72	0.32	6.88	11.57	16.06
	RS	17.36	10.61	4.90	3.99	3.73	8.10	9.43
MMC	1R	25.56	10.01	2.07	9.60	15.26	19.22	15.56
	RS	19.75	10.41	8.31	0.71	4.86	8.04	10.46
MER	1R	62.83	28.83	15.07	1.01	5.90	10.27	29.29
	RS	42.66	20.78	11.84	4.89	0.44	4.09	20.14
MET	1R	51.86	22.87	12.06	2.03	6.02	9.53	24.12
	RS	27.61	16.52	9.26	3.24	2.95	4.01	13.88
MWD	1R	51.70	25.87	13.57	2.15	6.47	10.30	24.76
	RS	29.54	15.23	7.35	2.75	3.83	6.09	14.25
PRU	1R	50.66	23.53	12.36	1.92	6.05	11.44	23.96
	RS	40.38	20.14	11.06	4.70	0.25	4.63	19.16

The table shows RRMSEs obtained by comparing model's prices to CDS data from January 2005 to December 2012. 1R stands for the “one-regime” equivalent model of Boudreault et al. (2013) and RS means the regime-switching model.

premiums, the RS model yields a RRMSE of 4.92%, which is 1.56% lower than the one obtained with the one-regime equivalent (i.e. a relative gain of about 32%).

Table 3.9: Relative root mean square errors in percentage: comparison between the regime-switching model and the “one-regime” equivalent of Boudreault et al. (2013), continued.

		1-year	2-year	3-year	5-year	7-year	10-year	All
SAFC	1R	25.29	11.81	1.78	9.45	13.17	17.54	15.02
	RS	20.59	14.63	8.28	4.13	3.36	6.15	11.35
STI	1R	30.58	30.38	10.36	6.46	10.39	8.35	19.08
	RS	32.22	30.50	10.45	8.18	9.16	6.87	19.47
USB	1R	26.93	21.91	12.34	8.21	9.56	13.31	16.80
	RS	23.95	20.81	12.99	9.73	7.76	11.63	15.63
WM	1R	22.11	11.02	3.49	9.10	11.14	14.78	13.21
	RS	17.57	16.33	7.21	6.24	1.20	4.75	10.73
WB	1R	22.63	11.33	5.52	3.38	6.05	8.77	11.52
	RS	21.28	13.82	7.77	5.98	1.99	4.39	11.28
WFC	1R	28.23	13.03	0.16	11.88	15.58	18.08	16.72
	RS	20.15	14.91	7.84	4.15	5.38	9.12	11.68
XL	1R	24.41	14.79	8.77	2.32	4.24	8.32	12.81
	RS	18.74	12.45	7.70	2.84	3.54	6.26	10.21
All firms	1R	35.91	18.45	9.21	6.48	9.79	13.23	18.39
	RS	27.79	16.77	9.34	4.92	3.75	6.38	14.26

The table shows RRMSEs obtained by comparing model's prices to CDS data from January 2005 to December 2012. 1R stands for the “one-regime” equivalent model of Boudreault et al. (2013) and RS means the regime-switching model.




Figure 3.11: Average one-year default probabilities computed using the regime-switching credit risk model and its “one-regime” equivalent.

This figure shows the time series of 1-year average default probabilities across the portfolio over the 2005–2012 period. Model time series are inferred from CDS premium market data. The different letters correspond to major events during the crisis: (A) The credit crunch begins in earnest (August 1, 2007). (B) The Federal Reserve Board approves the financing arrangement between JPM and BSC (March 14, 2008). (C) LEH files for Chapter 11 bankruptcy protection. MER is taken over by the BACORP. AIG almost defaults the day after (September 15, 2008). (D) Three large U.S. life insurance companies seek TARP funding: LNC, HIG and GNWTH (November 17, 2008). (E) The U.S. Treasury Department, Federal Reserve, and FDIC announce a package of guarantees, liquidity access, and capital for BACORP (January 16, 2009).

3.D.3 Default Probabilities

Figure 3.11 exhibits the average one-year default probabilities computed using the regime-switching credit risk model and its “one-regime” equivalent. In general, probabilities computed from the regime-switching framework are slightly higher during the crisis and during the European debt crisis for the banking subsector.

3.D.4 Recovery Rates

Recovery rates given by both models are highly correlated. The time series of 1-year expected recovery rates have a correlation of 96.37% in between both models, on average. Figure 3.12 shows comparison between “one-regime” and regime-switching one-year expected recovery rates. It seems that the one given by the “one-regime” equivalent are higher (and inconsistent with the current literature). The “one-regime” equivalent is not flexible enough to allow changes in both the level and the slope of the CDS term structure. Therefore, the filtered leverage is rather imprecise as it tries to account for these two effects simultaneously.

Indeed, Altman et al. (2005) find an average recovery at default of 53% and 35% for senior secured and unsecured bonds, respectively. Also, in Vazza and Gunter (2012), senior secured and unsecured bonds have an average discounted recovery rate of 56.4 and 42.9%, respectively, during the 1987-2012 period. These numbers are much more closer to the averages obtained with the regime-switching framework, when compared to what is obtained with the “one-regime” equivalent.

This difference could be explained by the fact that the regime-switching framework is much more flexible as it has two latent variables instead of only one. The additional flexibility allows us to capture more adequately the term structure of recovery rates (along with expected recovery over various horizons).

3.D.5 Dependence

The heat maps of Figure 3.13 summarize ρ for each firm. It also shows the correlation coefficients if both firms are in regime 1 and 2, respectively (while using the RS model). Generally speaking, it seems that $\rho_{1,1} < \rho < \rho_{2,2}$. Figure 3.14 exhibits the time series of median leverage correlation coefficients across firms and for the three categories of subsectors. The slight variation observed in the “one-regime” median correlation arises as the set of available firms change through time (i.e. defaults or additions). The variation observed for the regime-switching median correlation reflects the regime variations as well. Figure 3.15 shows the histogram of ρ for the three categories: Insurance/Insurance, Banking/Banking and Insurance/Banking. A comparison with Figure 7 of the paper shows that important variations in the correlation dynamics cannot be captured by the “one-regime” equivalent model.

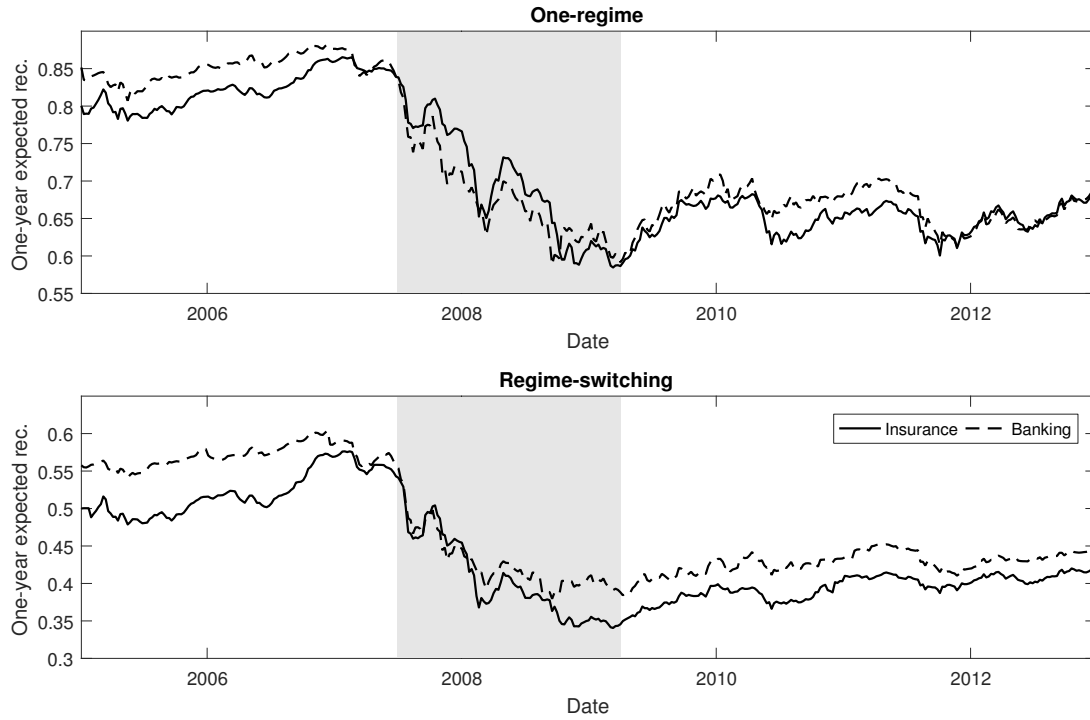


Figure 3.12: Average one-year expected recovery rates using the regime-switching credit risk model and its “one-regime” equivalent.

This figure shows the time series of one-year expected recovery rates across the portfolio over the 2005–2012 period. Model time series are inferred from CDS premium market data.




Figure 3.13: Heat map of ρ for the 35 firms using “one-regime” equivalent model of Boudreault et al. (2014) along with the correlation coefficients if both firms are in regime 1 and 2, respectively.

This figure shows ρ . The first 16 rows and columns correspond to the insurance subsector and last 19 ones to banks.




Figure 3.14: Time series of median leverage correlation coefficients for three groups: correlations across insurance firms (Insurance/Insurance), correlations across banking companies (Banking/Banking), and correlations between insurance and banking firms (Insurance/Banking).

The correlation coefficients are estimated for each pair of firms using the two-step approach based on the DEA-UKF methodology. Based on filtered regimes of both regime-switching and “one-regime” equivalent models, we compute the median correlation coefficients across the three groups. Note that median correlation in “one-regime” equivalent model slightly fluctuates over time as firms enter and leave our sample.




Figure 3.15: Histogram of ρ for three categories: correlations across insurance firms' leverages (Insurance/Insurance), correlations across banking companies' leverages (Banking/Banking), and correlations between insurance and banking firms' leverages (Insurance/Banking).

The correlation coefficients are estimated for the “one-regime” equivalent model of Boudreault et al. (2014). These figures show the empirical distribution of the ρ for three groups. The horizontal bar represents the sample mean.

The regime-switching model is able to capture the changing nature of correlation, especially during the past financial crisis. The “one-regime” equivalent model has a median correlation that remains virtually constant in time.³⁵

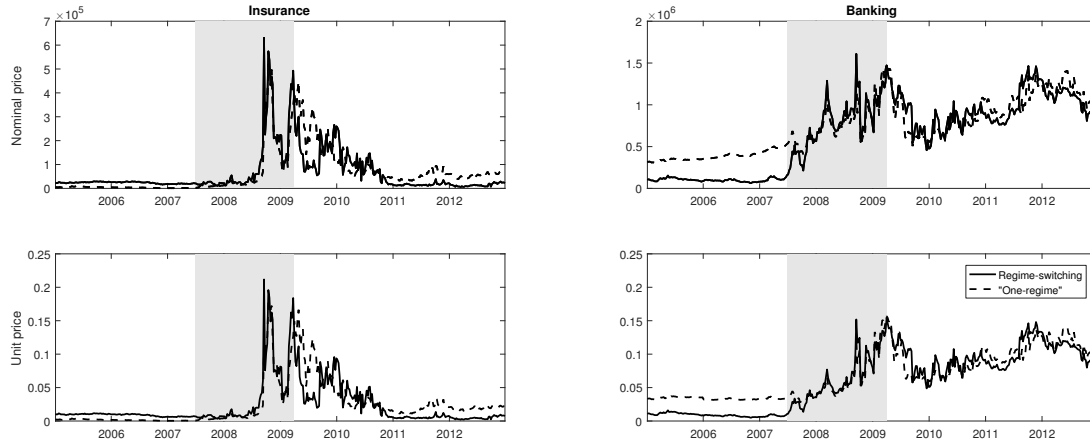


Figure 3.16: Time series of the systemic risk measure contribution and the relative systemic risk measure contribution using both regime-switching and “one-regime” equivalent models.

The systemic risk measures are computed using Monte Carlo methods and 5×10^5 paths over a span of three months. The theoretical debt portfolio that includes the total liabilities of each financial.

³⁵Note that median correlation in “one-regime” equivalent model slightly fluctuates over time as firms enter and leave our sample.

3.D.6 Systemic Risk Measures

We display the time series of nominal and unit price contributions in Figure 3.16 for both regime-switching and “one-regime” equivalent models. It seems that pre- and post-crisis contributions are high for the “one-regime” model. Not capturing adequately the changing behaviour of both volatility and correlation have an impact on the systemic risk measures.

References

- Acharya, V., R. Engle, and M. Richardson. 2012. Capital shortfall: A new approach to ranking and regulating systemic risks. *American Economic Review* 102:59–64.
- Acharya, V., L. H. Pedersen, T. Philippon, and M. P. Richardson. 2010. Measuring systemic risk. *Working paper*.
- Adrian, T., and M. K. Brunnermeier. 2009. CoVaR. *Working paper*.
- Allen, M., C. B. Rosenberg, C. Keller, B. Setser, and N. Roubini. 2002. A balance sheet approach to financial crisis. *Working paper*.
- Altman, E. I., B. Brady, A. Resti, and A. Sironi. 2005. The link between default and recovery rates: Theory, empirical evidence, and implications. *Journal of Business* 78:2203–2228.
- Baek, E., and W. Brock. 1992. A general test for nonlinear Granger causality: Bivariate model. *Working paper*.
- Bégin, J.-F., M. Boudreault, and G. Gauthier. 2016. Firm-specific credit risk modelling in the presence of statistical regimes and noisy prices. *Working paper*.
- Bharath, S. T., and T. Shumway. 2008. Forecasting default with the Merton distance to default model. *Review of Financial Studies* 21:1339–1369.
- Bielecki, T., D. Brigo, and F. Patras. 2011. *Credit risk frontiers: Subprime crisis, pricing and hedging, CVA, MBS, ratings, and liquidity*. John Wiley and Sons.
- Billio, M., M. Getmansky, A. W. Lo, and L. Pelizzon. 2012. Econometric measures of connectedness and systemic risk in the finance and insurance sectors. *Journal of Financial Economics* 104:535–559.

- Boudreault, M., G. Gauthier, and T. Thomassin. 2013. Recovery rate risk and credit spreads in a hybrid credit risk model. *Journal of Credit Risk* 9:3–39.
- Boudreault, M., G. Gauthier, and T. Thomassin. 2014. Contagion effect on bond portfolio risk measures in a hybrid credit risk model. *Finance Research Letters* 11:131–139.
- Boudreault, M., G. Gauthier, and T. Thomassin. 2015. Estimation of correlations in portfolio credit risk models based on noisy security prices. *Journal of Economic Dynamics and Control* 61:334–349.
- Chen, H., J. D. Cummins, K. S. Viswanathan, and M. A. Weiss. 2014. Systemic risk and the interconnectedness between banks and insurers: An econometric analysis. *Journal of Risk and Insurance* 81:623–652.
- Cummins, J. D., and M. A. Weiss. 2014. Systemic risk and the U.S. insurance sector. *Journal of Risk and Insurance* 81:489–528.
- Davis, M., and V. Lo. 2001. Infectious defaults. *Quantitative Finance* 1:382–387.
- Diks, C., and V. Panchenko. 2006. A new statistic and practical guidelines for non-parametric Granger causality testing. *Journal of Economic Dynamics and Control* 30:1647–1669.
- Dionne, G., G. Gauthier, K. Hammami, M. Maurice, and J.-G. Simonato. 2010. Default risk in corporate yield spreads. *Financial Management* 39:707–731.
- Duan, J.-C., and A. Fulop. 2009. Estimating the structural credit risk model when equity prices are contaminated by trading noises. *Journal of Econometrics* 150:288–296.
- Duffie, D., and N. Gârleanu. 2001. Risk and valuation of collateralized debt obligations. *Financial Analysis Journal* 57:41–59.
- Duffie, D., and D. Lando. 2001. Term structures of credit spreads with incomplete accounting information. *Econometrica* 69:633–664.
- Dumas, B., J. Fleming, and R. E. Whaley. 1998. Implied volatility functions: Empirical tests. *Journal of Finance* 53:2059–2106.
- Elliott, R. J., L. Aggoun, and J. B. Moore. 1995. *Hidden Markov models: Estimation and control*. Springer Science & Business Media.

- Engle, R. 2002. Dynamic conditional correlation: A simple class of multivariate generalized autoregressive conditional heteroskedasticity models. *Journal of Business & Economic Statistics* 20:339–350.
- Ericsson, J., K. Jacobs, and R. Oviedo. 2009. The determinants of credit default swap premia. *Journal of Financial and Quantitative Analysis* 44:109–132.
- Çetin, U., R. Jarrow, P. Protter, and Y. Yildirim. 2004. Modeling credit risk with partial information. *The Annals of Applied Probability* 14:1167–1178.
- Frey, R., and A. J. McNeil. 2003. Dependent defaults in models of portfolio credit risk. *Journal of Risk* 6:59–92.
- Friewald, N., R. Jankowitsch, and M. G. Subrahmanyam. 2012. Illiquidity or credit deterioration: A study of liquidity in the US corporate bond market during financial crises. *Journal of Financial Economics* 105:18–36.
- Garzarelli, F. 2009. The 2007-09 credit crisis and its aftermath. *Working paper*.
- Giesecke, K. 2006. Default and information. *Journal of Economic Dynamics and Control* 30:2281–2303.
- Giesecke, K., and L. Goldberg. 2003. Forecasting default in the face of uncertainty. *Journal of Derivatives* 12:11–25.
- Granger, C. W. 1969. Investigating causal relations by econometric models and cross-spectral methods. *Econometrica* 37:424–438.
- Gupton, G. M., C. C. Finger, and M. Bhatia. 2007. *Creditmetrics: Technical document*. CreditMetrics.
- Hamilton, J. D. 1994. *Time series analysis*. Princeton University Press.
- Harrington, S. E. 2009. The financial crisis, systemic risk, and the future of insurance regulation. *Journal of Risk and Insurance* 76:785–819.
- Hiemstra, C., and J. D. Jones. 1994. Testing for linear and nonlinear Granger causality in the stock price-volume relation. *Journal of Finance* 49:1639–1664.
- Huang, S. J., and J. Yu. 2010. Bayesian analysis of structural credit risk models with microstructure noises. *Journal of Economic Dynamics and Control* 34:2259–2272.
- Huang, X., H. Zhou, and H. Zhu. 2009. A framework for assessing the systemic risk of major financial institutions. *Journal of Banking and Finance* 33:2036–2049.

- Huang, X., H. Zhou, and H. Zhu. 2012. Assessing the systemic risk of a heterogeneous portfolio of banks during the recent financial crisis. *Journal of Financial Stability* 8:193–205.
- Hull, J., M. Predescu, and A. White. 2010. The valuation of correlation-dependent credit derivatives using a structural model. *Journal of Credit Risk* 6:99–132.
- Joe, H. 2014. *Dependence modeling with copulas*. CRC Press.
- Julier, S., and J. Uhlmann. 1997. A new extension of the Kalman Filter to nonlinear systems. In *SPIE proceedings series*, pp. 182–193. Society of Photo-Optical Instrumentation Engineers.
- Kalman, R. E. 1960. A new approach to linear filtering and prediction problems. *Journal of Basic Engineering* 82:35–45.
- Li, D. X. 2000. On default correlation: A copula function approach. *Journal of Fixed Income* 9:43–54.
- Longstaff, F. A., S. Mithal, and E. Neis. 2005. Corporate yield spreads: default risk or liquidity? New evidence from the credit default swap market. *Journal of Finance* 60:2213–2253.
- Maalaoui Chun, O., G. Dionne, and P. Francois. 2014. Detecting regime shifts in credit spreads. *Journal of Financial and Quantitative Analysis* 49:1339–1364.
- Madan, D., and H. Unal. 2000. A two-factor hazard rate model for pricing risky debt and the term structure of credit spreads. *Journal of Financial and Quantitative Analysis* 35:43–66.
- Markose, S., S. Giansante, and A. R. Shaghaghi. 2012. Too interconnected to fail financial network of US CDS market: Topological fragility and systemic risk. *Journal of Economic Behavior and Organization* 83:627–646.
- Meine, C., H. Supper, and G. N. Weiß. 2016. Is tail risk priced in credit default swap premia? *Review of Finance* 20:287–336.
- Milne, A. 2014. Distance to default and the financial crisis. *Journal of Financial Stability* 12:26–36.
- Mueller, P. 2008. Credit spreads and real activity. *Working paper*.
- Nier, E., J. Yang, T. Yorulmazer, and A. Alentorn. 2007. Network models and financial stability. *Journal of Economic Dynamics and Control* 31:2033–2060.

- Qi, H., and D. Sun. 2006. A quadratically convergent Newton method for computing the nearest correlation matrix. *SIAM Journal on Matrix Analysis and Applications* 28:360–385.
- Saldías, M. 2013. Systemic risk analysis using forward-looking distance-to-default series. *Journal of Financial Stability* 9:498–517.
- Saunders, A., and L. Allen. 2010. *Credit risk management in and out of the financial crisis: New approaches to value at risk and other paradigms*. John Wiley & Sons.
- Tang, D. Y., and H. Yan. 2007. Liquidity and credit default swap spreads. *Working paper*.
- Tugnait, J. K. 1982. Detection and estimation for abruptly changing systems. *Automatica* 18:607–615.
- Vazza, D., and E. Gunter. 2012. Recovery study (U.S.): Recoveries come into focus as the speculative-grade cycle turns negative. Tech. rep., Standard & Poor's Rating Services.
- Weiß, G. N., and J. Mühlnickel. 2014. Why do some insurers become systemically relevant? *Journal of Financial Stability* 13:95–117.

Chapter 4

Idiosyncratic Jump Risk Matters: Evidence from Equity Returns and Options

Abstract^{*}

The recent literature provides conflicting empirical evidence on the relationship between idiosyncratic risk and equity returns. This paper sheds new light on this relationship by exploiting the richness of option data. We disentangle four risk factors that potentially contribute to the equity risk premium: systematic Gaussian risk, systematic jump risk, idiosyncratic Gaussian risk, and idiosyncratic jump risk. First, we find that while systematic risk factors explain the greater part of the risk premium on a stock, idiosyncratic factors explain more than 40% of the average premium. Second, we show that the contribution of idiosyncratic risk to the equity risk premium arises exclusively from the jump risk component. Tail risk thus plays a central role in the pricing of idiosyncratic risk.

Keywords: risk premiums; tail risk; idiosyncratic risk; systematic risk; option valuation; GARCH.

^{*}Joint work with Christian Dorion and Geneviève Gauthier. Dorion and Gauthier are both affiliated with HEC Montréal.

4.1 Introduction

An investor should be rewarded for bearing systematic risk. One of the key insights of Sharpe (1964) and Lintner's (1965b) CAPM, building on the seminal work of Markowitz (1952), is that idiosyncratic risk, however, should not carry a risk premium as it can be diversified away. Since the CAPM, numerous asset pricing models have been developed building on the premise that idiosyncratic risk is not priced.¹ However, the recent literature has strongly challenged this notion.² Although the channel through which idiosyncratic risk could be priced is still a matter of debate, it is now widely accepted that, given market incompleteness, idiosyncratic risk can be priced.³ While previous studies are informative about the relative importance of idiosyncratic risk in explaining expected stock returns, they do not attempt to identify whether the importance of idiosyncratic risk arises from its diffusive or tail components. Thus, little is known on the relative contribution of systematic and idiosyncratic diffusive and tail risk in explaining the equity premium.

Our study departs from existing work by decomposing stocks' systematic and idiosyncratic shocks into a Gaussian and a jump components. Our approach offers an ideal framework to study the relative importance of each factor in explaining expected excess returns on equity. In particular, our study is the first to uncover the central role of idiosyncratic tail risk in explaining expected stock returns. Indeed, we find that idiosyncratic risk explains more than 40% of expected excess equity returns and, more importantly, that this is exclusively due to the jump risk component. Idiosyncratic Gaussian risk is not priced. This finding is consistent with the idea that investors have

¹Notably, Merton's (1973) ICAPM extends the insights of the CAPM to an intertemporal setup. The arbitrage pricing theory of Ross (1976) shows that any common return factor is a potential asset pricing factor. Fama and French (1992, 1993, 2015) and Carhart (1997), for instance, identify such potential factors, but diversifiable idiosyncratic risk is still assumed not to carry any risk premium.

²Concerns about the pricing of idiosyncratic risk dates back to Douglas (1969) and Lintner (1965a). Goyal and Santa-Clara (2003) contributed to putting this debate back at the forefront of the asset pricing literature by providing empirical evidence that idiosyncratic matters. Among others, Ang et al. (2006) find that stocks with high idiosyncratic volatilities had "abysmally" low average returns, lower than what could be explained by their exposure to aggregate volatility.

³Goyal and Santa-Clara (2003) highlight that a possible channel is background risk; investors hold nontraded assets (e.g. human capital or private businesses) which add background risk to their traded portfolio decisions. Jacobs and Wang (2004) provide evidence that idiosyncratic consumption risk is a priced factor in the cross section of stock returns. Hence, the average idiosyncratic stock variance being a proxy for idiosyncratic consumption risk could explain why idiosyncratic risk is priced. Consistent with this insight, Herskovic et al. (2016) provide evidence linking the average idiosyncratic volatility to income risk faced by households. Alternatively, Stambaugh et al. (2015) argue that the negative relationship between idiosyncratic volatility and stock returns could be driven by arbitrage asymmetry, as buying could be easier than shorting for many equity investors.

a hard time to hedge idiosyncratic tail risk and, thus, require a premium to bear their exposure to this risk.

We exploit the richness of stock option data to extract the expected risk premium associated with each risk factor, thereby avoiding the exclusive use of noisy realizations of historical equity returns. To this end, we develop a GARCH-jump model in which a firm's systematic and idiosyncratic risk have both a Gaussian and a tail component.

Our pricing kernel is such that each risk factor can potentially be priced. The model offers quasi-closed form solutions for the price of European options. We estimate the model on 260 firms that are or were part the S&P 500 index between 1996 and 2015, using equity returns and option prices of the market index and of each individual firm.⁴ To our knowledge, this is the most comprehensive joint estimation analysis of option-pricing models conducted in the literature.

Our empirical analysis highlights three new results. First, systematic risk accounts for only 59.8% of the average equity risk premium (ERP) on a stock, only one third of which is due to systematic normal risk.

Second, and most importantly, we find that the 40.2% contribution of idiosyncratic risk to the ERP is essentially due to idiosyncratic jump risk only. That is, the Gaussian component of idiosyncratic risk, which is easily diversifiable, is not priced once other sources of risk are accounted for.⁵ Consistent with Bates (2008), jump and normal risks are priced differently by investors. While the results of Christoffersen et al. (2012) and Ornathanalai (2014) already supported this view at the market level, our results document that both sources of risk have drastically different impact on the expected return of individual stocks.

When estimating a nested version of the model in which idiosyncratic jump risk is omitted, idiosyncratic normal risk appears to be priced. For the great majority of stocks, the nested variant of the model appears to be misspecified, however, because it offers a significantly worse fit to equity returns and options than the model with idiosyncratic jumps.

This result is of significant interest as most of the literature on idiosyncratic risk assumes conditional normality. Our third empirical finding is that idiosyncratic jump

⁴We considered all 1,000 stocks that were part of the index during this period; neglected stocks were set aside only because not enough options were liquidly traded over at least a consecutive 5-year window.

⁵Note that, while the expected stock return is not affected by idiosyncratic Gaussian risk, an option's vega is still positive and affected by total volatility.

risk shares a strong commonality across firms. Herskovic et al. (2016) document that idiosyncratic (total) variances have a strong factor structure. Based on the idiosyncratic volatilities of 20,000 CRSP stocks over 85 years, they document that a single factor explains 35% of the time variation firm-level idiosyncratic risk. In light of these results, our model of stock variance allows for two sources of commonality: one arising from commonality in idiosyncratic normal risk, the other from commonality in idiosyncratic jump risk. Over the 20 years in our sample, 260 firm-by-firm regressions of total idiosyncratic variance on these two sources of commonality yield an average R^2 of 73.4%; regressing on the commonality in idiosyncratic jump risk yields an average R^2 of 31.8%.

Our results thus extend the findings of Herskovic et al. (2016) in that we document that idiosyncratic tail risk explains a large fraction of the commonality in idiosyncratic variance. Already hard to hedge by nature, tail risk becomes virtually undiversifiable in times of turmoil, which justifies the risk premium attached to it.

Our study is the first to conduct a joint estimation, based on equity returns and options, of an option-valuation model to disentangle the four risk premiums associated with systematic and idiosyncratic, normal and tail (jump) risk. It is, however, closely related to several contemporaneous papers.

Christoffersen et al. (2013) document a strong factor structure in equity options. Consequently, building on Heston (1993), they develop a stochastic volatility model in which a firm's total variance is decomposed into a systematic and an idiosyncratic component. The authors study the effect of firm beta and market variance to explain the cross-sectional variations of equity options. Among others, their model predicts that stocks with higher betas have higher implied volatilities and steeper smiles, consistent with the empirical findings of Duan and Wei (2009).

Our framework extends that of Christoffersen et al. (2013) in that we allow for a jump component both in market returns and in the idiosyncratic part of stock returns. Moreover, our joint estimation methodology builds on those of Christoffersen et al. (2012) and Ornthanalai (2014), and allows us to quantify how the equity risk premium is affected by the four sources of risk affecting stocks in our setup. Our model and pricing kernel nest those of Elkamhi and Ornthanalai (2010) who complement the analysis in Christoffersen et al. (2013) and quantify the impact of market jump risk on equity options. They find that firms with a larger return compensation for systematic normal risk have a higher option-implied volatility level, while firms with a larger return compensation for systematic jump risk have steeper option-implied volatility slope.

However, stocks in their framework do not exhibit idiosyncratic jump risk, and they do not study the pricing of idiosyncratic risk. Along the same lines, Babaoğlu (2015) further document that a “jump beta” is needed to adequately explain equity returns, market risk exposures, and equity option prices.

Boloorforoosh (2014) extends the Christoffersen et al. (2013) model by allowing for idiosyncratic normal risk to be priced. He finds strong empirical support for the hypothesis that idiosyncratic risk is indeed priced. Boloorforoosh (2014) also documents that idiosyncratic volatilities exhibit a factor structure virtually as strong as that of total volatilities, consistent with Herskovic et al. (2016). Using a similar model, Xiao and Zhou (2014) study the same four risk factors as we do, but relying exclusively on realizations of historical equity returns.

Closer to our study, Gourié (2014) further extends Christoffersen et al. (2013) and estimates a continuous-time jump-diffusion model using a two-stage estimation procedure based on equity returns, options and intraday data observed on 29 stocks between 2006 and 2012. Gourié’s (2014) framework allows her to study the important role played by total (normal and jump) idiosyncratic risk in the equity and, most importantly, the variance risk premium. She finds that compensation for idiosyncratic risk represents, on average, 50% of the equity risk premium and 80% of the variance risk premium.

Although the models, datasets and estimation methods in our studies differ along several dimensions, the results that are common to our two studies are consistent and our analyses complement one another. In particular, Gourié (2014) provides strong empirical evidence that idiosyncratic risk is a key determinant of the equity risk premium; we provide strong empirical evidence that tail risk is actually at the core of the relationship between idiosyncratic risk and the equity risk premium.

This paper is organized as follows. Section 4.2 presents our model for the market and the individual stocks. Section 4.3 presents the data and discusses the estimation methodology. Then, Section 4.4 presents our empirical analysis. Section 4.5 concludes.

4.2 Model

We develop a model in which, in the spirit of the CAPM, stocks are exposed to systematic risk. Unlike the traditional one-factor CAPM, however, market and stock return

are not solely driven by a diffusive component. The market can crash, or more generally jump, and the stocks in our model are exposed to this systematic jump risk, as well as to idiosyncratic normal and jump risk. As such, our model falls under the framework of Kraus and Litzenberger (1976), but extends it in various directions.

4.2.1 Stock Returns

Returns on the market index, M_t , and a given stock, S_t , are modeled as follows:

$$R_{M,t+1} \equiv \log \left(\frac{M_{t+1}}{M_t} \right) = \mu_{M,t+1} - \xi_{M,t+1}^{\mathbb{P}} + z_{M,t+1} + y_{M,t+1}, \quad (4.1)$$

$$R_{S,t+1} \equiv \log \left(\frac{S_{t+1}}{S_t} \right) = \mu_{S,t+1} - \xi_{S,t+1}^{\mathbb{P}} + \beta_{S,z} z_{M,t+1} + \beta_{S,y} y_{M,t+1} + z_{S,t+1} + y_{S,t+1} \quad (4.2)$$

where stock returns are driven by the stock's exposure to systematic Gaussian and jump risk, $z_{M,t+1}$ and $y_{M,t+1}$, as well as stock-specific innovations $z_{S,t+1}$ and $y_{S,t+1}$, respectively capturing idiosyncratic normal and jump risk.

For $u \in \{M, S\}$, $\xi_{u,t+1}^{\mathbb{P}}$ is the convexity correction associated with the Gaussian, $z_{u,t+1}$, and the normal-inverse Gaussian (NIG) innovations, $y_{u,t+1}$.⁶ Hence,⁷

$$\mathbb{E}_t^{\mathbb{P}} [M_{t+1}] = M_t \exp(\mu_{M,t+1}) \text{ and } \mathbb{E}_t^{\mathbb{P}} [S_{t+1}] = S_t \exp(\mu_{S,t+1}).$$

That is, $\mu_{u,t+1} - r_{t+1}$ can be interpreted as the instantaneous equity risk premium on the index and the stock, given the risk-free rate r_{t+1} .⁸

Before discussing the exact form of the instantaneous risk premiums (Subsection 4.2.2), we further characterize the distribution of the shock processes. The Gaussian innovations are given by

$$z_{u,t+1} = \sqrt{h_{u,z,t+1}} \varepsilon_{u,t+1}, \quad u \in \{M, S\}$$

⁶ The convexity correction, $\xi_{M,t}^{\mathbb{P}} = \xi_{z_{M,t}}^{\mathbb{P}}(1) + \xi_{y_{M,t}}^{\mathbb{P}}(1)$, is based on the cumulant generating function of z_M and y_M (cf. Section 4.A). The same holds for $\xi_{S,t+1}^{\mathbb{P}} = \xi_{z_{M,t}}^{\mathbb{P}}(\beta_{S,z}) + \xi_{y_{M,t}}^{\mathbb{P}}(\beta_{S,y}) + \xi_{z_{S,t}}^{\mathbb{P}}(1) + \xi_{y_{S,t}}^{\mathbb{P}}(1)$.

⁷ The filtration is generated by the market noise terms as well as the stock noise terms, that is $\mathcal{F}_t^{\mathbb{S}} = \sigma\{z_{M,\tau}, y_{M,\tau}, z_{S,\tau}, y_{S,\tau}; S \in \mathbb{S}\}_{\tau=1}^t$. $\mathbb{E}_t^{\mathbb{P}} [S_{t+1}]$ is a shorthand for $\mathbb{E}^{\mathbb{P}} [S_{t+1} | \mathcal{F}_t^{\mathbb{S}}]$. Since all innovation time series are independent, $\mathbb{E}^{\mathbb{P}} [M_{t+1} | \mathcal{F}_t^{\mathbb{S}}] = \mathbb{E}^{\mathbb{P}} [M_{t+1} | \mathcal{F}_t^{\mathbb{M}}]$ where $\mathcal{F}_t^{\mathbb{M}} = \sigma\{z_{M,\tau}, y_{M,\tau}\}_{\tau=1}^t$ and we still use $\mathbb{E}_t^{\mathbb{P}} [\cdot]$ to represent both conditional expectations.

⁸ Over a short period of time, $\mu_{M,t+1}$ and r_{t+1} are close to zero, such that

$$\mathbb{E}_t^{\mathbb{P}} [M_{t+1}/M_t] - \mathbb{E}_t^{\mathbb{Q}} [M_{t+1}/M_t] = \exp(\mu_{M,t+1}) - \exp(r_{t+1}) \simeq \mu_{M,t+1} - r_{t+1}.$$

where the $\varepsilon_{u,t+1}$ are serially independent standard normal random variables and the conditional variance of $z_{u,t+1}$ follows GARCH dynamics. Indeed, the market conditional variance is

$$\begin{aligned} h_{M,z,t+1} &= w_{M,z} + b_{M,z} h_{M,z,t} + \frac{a_{M,z}}{h_{M,z,t}} (z_{M,t} - c_{M,z} h_{M,z,t})^2 \\ &= \sigma_{M,z}^2 + b'_{M,z} (h_{M,z,t} - \sigma_{M,z}^2) + \frac{a_{M,z}}{h_{M,z,t}} (z_{M,t}^2 - h_{M,z,t} - 2c_{M,z} h_{M,z,t} z_{M,t}), \end{aligned} \quad (4.3)$$

where $\sigma_{M,z}^2 = \frac{w_{M,z} + a_{M,z}}{1 - b'_{M,z}}$, is the unconditional level of the market variance, and $b'_{M,z} = b_{M,z} + a_{M,z} c_{M,z}^2$, is the variance persistence.

The specification of the stock's conditional variance is inspired from the literature on component volatility models,⁹ that is

$$\begin{aligned} h_{S,z,t+1} &= q_{S,z,t+1} + b_{S,z} (h_{S,z,t} - q_{S,z,t}) + \frac{a_{S,z}}{h_{S,z,t}} (z_{S,t}^2 - h_{S,z,t} - 2c_{S,z} h_{S,z,t} z_{S,t}) \\ &= \kappa_{S,z} h_{M,z,t+1} + b_{S,z} (h_{S,z,t} - \kappa_{S,z} h_{M,z,t}) \\ &\quad + \frac{a_{S,z}}{h_{S,z,t}} (z_{S,t}^2 - h_{S,z,t} - 2c_{S,z} h_{S,z,t} z_{S,t}). \end{aligned} \quad (4.4)$$

However, rather than varying around a long-run volatility component of its own, the conditional variance of a particular stock loads on market variance through $\kappa_{S,z} h_{M,z,t+1}$, while the idiosyncratic variance in excess of this central tendency, $h'_{S,z,t+1} = h_{S,z,t+1} - \kappa_{S,z} h_{M,z,t+1}$, has a GARCH structure. In the spirit of Martin and Wagner (2016), we refer to $h'_{S,z,t+1}$ as the excess idiosyncratic variance.¹⁰

The jumps, $y_{u,t+1}$, have a NIG distribution with location parameter set at 0, a tail heaviness parameter α_u and an asymmetry parameter δ_u . Following Ornathanalai (2014), the time-homogeneous scale parameter of the distribution is allowed to vary and is denoted by $h_{u,y,t+1}$.¹¹ We refer to $h_{u,y,t+1}$ as the jump intensity process.¹² The jump

⁹On GARCH component models, see, among others, Engle and Lee (1999), Christoffersen et al. (2008), Engle and Rangel (2008), and Engle et al. (2013).

¹⁰The variance process in Gourié (2014) has a similar structure and she refers to the analogue of $h'_{S,z,t+1}$ as residual idiosyncratic variance.

¹¹Earlier drafts of this paper featured Poisson rather than NIG jumps. While the main results were qualitatively similar, the estimated jump parameters were much less stable. In particular, the Poisson-jump version of the model had a harder time accommodating the positive jumps during and after the Great Recession.

¹²Strictly speaking, $h_{u,y,t+1}$ is not an intensity as it does not parameterize the number of jumps observed over a period Δt . However, the normal-inverse Gaussian distributions is closed under convolution in the sense that, given α_u and δ_u , the sum of two NIG shocks with scale parameters h_1 and h_2 would have a scale parameter of $h_1 + h_2$. Hence, the NIG jump as specified here is observationally equivalent to a compound Poisson process with i.i.d. NIG increments whose intensity would be time-varying (cf. Appendix A.1).

intensities of the market and the stock exhibit GARCH dynamics like those of their variance counterparts, but with separate parameters:¹³

$$h_{M,y,t+1} = w_{M,y} + b_{M,y}h_{M,y,t} + \frac{a_{M,y}}{h_{M,z,t}} \left(z_{M,t} - c_{M,y}h_{M,z,t} \right)^2, \quad (4.5)$$

$$\begin{aligned} h_{S,y,t+1} = & \kappa_{S,y}h_{M,y,t+1} + b_{S,y} \left(h_{S,y,t} - \kappa_{S,y}h_{M,y,t} \right) \\ & + \frac{a_{S,y}}{h_{S,z,t}} \left(z_{S,t}^2 - h_{S,z,t} - 2c_{S,y}h_{S,z,t}z_{S,t} \right). \end{aligned} \quad (4.6)$$

As for the variance of Gaussian shocks, idiosyncratic jump intensity has a central tendency $\kappa_{S,y}h_{M,y,t+1}$ and excess idiosyncratic intensity $h'_{S,y,t+1} = h_{S,y,t+1} - \kappa_{S,y}h_{M,y,t+1}$.

Conditional moments of the market and stock returns are derived in the Appendix A.1.2. In particular, total market and stock variances are given by

$$\text{Var}_t^{\mathbb{P}} [R_{M,t+1}] = h_{M,z,t+1} + \left(\frac{a_M^2}{(\alpha_M^2 - \delta_M^2)^{3/2}} \right) h_{M,y,t+1}, \quad (4.7)$$

$$\text{Var}_t^{\mathbb{P}} [R_{S,t+1}] = \underbrace{\beta_{S,z}^2 h_{M,z,t+1} + \beta_{S,y}^2 \left(\frac{a_M^2}{(\alpha_M^2 - \delta_M^2)^{3/2}} \right) h_{M,y,t+1}}_{\text{Total systematic variance}} + \underbrace{h_{S,z,t+1} + \left(\frac{a_S^2}{(\alpha_S^2 - \delta_S^2)^{3/2}} \right) h_{S,y,t+1}}_{\text{Total idiosyncratic variance}}. \quad (4.8)$$

Following the literature, we define idiosyncratic variance as the variance of the residuals obtained after accounting for systematic risk factors, here normal and jump market risk. In sum, our model of market returns is essentially the NIG variant of the model considered in Ornathanalai (2014).¹⁴ We simply extend his framework to allow stocks (i) to have systematic normal and jump risk exposure and (ii) to exhibit idiosyncratic normal and jump risk. In particular, our model remains in the affine class of models, which is key to obtaining a closed-form solution for the price of European options on the market index and individual stocks (cf. Subection 4.2.3). This solution generalizes those of Elkamhi and Ornathanalai (2010) and Babaoğlu (2015) by adding idiosyncratic jumps in stock returns.

4.2.2 Pricing Kernel and Risk Premiums

In an incomplete market setup, the pricing kernel, m_{t+1} , is potentially affected by untraded sources of risk. As highlighted in the literature on modelling the pricing kernel,

¹³ Christoffersen et al. (2012) compare, on market data, a model in which a single factor drives normal variance and jump intensity to a model akin to ours. They find the model with separate variance and intensity dynamics to dominate its counterpart in terms of fitting the data.

¹⁴In Ornathanalai's (2014) study, the NIG variant of the model offers the best fit to market data when compared to variants with Merton jumps, variance gamma jumps, or CGMY jumps (Carr et al., 2002).

in our context, it suffices to work with the projection of the pricing kernel on the observed sources of risk. Indeed, if p_t is the time t price of an asset with a time $t + 1$ cash flow x_{t+1} that depends on the realization of $\{z_{M,t+1}, y_{M,t+1}, z_{S,t+1}, y_{S,t+1}\}$,¹⁵ then

$$p_t = \mathbb{E}_t^{\mathbb{P}} [m_{t+1} x_{t+1}] = \mathbb{E}_t^{\mathbb{P}} \left[\mathbb{E}_t^{\mathbb{P}} [m_{t+1} x_{t+1} \mid z_{M,t+1}, y_{M,t+1}, z_{S,t+1}, y_{S,t+1}] \right] = \mathbb{E}_t^{\mathbb{P}} [\tilde{m}_{t+1} x_{t+1}]$$

where $\tilde{m}_{t+1} = \mathbb{E}_t^{\mathbb{P}} [m_{t+1} \mid z_{M,t+1}, y_{M,t+1}, z_{S,t+1}, y_{S,t+1}]$. If $z_{S,t+1}$ and $y_{S,t+1}$ are orthogonal to the pricing kernel, then they do not matter in the pricing and the projection is simply $\tilde{m}_{t+1} = \mathbb{E}_t^{\mathbb{P}} [m_{t+1} \mid z_{M,t+1}, y_{M,t+1}]$.

The recent literature, however, highlights that firm-specific (or idiosyncratic) risk can be correlated with risk factors that do enter the pricing kernel. In line with much of the option pricing literature, we take a reduced-form approach to modelling the pricing kernel and assume an exponentially affine Radon-Nikodym derivative (RND)

$$\begin{aligned} e^{r_{t+1}} \tilde{m}_{t+1} &= \frac{\frac{dQ}{dP} \big|_{\mathcal{F}_{t+1}^S}}{\frac{dQ}{dP} \big|_{\mathcal{F}_t^S}} \\ &= \frac{\exp(-\Lambda_M z_{M,t+1} - \Gamma_M y_{M,t+1} - \sum_{S \in \mathbb{S}} \Lambda_S z_{S,t+1} - \sum_{S \in \mathbb{S}} \Gamma_S y_{S,t+1})}{\mathbb{E}_{\mathcal{F}_t^S}^{\mathbb{P}} [\exp(-\Lambda_M z_{M,t+1} - \Gamma_M y_{M,t+1} - \sum_{S \in \mathbb{S}} \Lambda_S z_{S,t+1} - \sum_{S \in \mathbb{S}} \Gamma_S y_{S,t+1})]}, \end{aligned} \quad (4.9)$$

where \mathbb{S} is the set of firms in the economy. Implicitly, Λ_S and Γ_S are related to the projection of the pricing kernel on $z_{S,t+1}$ and $y_{S,t+1}$. In particular, if the firm-specific risk factors are not priced, that is $\Lambda_S = \Gamma_S = 0$, then our RND is equivalent to the one used by Christoffersen et al. (2012). As they point out, their RND is consistent with the pricing kernel studied by Bates (2008).¹⁶

As in Christoffersen et al. (2012) and Ornathanalai (2014), the pricing kernel in (4.9) yields an equity risk premium, $\mu_{M,t} - r_t$, which admits a decomposition in terms of a normal and a jump risk premium, that is

$$\mu_{M,t} - r_t = \lambda_M h_{M,z,t} + \gamma_M h_{M,y,t} \quad (4.10)$$

where the mappings between λ_M and γ_M and their pricing kernel counterparts Λ_M and Γ_M are given in Section 4.C. Note that if a market price of risk in the RND is zero (that is $\Gamma_M = 0$), then the associated risk premium is zero (e.g. $\gamma_M = 0$).

¹⁵See, for instance, Rubinstein (1975), Brennan (1979), Aït-Sahalia and Lo (1998), Jackwerth (2000) and Bakshi et al. (2010) for analyses of pricing kernels dependent on market returns.

¹⁶Similar pricing kernels are studied in continuous-time setups by, among others, Bates (1991, 2006), Liu et al. (2005), and Eraker (2008).

Section 4.C further establishes that the equity risk premium on a stock, $\mu_{S,t} - r_t$, can be decomposed in four risk premiums: the normal and jump market risk premiums, as well as the idiosyncratic normal and jump premiums:

$$\mu_{S,t} - r_t = \underbrace{\beta_{S,z} \lambda_M h_{M,z,t} + \gamma_{M,S}(\beta_{S,y}) h_{M,y,t}}_{\text{Systematic}} + \underbrace{\lambda_S h_{S,z,t} + \gamma_S h_{S,y,t}}_{\text{Idiosyncratic}} \quad (4.11)$$

Once again, if a market price of risk in the RND is zero (e.g. $\Gamma_S = 0$), then the associated risk premium is zero (e.g. $\gamma_S = 0$). Besides, note that although the model is affine, the premium associated with systematic jump depends non-linearly on the jump beta, $\beta_{S,y}$, and the market price of jump risk, γ_M , through function $\gamma_{M,S}(\cdot)$, which has a single root at 0. More details are provided in Section 4.C.

To illustrate the difference between our framework and a standard conditional CAPM framework, consider the familiar

$$\beta_{S,t+1}^{\text{CAPM}} = \frac{\text{Cov}_t^{\mathbb{P}}(r_{S,t+1}, r_{M,t+1})}{\text{Var}_t^{\mathbb{P}}(r_{M,t+1})} \equiv \frac{\text{Cov}_t^{\mathbb{P}}(e^{R_{S,t+1}-r_{t+1}}, e^{R_{M,t+1}-r_{t+1}})}{\text{Var}_t^{\mathbb{P}}(e^{R_{M,t+1}-r_{t+1}})}. \quad (4.12)$$

In the context of our model, a first-order approximation of this total beta yields

$$\begin{aligned} \beta_{S,t+1}^{\text{CAPM}} &\simeq \frac{\text{Cov}_t^{\mathbb{P}}(\beta_{S,z} z_{M,t+1} + \beta_{S,y} y_{M,t+1}, z_{M,t+1} + y_{M,t+1})}{\text{Var}_t^{\mathbb{P}}(z_{M,t+1} + y_{M,t+1})} \\ &= \frac{\beta_{S,z} h_{M,z,t+1} + \beta_{S,y} \left(\frac{\alpha_M^2}{(\alpha_M^2 - \delta_M^2)^{3/2}} \right) h_{M,y,t+1}}{h_{M,z,t+1} + \left(\frac{\alpha_M^2}{(\alpha_M^2 - \delta_M^2)^{3/2}} \right) h_{M,y,t+1}} \end{aligned} \quad (4.13)$$

and, in a CAPM setting, the risk premium on the stock would be

$$\mu_{S,t+1}^{\text{CAPM}} - r_{t+1} = \beta_{S,t+1}^{\text{CAPM}} (\mu_{M,t} - r_t) = \beta_{S,t+1}^{\text{CAPM}} (\lambda_M h_{M,z,t+1} + \gamma_M h_{M,y,t+1}). \quad (4.14)$$

Contrasting Equations (4.14) and (4.11) highlights two features of our model. First, in our model, stocks can have different sensitivities to normal and jump risk. Second, λ_S and γ_S are not assumed to be null, but are jointly estimated from past returns and option data.

4.2.3 Option Prices

The model, once risk-neutralized, remains within the affine class of models (see Section 4.D). Hence, we build on the work of Heston and Nandi (2000) and obtain

a closed-form solution for the price of European index and stock options.¹⁷ For $u_t \in \{M_t, S_t\}$, the price of an European call option is

$$C_t(u_t, K, T) = u_t P_{1,t,T} - K e^{-r_{t,T}(T-t)} P_{2,t,T} \quad (4.15)$$

where $r_{t,T} = \frac{1}{T-t} \sum_{j=1}^{T-t} r_{t+j}$, in which r_{t+j} is the deterministic risk-free rate at time $t+j$. The conditional probabilities $P_{1,t,T}$ and $P_{2,t,T}$ are given by

$$\begin{aligned} P_{1,t,T} &= \frac{1}{2} + \frac{1}{\pi} \int_0^\infty \operatorname{Re} \left[\frac{1}{\phi i} \exp(-i\phi \log \tilde{K}_{t,T}) \varphi_{t,T}^{\mathbb{Q}}(\phi i + 1) \right] d\phi \\ P_{2,t,T} &= \frac{1}{2} + \frac{1}{\pi} \int_0^\infty \operatorname{Re} \left[\frac{1}{\phi i} \exp(-i\phi \log \tilde{K}_{t,T}) \varphi_{t,T}^{\mathbb{Q}}(\phi i) \right] d\phi \end{aligned}$$

where i is the imaginary number,

$$\tilde{K}_{t,T} = \frac{K e^{-r_{t,T}(T-t)}}{u_t}$$

and the conditional moment generating function $\varphi_{t,T}^{\mathbb{Q}}(\phi) = \mathbb{E}_t^{\mathbb{Q}} \left[\exp \left(\phi \sum_{j=1}^{T-t} \tilde{R}_{u,t+j} \right) \right]$ of the aggregated excess returns $\sum_{j=1}^{T-t} \tilde{R}_{u,t+j} = \sum_{j=1}^{T-t} (R_{u,t+j} - r_{t+j})$ over the period $[t, T]$ satisfies

$$\varphi_{t,T}^{\mathbb{Q}}(\phi) = \exp \left(\begin{aligned} &\mathcal{A}_{u,T-t}(\phi) + \mathcal{B}_{u,T-1}(\phi) h_{M,z,t+1}^* + \mathcal{C}_{u,T-1}(\phi) h_{M,y,t+1}^* \\ &+ \mathcal{D}_{u,T-1}(\phi) h_{S,z,t+1}^* + \mathcal{E}_{u,T-1}(\phi) h_{S,y,t+1}^* \end{aligned} \right).$$

The deterministic functions $\mathcal{A}_u, \mathcal{B}_u, \mathcal{C}_u, \mathcal{D}_u, \mathcal{E}_u$ are calculated based on the recursion in Section 4.E. In particular, $\mathcal{D}_{M,T-1} = \mathcal{E}_{M,T-1} = 0$.

4.3 Joint Estimation Using Returns and Option Prices

Relying on a joint estimation procedure is of particular importance to our study. Indeed, the risk premium parameters we aim to study are relatively poorly identified under the physical measure. However, these parameters play a crucial role in the pricing kernel and, as such, are key to reconcile the price of the options and the underlying returns.¹⁸ Moreover, in the absence of jumps, a deep out-of-the-money option would be almost worthless, especially if the option is relatively short-dated. These options

¹⁷Heston and Nandi (2000) relies on an inversion similar to the one of Gil-Pelaez (1951).

¹⁸See, among others, Chernov and Ghysels (2000), Pan (2002), Chernov (2003), Eraker (2004), Santa-Clara and Yan (2010), Christoffersen et al. (2012), and Ornathanalai (2014).

will thus improve our ability to estimate the likelihood of jumps. Hence, the richness of stock option data plays a key role in allowing us to extract the expected risk premium associated with each risk factor.

4.3.1 Data

To estimate the model, we use the returns and prices of options on the S&P 500, as proxy for the market, and on 260 stocks that are or were part of the index since 1996. These stocks were selected based on whether their options had been actively traded over at least a consecutive 5-year window. Daily index and stock returns, from January 1996 to August 2015, are obtained from the Center for Research in Security Prices (CRSP).¹⁹ To compute the corresponding daily excess log-returns (henceforth, returns), we use the one-month Treasury bill rate (from Ibbotson Associates) as extracted from Kenneth French's data library.

The prices of options on the SPX and the stocks, between January 1996 to August 2015, are obtained from OptionMetrics.²⁰ We restrict our analysis to out-of-the-money monthly options with at least one week and at most one year to maturity. Observations for which the ask price is lower than the bid price are excluded. The price of the option is defined as the mid point between the ask and the bid, and options with a price lower than the bid-ask spread are excluded. Moreover, the open interest and the volume must be strictly positive. We further remove options that violate the common arbitrage conditions. For options on individual stocks, we follow Broadie et al. (2007) and de-Americanize the option prices.²¹ Finally, among the remaining options, we select the three most liquid puts and three most liquid calls on each Wednesday, for each maturity available.²² This leaves us with a total of 44,267 option prices on the

¹⁹In fact, we extract returns starting from January 1986. Returns between January 1986 and December 1995 are used to warmup the variance process; their likelihood, however, does not impact the estimation of the parameters. A similar procedure is used for individual stocks.

²⁰The zero-coupon term structure is also extracted from OptionMetrics and used for option pricing. The rate corresponding to an option's maturity is obtained through linear interpolation whenever necessary.

²¹Specifically, for each American option, OptionMetrics uses a Cox et al. (1979) binomial tree to derive the option's implied volatility, accounting for dividends. Given this implied volatility and dividends extracted from OptionMetrics, we compute the price of the corresponding European option.

²²We follow the literature and use Wednesday data because it is the least likely day to be a holiday and it is least likely to be affected by weekend effects. For more details on the advantages of using Wednesday data, see Dumas et al. (1998). If markets are closed on a given Wednesday (e.g. Christmas, January 1, Independence day or 9/11) we use the previous business day.

Table 4.1: Description of the SPX index option data (1996–2015).

Panel A: Number of option contracts.					
	DTM ≤ 30	30 < DTM ≤ 90	90 < DTM ≤ 180	180 < DTM ≤ 250	DTM > 250
0.80 < $K/F \leq 0.85$	256	1,666	1,046	715	487
0.85 < $K/F \leq 0.90$	505	1,999	1,258	861	639
0.90 < $K/F \leq 0.95$	789	2,220	1,410	1,031	828
0.95 < $K/F \leq 1.00$	1,278	2,505	1,530	1,193	884
1.00 < $K/F \leq 1.05$	2,234	4,379	2,008	1,101	973
1.05 < $K/F \leq 1.10$	484	2,432	1,517	908	723
1.10 < $K/F \leq 1.15$	83	972	871	624	512
1.15 < $K/F \leq 1.20$	18	288	380	373	287
All	5,647	16,461	10,020	6,806	5,333
Panel B: Average option prices.					
	DTM ≤ 30	30 < DTM ≤ 90	90 < DTM ≤ 180	180 < DTM ≤ 250	DTM > 250
0.80 < $K/F \leq 0.85$	1.19	4.18	12.00	21.45	29.92
0.85 < $K/F \leq 0.90$	1.61	7.06	18.45	30.56	41.24
0.90 < $K/F \leq 0.95$	2.93	13.05	27.94	43.18	55.69
0.95 < $K/F \leq 1.00$	8.71	26.05	45.48	62.61	78.04
1.00 < $K/F \leq 1.05$	7.29	19.61	38.88	60.88	77.47
1.05 < $K/F \leq 1.10$	1.92	6.78	16.99	31.11	44.74
1.10 < $K/F \leq 1.15$	1.34	3.68	9.16	18.44	28.45
1.15 < $K/F \leq 1.20$	1.22	3.17	6.17	10.57	17.08
All	5.65	13.50	25.84	39.90	53.11
Panel C: Average implied volatility.					
	DTM ≤ 30	30 < DTM ≤ 90	90 < DTM ≤ 180	180 < DTM ≤ 250	DTM > 250
0.80 < $K/F \leq 0.85$	0.4123	0.3094	0.2781	0.2586	0.2478
0.85 < $K/F \leq 0.90$	0.3252	0.2676	0.2509	0.2364	0.2294
0.90 < $K/F \leq 0.95$	0.2515	0.2320	0.2117	0.2186	0.2172
0.95 < $K/F \leq 1.00$	0.1830	0.1968	0.2049	0.2024	0.2018
1.00 < $K/F \leq 1.05$	0.1482	0.1574	0.1720	0.1844	0.1866
1.05 < $K/F \leq 1.10$	0.1910	0.1595	0.1609	0.1669	0.1701
1.10 < $K/F \leq 1.15$	0.2757	0.1879	0.1691	0.1679	0.1696
1.15 < $K/F \leq 1.20$	0.3725	0.2367	0.1884	0.1687	0.1657
All	0.2046	0.2057	0.2043	0.2024	0.1996

Moneyness is defined as K/F , where F is the forward price of the index and K is the option's strike price. DTM stands for days to maturity. Our final option dataset contains 44,267 observations.

Table 4.2: Description of firms option data (1996–2015).

Panel A: Number of option contracts.					
	DTM ≤ 30	30 < DTM ≤ 90	90 < DTM ≤ 180	180 < DTM ≤ 250	DTM > 250
0.80 < $K/F \leq 0.85$	6,791	53,859	75,543	46,305	14,036
0.85 < $K/F \leq 0.90$	16,416	100,294	101,743	56,700	15,120
0.90 < $K/F \leq 0.95$	42,747	168,293	120,217	61,422	15,549
0.95 < $K/F \leq 1.00$	100,267	216,362	123,431	60,425	14,575
1.00 < $K/F \leq 1.05$	103,111	245,827	160,306	85,404	20,169
1.05 < $K/F \leq 1.10$	36,508	169,034	146,104	80,884	19,519
1.10 < $K/F \leq 1.15$	13,119	90,264	110,095	67,035	17,543
1.15 < $K/F \leq 1.20$	5,390	47,073	73054	47,754	14,333
All	324,349	1,091,006	910,493	505,929	130,844
187,604					
2,962,621					
Panel B: Average option prices.					
	DTM ≤ 30	30 < DTM ≤ 90	90 < DTM ≤ 180	180 < DTM ≤ 250	DTM > 250
0.80 < $K/F \leq 0.85$	1.04	1.42	2.16	2.96	4.05
0.85 < $K/F \leq 0.90$	1.06	1.54	2.56	3.55	5.00
0.90 < $K/F \leq 0.95$	1.12	1.79	3.26	4.52	6.22
0.95 < $K/F \leq 1.00$	1.45	2.52	4.48	5.98	8.11
1.00 < $K/F \leq 1.05$	1.47	2.50	4.36	5.85	8.03
1.05 < $K/F \leq 1.10$	1.16	1.79	3.08	4.21	6.03
1.10 < $K/F \leq 1.15$	1.14	1.61	2.51	3.34	4.77
1.15 < $K/F \leq 1.20$	1.16	1.53	2.26	2.98	4.03
All	1.33	2.03	3.25	4.31	5.87
2.24					
2.44					
2.73					
3.34					
3.46					
2.77					
2.50					
2.36					
2.89					
Panel C: Average implied volatility.					
	DTM ≤ 30	30 < DTM ≤ 90	90 < DTM ≤ 180	180 < DTM ≤ 250	DTM > 250
0.80 < $K/F \leq 0.85$	0.7919	0.5436	0.4274	0.3872	0.3515
0.85 < $K/F \leq 0.90$	0.6392	0.4623	0.3851	0.3585	0.3352
0.90 < $K/F \leq 0.95$	0.4903	0.3869	0.3537	0.3398	0.3238
0.95 < $K/F \leq 1.00$	0.3618	0.3388	0.3386	0.3318	0.3169
1.00 < $K/F \leq 1.05$	0.3504	0.3258	0.3214	0.3162	0.3066
1.05 < $K/F \leq 1.10$	0.4675	0.3617	0.3217	0.3092	0.3017
1.10 < $K/F \leq 1.15$	0.5917	0.4244	0.3435	0.3168	0.2993
1.15 < $K/F \leq 1.20$	0.7147	0.4897	0.3802	0.3407	0.3072
All	0.4252	0.3819	0.3514	0.3334	0.3163
0.4569					
0.4183					
0.3785					
0.3418					
0.3268					
0.3453					
0.3703					
0.4017					
0.3661					

Moneyneess is defined as K/F , where F is the forward price of the underlying and K is the option's strike price. DTM stands for days to maturity. Our final option dataset contains 2,975,839 observations.

SPX, and 2,975,839 on the 260 stocks.²³ Tables 4.1 and 4.2 summarize the option data sets. Figure 4.1 provides an overview of how implied volatilities vary through time as the S&P 500 evolves. As evidenced in the lower panel of the figure, while implied volatilities on stocks comove with implied volatilities on S&P 500 options, the former are significantly larger than the latter.

4.3.2 Joint Estimation

Following Christoffersen et al. (2012) and Ornathanalai (2014), the model's parameters are estimated by maximizing the weighted joint log-likelihood function

$$L_u(\Theta_u) = \frac{T_u + N_u}{2} \left(\frac{L_{u,\text{returns}}(\Theta_u)}{T_u} + \frac{L_{u,\text{options}}(\Theta_u)}{N_u} \right), \quad (4.16)$$

where, $u \in \{M, S\}$, T_u is the number of index returns observed, N_u is the total number of option observations, and Θ_u represents the parameter set of the model.

We opt for a two-stage estimation approach. That is, we first maximize the joint likelihood L_M with respect to Θ_M and then, turn to maximizing L_S for each stock taking the results for the market as given. Although it has inconvenient, this approach is crucial to keeping the estimation procedure tractable in our settings. Indeed, in opposition to typical GARCH processes in which the noise term is fully determine once we condition on observed returns and the initial variance, the presence of jumps implies, focusing on the market model, that the Gaussian component $z_{M,t}$ and the jump part $y_{M,t}$ of time t innovation cannot be separated. Consequently, as pointed out by Durham et al. (2015), the conditional variance $h_{M,z,t}$ and intensity $h_{M,y,t}$ remain uncertain, even with the observed returns up to time t .²⁴ However, both $h_{M,z,t+1}$ and $h_{M,y,t+1}$ can be fully recovered from the initial conditions $h_{M,z,1}$, $h_{M,y,1}$, the returns $R_{M,1:t} = \{R_{M,s}\}_{s=1}^t$ and the jump innovations $y_{M,1:t} = \{y_{M,s}\}_{s=1}^t$. In this spirit, we propose a particle filter that infers the average (filtered) $z_{M,t}$, $y_{M,t}$, $h_{M,z,t}$ and $h_{M,y,t}$, while accounting for the uncertainty with respect to the conditional variance and the jump intensity. Section 4.F further describes the particle filter used to compute the log-likelihood $L_{M,\text{returns}}(\Theta_M)$.

²³In total, we considered options on the 1,000 different firms that were part of the S&P 500 at any point in our sample. Our selection procedure discarded 738 firms. Two additional firms (tickers BEN and NEE) were further discarded because they experimented very extreme returns that caused numerical problems in the particle filter; we are currently working on improving the importance sampling step in order to reintroduce these firms.

²⁴Technically speaking, $\mathcal{G}_t^M = \sigma\{R_{M,\tau}\}_{\tau=1}^{t-1}$ is the σ -field generated by the returns process which is coarser than the σ -field $\mathcal{F}_t^M = \sigma\{z_{M,\tau}, y_{M,\tau}\}_{\tau=1}^{t-1}$ generated by the innovation terms. The conditional variance $h_{M,z,t}$ and the jump intensity $h_{M,y,t}$ are both \mathcal{F}_{t-1}^M measurable, but they are not \mathcal{G}_{t-1}^M -measurable.




Figure 4.1: S&P 500 and ATM implied volatilities.

The upper panel of this figure presents the level of the S&P 500 index between January 1996 to August 2015; gray-shaded regions highlight NBER-dated recessions. The middle panel reports S&P 500 index excess returns over the same period. The lower panel reports the weekly at-the-money implied volatility from the nearest-to-maturity SPX options as extracted from OptionMetrics, along with the average of the weekly at-the-money implied volatility across the firms.

The same procedure is applied to each stock in a second estimation stage, keeping the market parameters and latent variables fixed. Conceptually, the particle filter could be extended to deal with a one-stage estimation of the market and the 260 firms; numerically, however, this would be absolutely intractable. Our alternative is computationally efficient and, given the richness of the index option data, we are confident that the two-stage estimation procedure yields parameter estimates for the market model that are more accurate than the ones that could be obtained from a poorly behaved one-stage procedure.

Following the option-pricing literature, the option fit's log-likelihood, $L_{u,\text{options}}(\Theta_u)$, is based on relative implied volatility pricing errors.²⁵ In particular, if $IV_{u,k}^{\text{MKT}}$ is the Black and Scholes (1973) implied volatility associated with the market price of option k on underlying $u \in \{M, S\}$ and $IV_{u,k}^{\text{MODEL}}$ the implied volatility inverted from the corresponding model price, then the relative implied volatility error is

$$e_{u,k} = \frac{IV_{u,k}^{\text{MODEL}} - IV_{u,k}^{\text{MKT}}}{IV_{u,k}^{\text{MKT}}}.$$

Assuming that the relative implied volatility error is normally distributed, and uncorrelated with shocks in returns, $e_{u,k} \sim \mathcal{N}(0, \sigma_e^2)$, we obtain

$$L_{u,\text{options}}(\Theta_u) = -\frac{1}{2} \sum_{k=1}^N \left(\log(2\pi\sigma_e^2) + \frac{e_{u,k}^2}{\sigma_e^2} \right).$$

Note that σ_e is identified using the sample standard deviation of $\{e_{u,k}\}_{k=1}^{N_u}$.

4.4 Empirical Results

4.4.1 Market

Although the focus of our study is the pricing of idiosyncratic risk, we first briefly discuss results obtained at the market level. Overall, these results are very close to those

²⁵This criterion, or variants thereof, is used by Bakshi et al. (2008), Christoffersen et al. (2012), and Ornathanalai (2014). Renault (1997) offers an interesting discussion on the benefits of using IVRMSE when comparing option pricing models. Alternatively, some authors will consider vega-weighted RMSE (VWRMSE) since VWRMSE and IVRMSE take very similar value, while the former have the advantage of being faster to compute than the latter. See for instance Carr and Wu (2007) and Trolle and Schwartz (2009). Note that using relative implied volatility errors has the advantage of not assigning excessive weighting to option prices observed during the financial crisis.

Table 4.3: Index parameters estimated using returns and option data.

	Normal	Jump
λ_M / γ_M	0.824 (2.40E-10)	0.701 (3.55E-10)
$w_{M,v}$	-1.63E-06 (2.50E-08)	-4.05E-07 (1.88E-08)
$a_{M,v}$	2.42E-06 (3.36E-09)	4.44E-06 (7.83E-09)
$b_{M,v}$	0.940 (6.94E-10)	0.934 (4.58E-09)
$c_{M,v}$	144.19 (2.86E-11)	140.50 (1.30E-11)
α_M		11.856 (1.31E-08)
δ_M		-7.018 (3.01E-10)
Average risk premium	2.33	3.85
Median risk premium	1.72	3.04
Persistence	0.991	
Percent of annual variance	74.0	26.0
Avg. $\text{Var}_{t-1}[h_{S,\cdot,t+1}]$	6.61E-11	3.79E-10
Average volatility (%)	18.01	
Average skewness	-6.25	
Average excess kurtosis	383.70	
Skewness of innovations, $\varepsilon_{M,t}$	-0.12	
Ex. kurtosis of innovations, $\varepsilon_{M,t}$	-0.07	
Return log-likelihood	78,395	
Option log-likelihood	12,788	
Total log-likelihood	91,183	
RIVRMSE	14.39	

The index parameters are estimated using daily index returns and weekly cross-sections of out-of-the-money options, from January 1996 to August 2015. Parameters are estimated using multiple simplex search method optimizations (fminsearch in Matlab). Robust standard errors are calculated from the outer product of the gradient at the optimal parameter values.

in the option pricing literature. In particular, our results are much in line with those reported by Ornathanalai (2014) for the NIG variant of his model, which is essentially our market model. Parameters, reported in Table 4.3, are largely similar, except maybe for $a_{M,v}$, the parameter governing the variance of jump intensity, which is much larger for us than it was for Ornathanalai. This difference could be due to our sample covering more of the Great Recession and its aftermath.

For each subset of option O , Table 4.4 reports two metrics

$$\begin{aligned} \text{IVRMSE} &= \sqrt{\frac{1}{N} \sum_{k \in O} (IV_k^{\text{MODEL}} - IV_k^{\text{MKT}})^2} \text{ and} \\ \text{RIVRMSE} &= \sqrt{\frac{1}{N} \sum_{k \in O} \left(\frac{IV_k^{\text{MODEL}} - IV_k^{\text{MKT}}}{IV_k^{\text{MKT}}} \right)^2}. \end{aligned} \quad (4.17)$$

The first, IVRMSE, provides an absolute measure of the implied-volatility pricing errors. The latter, a relative measure that is probably more informative when comparing

pricing errors through time. By both measures, our market fit to the option data, as detailed in Panel A of Table 4.4 compares favourably to the results in the option-pricing literature. This is true through time and across maturities and moneyness levels. As documented by Ornathanalai, the NIG jumps in our model allow for particularly large levels of (negative) skewness and excess kurtosis (cf. Table 4.3). This theoretical feature of the model explains its particularly good fit across maturities and moneyness levels. Moreover, the NIG jumps properly capture, empirically, the nonnormal innovations in returns; consequently, the filtered conditionally standard normal innovations, $\varepsilon_{M,t}$, have skewness and excess kurtosis that are close to zero, as should be. Figure 4.2 plots the filtered normal innovations $z_{M,t}$ (top panel), jumps (middle panel) and volatility components (bottom panel). Again, results are qualitatively similar to those of Ornathanalai.

Table 4.3 also reports risk premiums based on the average and median level of normal, $\lambda_M h_{M,z,t}$, and jump, $\gamma_M h_{M,y,t}$ components of the conditional equity risk premium (ERP). The median levels of the premiums are respectively 1.72% and 3.04%, for a median ERP of 4.76%. These numbers are comparable to those of Ornathanalai, who reports an annualized normal risk premium of 1.43% and a jump risk premium of 3.22%, for a total of 4.65%, based on the unconditional level of variance and jump intensity. Hence, although Table 4.3 reports that the jump component of variance, $h_{M,y}$, explains on average only 26.7% of total variance, $h_{M,z} + \left(\frac{\sigma_M^2}{(\alpha_M^2 - \delta_M^2)^{3/2}} \right) h_{M,y}$, the jump risk premium outweighs its normal counterpart.

The average ERP is higher than the median at 6.17%, and decomposes into an average normal premium of 2.32% and an average jump premium of 3.85%. Naturally, the average is more sensitive than the median (or any measure based on unconditional GARCH levels) to extreme values of the premiums observed during periods of turmoil. The top panel of Figure 4.3 reports how the premium unfolds through time. At its peak, in November 2008, the estimated ERP reaches 40.16%. While this number may appear high, Martin (2016) measure of the ERP, as extracted from one-month-to-maturity options alone, rises to more than 50% around the same time, while its three-month counterpart flirts with the 40% level. Using a panel of options with median time-to-expiration of 14 business days, Bollerslev and Todorov (2011) find that the jump component of the ERP rises above 40% during the same period.

The bottom panel of Figure 4.3 reports, on a daily basis, the ratio of the ERP that is explained by the jump component. This ratio is at its lowest during periods of turmoil, when the normal risk carries a higher than usual premium. When the ERP is particularly low, which coincides with periods of low volatility on the market, the




Figure 4.2: Filtered innovations and variances for the market model.

This figure presents the states for the market model, filtered using the parameters in Table 4.3. The top panel reports the filtered Gaussian innovations. The middle panel reports the filtered jump components (superposed on returns). The dashed line in the bottom panel reports, in annualized terms, the contribution of the jump component to the returns' conditional volatility, $\sqrt{252 \left(\frac{\sigma_M^2}{(\sigma_M^2 - \sigma_J^2)^2} \right) h_{M,Y,t}}$; the solid line reports the total volatility, $\sqrt{252 \left(h_{M,Z,t} + \left(\frac{\sigma_M^2}{(\sigma_M^2 - \sigma_J^2)^2} \right) h_{M,Y,t} \right)}$.

Table 4.4: Valuation errors on the options used for the estimation.

Panel A: Valuation errors on the options used for the estimation of the market model					
Overall IVRMSE and RIVRMSE			Sorting by year		
All	IVRMSE 3.086	RIVRMSE 14.391	1996	IVRMSE 2.326	RIVRMSE 15.834
Sorting by maturity	IVRMSE	RIVRMSE	1997	3.844	17.178
			1998	5.318	20.129
			1999	4.625	17.786
			2000	2.419	10.508
			2001	2.756	11.757
DTM ≤ 30	3.364	16.185	2002	2.192	9.345
30 < DTM ≤ 90	3.108	14.606	2003	2.313	11.214
90 < DTM ≤ 180	2.936	13.590	2004	1.260	8.067
180 < DTM ≤ 270	3.016	13.607	2005	2.307	17.593
270 < DTM ≤ 365	3.076	14.137	2006	2.327	17.721
Sorting by moneyness	IVRMSE	RIVRMSE	2007	2.002	14.031
			2008	5.086	13.635
			2009	4.018	15.546
			2010	3.724	17.210
			2011	3.416	13.978
0.80 < $K/F \leq 0.85$	3.815	11.663	2012	2.615	13.138
0.85 < $K/F \leq 0.90$	3.533	11.818	2013	1.255	8.961
0.90 < $K/F \leq 0.95$	3.334	12.543	2014	1.543	11.914
0.95 < $K/F \leq 1.00$	3.006	13.421	2015	2.018	16.756
1.00 < $K/F \leq 1.05$	2.684	16.316			
1.05 < $K/F \leq 1.10$	2.729	16.577			
1.10 < $K/F \leq 1.15$	2.808	15.384			
1.15 < $K/F \leq 1.20$	3.037	15.134			
Panel B: Average valuation errors on the options used for the estimation of the firm model					
Overall average IVRMSE and RIVRMSE			Sorting by year		
All	IVRMSE 6.241	RIVRMSE 13.483	1996	IVRMSE 3.914	RIVRMSE 11.750
Sorting by maturity	IVRMSE	RIVRMSE	1997	4.492	11.879
			1998	5.669	11.865
			1999	5.627	11.229
			2000	7.343	12.146
			2001	6.411	11.793
DTM ≤ 30	9.195	17.130	2002	6.638	13.457
30 < DTM ≤ 90	6.296	13.720	2003	4.624	11.786
90 < DTM ≤ 180	5.253	12.125	2004	3.131	9.881
180 < DTM ≤ 270	5.201	12.355	2005	3.473	12.214
270 < DTM ≤ 365	5.307	12.719	2006	3.482	11.293
Sorting by moneyness	IVRMSE	RIVRMSE	2007	4.315	12.787
			2008	10.121	14.176
			2009	8.138	14.595
			2010	4.097	11.882
			2011	4.656	11.911
0.80 < $K/F \leq 0.85$	7.359	12.184	2012	4.149	11.888
0.85 < $K/F \leq 0.90$	6.837	12.252	2013	3.062	10.215
0.90 < $K/F \leq 0.95$	6.249	12.670	2014	3.595	12.737
0.95 < $K/F \leq 1.00$	5.948	14.023	2015	3.718	12.907
1.00 < $K/F \leq 1.05$	5.613	14.239			
1.05 < $K/F \leq 1.10$	5.925	13.469			
1.10 < $K/F \leq 1.15$	6.318	13.413			
1.15 < $K/F \leq 1.20$	6.982	13.826			

We use the joint MLE estimates of Tables 4.3 and 4.5 to compute implied volatility root mean squared errors (IVRMSE) and relative implied volatility root mean squared errors (RIVRMSE) for various moneyness, maturity, and year bins. We then average IVRMSE and RIVRMSE for each moneyness, maturity and year bin across firms. IVRMSEs and RIVRMSEs are given in percentage.

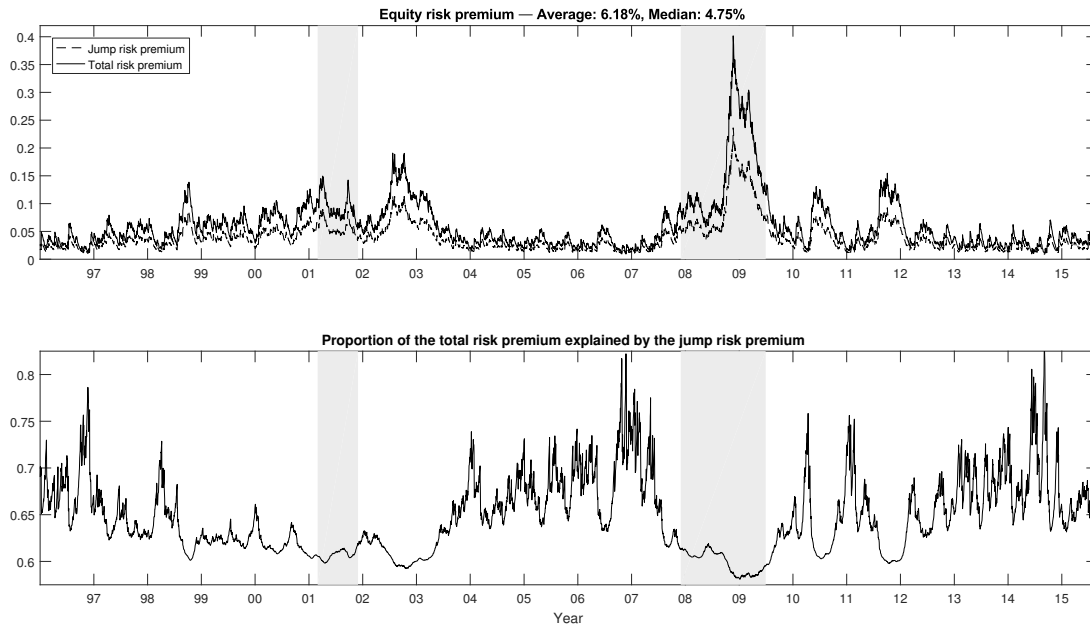


Figure 4.3: Annualized normal and jump risk premiums for the market model.

The top panel of this figure reports the annualized equity risk premium, $252(\lambda_M h_{M,z,t} + \gamma_M h_{M,y,t})$, and the component due to jump risk, $252\gamma_M h_{M,y,t}$. The lower panel reports, on a daily basis, the proportion of the total premium explained by the jump component.

jump risk premium explains up to 80% of the total ERP. Hence, while Figure 4.2 documents that as Bates (2008) pointed out, jump risk is countercyclical, the relative importance of jump risk in the ERP appears to be mildly cyclical.

In sum, our results at the market level are consistent with the literature.

4.4.2 Idiosyncratic *Jump* Risk Matters

We now turn to our paper's main empirical contribution. Namely, while our results are consistent with the literature highlighting that idiosyncratic does matter for the equity risk premium, we provide evidence that idiosyncratic jump risk is at the center of this empirical phenomenon.

Table 4.5 reports summary statistics on the parameters associated with the 260 stocks under consideration. While there is substantial cross-sectional variation, the average value of the parameters of the variance and intensity processes are comparable to the parameters obtained for the market model. Remarkably, more than 75% of the firms exhibit less negative skewness and excess kurtosis than the market. This is consistent

with Bakshi et al. (2003), who document that the option-implied skewness of individual stocks is typically much less negative than that of the market index.²⁶

Of particular interest, the normal and jump betas are on average 0.914 and 1.090 respectively. The normal beta ranges from 0.206 to 1.770, but 50% of firms under consideration have a normal beta between 0.688 and 1.120. In comparison, the jump beta ranges from 0.202 to 3.942, while 50% of firms under consideration have a normal beta between 0.815 and 1.291. Interestingly, the correlation between firms normal and jump beta is only of 0.195 (cf. Table 4.6). That is, there is a positive correlation, but firms with large normal betas do not necessarily have large jump betas, and the other way around.

Table 4.6 further reports the correlation between $\beta_{S,v}$, $v \in \{z, y\}$, and the firm-by-firm time series average of the systematic normal, $\beta_{S,z} \lambda_M h_{M,z,t}$, and jump, $\gamma_{M,S}(\beta_{S,y}) h_{M,y,t}$, risk premiums. Unsurprisingly, the correlation between $\beta_{S,v}$ and the corresponding systematic premium is high but imperfect.²⁷ Consistent with the modest correlation between $\beta_{S,z}$ and $\beta_{S,y}$, the correlation between the normal (jump) beta and the systematic jump (normal) premium are positive but modest at 0.191 (0.190). These results highlight the importance of accounting for separate systematic premiums on both types of risk, as emphasized by Elkamhi and Ornathanalai (2010) and Babaoğlu (2015).

Figure 4.4 decomposes, for each of the 260 stocks in our sample, the stock's equity risk premium in terms of the premiums associated with the four different risk factors in the model: (i) systematic normal, $\beta_{S,z} \lambda_M h_{M,z,t}$, (ii) systematic jump, $\gamma_{M,S}(\beta_{S,y}) h_{M,y,t}$, (iii) idiosyncratic normal, $\lambda_S h_{S,z,t}$, and (iv) idiosyncratic jump, $\gamma_S h_{S,y,t}$. Making this decomposition possible is the key econometric contribution of our paper. The empirical results are striking. First, consistent with financial theory, we find that systematic risk is priced and explains an important part (59.8%) of the equity risk premium. Normal systematic risk explains 20.3% of the total equity risk premium; systematic jump risk, 39.5%. Consistent with the discussion on the betas, the proportion of the systematic premium explained by its jump component ($\frac{(ii)}{(i)+(ii)}$), varies largely, from 14.7% to 67.7%.

²⁶Albuquerque (2012) develops and empirically supports a model in which conditional asymmetric stock return correlations and negative skewness in aggregate returns are caused by cross-sectional heterogeneity in firm announcement events.

²⁷ For the normal premium, the time series average

$$\frac{1}{T_S} \sum_{t \in T_S} \beta_{S,z} \lambda_M h_{M,z,t} = \beta_{S,z} \lambda_M \bar{h}_{M,z,T_S}, \quad S \in \mathbb{S}, \quad (4.18)$$

is linear in $\beta_{S,z}$ which makes the imperfect correlation puzzling at first sight. However, the firm-specific window of available data, T_S , introduces cross-sectional variation in \bar{h}_{M,z,T_S} .

Table 4.5: Firm parameters estimated using returns and option data.

	Average	SD	Min	Q1	Median	Q3	Max
$\beta_{S,z}$	0.908	0.297	0.206	0.676	0.925	1.116	1.756
$\beta_{S,y}$	1.085	0.417	0.202	0.809	1.034	1.283	3.926
$(\beta_{S,z} + \beta_{S,y})/2$	0.997	0.278	0.433	0.816	0.977	1.158	2.841
λ_S	0.000	0.002	-0.015	0.000	0.000	0.000	0.010
γ_S	0.953	0.300	0.083	0.791	0.974	1.086	2.092
$\kappa_{S,z}$	0.957	0.705	0.083	0.509	0.776	1.200	4.665
$a_{S,z}$	2.06E-06	1.23E-06	2.25E-07	1.31E-06	1.73E-06	2.55E-06	9.07E-06
$b_{S,z}$	0.992	0.003	0.978	0.991	0.993	0.994	0.999
$c_{S,z}$	103.90	46.85	-90.17	69.23	115.46	139.98	207.42
$\kappa_{S,y}$	0.518	0.291	0.091	0.329	0.471	0.627	2.212
$a_{S,y}$	4.82E-06	3.41E-06	4.54E-07	2.25E-06	4.22E-06	6.88E-06	2.43E-05
$b_{S,y}$	0.927	0.057	0.303	0.907	0.927	0.961	0.997
$c_{S,y}$	125.08	51.86	-252.57	107.15	132.38	147.54	325.27
α_S	10.757	4.880	0.744	8.353	10.587	12.294	43.452
δ_S	-6.239	2.209	-15.255	-7.316	-6.572	-5.128	-0.302
Avg. volatility (%)	39.21	13.57	18.90	30.56	36.39	45.67	162.41
Avg. skewness	-6.49	15.03	-132.05	-4.58	-2.81	-1.61	-0.11
Avg. excess kurtosis	1024.24	4359.87	1.28	52.63	114.23	231.93	42494.19
Skewness of innovations, $\varepsilon_{S,t}$	0.10	0.35	-3.60	0.06	0.12	0.18	0.86
Ex. kurtosis of innovations, $\varepsilon_{S,t}$	1.68	8.06	-0.07	0.49	0.74	1.04	100.40
RIVRMSE	13.48	2.73	8.87	11.93	12.96	14.59	39.63

The index parameters are estimated using daily index returns and weekly cross-sections of out-of-the-money options, from January 1996 to August 2015. Parameters are estimated using multiple simplex search method optimizations (fminsearch in Matlab). For firms, we report statistics on the joint MLE estimates obtained for the 260 individual stocks in our sample. Q1 and Q3 report the 25th and 75th percentiles of the estimates.

Table 4.6: Correlation between the parameters of stocks.

	$\beta_{S,y}$	γ_S	$\kappa_{S,z}$	$\kappa_{S,y}$	IVAR	$\widehat{\text{IVAR}}_{S,z}$	$\widehat{\text{IVAR}}_{S,y}$	$\text{RP}_{M,z}$	$\text{RP}_{M,y}$	$\text{RP}_{S,y}$
$\beta_{S,z}$										
$\beta_{S,y}$	0.195	-0.105	0.042	0.332	0.270	0.370	0.166	0.989	0.193	0.239
γ_S		-0.058	0.015	-0.016	0.066	0.159	0.216	0.190	0.982	0.069
$\kappa_{S,z}$			-0.006	0.002	-0.070	0.051	-0.114	-0.106	-0.059	0.489
$\kappa_{S,y}$				0.429	0.347	0.336	0.266	0.048	0.014	0.367
IVAR					0.280	0.504	0.145	0.346	-0.012	0.609
$\widehat{\text{IVAR}}_{S,z}$						0.480	0.375	0.273	0.076	0.314
$\widehat{\text{IVAR}}_{S,y}$							0.490	0.401	0.188	0.609
								0.177	0.222	0.576

This table reports the correlation between a subset of the parameters associated with the 260 stocks under consideration. IVAR stands for idiosyncratic variance; $\widehat{\text{IVAR}}_{S,z}$, and $\widehat{\text{IVAR}}_{S,y}$ are respectively the firm-by-firm average of the following time series: $h_{S,z,t} + \left(\frac{\sigma_S^2}{\sigma_S^2 - \sigma_S^2 y^2}\right) h_{S,y,t}$, $h_{S,z,t} - \kappa_{S,z} h_{M,z,t}$, and $h_{S,y,t} - \kappa_{S,y} h_{M,y,t}$. The $\text{RP}_{u,y}$ are the firm-by-firm average of the risk premiums associated with source of risk v_u , $v \in \{z, y\}$, $u \in \{M, S\}$.

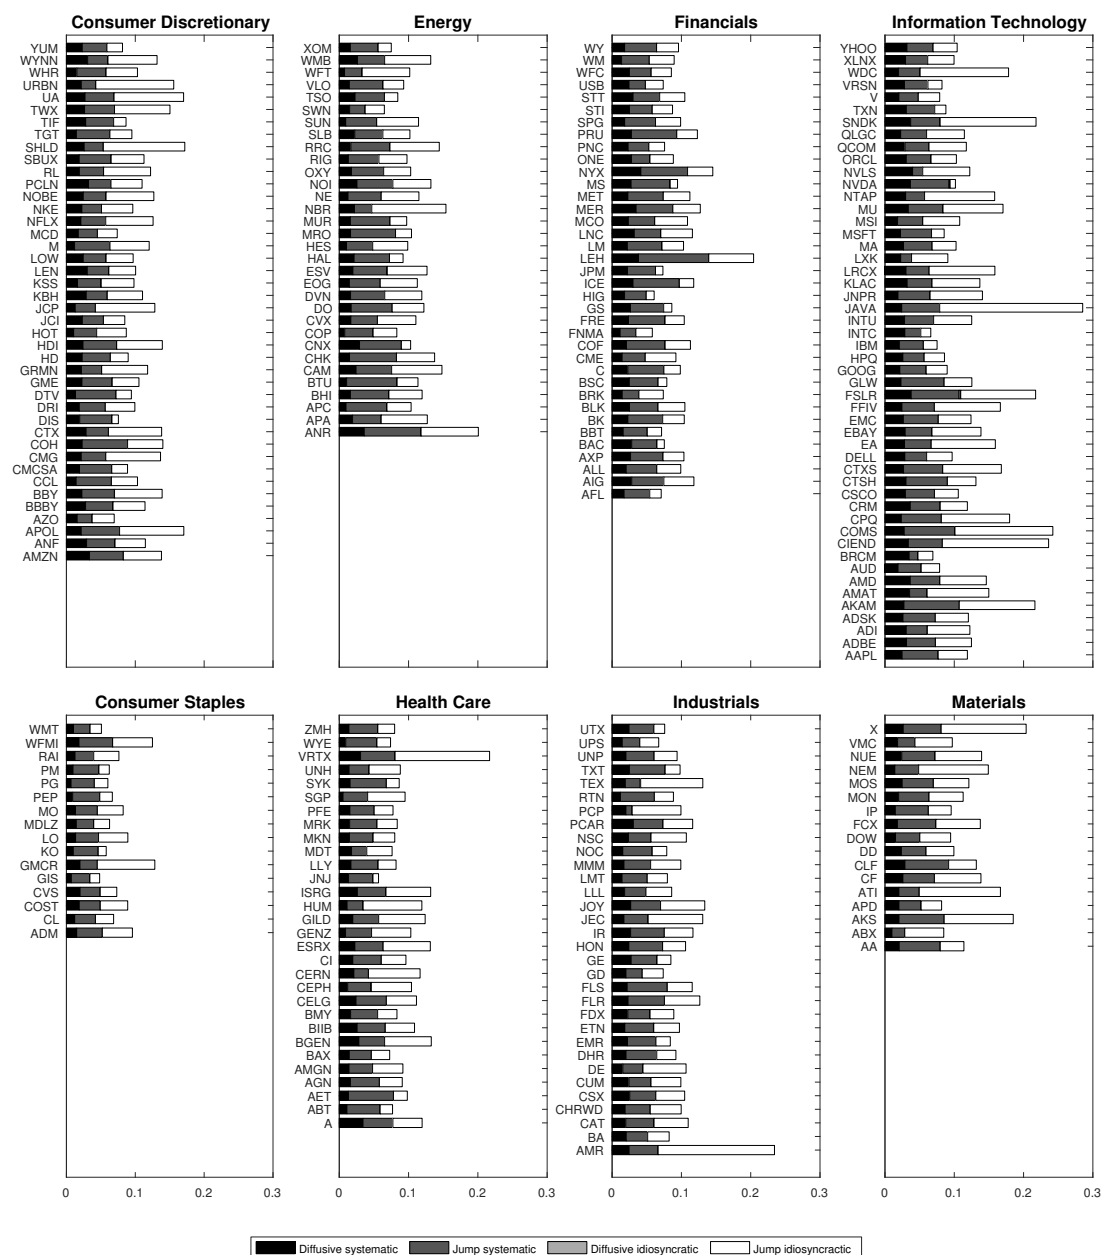


Figure 4.4: Decomposition of the equity risk premium by firm.

This figure presents, for each of the 260 stocks, the decomposition of its equity risk premium in terms of the premiums associated with the four different risk factors in the model: (i) systematic normal, (ii) systematic jump, (iii) idiosyncratic normal, and (iv) idiosyncratic jump. On average, systematic normal risk accounts for 20.3% of the total premium, systematic jump risk for 39.5%, and idiosyncratic jump risk 40.2%. Firms are grouped by industry, based on the Global Industry Classification Standard (GICS). Results for telecommunication services and utilities are not reported since they concern only five firms.

However, the most striking result illustrated in Figure 4.4 is that the premium associated with idiosyncratic risk accounts for a 40.2% fraction of the total premium ($\frac{(iii)+(iv)}{(i)+(ii)+(iii)+(iv)}$). Moreover, and perhaps most importantly, the normal component of idiosyncratic risk, which is easily diversifiable, is not priced once other sources of risk are accounted for. This is consistent with the average value of λ_S being very small at 5.079×10^{-5} ; Figure 4.4 shows that this leads to idiosyncratic normal risk premium that are economically insignificant. While it is now widely accepted that, given market incompleteness, idiosyncratic risk can be priced, we find that idiosyncratic jump risk, alone, matters in the equity risk premium. As shown in Figure 4.4, the proportion of the equity risk premium explained by the premium on the jump idiosyncratic risk factor varies significantly from firm to firm, but idiosyncratic normal risk virtually does not matter for any of the 260 firms in our sample.

4.4.2.1 Averages Across Industries

Figure 4.5 presents, for the eight largest Global Industry Classification Standard industries (GICS) covered by our sample, the evolution through time of the component of the industry's average firm's equity risk premium that is due to exposure to idiosyncratic jumps.²⁸ Note that all firms load, through their normal and jump betas, on the systematic risk premiums reported in Figure 4.3. Hence, the idiosyncratic jump risk premium (solid line) reported in Figure 4.5 adds to the premium arising from the firms' exposure to systematic risk factors (grey '+' marks).

For all industries, jump risk premiums increase around both recessions in our sample. In fact, the increase is relatively mild around the first recession for all industries, except Information Technology who had just been hit by the burst of the dot-com bubble. On the other hand, idiosyncratic jump risk premiums increase markedly for all industries around the Great Recession. Interestingly, the crisis peak in idiosyncratic jump risk premium for Financials is not as high as that experienced by Consumer Discretionary or Materials, for instance. However, as reported Table 4.7, firms from the Financial sector are, on average, the ones exhibiting the second highest normal beta and the second highest jump beta. Hence their total premium (summing the solid line with the grey '+') raises significantly during the crisis.

²⁸We do not report results for Telecommunication Services (2 firms) and Utilities (3 firms) as we do not have enough firms from these sectors.

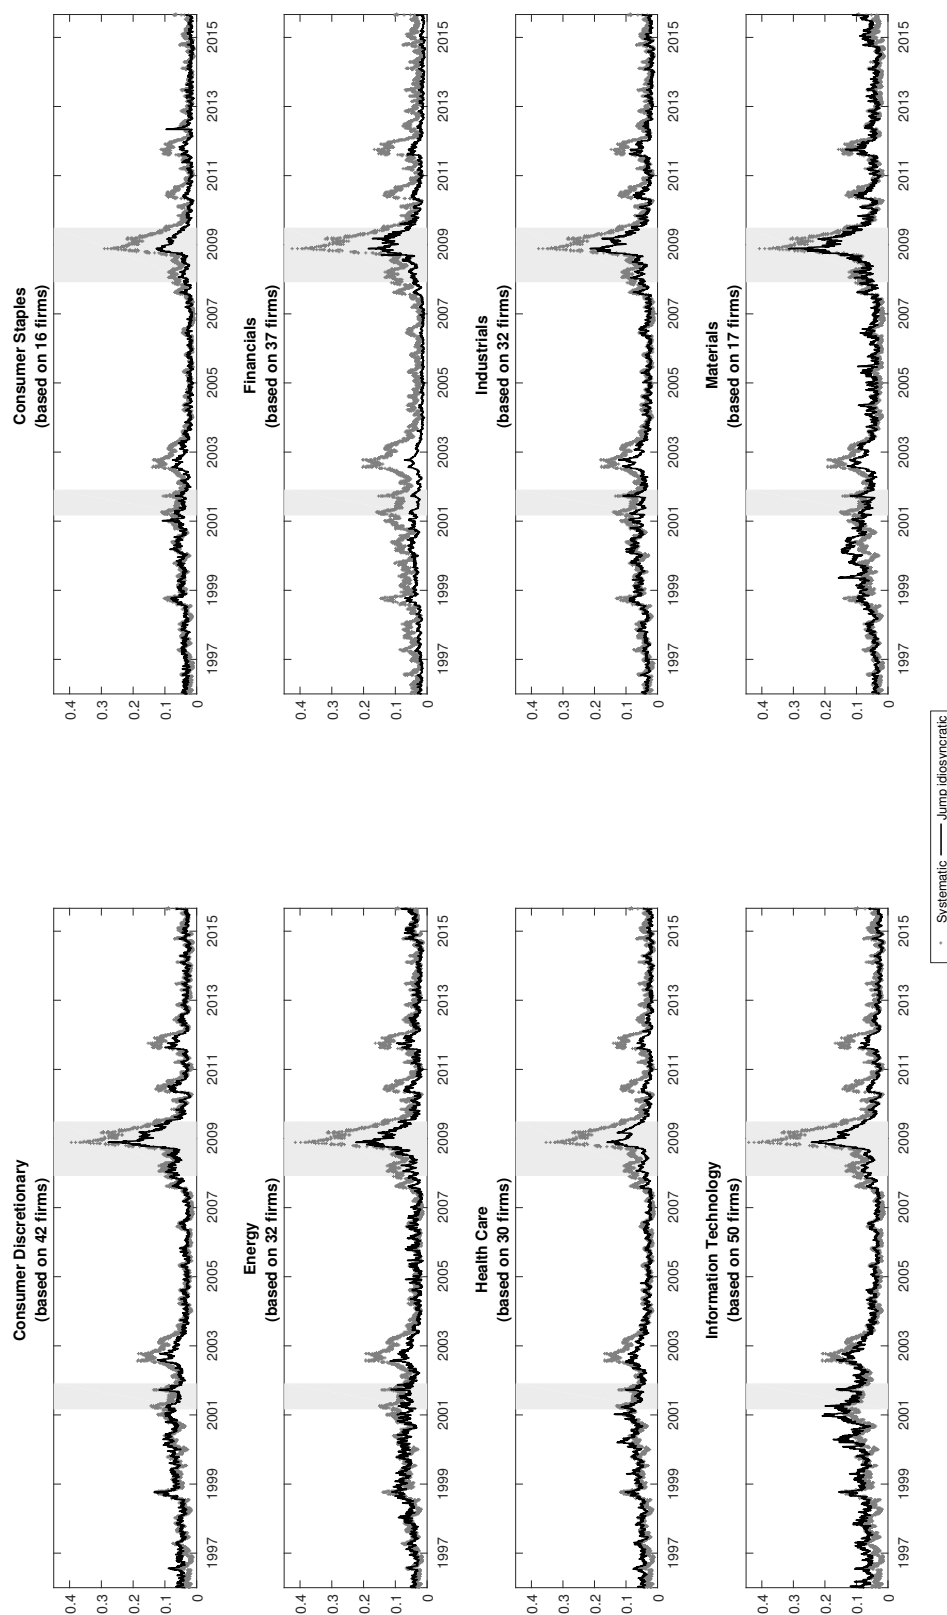


Figure 4.5: Time series decomposition of the average equity risk premium by industry.

This figure presents, for each of the eight largest GICS industries covered by our sample, the evolution through time of (i) the component of the industry's average equity risk premium that is due to exposure to idiosyncratic jumps (solid line), and (ii) the premium arising from the firms' exposure to systematic risk factors (grey '+' marks).

Table 4.7: Firm parameters estimated using returns and option data for each of the eight largest GICS industries.

	Average	S.D.	Min	Q1	Median	Q3	Max
Consumer Discretionary (based on 42 firms)							
$\beta_{S,z}$	0.937	0.241	0.477	0.762	0.963	1.092	1.421
$\beta_{S,y}$	1.025	0.321	0.500	0.779	0.983	1.259	1.945
$\kappa_{S,z}$	1.095	0.727	0.226	0.616	0.847	1.411	4.062
$\kappa_{S,y}$	0.617	0.379	0.185	0.378	0.511	0.754	2.212
Consumer Staples (based on 16 firms)							
$\beta_{S,z}$	0.567	0.181	0.289	0.435	0.541	0.711	0.851
$\beta_{S,y}$	0.832	0.215	0.582	0.647	0.780	0.946	1.338
$\kappa_{S,z}$	0.920	1.040	0.167	0.477	0.639	0.877	4.665
$\kappa_{S,y}$	0.341	0.156	0.143	0.201	0.300	0.482	0.575
Energy (based on 32 firms)							
$\beta_{S,z}$	0.735	0.263	0.326	0.576	0.695	0.930	1.472
$\beta_{S,y}$	1.331	0.442	0.536	1.069	1.248	1.619	2.503
$\kappa_{S,z}$	1.355	0.642	0.360	0.891	1.253	1.645	3.228
$\kappa_{S,y}$	0.529	0.230	0.097	0.422	0.529	0.650	1.155
Financials (based on 36 firms)							
$\beta_{S,z}$	1.015	0.266	0.445	0.849	1.007	1.176	1.756
$\beta_{S,y}$	1.206	0.578	0.495	0.898	1.179	1.378	3.926
$\kappa_{S,z}$	0.713	0.391	0.124	0.466	0.649	0.865	1.886
$\kappa_{S,y}$	0.441	0.152	0.175	0.351	0.442	0.500	1.011
Health Care (based on 30 firms)							
$\beta_{S,z}$	0.734	0.282	0.206	0.576	0.662	0.904	1.400
$\beta_{S,y}$	0.996	0.283	0.500	0.862	1.018	1.090	1.913
$\kappa_{S,z}$	0.660	0.510	0.130	0.326	0.584	0.832	2.880
$\kappa_{S,y}$	0.419	0.253	0.091	0.273	0.345	0.512	1.352
Industrials (based on 32 firms)							
$\beta_{S,z}$	0.896	0.189	0.516	0.749	0.868	1.039	1.316
$\beta_{S,y}$	0.985	0.304	0.202	0.802	0.984	1.146	1.661
$\kappa_{S,z}$	0.785	0.483	0.197	0.483	0.682	1.011	2.766
$\kappa_{S,y}$	0.472	0.273	0.140	0.329	0.419	0.542	1.444
Information Technology (based on 50 firms)							
$\beta_{S,z}$	1.184	0.218	0.765	0.999	1.199	1.326	1.568
$\beta_{S,y}$	1.083	0.423	0.275	0.814	0.987	1.282	2.416
$\kappa_{S,z}$	1.018	0.869	0.083	0.389	0.706	1.409	3.503
$\kappa_{S,y}$	0.635	0.330	0.125	0.423	0.565	0.777	1.674
Materials (based on 17 firms)							
$\beta_{S,z}$	0.864	0.208	0.430	0.732	0.864	1.035	1.236
$\beta_{S,y}$	1.173	0.440	0.435	0.856	1.159	1.545	1.943
$\kappa_{S,z}$	1.215	0.661	0.463	0.743	1.013	1.615	3.086
$\kappa_{S,y}$	0.550	0.290	0.091	0.359	0.569	0.698	1.134

The index parameters are estimated using daily index returns and weekly cross-sections of out-of-the-money options, from January 1996 to August 2015. Parameters are estimated using multiple simplex search method optimizations (fminsearch in Matlab). Robust standard errors are calculated from the outer product of the gradient at the optimal parameter values. For firms, we report statistics on the joint MLE estimates obtained for the 260 individual stocks in our sample across the eight largest GICS industries covered in our sample. Q1 and Q3 report the 25th and 75th percentiles of the estimates.

4.4.3 Commonality in Idiosyncratic Jump Risk

Following the literature, we defined idiosyncratic variance as the variance of the residuals obtained after accounting for risk factors, here normal and jump market risk. In particular, a stock's idiosyncratic variance and jump intensity are defined as (Equations (4.4) and (4.6))

$$\begin{aligned}
 h_{S,z,t+1} &= \kappa_{S,z} h_{M,z,t+1} + b_{S,z} (h_{S,z,t} - \kappa_{S,z} h_{M,z,t}) + \frac{a_{S,z}}{h_{S,z,t}} (z_{S,t}^2 - h_{S,z,t} - 2c_{S,z} h_{S,z,t} z_{S,t}), \\
 h_{S,y,t+1} &= \kappa_{S,y} h_{M,y,t+1} + b_{S,y} (h_{S,y,t} - \kappa_{S,y} h_{M,y,t}) + \frac{a_{S,y}}{h_{S,z,t}} (z_{S,t}^2 - h_{S,z,t} - 2c_{S,y} h_{S,z,t} z_{S,t}).
 \end{aligned}$$

In order to account for the documented strong commonality in idiosyncratic variances (Herskovic, Kelly, Lustig, and Van Nieuwerburgh (2014), henceforth HKLV), each process evolves around a dynamic level $\kappa_{S,v}h_{M,v,t+1}$, $v \in \{z, y\}$. Table 4.5 report that the normal kappa is on average 0.971, further supporting the commonality documented in HKLV. Our results extend those of HKLV by documenting a strong commonality in jump risk: the average jump kappa is 0.513. Hence, the commonality in jump risk is less important than that documented in variances, but is still sizable.

Firm-by-firm regressions (untabulated) of total idiosyncratic variance (cf. Equation (4.8)) on $\kappa_{S,z}h_{M,z,t}$ and $\kappa_{S,y}h_{M,y,t}$ yield an average R^2 of 73.4%; regressing on the $\kappa_{S,y}h_{M,y,t}$ alone yields an average R^2 of 31.8%. Our results thus extend the finding of Herskovic et al. (2016) in that we document that idiosyncratic tail risk explains a large fraction of the commonality in idiosyncratic variance. As such, tail risk, which is already hard to hedge by nature, becomes virtually undiversifiable in times of turmoil, which justifies the risk premium attached to it.

Interestingly, Table 4.6 reports that the three significant components of risk premiums are positively correlated with the firms' average level of total idiosyncratic variance (IVAR). The same holds for the excess idiosyncratic components of normal, $\widetilde{\text{IVAR}}_{S,z}$, and jump risk, $\widetilde{\text{IVAR}}_{S,y}$, which are respectively the average of the following time series: $h_{S,z,t} - \kappa_{S,z}h_{M,z,t}$, and $h_{S,y,t} - \kappa_{S,y}h_{M,y,t}$. This result is consistent with the results of Martin and Wagner (2016), who find that stocks exhibiting higher (lower) than average idiosyncratic volatility command higher (lower) expected excess returns.

4.4.4 On the Importance of Accounting for Equity-Specific Jumps

Financial theory tells us that diversifiable risk should not be priced. In most models, idiosyncratic risk is simply normal risk. As this normal risk should be easily diversified away, conditionally normal models imply that idiosyncratic risk should not be priced. Our results confirm that idiosyncratic normal risk indeed is not priced. Idiosyncratic jump risk, on the other hand, is difficult to diversify by nature. As such, it can bear a risk premium, and it does.

When one estimates a conditionally normal model on actual returns, the filtered “normal” innovations are all but normal. They typically have a very large kurtosis; in a misspecified normal model, the supposedly normal innovations are also capturing jumps. Given the importance of the premium on these idiosyncratic jumps in our

model, we conjecture that, in conditionally normal models, the risk premium on idiosyncratic normal risk originates from the model's misspecification.

To validate this conjecture, we estimate a nested version of our model in which the market model remains unchanged, but the stock model does not exhibit idiosyncratic jump risk. That is,

$$\begin{aligned} R_{S,t+1} &= \mu_{S,t+1} - \xi_{S,t+1}^{\mathbb{P}} + \beta_{S,z} z_{M,t+1} + \beta_{S,y} y_{M,t+1} + z_{S,t+1}, \\ \mu_{S,t+1} - r_{t+1} &= \beta_{S,z} \lambda_M h_{M,z,t+1} + \gamma_{M,S}(\beta_{S,y}) h_{M,y,t+1} + \lambda_S h_{S,z,t+1}, \\ \xi_{S,t+1}^{\mathbb{P}} &= \xi_{z_{M,t}}^{\mathbb{P}}(\beta_{S,z}) + \xi_{y_{M,t}}^{\mathbb{P}}(\beta_{S,y}) + \xi_{z_{S,t}}^{\mathbb{P}}(1), \end{aligned} \quad (4.19)$$

where market innovations have separate normal variance and jump intensity, as specified in Section 4.2, and $z_{S,t+1}$ is simply assumed to be conditionally normal, with GARCH variance as specified in Equation (4.4).

Figure 4.6 shows that the composition of the total risk premium is drastically different once idiosyncratic jumps are neglected. The systematic components are very similar to those reported in Figure 4.4. The premium associated with idiosyncratic normal risk, however, is now more important than the sum of premiums associated with systematic risk. The expected excess return on an average stock (not tabulated) rises significantly, from 10.3% (5.9% systematic and 4.4% idiosyncratic) in the full model to 13.3% (6.8% and 6.5%, respectively) in the nested model of Equation (4.19). Both Christoffersen et al. (2012) and Ornathanalai (2014) document, at the market level, that the equity risk premium levels implied by conditionally normal model are unreasonably high.²⁹ The difference is not as marked at the stock level, perhaps due to the presence of systematic jump risk. Yet, it appears that ignoring idiosyncratic jump risk also leads to a severe misspecification at the stock level.

Panel A of Table 4.8 provides further evidence of this misspecification. In particular, the filtered $\varepsilon_{S,t}$, which are supposed to be conditionally standard normal innovations under the model of Equation (4.19), exhibit levels of excess kurtosis that are much too high. While the theoretical level should be 0, the median level reported in Panel A of Table 4.8 is 9.14. In comparison, the corresponding median is 0.74 in Table 4.5. While the filtered “normal” innovations exhibit skewness and kurtosis, they theoretically don't, which reduces the ability of the stock model to properly fit the stocks' implied volatility smile through time. Indeed, the entire RIVRMSE distribution reported in Panel A of Table 4.8 is shifted to the right when compared with the full

²⁹Christoffersen et al. (2012) find 22.15% (Table 6) and Ornathanalai (2014) obtains 15.50% (Table 3).

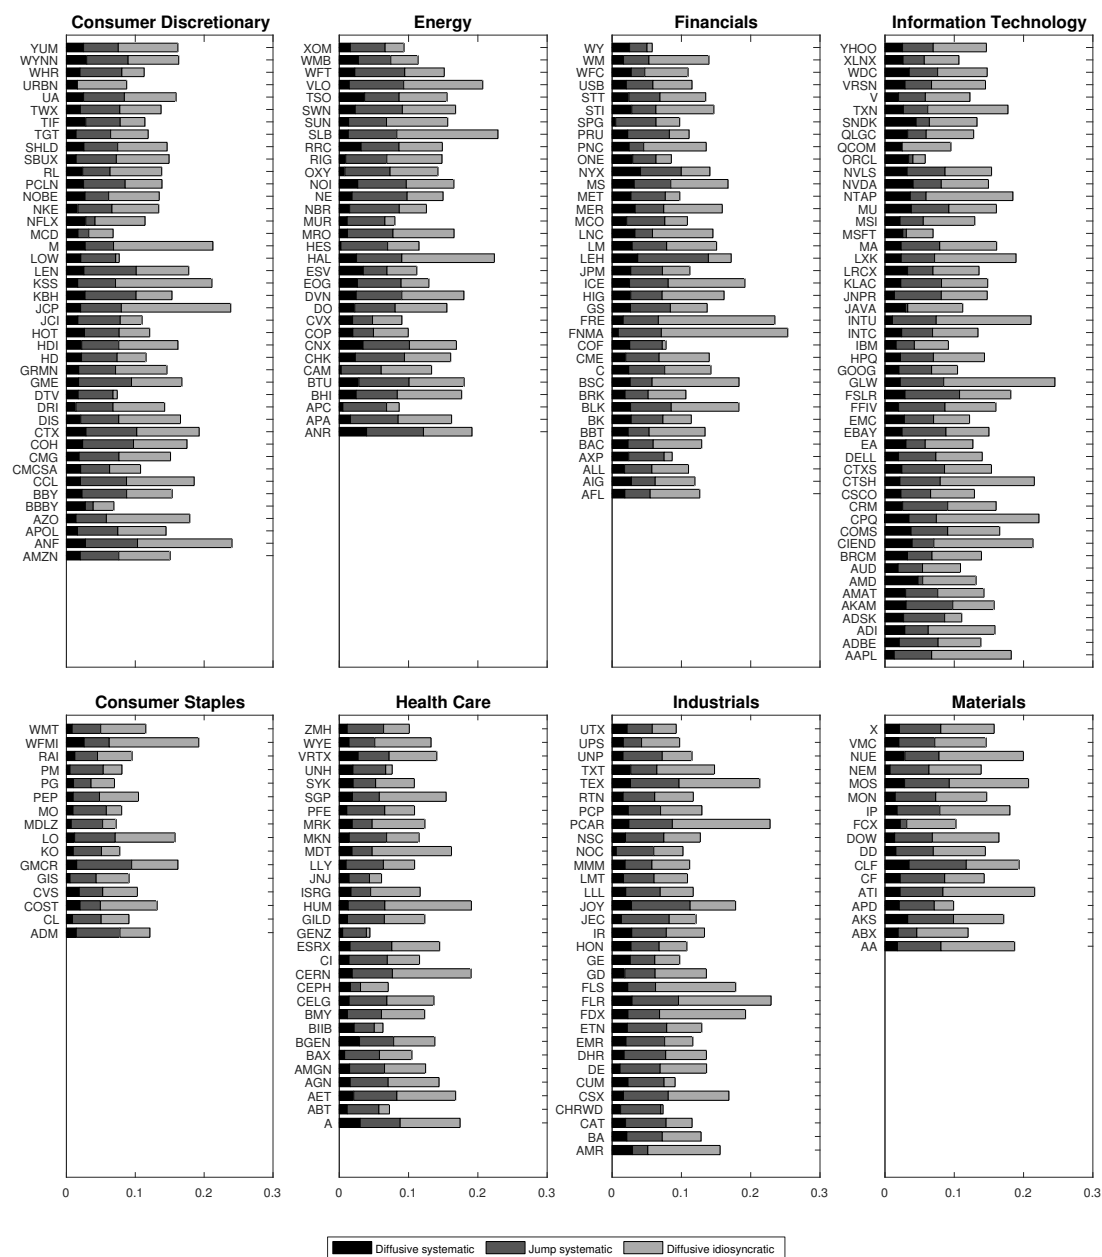


Figure 4.6: Decomposition of the equity risk premium by firm.

This figure presents, for each of the 260 stocks, the decomposition of its equity risk premium in terms of the premiums associated with the three different risk factors in the nested model of Equation (4.19): (i) systematic normal, (ii) systematic jump, and (iii) idiosyncratic normal. On average, systematic normal risk accounts for 15.0% of the total premium, systematic jump risk for 36.1%, and idiosyncratic normal risk 48.9%. Firms are grouped by industry, based on the Global Industry Classification Standard (GICS). Results for telecommunication services and utilities are not reported since they concern only five firms.

model's RIVRMSE in Table 4.5. That is, although the model of Equation (4.19) still allows for relatively high levels of (negative) skewness and kurtosis, thanks to systematic jumps, the normality assumption at the idiosyncratic level clearly deteriorates the fit to options.

In sum, the contrast between the results obtained when considering or neglecting idiosyncratic jumps highlights the importance of accounting for these equity-specific jumps.

4.4.5 Realized Premium based Portfolios of Stocks

In a typical study of factor models, the market prices associated with the different risk factors are estimated from the panel of returns. Here, they are identified from returns and option prices. A potential issue, if stock and option markets are partly segmented, is that the estimated risk premiums might reflect, for instance, option market makers' shadow price of equity. While this concern is partly mitigated by the use of stock returns in the joint estimation, it still might affect the estimated magnitude of the premium. This criticism applies to any study using options prices to learn about the equity risk premium or even the physical distribution.³⁰

To appraise whether this criticism indeed points to a weakness of our framework, we here analyze portfolios formed according to the model-implied risk premium associated with idiosyncratic jump risk. First, on each day t out of the 4,951 days in our sample, we sort available stocks according to the expected excess return associated to idiosyncratic jumps:

$$\text{RP}_{S,y,t} = \mathbb{E}_t^{\mathbb{P}} \left[\exp \left(\gamma_S h_{S,y,t+1} - \xi_{yS,t+1}^{\mathbb{P}} + y_{S,t+1} \right) \right] = e^{\gamma_S h_{S,y,t+1}}. \quad (4.20)$$

We then divide these stocks in five quintile portfolios P1 to P5, from lowest to highest expected return; stocks within a portfolio are weighted according to their market capitalization on day t . Finally, we create a long-short portfolio with a long position in the portfolio with the highest expected return, P5, and a short position in that with the lowest expected return, P1. For comparison, the same procedure is used to create a long-short portfolio on the basis of $\text{RP}_{S,z,t} = e^{\lambda_S h_{S,z,t+1}}$.

³⁰For instance, Martin (2016) and Martin and Wagner (2016) infer the equity risk premium directly from option prices. Ross (2015) and Jensen et al. (2016) infer the \mathbb{P} distribution from the evolution of the \mathbb{Q} distribution.

Table 4.8: **Firm parameters estimated using returns and option data: Neglecting idiosyncratic jumps.**

	Average	S.D.	Min	Q1	Median	Q3	Max
$\beta_{S,z}$	0.908	0.335	0.101	0.684	0.908	1.141	2.051
$\beta_{S,y}$	1.365	0.550	0.000	0.985	1.398	1.698	3.961
λ_S	0.765	0.440	0.030	0.397	0.727	1.016	2.242
$\kappa_{S,z}$	0.874	0.874	0.045	0.309	0.649	1.200	5.690
$a_{S,z}$	2.66E-06	2.25E-06	9.29E-08	1.12E-06	2.12E-06	3.63E-06	2.30E-05
$b_{S,z}$	0.995	0.014	0.906	0.997	0.998	0.999	1.000
$c_{S,z}$	53.30	84.76	-282.82	7.99	52.37	95.89	437.76
Avg. volatility (%)	36.56	9.50	18.33	29.41	35.10	43.24	64.32
Avg. skewness	-2.92	2.33	-13.32	-4.26	-2.45	-1.03	0.00
Avg. excess kurtosis	149.03	150.39	0.00	35.04	107.46	216.51	1014.66
Skewness of innovations, $\varepsilon_{S,t}$	-0.37	1.20	-9.34	-0.50	-0.08	0.12	2.97
Ex. kurtosis of innovations, $\varepsilon_{S,t}$	16.78	31.96	0.69	5.00	9.14	15.34	349.03
RIVRMSE	15.57	6.42	9.94	12.24	14.19	17.01	80.62

The index parameters are estimated using daily index returns and weekly cross-sections of out-of-the-money options, from January 1996 to August 2015. Parameters are estimated using multiple simplex search method optimizations (fminsearch in Matlab). Robust standard errors are calculated from the outer product of the gradient at the optimal parameter values. For firms, we report statistics on the joint MLE estimates obtained for the 260 individual stocks in our sample. Q1 and Q3 report the 25th and 75th percentiles of the estimates.

Table 4.9 reports regressions of the returns of these long-short portfolios on some of the most prevalent factors in empirical asset pricing. First, the regression labelled FF3 is based on the Fama and French (1993) 3-factor models: market (MKT), small minus big (SMB), high minus low (HML). The regression labelled FF5 extends the set of regressors to those of weak (RMW), and conservative minus aggressive (CMA) factors. The regression labelled CF4 considers the Carhart (1997) 4-factor model, essentially adding a momentum (MOM) factor to FF3. The regression labelled AHXZ is inspired by the work of Ang et al. (2006) and extends FF3 by adding the innovation on the CBOE Volatility Index (VIX) to the set of regressors.³¹ Finally, we consider a kitchen sink regression, labelled All, in which we control for all the aforementioned factors, that is:

$$r_t^{P5-P1} = \alpha + \beta_{\text{MKT}} \text{MKT}_t + \beta_{\text{SMB}} \text{SMB}_t + \beta_{\text{HML}} \text{HML}_t + \beta_{\text{RMW}} \text{RMW}_t + \beta_{\text{CMA}} \text{CMA}_t + \beta_{\text{MOM}} \text{MOM}_t + \beta_{\Delta \text{VIX}} \Delta \text{VIX}_t + \epsilon_t, \quad (4.21)$$

where r_t^{P5-P1} is the day- t simple excess return on the long-short portfolio formed on the basis of the model-implied risk premiums, $\text{RP}_{S,\cdot,t}$, and $\Delta \text{VIX}_t = \text{VIX}_t - \text{VIX}_{t-1}$ is the innovation on the VIX. The alphas of the regressions are reported in annualized percentage terms.

Table 4.9 is divided in two panels. In the first five columns of Panel A, the long-short portfolio is formed according to the quintiles of the risk premium associated with idiosyncratic Gaussian risk in our model, $\text{RP}_{S,z,t}$. Regression coefficients are in bold whenever they are significant at the 5% level or in italics if they are at the 10% level; t -statistics are based on robust Newey and West (1987) standard errors. In particular, only two of the alphas in these five regressions are significant at the 10% level. Panel A contains a second set of regressions. In these regressions, the long-short portfolio is constructed considering the risk-premium in excess of the common component. That is, to obtain the long-short portfolio used in the last five columns of Panel A, we sort stocks into quintiles of $\text{RP}'_{S,z,t} = e^{\lambda_S h'_{S,z,t+1}}$, where $h'_{S,z,t+1}$ is the predicted excess idiosyncratic Gaussian variance introduced in Subsection 4.2.1. None of the alphas in these five regressions are significant. Consistent with our results in Subsection 4.4.2, idiosyncratic Gaussian risk does not appear to be priced.

In Panel B, we repeat the same analysis, but forming the long-short portfolio on the basis of the risk premiums associated with (excess) idiosyncratic jump risk, $\text{RP}_{S,y,t}$

³¹This regression is akin to their ex post Regression (6), with the difference that we use ΔVIX rather than a factor mimicking aggregate volatility risk.

($RP'_{S,y,t}$). The alphas of the first five regressions are highly significant, both statistically and economically. They vary between 8.3% to 17.8% annually, depending on the regression considered. It is worth noting that, while the metric used to sort stocks into portfolios is inferred in part from option prices, the alphas reported here are obtained trading only stocks. These results thus confirm that idiosyncratic jump risk is priced in stock markets.

Besides, the alphas on the $RP'_{S,y,t}$ long-short portfolios are just as significant, statistically and economically, as the alphas on the portfolio based on the total idiosyncratic jump risk premium. These results show that idiosyncratic jump risk is not just priced through its common component; excess idiosyncratic jump risk matters. This result is consistent with the intuition one could build from the theoretical analysis in Martin and Wagner (2016). Unfortunately, to the best of our knowledge, it is not possible to assign different prices of risk to the common and excess components of idiosyncratic jump risk and remain in an affine option-pricing framework.

4.4.5.1 Characteristics of the Quintile Portfolios and Double-Sort Portfolios

In our model, the only systematic risk factors are the Gaussian and jump innovations on the market. Accounting for more factors, like the Fama and French (2015) or Carhart (1997) factors in the option-pricing model would have required postulating a dynamics for each of the factors, introducing many more parameters. It further would have forced estimation only from returns, at least for these factors, as options are not traded on these factors.

Yet, in Table 4.9, the loadings on these factors are almost all significant. One can thus conjecture that neglecting these factors in the model led to idiosyncratic risk proxies that are partly driven by these factors. Consistent with this intuition, Table 4.10 shows that the quintile portfolios obtained based on the idiosyncratic jump risk premium display near monotonic market betas, market capitalizations (ME), book-to-market ratios (BE/ME), operating profitability (OP) investment levels, trailing 12-month returns, and the volatility betas.³² See Section 4.G for more details on these variables.

³² The market and volatility betas here are obtained by performing the following regression

$$r_{S,t-k} = \alpha + \beta_{\text{MKT},t} \text{MKT}_{t-k} + \beta_{\Delta\text{VIX},t} \Delta\text{VIX}_{t-k} + \epsilon_{t-k}, \quad k = 0, \dots, 252.$$

This regression corresponds to the pre-formation regression of Ang et al. (2006), over the past year (rather than month, as in AHXZ) of data.

The alphas in Panel B of Table 4.9 are significant even after linearly controlling for the factors corresponding to these variables. The patterns observed in Table 4.10 could nonetheless raise concerns that the alphas are somehow nonlinearly related to the fundamentals of the stocks in the top and bottom quintile portfolios used in the long-short strategy. To alleviate these concerns, we perform a double sort. On each day t , we first sort stocks into quintiles (Q1 to Q5) based on their market beta over the past year. Then, within each market-beta quintile, we sort stocks into terciles based on $RP_{S,y,t}$ (first five columns; the last five columns are based on $RP'_{S,y,t}$). We then take a long position in a cap-weighted portfolio of the top-tercile stocks, and a short position in a cap-weighted portfolio of the bottom-tercile stocks. This leaves us with five long-short portfolios, each of which is composed of stocks with homogenous market betas. The first row of Table 4.11 reports the alphas of Regression (4.21) for each of these long-short portfolios. The procedure is repeated for the other six variables in Table 4.10.

The alphas on the long-short portfolios built from (excess) idiosyncratic jump risk premiums are positive and significant, both statistically and economically, in 30 (31) of the 35 regressions. When they are not statistically significant, they are still positive; the lack of significance is mainly due to the large standard errors on some of these portfolios. There are no clear patterns in the alphas across the quintiles of most variables, the exception potentially being volatility betas. In sum, no single one of these seven variables appears to be driving the main result of our paper: idiosyncratic jump risk carries a positive risk premium.

4.5 Conclusion

In this study, we shed new light on the relationship between idiosyncratic risk and equity returns. We develop a model allowing us to disentangle the contribution of four different risk factors to the equity risk premium: systematic and idiosyncratic risk are both decomposed in their normal and jump components. Using 20 years of returns and options on the S&P 500 and more than 250 stocks, we find that normal and jump risk have a drastically different impact on the expected return on individual stocks.

While our pricing kernel is such that each risk factor can potentially be priced, we find that the normal component of idiosyncratic risk, which is easily diversifiable, is not priced once other sources of risk are accounted for. Firm-specific jump risk, however, is priced and justifies more than 40% of the expected excess return on an average

Table 4.9: **Excess returns of portfolios based on idiosyncratic jump risk premiums.**

Panel A: Idiosyncratic jump risk premium										
	Idiosyncratic diffusive risk premium					In excess of the common component				
	FF3	FF5	MOM	AHXZ	All	FF3	FF5	MOM	AHXZ	All
Cst	1.748 (0.71)	4.220 (1.74)	2.599 (1.05)	1.534 (0.63)	4.449 (1.83)	1.287 (0.53)	3.559 (1.46)	2.082 (0.84)	1.011 (0.41)	3.657 (1.50)
MKT	0.119 (6.80)	0.049 (2.50)	0.097 (5.57)	0.144 (6.28)	0.048 (2.05)	0.089 (5.10)	0.024 (1.22)	0.069 (3.89)	0.122 (5.24)	0.034 (1.39)
SMB	0.095 (3.72)	0.053 (2.25)	0.103 (3.88)	0.093 (3.63)	0.058 (2.46)	0.085 (3.25)	0.051 (2.02)	0.092 (3.44)	0.082 (3.13)	0.055 (2.18)
HML	-0.542 (-18.22)	-0.374 (-12.03)	-0.587 (-18.76)	-0.538 (-18.32)	-0.405 (-11.59)	-0.561 (-18.93)	-0.396 (-11.37)	-0.603 (-18.97)	-0.556 (-19.07)	-0.422 (-11.14)
RMW		-0.202 (-4.97)			-0.192 (-4.67)		-0.167 (-4.14)			-0.157 (-3.88)
CMA		-0.424 (-7.46)			-0.399 (-6.95)		-0.426 (-6.99)			-0.404 (-6.57)
MOM			-0.094 (-4.29)		-0.045 (-2.31)			-0.087 (-3.96)		-0.039 (-1.93)
ΔVIX				0.023 (1.76)	0.005 (0.34)				0.030 (2.22)	0.013 (0.95)
Adj. R^2	22.0%	26.2%	22.9%	22.0%	26.4%	21.1%	25.0%	21.9%	21.2%	25.1%
Panel B: Idiosyncratic jump risk premium in excess of the common component										
	Idiosyncratic jump risk premium					In excess of the common component				
	FF3	FF5	MOM	AHXZ	All	FF3	FF5	MOM	AHXZ	All
Cst	8.412 (2.28)	16.317 (4.54)	11.070 (3.06)	8.294 (2.25)	17.840 (4.97)	9.513 (2.50)	16.201 (4.35)	12.148 (3.31)	9.488 (2.52)	17.807 (4.91)
MKT	0.473 (15.59)	0.259 (11.92)	0.406 (15.76)	0.487 (10.79)	0.179 (5.96)	0.487 (14.76)	0.299 (12.29)	0.421 (14.80)	0.490 (9.72)	0.214 (6.56)
SMB	0.401 (10.16)	0.231 (6.09)	0.424 (10.84)	0.399 (10.06)	0.253 (6.95)	0.285 (6.76)	0.173 (4.38)	0.309 (7.11)	0.285 (6.75)	0.196 (5.05)
HML	-0.428 (-5.25)	0.044 (1.00)	-0.568 (-7.04)	-0.426 (-5.28)	-0.073 (-1.56)	-0.115 (-1.34)	0.346 (5.87)	-0.254 (-3.38)	-0.115 (-1.34)	0.223 (4.06)
RMW		-0.756 (-12.90)			-0.731 (-13.87)		-0.536 (-9.41)			-0.510 (-8.71)
CMA		-1.142 (-11.39)			-1.057 (-11.96)		-1.170 (-12.70)			-1.082 (-12.14)
MOM			-0.292 (-5.52)		-0.161 (-4.01)			-0.290 (-5.92)		-0.168 (-4.17)
ΔVIX				0.013 (0.52)	-0.049 (-2.41)				0.003 (0.10)	-0.053 (-2.21)
Adj. R^2	29.1%	44.2%	32.6%	29.1%	45.2%	24.1%	37.4%	27.7%	24.0%	38.6%

Each day, the model-implied risk premium associated with idiosyncratic jump risk, $RP_{S,y}$, is computed for each firm. Firms are then sorted into quintile portfolios from the lowest (1) to the highest (5) level of $RP_{S,y}$; portfolios are weighted according to market capitalization. A “top minus bottom” portfolio is created from taking long position in (5) and a short position in (1). The daily returns of the quintile portfolios and of the long-short portfolio are then regressed on (subsets of) the following seven variables: the Fama-French market (MKT), small minus big (SMB), high minus low (HML), robust minus weak (RMW), and conservative minus aggressive (CMA) factors, the momentum (MOM) factor, and returns on the CBOE volatility index (ΔVIX). The regression constant (Cst) is reported in annualized percentage points ($\Delta t = 1/252$).

Table 4.10: **Description of quintile portfolios.**

	Idiosyncratic Jump Risk Premium					In Excess of the Common Component				
	P1	P2	P3	P4	P5	P1	P2	P3	P4	P5
Market beta	0.97 [0.11]	1.00 [0.09]	1.06 [0.12]	1.11 [0.15]	1.26 [0.25]	0.98 [0.13]	0.99 [0.12]	1.05 [0.11]	1.09 [0.14]	1.21 [0.23]
log(ME)	25.40 [0.28]	24.96 [0.38]	24.61 [0.46]	24.30 [0.52]	23.79 [0.54]	25.14 [0.38]	25.17 [0.40]	24.86 [0.42]	24.49 [0.45]	23.99 [0.51]
BE/ME	1.76 [0.53]	1.42 [0.56]	1.47 [0.75]	1.24 [0.83]	1.20 [1.08]	1.43 [0.58]	1.54 [0.55]	1.53 [0.93]	1.55 [1.33]	1.56 [1.90]
OP (%)	5.74 [1.30]	5.46 [1.35]	5.93 [1.36]	6.10 [1.60]	5.42 [1.85]	5.80 [1.54]	5.79 [1.40]	5.58 [1.45]	5.74 [1.50]	5.70 [1.71]
Investment (%)	14.63 [18.44]	15.86 [19.06]	19.31 [19.78]	26.96 [45.46]	35.51 [77.07]	20.46 [33.20]	16.30 [15.94]	15.37 [16.53]	18.98 [23.66]	28.43 [49.08]
12-month return (%)	15.25 [15.54]	15.72 [16.84]	17.84 [18.21]	19.27 [19.74]	20.11 [29.61]	19.11 [16.43]	16.57 [15.81]	15.12 [17.32]	14.79 [19.20]	16.73 [27.06]
Volatility beta (%)	0.29 [3.07]	0.50 [3.44]	0.53 [4.30]	-1.11 [6.29]	-2.65 [9.71]	-0.32 [3.97]	0.46 [3.44]	0.83 [4.16]	-0.06 [5.44]	-2.37 [9.13]

This table describes the quintile portfolios obtained in Table 4.9. Each day, the following variables are recorded for the firms in each quintile portfolio: market beta, log of market capitalisation, book-to-market ratio, operating profitability (OP), investment, trailing 12-month return, and the volatility beta. This table reports the time-series average of these variables for each of the quintile portfolios obtained using the total idiosyncratic risk premium (first five columns) or its component in excess of the common component (last five columns). Standard deviations are reported within square brackets.

Table 4.11: **Double sort: Excess returns of portfolios based on idiosyncratic jump risk premiums.**

	Idiosyncratic Jump Risk Premium					In Excess of the Common Component				
	Q1	Q2	Q3	Q4	Q5	Q1	Q2	Q3	Q4	Q5
Market beta	10.408 (2.71)	8.957 (2.44)	16.837 (4.43)	15.242 (3.62)	10.024 (1.43)	11.783 (3.23)	11.368 (3.04)	20.495 (5.35)	15.627 (3.60)	10.759 (1.89)
log(ME)	11.629 (2.46)	7.435 (2.05)	8.791 (2.73)	17.906 (5.77)	10.027 (3.80)	11.283 (2.30)	9.990 (2.57)	9.429 (2.84)	18.620 (5.99)	8.780 (3.02)
BE/ME	20.747 (4.63)	13.477 (2.89)	9.252 (2.21)	14.811 (3.50)	14.381 (3.47)	15.653 (2.66)	18.488 (3.92)	12.847 (3.12)	17.991 (4.33)	15.290 (3.72)
OP (%)	9.027 (1.61)	10.796 (2.26)	11.430 (2.46)	7.035 (1.47)	24.957 (5.02)	18.828 (3.46)	12.194 (2.71)	14.171 (3.03)	10.364 (2.15)	13.506 (2.70)
Investment (%)	13.785 (2.97)	12.780 (3.39)	15.409 (3.98)	23.275 (5.13)	11.985 (2.36)	16.072 (3.42)	11.650 (3.07)	22.329 (5.86)	16.914 (3.64)	5.305 (1.01)
12-month return (%)	17.567 (3.42)	5.529 (1.46)	12.164 (3.36)	14.953 (3.66)	12.720 (2.57)	14.197 (2.85)	13.406 (3.46)	13.994 (3.89)	13.975 (3.40)	4.232 (0.88)
Volatility beta (%)	3.451 (0.68)	12.885 (2.97)	14.973 (3.95)	14.183 (3.67)	15.743 (3.18)	7.568 (1.49)	14.425 (3.27)	12.403 (3.39)	16.668 (4.35)	16.876 (3.21)

On each day t , we first sort stocks into quintiles (Q1 to Q5) based on their market beta over the past year. Then, within each market-beta quintile, we sort stocks into terciles based on $RP_{S,y}$ (first five columns; the last five columns are based on $RP'_{S,y}$). We then take a long position in a cap-weighted portfolio of the top-tercile stocks, and a short position in a cap-weighted portfolio of the bottom-tercile stocks. This leaves us with five long-short portfolios, each of which is composed of stocks with homogenous market betas. The first row of this table reports the alphas of Regression (4.21) for each of these long-short portfolios. The procedure is repeated for six other variables: log of market capitalisation, book-to-market ratio, operating profitability (OP), investment, trailing 12-month return, and the volatility beta.

stock. Given the recent conflicting empirical evidence regarding how idiosyncratic risk affects expected returns, these findings might provide new guidance for future studies.

Our focus in this paper is on the relationship between jump risk and the equity risk premium. Given the strong links between the equity risk premium and the variance risk premium, it is natural to wonder whether our findings extend to the variance risk premium; the results of Gouriéroux (2014) certainly suggest they do. Hence, it appears that properly accounting for jump risk is crucial in any attempts to study the risk premiums associated with idiosyncratic risk.

4.A Innovations' Cumulant Generating Functions

4.A.1 Continuous Component

For any $z_{u,t} \in \{z_{M,t}, z_{S,t} : S \in \mathbb{S}\}$, the conditional cumulant generating function of $z_{u,t}$ satisfies

$$\xi_{z_{u,t}}^{\mathbb{P}}(\phi) = \log \mathbb{E}_{\mathcal{F}_{t-1}^{\mathbb{S}}}^{\mathbb{P}} [\exp(\phi z_{u,t})] = \frac{\phi^2}{2} h_{u,z,t}.$$

4.A.2 Jump Component

The conditional cumulant generating function of $y_{u,t} \in \{y_{M,t}, y_{S,t} : S \in \mathbb{S}\}$ is

$$\xi_{y_{u,t}}^{\mathbb{P}}(\phi) = \log \mathbb{E}_{\mathcal{F}_{t-1}^{\mathbb{S}}}^{\mathbb{P}} [\exp(\phi y_{u,t})] = \Pi_u(\phi) h_{u,y,t}$$

where

$$\Pi_u(\phi) = \left(\sqrt{\alpha_u^2 - \delta_u^2} - \sqrt{\alpha_u^2 - (\delta_u + \phi)^2} \right). \quad (4.22)$$

4.B Innovations' Risk Neutral Cumulant Generating Functions

Lemma 4.1. *For any $\varepsilon_{u,t} \in \{\varepsilon_{M,t}, \varepsilon_{S,t} : S \in \mathbb{S}\}$, the conditional cumulant generating function of $\varepsilon_{u,t}$ under \mathbb{Q} is*

$$\xi_{\varepsilon_{u,t}}^{\mathbb{Q}}(\phi) = \log \mathbb{E}_{t-1}^{\mathbb{Q}} [\exp(\phi \varepsilon_{u,t})] = \frac{1}{2} \phi^2 - \Lambda_u \sqrt{h_{u,z,t}} \phi$$

which corresponds to the cumulant generating function of a Gaussian random variable of expectation $-\Lambda_u \sqrt{h_{u,z,t}}$ and variance 1. To obtain a risk neutral sequence of standard normal innovations, we must set

$$\varepsilon_{u,t}^* = \varepsilon_{u,t} + \Lambda_u \sqrt{h_{u,z,t}}. \quad (4.23)$$

Sketch of the proof.

$$\begin{aligned} \xi_{\varepsilon_{u,t}}^{\mathbb{Q}}(\phi) &= \log \mathbb{E}_{t-1}^{\mathbb{Q}} [\exp(\phi \varepsilon_{u,t})] \\ &= \log \mathbb{E}_{t-1}^{\mathbb{P}} \left[\frac{\exp(-\Lambda_M z_{M,t} - \Gamma_M y_{M,t} - \sum_{S \in \mathbb{S}} \Lambda_S z_{S,t} - \sum_{S \in \mathbb{S}} \Gamma_S y_{S,t})}{\mathbb{E}_{t-1}^{\mathbb{P}} [\exp(-\Lambda_M z_{M,t} - \Gamma_M y_{M,t} - \sum_{S \in \mathbb{S}} \Lambda_S z_{S,t} - \sum_{S \in \mathbb{S}} \Gamma_S y_{S,t})]} \exp(\phi \varepsilon_{u,t}) \right] \\ &= \xi_{\varepsilon_{u,t}}^{\mathbb{P}}(\phi - \Lambda_u \sqrt{h_{u,z,t}}) - \xi_{\varepsilon_{u,t}}^{\mathbb{P}}(-\Lambda_u \sqrt{h_{u,z,t}}). \end{aligned}$$

Note that, given that $h_{u,z,t}$ is $\mathcal{F}_{t-1}^{\mathbb{S}}$ -measurable, the conditional cumulant generating function of $\varepsilon_{u,t}$ under \mathbb{Q} can easily be shown (by replicating the above) to be that of a normal variable of expectation $-\Lambda_u \sqrt{h_{u,z,t}}$ and variance $h_{u,z,t}$. In other words, consistent with the results in Christoffersen et al. (2010), the risk-neutral $\mathcal{F}_{t-1}^{\mathbb{S}}$ -conditional variance of $\varepsilon_{u,t}$, $h_{u,z,t}^*$, is equal to its physical counterpart, $h_{u,z,t}$.

Lemma 4.2. *For any $y_{u,t} \in \{y_{M,t}, y_{S,t} : S \in \mathbb{S}\}$, the conditional cumulant generating function of $y_{u,t}$ under \mathbb{Q} is*

$$\xi_{y_{u,t}}^{\mathbb{Q}}(\phi) = \log \mathbb{E}_{t-1}^{\mathbb{Q}} [\exp(\phi y_{u,t})] = \Pi_u^*(\phi) h_{u,y,t}^* \quad (4.24)$$

where

$$\Pi_u^*(\phi) = \sqrt{\alpha_u^2 - (\delta_u - \Gamma_u)^2} - \sqrt{\alpha_u^2 - (\delta_u - \Gamma_u + \phi)^2} \quad (4.25)$$

The proof uses a similar argument as for $\xi_{\varepsilon_{u,t}}^{\mathbb{Q}}(\phi)$. Details are provided in the Appendix A.2.

The risk neutral jump component is still a NIG random variable with no location parameter, the tail heaviness parameter $\alpha_u^* = \alpha_u$ is not affected by the change of measure, the asymmetry parameter becomes $\delta_u^* = \delta_u - \Gamma_u$ and the scale variable is $h_{u,y,t}^* = h_{u,y,t}$.

4.C Risk Premiums

Lemma 4.3. *The mappings between λ_M and γ_M and their pricing kernel counterparts Λ_M and Γ_M are*

$$\lambda_M = \Lambda_M \text{ and } \gamma_M = \Pi_u(1) - \Pi_u^*(1).$$

For the stock parameters λ_S and γ_S , the relation is

$$\lambda_S = \Lambda_S, \gamma_{M,S}(\beta_{S,y}) = \Pi_M(\beta_{S,y}) - \Pi_M^*(\beta_{S,y}), \gamma_S = \Pi_S(1) - \Pi_S^*(1)$$

where $\Pi_u(\cdot)$ and $\Pi_u^*(\cdot)$ are defined at Equations (4.22) and (4.25).

Proof of Lemma 4.3. Since the proof for the market component is similar, the focus is put on the stock specific parameters. More details are available in Appendix A.

Since the discounted stock price should behave as a \mathbb{Q} -martingale,

$$\begin{aligned} 1 &= \mathbb{E}_{t-1}^{\mathbb{Q}} \left[\frac{\exp(-r_t) S_t}{S_{t-1}} \right] \\ &= \mathbb{E}_{t-1}^{\mathbb{P}} \left[\frac{\exp(-\Lambda_M z_{M,t} - \Gamma_M y_{M,t} - \sum_{S \in \mathbb{S}} \Lambda_S z_{S,t} - \sum_{S \in \mathbb{S}} \Gamma_S y_{S,t})}{\mathbb{E}_{t-1}^{\mathbb{P}} [\exp(-\Lambda_M z_{M,t} - \Gamma_M y_{M,t} - \sum_{S \in \mathbb{S}} \Lambda_S z_{S,t} - \sum_{S \in \mathbb{S}} \Gamma_S y_{S,t})]} \exp(R_{S,t} - r_t) \right]. \end{aligned}$$

Replacing the excess return using (4.2) and the cumulant generating functions, we get

$$1 = \exp \left(\begin{aligned} &\mu_{S,t}^{\mathbb{P}} - r_t - \xi_{z_{M,t}}^{\mathbb{P}}(\beta_{S,z}) - \xi_{y_{M,t}}^{\mathbb{P}}(\beta_{S,y}) - \xi_{z_{S,t}}^{\mathbb{P}}(1) - \xi_{y_{S,t}}^{\mathbb{P}}(1) \\ &+ \xi_{z_{M,t}}^{\mathbb{P}}(\beta_{S,z} - \Lambda_M) + \xi_{y_{M,t}}^{\mathbb{P}}(\beta_{S,y} - \Gamma_M) + \xi_{z_{S,t}}^{\mathbb{P}}(1 - \Lambda_S) + \xi_{y_{S,t}}^{\mathbb{P}}(1 - \Gamma_S) \\ &- \xi_{z_{M,t}}^{\mathbb{P}}(-\Lambda_M) - \xi_{y_{M,t}}^{\mathbb{P}}(-\Gamma_M) - \xi_{z_{S,t}}^{\mathbb{P}}(-\Lambda_S) - \xi_{y_{S,t}}^{\mathbb{P}}(-\Gamma_S) \end{aligned} \right).$$

Because,

$$\begin{aligned} -\xi_{z_{M,t}}^{\mathbb{P}}(\beta_{S,z}) + \xi_{z_{M,t}}^{\mathbb{P}}(\beta_{S,z} - \Lambda_M) - \xi_{z_{M,t}}^{\mathbb{P}}(-\Lambda_M) &= -\Lambda_M \beta_{S,z} h_{M,z,t}, \\ -\xi_{z_{S,t}}^{\mathbb{P}}(1) + \xi_{z_{S,t}}^{\mathbb{P}}(1 - \Lambda_S) - \xi_{z_{S,t}}^{\mathbb{P}}(-\Lambda_S) &= -\Lambda_S h_{S,z,t}, \\ -\xi_{y_{M,t}}^{\mathbb{P}}(\beta_{S,y}) + \xi_{y_{M,t}}^{\mathbb{P}}(\beta_{S,y} - \Gamma_M) - \xi_{y_{M,t}}^{\mathbb{P}}(-\Gamma_M) &= -h_{M,y,t} \gamma_{M,S}(\beta_{S,y}) \\ -\xi_{y_{S,t}}^{\mathbb{P}}(1) + \xi_{y_{S,t}}^{\mathbb{P}}(1 - \Gamma_S) - \xi_{y_{S,t}}^{\mathbb{P}}(-\Gamma_S) &= -h_{S,y,t} \gamma_S, \end{aligned}$$

we conclude that

$$1 = \exp\left(\mu_{S,t}^{\mathbb{P}} - r_t - \Lambda_M \beta_{S,z} h_{M,z,t} - h_{M,y,t} \gamma_{M,S}(\beta_{S,y}) - \Lambda_S h_{S,z,t} - h_{S,y,t} \gamma_S\right).$$

Therefore,

$$\mu_{S,t}^{\mathbb{P}} = r_t + \Lambda_M \beta_{S,z} h_{M,z,t} + h_{M,y,t} \gamma_{M,S}(\beta_{S,y}) + \Lambda_S h_{S,z,t} + h_{S,y,t} \gamma_S.$$

□

4.D Risk Neutral Conditional Variances and Jump Intensities

Lemma 4.4. *Let*

$$\eta_t^* = \left[1 \quad h_{M,z,t}^* \quad h_{M,y,t}^* \quad h_{S,z,t}^* \quad h_{S,y,t}^* \quad (\varepsilon_{M,t}^*)^2 \quad \sqrt{h_{M,z,t}^*} \varepsilon_{M,t}^* \quad (\varepsilon_{S,t}^*)^2 \quad \sqrt{h_{S,z,t}^*} \varepsilon_{S,t}^* \right]'$$

Then, for any $u \in \{M, S\}$ and $v \in \{z, y\}$,

$$h_{u,v,t+1}^* = \pi_{u,v} \eta_t^* \tag{4.26}$$

where $\pi_{u,v}$ is a 1×9 vector of constants satisfying

$$\begin{aligned}
 \pi_{M,z,1} &= w_{M,z} & \pi_{M,z,7} &= -2a_{M,z}(c_{M,z} + \Lambda_M) \\
 \pi_{M,z,2} &= b_{M,z} + a_{M,z}(c_{M,z} + \Lambda_M)^2 & \pi_{M,z,i} &= 0 \text{ for } i \in \{3, 4, 5, 8, 9\} \\
 \pi_{M,z,6} &= a_{M,z} \\
 \\
 \pi_{S,z,1} &= \kappa_{S,z}\pi_{M,z,1} - a_{S,z} & \pi_{S,z,6} &= \kappa_{S,z}\pi_{M,z,6} \\
 \pi_{S,z,2} &= \kappa_{S,z}(\pi_{M,z,2} - b_{S,z}) & \pi_{S,z,7} &= \kappa_{S,z}\pi_{M,z,7} \\
 \pi_{S,z,4} &= b_{S,z} + a_{S,z}(2c_{S,z} + \Lambda_S)\Lambda_S & \pi_{S,z,8} &= a_{S,z} \\
 \pi_{S,z,i} &= 0 \text{ for } i \in \{3, 5\} & \pi_{S,z,9} &= -2a_{S,z}(c_{S,z} + \Lambda_S) \\
 \\
 \pi_{M,y,1} &= w_{M,y} & \pi_{M,y,6} &= a_{M,y} \\
 \pi_{M,y,2} &= a_{M,y}(c_{M,y} + \Lambda_M)^2 & \pi_{M,y,7} &= -2a_{M,y}(c_{M,y} + \Lambda_M) \\
 \pi_{M,y,3} &= b_{M,y} & \pi_{M,y,i} &= 0 \text{ for } i \in \{4, 5, 8, 9\} \\
 \\
 \pi_{S,y,1} &= \kappa_{S,y}\pi_{M,y,1} - a_{S,y} & \pi_{S,y,6} &= \kappa_{S,y}\pi_{M,y,6} \\
 \pi_{S,y,2} &= \kappa_{S,y}\pi_{M,y,2} & \pi_{S,y,7} &= \kappa_{S,y}\pi_{M,y,7} \\
 \pi_{S,y,3} &= \kappa_{S,y}(\pi_{M,y,3} - b_{S,y}) & \pi_{S,y,8} &= a_{S,y} \\
 \pi_{S,y,4} &= a_{S,y}(2c_{S,y} + \Lambda_S)\Lambda_S & \pi_{S,y,9} &= -2a_{S,y}(c_{S,y} + \Lambda_S) \\
 \pi_{S,y,5} &= b_{S,y}
 \end{aligned}$$

Proof of Lemma 4.4. The risk neutral market conditional variance $h_{M,z,t+1}^*$ and jump intensity variable $h_{M,y,t+1}^*$ are obtained by replacing (4.23) in (4.3) and (4.5).

In the case of the stocks, for any $v \in \{z, y\}$,

$$\begin{aligned}
 & (\varepsilon_{S,t}^2 - 1 - 2c_{S,v}\sqrt{h_{S,z,t}}\varepsilon_{S,t}) \\
 &= \left(\varepsilon_{S,t}^* - \Lambda_S\sqrt{h_{S,z,t}^*}\right)^2 - 1 - 2c_{S,v}\sqrt{h_{S,z,t}^*}\left(\varepsilon_{S,t}^* - \Lambda_S\sqrt{h_{S,z,t}^*}\right) \\
 &= (2c_{S,v} + \Lambda_S)\Lambda_S h_{S,z,t}^* + \left((\varepsilon_{S,t}^*)^2 - 1 - 2(c_{S,v} + \Lambda_S)\sqrt{h_{S,z,t}^*}\varepsilon_{S,t}^*\right). \quad (4.27)
 \end{aligned}$$

where the first equality arises from (4.23). Replacing back in the conditional variance (4.4) and the jump intensity process (4.6) leads to their risk neutral versions. \square

4.E Moment Generating Function of Risk-Neutral Excess Returns

Lemma 4.5. *For $u \in \{M, S\}$, the conditional moment generating function of the excess returns satisfies*

$$\begin{aligned} \varphi_{\tilde{R},t,T}^{\mathbb{Q}}(\phi) = & \exp \left(\mathcal{A}_{u,T-t}(\phi) + \mathcal{B}_{u,T-t}(\phi) h_{M,z,t+1}^* \right. \\ & \left. + \mathcal{C}_{u,T-t}(\phi) h_{M,y,t+1}^* + \mathcal{D}_{u,T-t}(\phi) h_{S,z,t+1}^* + \mathcal{E}_{u,T-t}(\phi) h_{S,y,t+1}^* \right) \end{aligned}$$

where the coefficients are found using a backward recursion over time. Indeed, $\varphi_{u,0}^{\mathbb{Q}}(\phi) = 1$ implies that

$$\mathcal{A}_{u,0}(\phi) = \mathcal{B}_{u,0}(\phi) = \mathcal{C}_{u,0}(\phi) = \mathcal{D}_{u,0}(\phi) = \mathcal{E}_{u,0}(\phi) = 0.$$

For $i \in \{0, 1, \dots, 9\}$, let

$$\zeta_{u,T-t-1,i}(\phi) = \mathcal{B}_{u,T-t-1}(\phi) \pi_{M,z,i} + \mathcal{C}_{u,T-t-1}(\phi) \pi_{M,y,i} + \mathcal{D}_{u,T-t-1}(\phi) \pi_{S,z,i} + \mathcal{E}_{u,T-t-1}(\phi) \pi_{S,y,i},$$

where the π . are as provided in Section 4.D. If $\zeta_{s,6}(\phi) < \frac{1}{2}$ and $\zeta_{s,8}(\phi) < \frac{1}{2}$ for any $s \in \{t+1, \dots, T\}$, then

$$\begin{aligned} \mathcal{A}_{u,T-t}(\phi) &= \mathcal{A}_{u,T-t-1}(\phi) + \zeta_{u,T-t-1,1}(\phi) - \frac{1}{2} \log(1 - 2\zeta_{u,T-t-1,6}(\phi)) \\ &\quad - \frac{1}{2} \log(1 - 2\zeta_{u,T-t-1,8}(\phi)), \\ \mathcal{B}_{u,T-t}(\phi) &= \zeta_{u,T-t-1,2}(\phi) - \frac{1}{2} \beta_{u,z}^2 \phi + \frac{1}{2} \frac{(\zeta_{u,T-t-1,7}(\phi) + \beta_{u,z} \phi)^2}{1 - 2\zeta_{u,T-t-1,6}(\phi)}, \\ \mathcal{C}_{u,T-t}(\phi) &= \zeta_{u,T-t-1,3}(\phi) - \Pi_M^*(\beta_{u,y}) \phi + \Pi_M^*(\beta_{u,y} \phi), \\ \mathcal{D}_{u,T-t}(\phi) &= \zeta_{u,T-t-1,4}(\phi) - \frac{1}{2} (\beta'_{u,z})^2 \phi + \frac{1}{2} \frac{(\zeta_{u,T-t-1,9}(\phi) + \beta'_{u,z} \phi)^2}{1 - 2\zeta_{u,T-t-1,8}(\phi)}, \\ \mathcal{E}_{u,T-t}(\phi) &= \zeta_{u,T-t-1,5}(\phi) - \Pi_S^*(\beta'_{u,y}) \phi + \Pi_S^*(\beta'_{u,y} \phi). \end{aligned}$$

where for the market case, $\beta_{M,z} = \beta_{M,y} = 1$ and $\beta'_{M,z} = \beta'_{M,y} = 0$ while for the stock, $\beta'_{S,z} = \beta'_{S,y} = 1$.

As the proof is strongly inspired from the existing literature, we refer the reader to Appendix A.5.

4.F Particle Filter

In the following, whenever the subscript M and S have been dropped, the approach is applicable to both market and stock data.

The filter is based on pure jump particle paths $y_{1:T}^{(i)} = \{y_1^{(i)}, y_2^{(i)}, \dots, y_T^{(i)}\}$, $i \in \{1, \dots, N\}$ and the sequential importance resampling (SIR) of Gordon et al. (1993) is implemented.³³ A single step of the SIR is now described.

Assume that N jump paths $y_{1:t-1}^{(i)}$, $i \in \{1, 2, \dots, N\}$ are available up to time $t - 1$. As a by-product, the conditional variance $h_{z,t}^{(i)}$ and the jump scale variable $h_{y,t}^{(i)}$ are recovered.

1. For $i \in \{1, 2, \dots, N\}$, the time t jump $y_t^{(i)}$ is simulated from the proposal distribution³⁴

$$f(\cdot \mid y_{1:t-1}^{(i)}, R_{1:t-1}) = f_{NIG}(\cdot; \alpha, \delta, h_{y,t}^{(i)}).$$

2. For $i \in \{1, 2, \dots, N\}$, update the importance weights (up to a normalizing constant) to reflect how likely the simulated particles are with respect to the time t information R_t :

$$\bar{\omega}_t^{(i)} = f(R_t \mid R_{1:t-1}, y_{1:t}^{(i)}).$$

More precisely, from Equations (4.1) and (4.10), the market returns satisfy³⁵

$$R_{M,t} = r_t + \left(\lambda_M - \frac{1}{2} \right) h_{M,z,t} + (\gamma_M - \Pi_M(1)) h_{M,y,t} + z_{M,t} + y_{M,t}.$$

³³Throughout the paper, $N = 25,000$ particles are used.

³⁴More precisely,

$$\begin{aligned} f_{NIG}(x; \alpha, \delta, h) &= \frac{\alpha h K_1(\alpha \sqrt{h^2 + x^2})}{\pi \sqrt{h^2 + x^2}} \exp(h \sqrt{\alpha^2 - \delta^2} + \delta x) \\ K_1(x) &= \int_0^\infty \exp(-x \cosh(t)) \cosh(t) dt. \end{aligned}$$

³⁵Similarly, for the stock returns, we have

$$f(R_{S,t} \mid R_{S,1:t-1}, y_{S,1:t}^{(i)}, \tilde{h}_{M,z,1:t}, \tilde{h}_{M,y,1:t}, \tilde{z}_{M,1:t+1}, \tilde{y}_{M,1:t+1}) = \phi(R_{S,t}; m_{S,t}^{(i)}, h_{S,z,t}^{(i)})$$

with

$$\begin{aligned} m_{S,t}^{(i)} &= r_t + \left(\beta_{S,z} \lambda_M - \frac{1}{2} \beta_{S,z}^2 \right) \tilde{h}_{M,z,t} + [\gamma_{M,S}(\beta_{S,y}) - \Pi_M(\beta_{S,y})] \tilde{h}_{M,y,t} + \beta_{S,z} \tilde{z}_{M,t+1} + \beta_{S,y} \tilde{y}_{M,t+1} \\ &\quad + \left(\lambda_S - \frac{1}{2} \right) h_{S,z,t}^{(i)} + [\gamma_S - \Pi_S(1)] h_{S,y,t}^{(i)} + y_{S,t+1}^{(i)} \end{aligned}$$

and variance $h_{S,z,t}^{(i)}$.

Therefore, conditionally on a simulated path $y_{M,1:t}^{(i)}$ and on the past returns $R_{M,1:t-1}$, the time t market return $R_{M,t}$ is normally distributed with expectation

$$m_{M,t}^{(i)} = r_t + \left(\lambda_M - \frac{1}{2} \right) h_{M,z,t}^{(i)} + (\gamma_M - \Pi_M(1)) h_{M,y,t}^{(i)} + y_{M,t}^{(i)}$$

and variance $h_{M,z,t}^{(i)}$,

$$f(R_{M,t} \mid R_{M,1:t-1}, y_{M,1:t}^{(i)}) = \frac{1}{\sqrt{2\pi h_{M,z,t}^{(i)}}} \exp \left\{ -\frac{1}{2} \frac{(R_{M,t} - m_{M,t}^{(i)})^2}{h_{M,z,t}^{(i)}} \right\}.$$

3. For $i \in \{1, 2, \dots, N\}$, compute the normalized weights

$$\omega_t^{(i)} = \frac{\bar{\omega}_t^{(i)}}{\sum_{k=1}^N \bar{\omega}_t^{(k)}}.$$

4. For $i \in \{1, 2, \dots, N\}$, update the conditional variance and the jump scale variable. For the market, based on (4.3) and (4.5),

$$\begin{aligned} h_{M,z,t+1}^{(i)} &= w_{M,z} + b_{M,z} h_{M,z,t}^{(i)} + \frac{a_{M,z}}{h_{M,z,t}^{(i)}} \left(z_{M,t}^{(i)} - c_{M,y} h_{M,z,t}^{(i)} \right)^2 \\ h_{M,y,t+1}^{(i)} &= w_{M,y} + b_{M,y} h_{M,y,t}^{(i)} + \frac{a_{M,y}}{h_{M,z,t}^{(i)}} \left(z_{M,t}^{(i)} - c_{M,y} h_{M,z,t}^{(i)} \right)^2 \end{aligned}$$

where $z_{M,t}^{(i)} = R_{M,t} - m_{M,t}^{(i)}$.³⁶

5. From normalized importance weights, compute the filtered variables

$$\begin{aligned} \tilde{z}_{M,t} &= \sum_{i=1}^N z_{M,t}^{(i)} \omega_t^{(i)}, & \tilde{h}_{M,z,t+1} &= \sum_{i=1}^N h_{M,z,t+1}^{(i)} \omega_t^{(i)}, \\ \tilde{y}_{M,t} &= \sum_{i=1}^N y_{M,t}^{(i)} \omega_t^{(i)}, & \tilde{h}_{M,y,t+1} &= \sum_{i=1}^N h_{M,y,t+1}^{(i)} \omega_t^{(i)}. \end{aligned}$$

6. Resample the particles using the continuous sampling of Malik and Pitt (2011).³⁷

³⁶For the stock,

$$\begin{aligned} h_{S,z,t+1}^{(i)} &= \kappa_{S,z} \tilde{h}_{M,z,t+1} + b_{S,z} \left(h_{S,z,t}^{(i)} - \kappa_{S,z} \tilde{h}_{M,z,t} \right) + a_{S,z} \left(\left(h_{S,z,t}^{(i)} \right)^{-1} \left(z_{S,t}^{(i)} \right)^2 - 1 - 2c_{S,z} z_{S,t}^{(i)} \right) \\ h_{S,y,t+1}^{(i)} &= \kappa_{S,y} \tilde{h}_{M,y,t+1} + b_{S,y} \left(h_{S,y,t}^{(i)} - \kappa_{S,y} \tilde{h}_{M,y,t} \right) + a_{S,y} \left(\left(h_{S,z,t}^{(i)} \right)^{-1} \left(z_{S,t}^{(i)} \right)^2 - 1 - 2c_{S,y} z_{S,t}^{(i)} \right) \end{aligned}$$

where $z_{S,t}^{(i)} = R_{S,t} - m_{S,t}^{(i)}$.

³⁷As argued in Creal (2012), basic resampling methods are ill-suited for maximum likelihood estimation.

- (a) Draw N particles from the current particle set from a smoothed empirical cdf as proposed in Malik and Pitt (2011) and let $\{h_{M,z,t+1}^{(ji)}\}_{i=1}^N$ and $\{h_{M,y,t+1}^{(ji)}\}_{i=1}^N$ denotes the resulting conditional variances and the jump intensity variables once the resampling is accomplished.³⁸
- (b) Replace the current conditional variance and jump intensity with their re-sampled values:

$$h_{M,z,t+1}^{(i)} \leftarrow h_{M,z,t+1}^{(ji)}, \quad \text{and} \quad h_{M,y,t+1}^{(i)} \leftarrow h_{M,y,t+1}^{(ji)}.$$

The log-likelihood is obtained as a by-product of the particle filter. Indeed,

$$L_{M,\text{returns}}(\Theta_M) = \sum_{t=1}^T \log \left(\sum_{i=1}^N \bar{\omega}_t^{(i)} \right).$$

4.G Stock Fundamentals

The market and the volatility betas are obtained by regressing a stock's excess returns on the S&P 500 excess returns and daily changes on the VIX using the past year of data:

$$r_{S,t-k} = \alpha + \beta_{\text{MKT},t} \text{MKT}_{t-k} + \beta_{\Delta\text{VIX},t} \Delta\text{VIX}_{t-k} + \epsilon_{t-k}, \quad k = 0, \dots, 252.$$

The betas are considered missing if less than 63 data points are available over the past year.

The market equity (ME) is obtained by multiplying the number of outstanding shares by the close price for each stock.

The book equity (BE) is computed as the difference between the total assets of a firm (ATQ in Compustat) and its liabilities. The latter are defined as the sum of the debt in current liabilities (DLCQ) and half of the long-term debt (DLTTQ) as in Bharath and Shumway (2008). Both the debt in current liabilities and the long-term debt are linearly interpolated between quarterly data points to obtain daily estimates. BE is considered missing when negative.

³⁸Note that when the number of resampled particles is small, we use importance sampling to increase it. To this end, the jump intensity variable is artificially increased and a weight correction is applied accordingly.

The operating profitability (OP) is defined as the quarterly revenue at time t (REVTQ), minus the cost of goods sold at time t (COGSQ), the interest expense at time t (XINTQ) and selling, general, and administrative expenses at time t (XSGAQ), divided by book equity for the last year (i.e. at t minus 1 year). All the fundamental values used to compute OP were linearly interpolated from quarterly data.

The investment level is obtained from the book value of assets. Specifically, it is computed as the change in total assets over the previous year (from t minus 1 year to t), divided by the total assets at the end of the previous year (i.e. at time t minus 1 year). The values of the assets are also linearly interpolated from quarterly data to obtain daily estimates.

Finally, the trailing twelve-month return is obtained by taking the sum of daily excess returns over the last year (i.e. 252 previous business days, when available).

References

- Aït-Sahalia, Y., and A. Lo. 1998. Nonparametric estimation of state-price densities implicit in financial asset prices. *Journal of Finance* 53:499–547.
- Albuquerque, R. 2012. Skewness in stock returns: Reconciling the evidence on firm versus aggregate returns. *Review of Financial Studies* 25:1630–1673.
- Ang, A., R. J. Hodrick, Y. Xing, and X. Zhang. 2006. The cross-section of volatility and expected returns. *Journal of Finance* 61:259–299.
- Babaoğlu, K. 2015. The pricing of market jumps in the cross-section of stocks and options. *Working paper*.
- Bakshi, G., P. Carr, and L. Wu. 2008. Stochastic risk premiums, stochastic skewness in currency options, and stochastic discount factors in international economies. *Journal of Financial Economics* 87:132–156.
- Bakshi, G., N. Kapadia, and D. Madan. 2003. Stock returns characteristics, skew laws, and the differential pricing of individual equity options. *Review of Financial Studies* 16:101–143.
- Bakshi, G., D. Madan, and G. Panayotov. 2010. Returns of claims on the upside and the viability of U-shaped pricing kernels. *Journal of Financial Economics* 97:130–154.

- Bates, D. S. 1991. The Crash of '87: Was it expected? The evidence from options markets. *Journal of Finance* 46:1009–1044.
- Bates, D. S. 2006. Maximum likelihood estimation of latent affine processes. *Review of Financial Studies* 19:909–965.
- Bates, D. S. 2008. The market for crash risk. *Journal of Economic Dynamics and Control* 32:2291–2321.
- Bharath, S. T., and T. Shumway. 2008. Forecasting default with the Merton distance to default model. *Review of Financial Studies* 21:1339–1369.
- Black, F., and M. Scholes. 1973. The pricing of options and corporate liabilities. *Journal of Political Economy* 81:637–654.
- Bollerslev, T., and V. Todorov. 2011. Tails, fears, and risk premia. *Journal of Finance* 66:2165–2211.
- Bolorforoosh, A. 2014. Is idiosyncratic volatility risk priced? Evidence from the physical and risk-neutral distributions. *Working paper*.
- Brennan, M. 1979. The pricing of contingent claims in discrete time models. *Journal of Finance* 34:53–68.
- Broadie, M., M. Chernov, and M. Johannes. 2007. Model specification and risk premia: Evidence from futures options. *Journal of Finance* 62:1453–1490.
- Carhart, M. 1997. On persistence in mutual fund performance. *Journal of Finance* 52:57–82.
- Carr, P., H. Geman, D. Madan, and M. Yor. 2002. The fine structure of asset returns: An empirical investigation. *Journal of Business* 75:305–332.
- Carr, P., and L. Wu. 2007. Stochastic skew in currency options. *Journal of Financial Economics* 86:213–247.
- Chernov, M. 2003. Empirical reverse engineering of the pricing kernel. *Journal of Econometrics* 116:329–364.
- Chernov, M., and E. Ghysels. 2000. A study towards a unified approach to the joint estimation of objective and risk neutral measures for the purpose of options valuation. *Journal of Financial Economics* 56:407–458.

- Christoffersen, P., R. Elkamhi, B. Feunou, and K. Jacobs. 2010. Option valuation with conditional heteroskedasticity and nonnormality. *Review of Financial Studies* 23:2139–2183.
- Christoffersen, P., M. Fournier, and K. Jacobs. 2013. The factor structure in equity options. *Working paper*.
- Christoffersen, P., K. Jacobs, and C. Ornathanalai. 2012. Dynamic jump intensities and risk premiums: Evidence from S&P 500 returns and options. *Journal of Financial Economics* 106:447–472.
- Christoffersen, P. F., K. Jacobs, C. Ornathanalai, and Y. Wang. 2008. Option valuation with long-run and short-run volatility components. *Journal of Financial Economics* 90:272–297.
- Cox, J. C., S. A. Ross, and M. Rubinstein. 1979. Option pricing: A simplified approach. *Journal of Financial Economics* 7:229–263.
- Creal, D. 2012. A survey of sequential Monte Carlo methods for economics and finance. *Econometric Reviews* 31:245–296.
- Douglas, G. W. 1969. Risk in the equity markets: An empirical appraisal of market efficiency. *Yale Economic Essays* 9:3–45.
- Duan, J.-C., and J. Wei. 2009. Systematic risk and the price structure of individual equity options. *Review of Financial Studies* 22:1981–2206.
- Dumas, B., J. Fleming, and R. E. Whaley. 1998. Implied volatility functions: Empirical tests. *Journal of Finance* 53:2059–2106.
- Durham, G., J. Geweke, and P. Ghosh. 2015. A comment on Christoffersen, Jacobs, and Ornathanalai (2012), “Dynamic jump intensities and risk premiums: Evidence from S&P 500 returns and options”. *Journal of Financial Economics* 115:210–214.
- Elkamhi, R., and C. Ornathanalai. 2010. Market jump risk and the price structure of individual equity options. *Working paper*pp. 1–55.
- Engle, R., E. Ghysels, and B. Sohn. 2013. Stock market volatility and macroeconomic fundamentals. *Review of Economics and Statistics* 95:776–797.
- Engle, R. F., and G. Lee. 1999. A permanent and transitory component model of stock return volatility. In *Cointegration, causality, and forecasting: A festschrift in honor of Clive W.J. Granger*, pp. 475–497. Oxford University Press.

- Engle, R. F., and J. G. Rangel. 2008. The spline-GARCH model for low-frequency volatility and its global macroeconomic causes. *Review of Financial Studies* 21:1187–1222.
- Eraker, B. 2004. Do stock prices and volatility jump? Reconciling evidence from spot and option prices. *Journal of Finance* 59:1367–1404.
- Eraker, B. 2008. Affine general equilibrium models. *Management Science* 54:2068–2080.
- Fama, E. F., and K. French. 1992. The cross-section of expected stock returns. *Journal of Finance* 47:427–465.
- Fama, E. F., and K. French. 1993. Common risk factors in the returns on stocks and bonds. *Journal of Financial Economics* 33:3–56.
- Fama, E. F., and K. French. 2015. A five-factor asset pricing model. *Journal of Financial Economics* 116:1–22.
- Gil-Pelaez, J. 1951. Note on the inversion theorem. *Biometrika* 38:481–482.
- Gordon, N. J., D. J. Salmond, and A. F. Smith. 1993. Novel approach to nonlinear/non-Gaussian Bayesian state estimation. In *IEE Proceedings F (Radar and Signal Processing)*, vol. 140, pp. 107–113.
- Gourier, E. 2014. Pricing of idiosyncratic equity and variance risks. *Working paper*.
- Goyal, A., and P. Santa-Clara. 2003. Idiosyncratic risk matters! *Journal of Finance* 58:975–1008.
- Herskovic, B., B. Kelly, H. Lustig, and S. Van Nieuwerburgh. 2016. The common factor in idiosyncratic volatility: Quantitative asset pricing implications. *Journal of Financial Economics* 119:249–283.
- Heston, S. 1993. A closed-form solution for options with stochastic volatility with applications to bond and currency options. *Review of Financial Studies* 6:327.
- Heston, S., and S. Nandi. 2000. A closed-form GARCH option valuation model. *Review of Financial Studies* 13:585–625.
- Jackwerth, J. 2000. Recovering risk aversion from option prices and realized returns. *Review of Financial Studies* 13:433–451.

- Jacobs, K., and K. Q. Wang. 2004. Idiosyncratic consumption risk and the cross section of asset returns. *Journal of Finance* 59:2211–2252.
- Jensen, C. S., D. Lando, and L. H. Pedersen. 2016. Generalized recovery. *Working paper*.
- Kraus, A., and R. H. Litzenberger. 1976. Skewness preference and the valuation of risk assets. *Journal of Finance* 31:1085–1100.
- Lintner, J. 1965a. Security prices and risk: The theory and comparative analysis of A.T.&T. and leading industrials. *Presented at the conference on "The Economics of Regulated Public Utilities" at the University of Chicago Business School*.
- Lintner, J. 1965b. The valuation of risk assets and the selection of risky investments in stock portfolios and capital budgets. *Review of Economics and Statistics* 47:13–37.
- Liu, J., J. Pan, and T. Wang. 2005. An equilibrium model of rare-event premia and its implication for option smirks. *Review of Financial Studies* 18:131–164.
- Malik, S., and M. K. Pitt. 2011. Particle filters for continuous likelihood evaluation and maximisation. *Journal of Econometrics* 165:190–209.
- Markowitz, H. 1952. Portfolio selection. *Journal of Finance* 7:77–91.
- Martin, I. 2016. What is the expected return on the market? *Working paper*.
- Martin, I., and C. Wagner. 2016. What is the expected return on a stock? *Working paper*.
- Merton, R. 1973. An intertemporal capital asset pricing model. *Econometrica* 41:867–887.
- Newey, W., and K. West. 1987. A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix. *Econometrica* 55:703–708.
- Ornthanalai, C. 2014. Levy jump risk: Evidence from options and returns. *Journal of Financial Economics* 112:69–90.
- Pan, J. 2002. The jump-risk premia implicit in options: Evidence from an integrated time-series study. *Journal of Financial Economics* 63:3–50.
- Renault, E. 1997. Econometric models of option pricing errors. *Econometric Society Monographs* 28:223–278.

- Ross, S. 1976. The arbitrage theory of capital asset pricing. *Journal of Economic Theory* 13:341–360.
- Ross, S. 2015. The recovery theorem. *Journal of Finance* 70:615–648.
- Rubinstein, M. 1975. The strong case for the generalized logarithmic utility model of financial markets. *Journal of Finance* 31:551–571.
- Santa-Clara, P., and S. Yan. 2010. Crashes, volatility, and the equity premium: lessons from S&P 500 options. *Review of Economics and Statistics* 92:435–451.
- Sharpe, W. F. 1964. Capital asset prices: A theory of market equilibrium under conditions of risk. *Journal of Finance* 19:425–442.
- Stambaugh, R. F., J. Yu, and Y. Yuan. 2015. Arbitrage asymmetry and the idiosyncratic volatility puzzle. *Journal of Finance* 70:1903–1948.
- Trolle, A., and E. Schwartz. 2009. Unspanned stochastic volatility and the pricing of commodity derivatives. *Review of Financial Studies* 22:4423–4461.
- Xiao, X., and C. Zhou. 2014. Systematic and idiosyncratic jump risks in the expected stock returns. *Working paper*.

Chapter 5

Extracting Latent States with High Frequency Option Prices

Abstract^{*}

We propose the option realized variance as a new observable covariate that integrates high frequency option prices in the inference of option pricing models. Using several simulation and empirical studies, this paper documents the incremental information offered by this variable. Our empirical results show that the information contained in this new covariate improves the inference of model variables such as the instantaneous variance, return jumps, and variance jumps of the S&P 500 index. Parameter estimates indicate that the risk premium breakdown between jump and diffusive risks is affected by the omission of this information.

Keywords: high frequency data; option realized variance; options; jump-diffusions; particle filter.

^{*}Joint work with Diego Amaya and Geneviève Gauthier. Amaya is affiliated with Wilfrid Laurier University and Gauthier with HEC Montréal.

5.1 Introduction

The existing literature on option pricing has almost exclusively focused on model specification¹ and estimation², but has not given much regard to the incremental information offered by observable variables. This paper addresses this issue by enlarging the set of observable covariates with the option realized variance (*ORV*). This novel source of information parsimoniously integrates high frequency option prices into the list of observable covariates.

Identification of model components relies on time series of observable variables. The first set of observable covariates corresponds to the time series of underlying asset prices, which contain information about physical measure dynamics. Traditional discrete-time models such as those studied in Engle (1982) and Bollerslev (1986) employ these prices directly to infer volatility process dynamics. In continuous-time models, these prices have been used to analyze the role of jumps on returns and their volatility (Eraker et al., 2003). Advances in econometrics have allowed researchers to estimate different characteristics of the return variation using measures such as the realized variance and bipower variation (Andersen et al., 2001, Barndorff-Nielsen and Shephard, 2004, among others). In the context of option pricing, Christoffersen et al. (2014) and Christoffersen et al. (2015) show that these realized measures reduce option pricing errors significantly, providing evidence of the economic value of high frequency asset prices.

A second important source of information comes from the rich cross section of option prices. These prices not only provide information about the dynamics under the risk-neutral measure, but also about the parameters that govern the conditional return distribution. For instance, Bates (2000) uses option contracts on the S&P 500 index to extract implicit distributions of competing jump-diffusion models. Combining both underlying asset returns and option prices, Eraker (2004) conducts an empirical analysis of jump-diffusion models by looking at their ability to simultaneously fit option and return data. Johannes et al. (2009) propose a filtering methodology based on return and option datasets to disentangle jumps and diffusive components. More

¹In the recent literature, few model features have been considered: jumps in volatility (Eraker et al., 2003, Pan, 2002, Todorov and Tauchen, 2011), co-jumps (Broadie et al., 2007, Chernov et al., 2003, Eraker, 2004), jump arrival intensities (Bates, 2000), and multifactor stochastic volatility with jumps (Bardgett et al., 2015), among others.

²A non-exhaustive list of methods for the estimation of these models includes Markov chain Monte Carlo methods (Eraker, 2001, Jones, 1998), nonparametric approaches (Aït-Sahalia and Lo, 1998), the simulated method of moments (Duffie and Singleton, 1993, Gallant and Tauchen, 2010), the generalized method of moments (Pan, 2002), and approximate maximum likelihood (Bates, 2006), among others.

recently, the emergence of derivative contracts on these options has allowed studies such as Cont and Kokholm (2013) and Bardgett et al. (2015) to extract information about the expected future realized variance.

To understand the information content of option realized variances, we employ a model of stock returns that exhibits stochastic volatility, jumps in returns with stochastic intensity, and independent jumps in variance using the general framework of Duffie et al. (2000). These characteristics are similar to those studied in Pan (2002), Eraker et al. (2003), Broadie et al. (2007) and Johannes et al. (2009), among others. We favour this specification over more sophisticated ones because it is flexible enough to capture dynamic properties of stock returns while providing analytical tractability to conduct empirical analysis with time series of option panels.

Jumps in the variance process are generally difficult to infer as the contributions to diffusive and jump components of the variance process are hard to disentangle. The inclusion of option realized variances in the inference process circumvents this issue because it brings in non-redundant information. Contrary to Andersen et al. (2015b), we use a parametric approach to characterize model components from high frequency option prices. A distinctive characteristic of *ORV* is that, depending on its moneyness, specific features of the underlying processes can be isolated. For instance, deep out-of-the-money options have very low deltas and vegas, so most of the variability in option prices comes from discontinuous sample paths generated by jumps in returns and volatility.

The paper conducts simulated and empirical studies on the incremental information content of several information sources. In addition to returns and a panel of daily option implied volatilities (*IV*), we add measures of the index return variance such as realized variance and bipower variation. We complement this sample with the newly proposed option realized variance, which is available every day for each option in the panel. An extensive Monte Carlo study that relies on an extension of Johannes et al.'s (2009) filter shows that option realized variances contribute significantly to the identification of variance jumps.

In our empirical study, we employ intraday data from options on the S&P 500 index and from futures prices on the E-mini S&P 500 to document several properties of the option realized variance.³ Using a sample of daily observations that extends from July 2004 to December 2012, we analyze the large cross section of option data by

³Recently, we have been made aware that Audrino and Fengler (2015) have developed, independently, a competing measure of the option realized variance similar to ours. Whereas they use option log-prices to compute their measure, we employ option prices. We privilege the use of option prices

constructing *ORV* surfaces across moneynesses and maturities. We find that there is an important level of commonality between *ORVs* and the realized variance of index returns, especially when large, sporadic shocks happen.⁴ This degree of commonality is expected as a result of intense trading activity in both markets (Stephan and Whaley, 1990). We also assess the economic relationship between *ORVs* and index return variations with predictive regressions: lagged values of selected *ORVs* predict the one-day ahead realized variance and jump variation activity.⁵

The paper then investigates several empirical implications of adding *ORVs* as additional observable covariates in the estimation of option pricing models. We first observe that the addition of *ORV* produces less frequent larger jumps in the variance process and more frequent smaller negative jumps in the price process. Next, we find that the addition of this new covariate increases the compensation for bearing price jump risk (i.e. 0.1% without *ORV* vs 1.6% with *ORV*, on average) and decreases the compensation for diffusive risk premium (i.e. 4.4% when *ORV* is excluded and 1.8% when *ORV* is used). Thus, in total, the average risk premium that results from using both information sets is similar, but the risk premium breakdown between diffusive and jump risks is different. Third, regarding the uncertainty around the estimation of latent quantities, posterior standard deviations show that adding *ORV* to the information set decreases this uncertainty. The most staggering decrease is observed for variance jumps, which experience approximately a fivefold decrease. Finally, in- and out-of-sample analyses show that the inclusion of all information sources produces smaller forecasting errors of implied volatilities.

Several studies in the literature on asset pricing employ non-parametrical methods to infer the underlying structure of stock returns and their volatilities from high frequency prices.⁶ We focus on a parametrical approach to study the incremental information

because the variance measure resulting from these prices does not depend on the price of the option itself, which is important for large-scale optimizations in empirical option pricing studies.

⁴Events such as the financial crisis of 2008, the flash crash episode of 2010, or the downgrade of the U.S. debt in 2011 have common spikes in all time series.

⁵The jump variation activity is measured as the positive difference between realized variance and bipower variation.

⁶Todorov and Tauchen (2011) find evidence of discontinuous co-movements between the instantaneous volatilities and returns using high frequency data of VIX and the S&P 500 index. Andersen et al. (2015a) employ a nonparametric framework to infer latent instantaneous volatilities and jump intensities with intraday Black-Scholes implied volatilities. Using the rich information embedded in option prices, the authors find that the dynamics for *IV* of deep OTM puts behave as a pure jump process, whereas those of near-the-money are better characterized by a diffusive process. Bandi and Reno (2016) use high frequency data for the S&P 500 index to estimate its realized variance over intraday intervals. Using a novel moment-based procedure, the authors find support for the existence of independent jumps in the instantaneous volatilities and co-jumps between the price process and that of the volatility.

offered by intraday prices. Our results suggest that not all sources give the same information and that there is unique information about specific features of the model in each source. Our paper also contributes to the literature with a parsimonious measure that summarizes intraday information from option prices. In line with Andersen et al. (2015a), we document that the cross section of option high frequency prices, summarized by the option realized variance in this study, contains relevant information about the discontinuous part of the variance process.

This paper is also related to studies that deal with the filtering of continuous-time jump-diffusion models.⁷ Our paper differs from this literature in several ways. First, we propose a filter based on the sequential importance resampling (SIR) algorithm (Gordon et al., 1993) that, in addition to index returns and option prices, employs realized variance, bipower variation, and option realized variances as observable sources. Second, contrary to Johannes et al. (2009), we estimate the proposed model using the filter, which requires special attention to the way the likelihood function is approximated and provides a less computationally intensive procedure than the one employed in Eraker (2004).

The rest of this paper is organized as follows. Section 5.2 describes the model and provides insights about the information content of observable variables. Section 5.3 presents the estimation methodology of the model. Simulated-based results are given in Section 5.4. Section 5.5 conducts a nonparametric study of option realized variances. Section 5.6 presents the option pricing implications of adding *ORVs* in the estimation set. Finally, Section 5.7 concludes.

5.2 Framework

The goal of this section is to present a model that captures different properties of asset prices. The building block is the square root stochastic volatility specification of Heston (1993). To adequately model fat tails of return distributions, a compound Poisson process with Gaussian jumps is added to the price process.⁸ Although both stochastic volatility and jumps in the price process can generate rich dynamics, Bakshi

⁷Johannes et al. (2009) and Bardgett et al. (2015) use an optimal filtering methodology that combines time-discretization schemes with Monte Carlo methods to obtain latent states. Eraker (2004) employs Markov chain Monte Carlo simulation to estimate the posterior distribution of parameters, as well as volatility and jump processes.

⁸This kind of jump process is used by Bates (1996), Bakshi et al. (1997), Duffie et al. (2000), Pan (2002), Eraker et al. (2003), Johannes et al. (2009) to name a few.

et al. (1997), Bates (2000) and Pan (2002) find that volatility dynamics resulting from this specification tend to be misspecified. As explained by Andersen et al. (2001) and Alizadeh et al. (2002), two factors generate volatility: a rapidly changing factor and a highly persistent slowly moving one. The latter could be modelled by jumps in variance process. Therefore, exponentially distributed jumps are also added to the variance process.⁹

We deliberately focus our attention on a model with the following latent factors: diffusive stochastic volatility, log-equity jumps, and instantaneous variance jumps. These three risk factors induce different behaviour in the asset price and hence can be identified from an econometric perspective, given the right amount of data (Eraker et al., 2003). This insight is explored in the second part of this section, in which we discuss the list of available data sources and link these quantities to different features of the data generating process.

5.2.1 Model

The model is part of the family of stochastic volatility models with jumps (SVJ). This affine jump-diffusion model yields semi-closed form solutions for option pricing as shown in Duffie et al. (2000). Under the objective measure \mathbb{P} , the process governing the dynamics of the log-equity price, Y , and its instantaneous stochastic variance, V , are

$$dY_t = \alpha_t^{\mathbb{P}} dt + \sqrt{V_t} dW_{Y,t} + dJ_{Y,t}, \quad (5.1)$$

$$dV_t = \kappa(\theta - V_t) dt + \sigma \sqrt{V_t} dW_{V,t} + dJ_{V,t}, \quad (5.2)$$

$$W_{Y,t} = \rho W_{V,t} + \sqrt{1 - \rho^2} W_{\perp,t},$$

$$J_{Y,t} = \sum_{n=1}^{N_{Y,t}} Z_{Y,n}, \quad Z_{Y,n} \sim \mathcal{N}(\mu_Y; \sigma_Y^2)$$

$$J_{V,t} = \sum_{n=1}^{N_{V,t}} Z_{V,n}, \quad Z_{V,n} \sim \text{Exp}(\mu_V)$$

⁹Among others, Duffie et al. (2000), Eraker et al. (2003), Johannes et al. (2009) and Andersen et al. (2015b) use exponentially distributed variance jumps. Note that Bates (2000), Duffie et al. (2000), Pan (2002), Eraker et al. (2003) and Todorov and Tauchen (2011) provide evidence for the presence of positive jumps in volatility.

where $\{W_{V,t}\}_{t \geq 0}$ and $\{W_{\perp,t}\}_{t \geq 0}$ are two independent standard \mathbb{P} -Brownian motions and $\{\alpha_{t^-}^{\mathbb{P}}\}_{t \geq 0}$ is a predictable process such that

$$\alpha_{t^-}^{\mathbb{P}} = r - q + \left(\eta_Y - \frac{1}{2} \right) V_{t^-} + \left(\gamma_Y - \left(\varphi_{Z_Y}^{\mathbb{P}}(1) - 1 \right) \right) \lambda_{Y,t^-}. \quad (5.3)$$

The risk-free interest rate is given by r and the instantaneous dividend yield by q . The parameters η_Y and γ_Y capture the diffusive and jump risk premiums, respectively.¹⁰ The function $\varphi_{Z_Y}^{\mathbb{P}}(1)$ is the cumulant generating function of Z_Y evaluated at 1. The drift process $\alpha_{t^-}^{\mathbb{P}}$ is a by-product of the Radon-Nikodym derivative used in this study. We refer the interested reader to Subsection 5.B.1 for more details.

Log-equity jumps are generated by a Poisson process $\{N_{Y,t}\}_{t \geq 0}$ with stochastic intensity that depends on the instantaneous variance: $\lambda_{Y,t^-} = \lambda_{Y,0} + \lambda_{Y,1} V_{t^-}$. The size of these jumps are given by Gaussian random variables with mean μ_Y and standard deviation σ_Y .¹¹ Regarding jumps in volatility, these are governed by a Poisson process $\{N_{V,t}\}_{t \geq 0}$ that has a constant intensity $\lambda_{V,t^-} = \lambda_{V,0}$. Variance jump sizes $\{Z_{V,n}\}_{n=1}^{\infty}$ are given by independent exponentially distributed random variables with mean μ_V .

5.2.2 Links Between Theoretical and Empirical Quantities

We now analyze variables commonly used to conduct inference of models displaying latent characteristics, such as those embedded in Equations (5.1) and (5.2).

The first group of observables corresponds to time series of index prices, which contain information under the physical measure. In principle, high frequency prices are sufficient to extract all the information required for inferring latent components. In practice, however, parametric estimation from these prices would require the use of a very large sample set. The realized variance and bipower variation help reduce this dimensionality problem by keeping important information about the price and variance processes.

¹⁰Our Radon-Nikodym derivative explained in Subsection 5.A.1 includes four equivalent martingale measure coefficients: Λ_{Y,u^-} , Λ_{V,u^-} , Γ_Y and Γ_V . The process Λ_{Y,u^-} is $\eta_Y \sqrt{V_{u^-}}$, as in Heston (1993) among others. Λ_{V,u^-} is defined analogously. Moreover, γ_Y is a nontrivial function of Γ_Y . Even though η_V and Γ_V are not involved directly in our \mathbb{P} -measure modelling, these two parameters deal with the change of measure of the variance diffusive and jump components.

¹¹For this model, the jump convexity correction is given by $\varphi_{Z_Y}^{\mathbb{P}}(1) - 1 = \exp\left(\mu_Y + \frac{\sigma_Y^2}{2}\right) - 1$.

The second group is composed of options written on the S&P 500 index.¹² Option implied volatilities capture the wedge between the physical and the risk-neutral measure, which should provide more precise estimates of model parameters. In addition to this variable, we consider the variance of the option price as a complementary information source.

5.2.2.1 Index Prices as a First Source of Information

Log-prices constitute the fundamental source of information available under the physical measure:

$$\text{Observable \#1: } Y_t. \quad (5.4)$$

One of the interesting properties of these prices, as pointed out in Merton (1980), is that information about volatility can be obtained by arbitrarily increasing the sample frequency. Moreover, as discussed in nonparametric studies (see Barndorff-Nielsen and Shephard, 2004), this source not only provides information about the volatility process, but also about jumps in asset prices.

The fundamental variable that links intraday log-prices with price volatility is the quadratic variation (QV). This variable employs intraday log-prices over the interval $[0, t]$ in the following way:

$$QV_t = [Y, Y]_t = \lim_{N \rightarrow \infty} \sum_{j=1}^N \left(Y_{tj/N} - Y_{t(j-1)/N} \right)^2$$

where N is the number of elements in an equidistant grid dividing the interval $[0, t]$. Under the proposed jump-diffusion model, the quadratic variation is composed of the integrated variance and the log-price jump induced variation:

$$QV_t = \int_0^t V_{s-} ds + \sum_{n=1}^{N_{Y,t}} (Z_{Y,n})^2.$$

In the model considered in Equation (5.1), the quadratic variation depends explicitly on log-equity price jumps sizes $Z_{Y,n}$, the number of price jumps $N_{Y,t}$, and the path

¹²An interesting extension would be to consider the VIX index and its derivatives within a joint framework as the one proposed in Bardgett et al. (2015). We leave for future research the question of identifying the value added by these derivatives and choose to explore in detail what options written on the S&P 500 contribute in terms of information.

followed by the volatility process through $\int_0^t V_{s^-} ds$. The impact of jumps on the instantaneous variance is implicitly captured in the integrated variance over time. Over a time interval of length Δ (e.g. one day), the increments of the quadratic variation are given by:

$$\Delta QV_{t-\Delta,t} = QV_t - QV_{t-\Delta} = \int_{t-\Delta}^t V_{s^-} ds + \sum_{n=N_{Y,t-\Delta}+1}^{N_{Y,t}} (Z_{Y,n})^2. \quad (5.5)$$

As argued by Andersen et al. (2001), estimates of these increments can be obtained with the realized variance (RV), which is computed from intraday log-prices in the following way:

$$\text{Observable \#2: } RV_{t-\Delta,t} = \sum_{j=1}^N \left(Y_{t-\Delta+j\Delta/N} - Y_{t-\Delta+(j-1)\Delta/N} \right)^2. \quad (5.6)$$

In general, if N is large enough, $\Delta QV_{t-\Delta,t}$ and $RV_{t-\Delta,t}$ should be similar.¹³

Whereas the quadratic variation provides an overall measure of price volatility, the integrated variance, defined by

$$I_t = \int_0^t V_{s^-} ds$$

depends only on the diffusive component of the variance process and its jumps. Increments of the integrated variance over a time period of length Δ are given by:

$$\Delta I_{t-\Delta,t} = I_t - I_{t-\Delta} = \int_{t-\Delta}^t V_{s^-} ds, \quad (5.7)$$

which can be estimated using the realized bipower variation:

$$\begin{aligned} \text{Observable \#3: } BV_{t-\Delta,t} = \frac{\pi}{2} \sum_{j=2}^N & \left| Y_{t-\Delta+j\Delta/N} - Y_{t-\Delta+(j-1)\Delta/N} \right| \\ & \times \left| Y_{t-\Delta+(j-1)\Delta/N} - Y_{t-\Delta+(j-2)\Delta/N} \right|. \end{aligned} \quad (5.8)$$

In general, if N is large enough, $\Delta I_{t-\Delta,t}$ and $BV_{t-\Delta,t}$ should be similar.¹⁴ Thus, in order to identify log-price jump activity, it is necessary to combine the information contained in RV and BV .

¹³Protter (2004) shows that $RV_{t-\Delta,t}$ converges uniformly in probability to $\Delta QV_{t-\Delta,t}$ as the number of intraday observations increases.

¹⁴In a similar class of models, Barndorff-Nielsen and Shephard (2004) prove that $BV_{t-\Delta,t}$ converges to $\Delta I_{t-\Delta,t}$ as the number of intraday observations increases.

5.2.2.2 Option-Based Information

Option prices provide an additional source of information because they are sensitive to changes in price and volatility. Under usual conditions, the price of an option is given by

$$O_t = \mathbb{E}_t^{\mathbb{Q}} \left[e^{-r(T-t)} F(Y_T, K) \mid Y_t, V_t \right],$$

where $F(Y_T, K)$ is the payoff at time T , K the strike price of the option, and \mathbb{Q} a risk-neutral probability measure. Several authors highlight different advantages of adding these prices to the information set. First, these prices are conditional functions of stock returns, which allows researchers to estimate parameters governing the shape of these distributions (Dumas et al., 1998, Jackwerth and Rubinstein, 1996). Second, option prices help capture the wedge between the measures \mathbb{P} and \mathbb{Q} , thus providing information about volatility and jump risk premiums (Chernov and Ghysels, 2000, Christoffersen et al., 2012, Eraker, 2004, Pan, 2002, Santa-Clara and Yan, 2010). Third, option prices are highly informative about the instantaneous variance level (Broadie et al., 2007).¹⁵

Rather than using option prices directly, we employ the implied volatility (*IV*) resulting from the Black-Scholes formula:¹⁶

$$\text{Observable \#4: } \sigma_{t,i}^{BS}, t \in \{1, 2, \dots, n_t\}. \quad (5.9)$$

This change of variable does not only provide measures that are invariant to price levels, but also ones that can be characterized with OTM call and put options of different maturities (the so-called *IV* surface). We digress from the pure common calibration approach and use a time series of cross sections of options to capture the links between physical and risk-neutral parameters.

The literature in option pricing has relied almost exclusively on end-of-day option prices. With the availability of high frequency option prices, we now explore how to construct variables that capture information from fine grids. To this end, we extend the well-established concept of realized variance of log-prices to option prices and compute what we call the option realized variance (*ORV*).

¹⁵Within the proposed framework, semi closed-forms for option prices exist and are provided in Subsection 5.A.3.

¹⁶There is indeed a bijection between option prices and implied volatilities. See Renault (1997) for the benefits of using implied volatilities over option prices.

To understand ORV , we need to first look at the quadratic variation of the option price, as the former quantity corresponds to an approximation of the latter. In the proposed framework, we can characterize this variation for a European option as¹⁷

$$[O, O]_t = \int_0^t \left(\left(\frac{\partial O_u}{\partial y}(Y_{u-}, V_{u-}) \right)^2 + 2\sigma\rho \frac{\partial O_u}{\partial y}(Y_{u-}, V_{u-}) \frac{\partial O_u}{\partial v}(Y_{u-}, V_{u-}) + \sigma^2 \left(\frac{\partial O_u}{\partial v}(Y_{u-}, V_{u-}) \right)^2 \right) V_{u-} du + \sum_{0 < u \leq t} \{O_u(Y_u, V_u) - O_u(Y_{u-}, V_{u-})\}^2. \quad (5.10)$$

The last term of Equation (5.10) is associated with the jump variation contribution and can be rewritten as:

$$\begin{aligned} & \sum_{0 < u \leq t} \{O_u(Y_u, V_u) - O_u(Y_{u-}, V_{u-})\}^2 \\ &= \sum_{0 < u \leq t} \{O_u(Y_u, V_u) - O_u(Y_{u-}, V_{u-})\}^2 + \sum_{0 < u \leq t} \{O_u(Y_u, V_u) - O_u(Y_u, V_{u-})\}^2, \end{aligned}$$

where the first term relates to log-equity jumps and the second one to variance jumps.

The realized equivalent of the change in option quadratic variation

$$\Delta OQV_{t-\Delta, t} = [O, O]_t - [O, O]_{t-\Delta}$$

is the option realized variance. This variable is computed from intraday option prices according to the following expression:

$$\text{Observable \#5: } ORV_{t-\Delta, t, i} = \sum_{j=1}^N \left(O_{t-\Delta+j\Delta/N} - O_{t-\Delta+(j-1)\Delta/N} \right)^2, \quad t \in \{1, \dots, n_t\}. \quad (5.11)$$

Several properties of option realized variances can be characterized from Equation (5.10). Notice first that the option realized variance depends on the delta of the option through

$$\frac{\partial O_t}{\partial y} = \frac{\partial O_t}{\partial S} S$$

and the variance vega $\frac{\partial O_t}{\partial v}$. Accordingly, the moneyness of the option will determine which aspects of the data generating process drive the option realized variance. If we consider a deep in-the-money (ITM) call option, its delta is close to one and the

¹⁷Since the option price is a smooth function of Y and V , Itô's lemma can be applied to determine its quadratic variation. Details are available in the Subsection 5.B.2. Appendix B.2.2.1 provides a description on how to compute the derivatives used in Equation (5.10).

sensitivity to the instantaneous variance close to zero, so the option quadratic variation for this option will be approximately equal to:

$$[O, O]_t^{\text{ITM}} \cong \int_0^t \exp(2Y_{u-}) V_{u-} du + \sum_{0 < u \leq t} (\Delta O_u)^2.$$

In this case, we would expect the quadratic variance of this option to be “perfectly elastic” to the integrated variance and jump induced variation (in a day without jumps we should expect the option quadratic variance for such a derivative to be proportional to the integrated variance of the underlying).¹⁸ Therefore, information obtained using ITM option quadratic variance is somewhat redundant when one has already included realized variance as a source of information. If we consider an ATM option, the variance vega is at its highest value and delta should be close to $\frac{1}{2}$, so Equation (5.10) is approximated to:

$$[O, O]_t^{\text{ATM}} \cong \int_0^t \left(\frac{\exp(2Y_{u-})}{4} + \rho\sigma \frac{\partial O_u}{\partial v} + \sigma^2 \left(\frac{\partial O_u}{\partial v} \right)^2 \right) V_{u-} du + \sum_{0 < u \leq t} (\Delta O_u)^2.$$

For this level of moneyness, we would expect that the option quadratic variation responds more than proportionally to changes in the integrated variance. Finally, for deep OTM options, delta and variance vega should be close to zero, so most of the option quadratic variation comes from jump-induced variation. In this case we have:

$$[O, O]_t^{\text{OTM}} \cong \sum_{0 < u \leq t} (\Delta O_u)^2.$$

The behaviour of this type of option is not new to the finance literature; Andersen et al. (2015a) show nonparametrically that OTM put options behave like a pure-jump process. In this paper we show that the quadratic variation of OTM and deep OTM options captures this information, which helps to identify jump activity in variance and log-price processes by bringing an alternative variable to the existing list of observable covariates.

We end this section by summarizing how realized moments and other observed quantities can be combined to identify different risk sources. First, the difference between

¹⁸Empirically, we could look at this by running log-regressions of ORV over RV for different moneyness (controlling with a variable related to jumps, or excluding days with detected jumps). We should observe that coefficients for ATM are higher than those of ITM, and that the latter should be higher than those of OTM. Subsection 5.4.2 shows examples of these regressions using simulated data.

the realized variance and the bipower variation,

$$RV_{t-\Delta,t} - BV_{t-\Delta,t} \cong \sum_{n=N_{Y,t-\Delta}+1}^{N_{Y,t}} (Z_{Y,n})^2,$$

provides information about the presence and importance of log-equity price jumps; however, this difference does not give enough information to identify the sign of the jump. To this end, deep out-of-the-money (OTM) call (put) options are contracts that capture, to some extent, the sign of price jumps since their prices change quite drastically as soon as there is a substantial increase (decrease) in the underlying price.

Second, OTM option realized variations, combined with realized measures of log-equity prices, help to identify jumps in the volatility process. As discussed before, the OTM option realized variance can be approximated by

$$ORV_{t-\Delta,t}^{\text{OTM}} \cong \sum_{t-\Delta < u \leq t} (O_u(Y_u, V_u) - O_u(Y_{u-}, V_u))^2 + \sum_{t-\Delta < u \leq t} (O_u(Y_u, V_u) - O_u(Y_u, V_{u-}))^2.$$

In the absence of log-equity price jumps, $RV_{t-\Delta,t} - BV_{t-\Delta,t}$ is close to zero, so $ORV_{t-\Delta,t}^{\text{OTM}}$ becomes

$$ORV_{t-\Delta,t}^{\text{OTM}} \cong \sum_{t-\Delta < u \leq t} (O_u(Y_u, V_u) - O_u(Y_u, V_{u-}))^2.$$

The previous equation shows that $ORV_{t-\Delta,t}^{\text{OTM}}$ encompasses information about the presence of variance jumps.

5.3 Filtering and Estimation

This section provides details about the implementation of a SIR-type filter following Gordon et al. (1993). Our approach differs from the current literature in two ways. First, we propose a filter based on the SIR that employs realized variance, bipower variation and option realized variances as observable sources. Second, we show how to estimate the model using the filter, which requires special attention to the way the likelihood function is computed.

5.3.1 Filtering Algorithm

The filter is constructed over samples of daily observations¹⁹, so we define the elapsed time between two consecutive time steps by $\Delta = 1/252$ and denote the observed sample by $\{z_{k\Delta}\}_{k=1}^T$, where

$$z_{k\Delta} = \left[Y_{k\Delta}, RV_{(k-1)\Delta, k\Delta}, BV_{(k-1)\Delta, k\Delta}, \sigma_{k\Delta}^{BS}, \mathbf{ORV}_{(k-1)\Delta, k\Delta} \right],$$

$\sigma_{k\Delta}^{BS} = [\sigma_{k\Delta,1}^{BS}, \dots, \sigma_{k\Delta, n_{k\Delta}}^{BS}]$ being the implied volatility vector of the $n_{k\Delta}$ options considered that day and $\mathbf{ORV}_{(k-1)\Delta, k\Delta} = [ORV_{(k-1)\Delta, k\Delta, 1}, \dots, ORV_{(k-1)\Delta, k\Delta, n_{k\Delta}}]$ the corresponding vector of option realized variances.

As it is the case for any particle filter, the first step is to simulate particles for the latent variables. The interconnection between instantaneous variance and log-price processes requires the generation of the complete model. More precisely, assuming that the log-price $Y_{(k-1)\Delta}$ and its instantaneous variance $V_{(k-1)\Delta}$ are known at the end of day $k - 1$, the variables are simulated based on a time discretization of Equations (5.1) and (5.2). Intraday steps are required to capture the possibility of having multiple jumps in a single day and to ensure that the Euler discretization scheme is not too far away from the original model. Subsection 5.B.2.1 provides a detailed description of the simulation step.

The calculation of option realized variances requires the generation of intraday latent quantities. In this regard, an aggregation step is proposed. Specifically, we use the simulation method of Subsection 5.B.2.1 to generate paths on a fine grid (i.e. M intraday values per day). To obtain end-of-day quantities, we aggregate the M simulated values as explained in Subsection 5.B.2.2.

The aggregation produces daily simulated particles:

$$x_{k\Delta} = \left[Y_{k\Delta}, V_{k\Delta}, \Delta I_{(k-1)\Delta, k\Delta}, \Delta QV_{(k-1)\Delta, k\Delta}, \mathbf{IV}_{k\Delta}(Y_{k\Delta}, V_{k\Delta}), \Delta \mathbf{OQV}_{(k-1)\Delta, k\Delta} \right]$$

where $\Delta I_{(k-1)\Delta, k\Delta}$ is the integrated variance generated as a by-product of the simulation stage, $\Delta QV_{(k-1)\Delta, k\Delta}$ is the simulated quadratic variation derived from the integrated variance and the simulated jumps, $\mathbf{IV}_{k\Delta}(Y_{k\Delta}, V_{k\Delta})$ represents the vector of the $n_{k\Delta}$ model implied volatilities based on the simulated log-price and variance values,

¹⁹There is a vast literature on how many intraday data points should be used to avoid the microstructure noise effect and what procedures can be implemented to get rid of it. This will be discussed in the empirical study section. For now, we assume that realized variance, bipower variation and option realized variance have been measured adequately.

and $\Delta \mathbf{OQV}_{(k-1)\Delta, k\Delta}$ is the vector of the corresponding option quadratic variation calculated from Equation (5.10).

Further distributional assumptions about the measurement errors are required to connect the observed variables to the state variables. Since the realized variance is based on a finite sample of returns, it has not converged to its limit. Hence, the relative error in the realized variance follows²⁰

$$\frac{\Delta QV_{(k-1)\Delta, k\Delta} - RV_{(k-1)\Delta, k\Delta}}{RV_{(k-1)\Delta, k\Delta}} \sim \mathcal{N}(0, \eta_1^2). \quad (5.12)$$

Similarly, the relative error between bipower variation and integrated variance is assumed to be normally distributed:

$$\frac{\Delta I_{(k-1)\Delta, k\Delta} - BV_{(k-1)\Delta, k\Delta}}{BV_{(k-1)\Delta, k\Delta}} \sim \mathcal{N}(0, \eta_2^2). \quad (5.13)$$

The relative implied volatility error of the i^{th} option of the sample is assumed to have a Gaussian distribution:

$$\frac{IV_{k\Delta, i}(Y_{k\Delta}, V_{k\Delta}) - \sigma_{k\Delta, i}^{BS}}{\sigma_{k\Delta, i}^{BS}} \sim \mathcal{N}(0, \eta_3^2) \quad (5.14)$$

where $IV_{k\Delta, i}(Y_{k\Delta}, V_{k\Delta})$ is the model IV and $\sigma_{k\Delta, i}^{BS}$ is the market implied volatility.

Finally, as shown in Protter (2004), $ORV_{(k-1)\Delta, k\Delta, i}$ converges to $\Delta OQV_{(k-1)\Delta, k\Delta, i}$. Again the relative error is presumed to be Gaussian:

$$\frac{\Delta OQV_{(k-1)\Delta, k\Delta, i} - ORV_{(k-1)\Delta, k\Delta, i}}{ORV_{(k-1)\Delta, k\Delta, i}} \sim \mathcal{N}(0, \eta_4^2). \quad (5.15)$$

All these relative errors are assumed to be independent.

5.3.2 Estimation

In the SIR method, the proposal distribution depends on the most recent values of state variables. To ensure that the joint estimation is not dominated by one particular source, the likelihood associated with each observation receives a weight inversely

²⁰We are aware of Barndorff-Nielsen and Shephard's (2002) asymptotic results as the number M of intraday observations tends to infinity. However, the time discretization of our empirical implementation is too coarse to pretend that the asymptotic distribution has been reached. In fact, the goal here is to get an estimation procedure as efficient as possible and a large M makes the simulation step very time consuming. The Monte Carlo study shows that good precision is attained with quite small M .

proportional to the number of sources of a given type. That is, log-equity price, RV , and BV receive a weight of 1, and σ^{BS} and ORV a weight of one over $n_{k\Delta}$.²¹ Based on Assumptions (5.12), (5.13), (5.14) and (5.15), the weighted contribution to the likelihood function at time $k\Delta$ for particle $x_{k\Delta}$ is

$$\begin{aligned} f(z_{k\Delta} | x_{k\Delta}) = & \phi\left(Y_{k\Delta}; \mu_{k\Delta}, \sigma_{k\Delta}^2\right) \phi\left(\frac{\Delta QV_{(k-1)\Delta, k\Delta} - RV_{(k-1)\Delta, k\Delta}}{RV_{(k-1)\Delta, k\Delta}}; 0, \eta_1^2\right) \\ & \times \phi\left(\frac{\Delta I_{(k-1)\Delta, k\Delta} - BV_{(k-1)\Delta, k\Delta}}{BV_{(k-1)\Delta, k\Delta}}; 0, \eta_2^2\right) \left(\prod_{i=1}^{n_{k\Delta}} \phi\left(\frac{IV_{k\Delta, i}(Y_{k\Delta}, V_{k\Delta}) - \sigma_{k\Delta, i}^{BS}}{\sigma_{k\Delta, i}^{BS}}; 0, \eta_3^2\right)\right)^{1/n_{k\Delta}} \\ & \times \left(\prod_{i=1}^{n_{k\Delta}} \phi\left(\frac{\Delta OQV_{(k-1)\Delta, k\Delta, i} - ORV_{(k-1)\Delta, k\Delta, i}}{ORV_{(k-1)\Delta, k\Delta, i}}; 0, \eta_4^2\right)\right)^{1/n_{k\Delta}}, \end{aligned}$$

where $\phi(\cdot; m, s^2)$ is the density function of a Gaussian variable with mean m and variance s^2 . The conditional expectation $\mu_{k\Delta}$ and standard deviation $\sigma_{k\Delta}$ are defined in Subsection 5.B.1. Let

$$f(z_{k\Delta}) = \frac{1}{N_x} \sum_{x_{k\Delta}} f(z_{k\Delta} | x_{k\Delta})$$

be the average likelihood at time $k\Delta$ across particles and the log-likelihood function up to time $k\Delta$ to be defined recursively:

$$\mathcal{L}(z_{\Delta:k\Delta}) \propto \mathcal{L}(z_{\Delta:(k-1)\Delta}) + \log f(z_{k\Delta}).$$

As usual in the SIR filter, particles are resampled according to their weights before handling the information of the following day. More precisely, the path $x_{\Delta:(k+1)\Delta}$ is resampled proportionally to its likelihood, that is,²²

$$\omega(x_{\Delta:k\Delta}) \propto \omega(x_{\Delta:(k-1)\Delta}) f(z_{k\Delta} | x_{k\Delta}).$$

Although the particle filter approximation of the likelihood function at any point is asymptotically consistent in the number of particles, the log-likelihood function is not a continuous function of the parameters. Indeed, particle filters are known for being ill-suited for maximum likelihood estimation when using naive resampling methodologies. The likelihood function is not a continuous function of the model parameters and this could cause problems for gradient-based optimizers (e.g. Hürzeler and Künsch, 2001). To circumvent the issue, we use the continuous resampling of Malik and

²¹This idea is similar to the weighted likelihood estimator. Hu and Zidek (2002) study the properties of the weighted likelihood estimator and show that the key asymptotic results continue to hold.

²²Because of our particular resampling strategy, $\omega(x_{\Delta:(k-1)\Delta}) = 1/N_x$, N_x being the number of particles, and the weight reduces to $\omega(x_{k\Delta}) \propto f(z_{k\Delta} | x_{k\Delta})$.

Pitt (2011) which allows for continuous likelihood as a function of the unknown parameters under rather general conditions.

5.4 Simulation-Based Results

This section provides simulated-based results using the model proposed in Section 5.2. Section 5.4.1 studies the performance of the filter introduced in Section 5.3 with respect to the information provided by different sources. Section 5.4.2 analyzes the information conveyed in *ORV*. Finally, Section 5.4.3 conducts a robustness test on the approximation employed by the filter in order to compute *ORV* from intraday option prices.

Table 5.1 shows the parameters employed in the tests of this section. These parameters are qualitatively consistent with the ones estimated in the current literature, e.g. Eraker et al. (2003) and Eraker (2004), among others.^{23,24}

Table 5.1: **Model parameters used in the simulation study.**

Log-price process		Variance process		Standard deviations of error terms	
η_Y	0.500	η_V	-0.500	η_1	0.050
γ_Y	0.005	Γ_V	1.000	η_2	0.050
r	0.010	κ	4.000	η_3	0.050
q	0.010	θ	0.020	η_4	0.050
ρ	0.000	σ	0.500		
$\lambda_{Y,0}$	15.000	$\lambda_{V,0}$	15.000		
$\lambda_{Y,1}$	40.000	μ_V	0.030		
μ_Y	-0.020	V_0	0.020		
σ_Y	0.020				
Y_0	log(1000)				

The long-term expected variance $\theta = 0.02$ corresponds to about $\sqrt{0.02} = 14\%$ annualized volatility. The parameter ρ is constrained to zero to reduce the signal to noise ratio, making it more difficult to estimate the volatility. The average log-price jump intensity is 15.8%, meaning that about 16 jumps per year are expected. The log-equity price jumps are on average negative, with an average magnitude of -2% and standard deviation of 2% . The average variance jump is about 3% . The theoretical average diffusive risk premium is given by $\eta_Y\theta = 1\%$. The jump risk premium is about $\gamma_Y(\lambda_{Y,0} + \lambda_{Y,1}\theta) = 7.9\%$. Our results are robust to other choices of η_Y and

²³They are also consistent with the ones estimated in Subsection 5.6.1.

²⁴A few other specifications have been tested and results were qualitatively robust to these other parameters.

γ_Y . Consistent with most econometric studies, η_V is negative: the \mathbb{Q} -measure variance persistence is thus lower than its \mathbb{P} counterpart (3.75 vs. 4, respectively) and the long-term expected variance is higher under the risk-neutral measure (0.0213 vs. 0.02, respectively). Finally, since Γ_V is higher than zero, the average variance jump size under the \mathbb{Q} -measure is slightly higher (0.0309) than the one under the physical measure (0.03). Throughout the simulation experiments, we use 50,000 particles in the filter.

5.4.1 Filter's Performance as a Function of Information

We simulate 100 sample paths of the data generating process (DGP) using the Broadie and Kaya (2006) drift-interpolation scheme of van Haastrecht and Pelsser (2010). The sampling frequency is set at $1/1,560$ of a day (one sample every 15 seconds during a 6.5 hour trading day) and the length of a path is set to one year. Based on a path of the DGP, quantities such as integrated variance and realized variance are computed over a daily frequency. Additional to a path of the DGP, we also compute samples of intraday option prices. Our sample is composed of short-term options (30 days to maturity) with call-equivalent deltas of 0.20, 0.35, 0.50, 0.65, and 0.80, and a sample of long-term options (90 days to maturity) with same call-equivalent deltas. Intraday option prices are used to compute option realized variances for each day, while implied volatilities are deduced from end-of-the-day option prices. To take into account measurement errors, all end-of-day quantities include error terms as defined in Assumptions (5.12) to (5.15). These observations constitute our set of observables, from which the filter is run for different data aggregation periods of $M = 1, 2, 3, 5$, and 10. We regard $M = 15$ as the *true* filter that avoids discretization errors.

We start with a graphical illustration of how different sources of information impact estimates from the filter. Figure 5.1 compares the true instantaneous variance with five different densities for a randomly selected day during which a variance jump occurred. Each density corresponds to the posterior instantaneous variance density function obtained from running the filter with a specific set of variables (i.e. the information set). When only returns are employed in the filter, we observe that the filter lacks precision as values are highly dispersed on the left of the true value. When BV is included, the distribution exhibits a mode closer to the true value with a heavy tail on the right. However, the main mode still exhibits a downward bias. This shift in the distribution is evidence of the valuable information that BV adds about the presence of return jumps in the total variation. An interesting impact in the distribution is observed when option prices are included in the data set, since the data points cluster closer to the




Figure 5.1: Example of instantaneous variance densities using different data sources.

Y means daily log-equity value, RV means realized variance, BV means bipower variation, IV means implied volatility here, and ORV means option realized variance. The parameters of Table 5.1 are used. The vertical line represents the true instantaneous variance value. The density are constructed by using Gaussian kernels along with simulated particles obtained using the SIR-based methodology.

true variance value producing a distribution that is noticeably more peaked, that is, the variability of the estimator is reduced, but the downward bias is still present. Finally, the inclusion of option realized variances centers the distribution around the true value.

We now turn to a more detailed analysis of the filter performance with different information sources. This time we employ the root mean square error (RMSE) to compare the approximate filtered mean with the true filtered one. We also look at the performance of the filter by comparing the approximate filtered mean with the true simulated variable.

Table 5.2 summarizes the performance of the information source and the filter for different quantities of interest. Consider first the case of the instantaneous variance in Panel A. Notice that no matter which information set is employed, increasing M has little effect on the performance of the filter for this variable. This effect comes as no surprise due to the low discretization bias present at daily frequencies. On the other hand, we do observe large differences depending on the source of information. The most important case corresponds to a fourfold decrease when RV is combined with log-returns. We still observe improvements of about 10% each time other observable variables are included in the information set, which goes in line with the patterns

Table 5.2: **RMSE for instantaneous variance, integrated variance, quadratic variance, log-equity jump, and instantaneous variance jump across 100 simulated paths.**

Panel A: Instantaneous variance										
<i>M</i>	RMSE (true)					RMSE (filtered)				
	1	2	3	5	10	1	2	3	5	10
<i>Y</i>	45.372	45.308	45.295	45.402	45.213	44.572	44.500	44.487	44.590	44.403
<i>Y</i> and <i>RV</i>	13.0213	12.750	12.688	12.806	12.780	10.739	10.376	10.317	10.425	10.405
<i>Y, RV</i> and <i>BV</i>	9.422	9.246	9.233	9.186	9.182	6.033	5.743	5.680	5.568	5.596
<i>Y, RV, BV</i> and <i>IV</i>	8.5055	7.853	7.531	7.516	7.478	5.389	4.626	4.012	4.001	3.940
<i>Y, RV, BV, IV</i> and <i>ORV</i>	7.496	6.256	6.131	6.085	6.078	4.819	2.147	1.828	1.773	1.713
Panel B: Quadratic variation increment										
<i>M</i>	RMSE (true)					RMSE (filtered)				
	1	2	3	5	10	1	2	3	5	10
<i>Y</i>	0.3121	0.3116	0.3120	0.3125	0.3118	0.3082	0.3074	0.3079	0.3085	0.3077
<i>Y</i> and <i>RV</i>	0.0304	0.0304	0.0311	0.0313	0.0310	0.0249	0.0245	0.0254	0.0255	0.0256
<i>Y, RV</i> and <i>BV</i>	0.0278	0.0282	0.0283	0.0291	0.0286	0.0210	0.0203	0.0209	0.0218	0.0212
<i>Y, RV, BV</i> and <i>IV</i>	0.0287	0.0278	0.0278	0.0287	0.0279	0.0214	0.0202	0.0202	0.0216	0.0212
<i>Y, RV, BV, IV</i> and <i>ORV</i>	0.0292	0.0203	0.0202	0.0216	0.0208	0.0259	0.0147	0.0157	0.0164	0.0148
Panel C: Integrated variance increment										
<i>M</i>	RMSE (true)					RMSE (filtered)				
	1	2	3	5	10	1	2	3	5	10
<i>Y</i>	0.1753	0.1752	0.1752	0.1758	0.1750	0.1738	0.1737	0.1737	0.1742	0.1734
<i>Y</i> and <i>RV</i>	0.0363	0.0351	0.0350	0.0353	0.0351	0.0325	0.0311	0.0310	0.0314	0.0311
<i>Y, RV</i> and <i>BV</i>	0.0173	0.0169	0.0170	0.0170	0.0169	0.0096	0.0091	0.0090	0.0090	0.0090
<i>Y, RV, BV</i> and <i>IV</i>	0.0209	0.0162	0.0159	0.0159	0.0158	0.0142	0.0085	0.0079	0.0079	0.0078
<i>Y, RV, BV, IV</i> and <i>ORV</i>	0.0271	0.0129	0.0119	0.0116	0.0114	0.0234	0.0060	0.0044	0.0042	0.0040
Panel D: Log-equity jump										
<i>M</i>	RMSE (true)					RMSE (filtered)				
	1	2	3	5	10	1	2	3	5	10
<i>Y</i>	6.2694	6.2618	6.2613	6.2577	6.2576	5.9798	5.9723	5.9721	5.9718	5.9724
<i>Y</i> and <i>RV</i>	2.2770	2.2232	2.2296	2.2608	2.2431	1.5404	1.4689	1.4931	1.5001	1.5047
<i>Y, RV</i> and <i>BV</i>	1.8560	1.8554	1.8715	1.9002	1.8427	0.9038	0.8786	0.9379	0.9998	0.9654
<i>Y, RV, BV</i> and <i>IV</i>	1.9777	1.8463	1.8800	1.8820	1.8203	1.0732	0.8666	0.9402	0.9799	0.9765
<i>Y, RV, BV, IV</i> and <i>ORV</i>	3.9788	1.9680	1.8765	1.9214	1.8785	3.5345	1.0410	1.0325	0.9730	0.9531
Panel E: Instantaneous variance jump										
<i>M</i>	RMSE (true)					RMSE (filtered)				
	1	2	3	5	10	1	2	3	5	10
<i>Y</i>	9.8003	9.7880	9.7880	9.7862	9.7837	9.3859	9.3742	9.3746	9.3737	9.3705
<i>Y</i> and <i>RV</i>	9.8022	9.1877	9.1002	9.0421	9.0619	9.3878	8.7586	8.6670	8.6477	8.6606
<i>Y, RV</i> and <i>BV</i>	9.8147	8.6952	8.4181	7.9956	7.9411	9.4020	8.2883	7.9502	7.5375	7.4418
<i>Y, RV, BV</i> and <i>IV</i>	9.0258	6.2434	5.8524	5.7860	5.7234	8.5814	5.7711	5.3274	5.2712	5.2238
<i>Y, RV, BV, IV</i> and <i>ORV</i>	4.5677	2.0562	1.9963	1.9173	1.9068	4.2681	0.9991	0.9616	0.9880	0.9240

Y means daily log-equity, *RV* means realized variance, *BV* means bipower variation, *IV* means implied volatility, and *ORV* means option realized variance. Quantities were multiplied by 1,000. Filtered values are computed as the mean of resampled particles obtained via a SIR particle filter with $M = 15$ and using *Y, RV, BV, IV* and *ORV*.

observed in Figure 5.1. This result complements the findings of Johannes et al. (2009), in which the authors report similar gains when daily option prices are included in addition to asset returns. The large improvement offered by RV is also observed when filtering other measures of variance such as the quadratic variation (Panel B) and the integrated variance (Panel C).

Concerning the performance of jump estimation, Panel D shows results regarding log-equity jump sizes and Panel E does so for volatility jumps. The results indicate that data augmentation benefits the identification of log-equity jump sizes when RV is added to the information set, but it has little effect when other sources are employed. Most of the gains in RMSE come from the inclusion of RV and to some extent from BV . There is little or no benefit when option prices and $ORVs$ are added. These results contrast with those of volatility jumps, where there is little benefit from including RV in the filter and the largest RMSE gains are observed when ORV is included, yielding an almost threefold decrease in RMSE for $M = 10$. Notice also how data augmentation helps reduce RMSE for volatility jump size estimation, which shows how difficult it is to estimate jumps in volatility at the daily level even if parameters are known. If we look at the sum of the RMSE for both types of jumps, we observe that the information content of ORV is very useful for volatility jump estimation, even with no data aggregation ($M = 1$).

Previous results about jump size estimation are complemented with jump detection results from Table 5.3. The focus of this set of results is on the ability of a given source to identify the occurrence of jumps. A jump is detected when the filtered jump probability is greater than 50%. This statistic, when compared with a true simulated jump, provides an idea of the general performance of the filter. On the other hand, when the statistic is compared with the filtered jump (a jump detected with the filter using $M = 15$), it quantifies the ability of a given source to identify the occurrence of a jump despite the discretization bias. According to Panel A, the log-equity price jump times are filtered adequately when log-equity values and realized variance are included, with a level of concordance of about 96.9% (98.7% with the discretization bias). This level is slightly improved when more sources are employed. As observed with jump sizes, the detection of volatility jumps is improved when ORV are employed, with concordance levels of 97.3% (99.7% with the discretization bias).

Table 5.3: Average jump times in percentages.

Panel A: Average log-equity jump times										
<i>M</i>	Hit (true)					Hit (filtered)				
	1	2	3	5	10	1	2	3	5	10
<i>Y</i>	93.61	93.63	93.64	93.65	93.65	95.30	95.33	95.34	95.35	95.35
<i>Y</i> and <i>RV</i>	96.85	96.97	96.97	96.97	96.96	98.65	98.77	98.78	98.76	98.76
<i>Y</i> , <i>RV</i> and <i>BV</i>	97.90	97.90	97.90	97.89	97.86	99.77	99.76	99.77	99.76	99.74
<i>Y</i> , <i>RV</i> , <i>BV</i> and <i>IV</i>	97.94	97.92	97.92	97.92	97.89	99.78	99.78	99.77	99.78	99.77
<i>Y</i> , <i>RV</i> , <i>BV</i> , <i>IV</i> and <i>ORV</i>	98.29	98.00	97.95	97.95	97.91	99.87	99.85	99.81	99.82	99.82
Panel B: Average log-equity jump times										
<i>M</i>	Hit (true)					Hit (filtered)				
	1	2	3	5	10	1	2	3	5	10
<i>Y</i>	94.08	94.08	94.08	94.08	94.08	96.53	96.53	96.53	96.53	96.53
<i>Y</i> and <i>RV</i>	94.08	94.15	94.29	94.41	94.48	96.53	96.60	96.73	96.88	96.94
<i>Y</i> , <i>RV</i> and <i>BV</i>	94.08	94.47	94.79	95.03	95.12	96.53	96.91	97.24	97.48	97.58
<i>Y</i> , <i>RV</i> , <i>BV</i> and <i>IV</i>	94.31	95.21	95.33	95.37	95.38	96.76	97.67	97.79	97.83	97.85
<i>Y</i> , <i>RV</i> , <i>BV</i> , <i>IV</i> and <i>ORV</i>	97.60	97.33	97.29	97.31	97.33	99.74	99.80	99.78	99.80	99.79

In this table, we show the average number of times a jump has been adequately filtered (i.e. whether the probability of having a jump on a given day is higher than 0.5). Hit reveals the proportion of time that jumps are adequately filtered. *Y* means daily log-equity, *RV* means realized variance, *BV* means bipower variation, *IV* means implied volatility, and *ORV* means option realized variance. Filtered values are computed as the mean of resampled particles obtained via a SIR particle filter with $M = 15$ and using *Y*, *RV*, *BV*, *IV* and *ORV*.

5.4.2 Information Embedded in *ORV*

The previous tests highlight the benefit of using *ORV* as a source for disentangling jumps and volatility in the filtering process. The tests conducted in this section study the link between *ORV* and information contained in *RV* and jumps.

The first experiment consists of simulating 1,000 paths of one day for which there are no jumps in the DGP. Each path is generated at a frequency of 1/1,560 of a day using the Broadie and Kaya (2006) drift-interpolation scheme of van Haastrecht and Pelsser (2010). Each time a path is simulated, the parameters of a 30-day European call option are simulated so that *ORV* can be computed along the path. The option's delta is uniformly simulated between 0.1 and 0.9 and the initial spot variance between 0.01 and 0.10.

The simulated data produces daily values of *RV* and *ORV* for a random sample of options, which allows us to measure the degree of redundancy between these two variables with the following regression:

$$ORV_i = \gamma_1 \left(\frac{\partial O_i}{\partial y} \right)^2 RV_i + \gamma_2 \left(2 \frac{\partial O_i}{\partial y} \frac{\partial O_i}{\partial v} \right) RV_i + \gamma_3 \left(\frac{\partial O_i}{\partial v} \right)^2 RV_i + \varepsilon_i. \quad (5.16)$$

This regression model comes from fixing the option derivatives in Equation (5.10)

to their beginning-of-the-day values, so that the impacts of changes in RV are only modulated by option parameters from sample to sample.

Table 5.4: **Information content of ORV when jumps are absent.**

	Coefficient	Standard error
γ_1	1.0136	0.0156
γ_2	-0.0234	0.0674
γ_3	0.4356	0.2552
R^2	0.9891	

The following regression is applied to ORV :

$$ORV_i = \gamma_1 \left(\frac{\partial O_i}{\partial y} \right)^2 RV_i + \gamma_2 \left(2 \frac{\partial O_i}{\partial y} \frac{\partial O_i}{\partial v} \right) RV_i + \gamma_3 \left(\frac{\partial O_i}{\partial v} \right)^2 RV_i + \varepsilon_i.$$

The derivatives are computed based on beginning-of-the-day information. Values in bold are statistically different (at a confidence level of 95%) from their theoretical values (i.e. 1, 0, and 0.25 respectively). We use a simulated sample of 1000 observations.

Table 5.4 reports estimate results of the regression. The R -squared close to one reveals that, in the absence of jumps, ORV is redundant with respect to the information embedded in RV . That is, ORV only amplifies or attenuates values of RV depending on the parameters of the option. Regarding regression estimates, we observe that the individual hypothesis that $\gamma_1 = 1$, $\gamma_2 = 2\sigma\rho = 0$, and $\gamma_3 = \sigma^2 = 0.25$ cannot be rejected with a confidence level of 95%, suggesting that this specification follows Equation (5.10) closely with the assumption of constant option derivatives. This last remark supports the sampling technique used in the filter to compute particles of ORV with few intraday samples, as option derivatives vary little between days.

We now proceed to analyze the impact that jumps in the underlying process might have on the information content of ORV . The experiment is run this time by taking into account the specificity of the option contract. To do this, we fix in advance the option type (call or put) and the call-equivalent delta (0.20, 0.35, 0.50, 0.65, and 0.8), and then the daily ORV is computed from the path generated by the simulation algorithm. The following regression is run separately for each type of contract and moneyness:

$$\log(ORV_i) = \beta_0 + \beta_1 \log(RV_i) + \beta_2 \Delta N_{Y,i} + \beta_3 \Delta N_{V,i} + \varepsilon_i \quad (5.17)$$

where $\Delta N_{Y,i}$ and $\Delta N_{V,i}$ are variables that capture the number of log-equity price jumps and variance jumps during day i , respectively.

Table 5.5 presents two specifications of the previous regression model. The first one includes information about log-return jumps and the second one adds to this specification information about volatility jumps. Taking R -squared as a measure of redundancy,

we observe that RV and ORV are more redundant when the option is in-the-money (delta higher than 0.5 for calls and lower than 0.5 for puts). As the option becomes out-of-the-money, redundancy decreases. A closer look at the coefficients in this regression show the sensitiveness of ORV to different types of information. Whereas the coefficient associated to RV increases with the moneyness of the option, those associated with jump activity do so (in absolute value) when the moneyness decreases. These two pieces of evidence show that the ORV of OTM options provides complementary information not contained in RV .

Table 5.5: Information content of ORV when jumps are present.

Δ^e	Call					Put				
	(1)		(2)		(1)		(2)			
	Coefficient	SE	Coefficient	SE	Coefficient	SE	Coefficient	SE		
0.20	β_0	13.945	(0.383)	11.161	(0.327)	β_0	14.149	(0.080)	13.811	(0.080)
	β_1	1.333	(0.041)	1.045	(0.035)	β_1	1.079	(0.008)	1.044	(0.008)
	β_2	-1.064	(0.086)	-0.582	(0.071)	β_2	-0.111	(0.018)	-0.052	(0.017)
	β_3			1.067	(0.045)	β_3			0.129	(0.011)
	R^2	0.627		0.763		R^2	0.978		0.981	
0.35	β_0	13.796	(0.285)	11.805	(0.249)	β_0	14.757	(0.138)	14.018	(0.132)
	β_1	1.199	(0.030)	0.994	(0.026)	β_1	1.187	(0.015)	1.111	(0.014)
	β_2	-0.726	(0.064)	-0.382	(0.054)	β_2	-0.265	(0.031)	-0.137	(0.029)
	β_3			0.763	(0.034)	β_3			0.283	(0.018)
	R^2	0.742		0.829		R^2	0.943		0.954	
0.50	β_0	13.681	(0.215)	12.278	(0.194)	β_0	15.303	(0.192)	14.177	(0.179)
	β_1	1.119	(0.023)	0.974	(0.020)	β_1	1.297	(0.020)	1.181	(0.019)
	β_2	-0.497	(0.048)	-0.255	(0.042)	β_2	-0.426	(0.043)	-0.231	(0.039)
	β_3			0.538	(0.026)	β_3			0.432	(0.024)
	R^2	0.833		0.882		R^2	0.903		0.926	
0.65	β_0	13.548	(0.141)	12.740	(0.133)	β_0	15.910	(0.258)	14.272	(0.235)
	β_1	1.048	(0.015)	0.964	(0.014)	β_1	1.448	(0.027)	1.278	(0.025)
	β_2	-0.282	(0.031)	-0.142	(0.029)	β_2	-0.648	(0.058)	-0.364	(0.051)
	β_3			0.310	(0.018)	β_3			0.628	(0.032)
	R^2	0.922		0.940		R^2	0.853		0.894	
0.80	β_0	13.470	(0.083)	13.079	(0.082)	β_0	16.273	(0.323)	14.136	(0.288)
	β_1	1.004	(0.009)	0.964	(0.009)	β_1	1.596	(0.034)	1.375	(0.030)
	β_2	-0.131	(0.019)	-0.064	(0.018)	β_2	-0.883	(0.072)	-0.514	(0.063)
	β_3			0.150	(0.011)	β_3			0.819	(0.039)
	R^2	0.972		0.976		R^2	0.806		0.865	

The following regression is applied to ORV :

$$\log(ORV_i) = \beta_0 + \beta_1 \log(RV_i) + \beta_2 \Delta N_{Y,i} + \beta_3 \Delta N_{V,i} + \varepsilon_i$$

where $\Delta N_{Y,i}$ is the number of log-equity price jumps during day i and $\Delta N_{V,i}$ is the number of variance jumps during day i . We estimate two versions of Equation (5.17): Regression (1) we force $\beta_3 = 0$, and Regression (2) we let β_3 be different than zero. R^2 and Newey and West (1987) standard errors (SE, in parentheses) are reported. All coefficients estimated are statistically significant at a confidence level of 95%.

5.4.3 Option Quadratic Variation Increments Approximation

One of the key features of the filter proposed in Section 5.3 is that very few number of intraday prices are required to compute ORV . We test this approximation by studying the impact of M on the computation of ORV .

Using the parameters given in Table 5.1, we simulate 500 paths at a frequency of $1/1,560$ over a day and compute the option realized variance for a call along each path. Next, the simulated ORV is compared to approximations of this quantity, denoted by ΔOQV , which are computed with a lower number of intraday observations $M \in \{1; 2; 3; 5; 10; 1,560\}$. The maturity of the options used in this exercise is 30 days. As a measure of the quality of the approximation, we compute the RMSE and relative RMSE that result from comparing the ORV computed at the highest frequency with that of the approximation. We also regress ΔOQV on ORV and employ the R -squared of this regression as a measure of the quality of the fit.²⁵

Table 5.6 shows the root mean square errors (RMSE), relative RMSE, and regression R -squareds for different levels of moneyness. As expected, we find that errors decrease as M becomes larger and that with as few as $M = 5$ observations per day the approximation provides satisfactory results. Note that the presence of jumps has an apparent impact on the performance of the approximation, which is not surprising as these rare events introduce more variation in the estimation.

To visualize the performance of the approximation, Figure 5.2 plots ORV and their approximations ΔOQV for days without (top panels) and with jumps (bottom panels), respectively. If ΔOQV constitutes a good approximation of ORV , all the points should be aligned on the diagonal, as it is the case for $M = 1,560$. Note that ΔOQV approaches to ORV for values of $M = 3$ and higher.

5.5 Exploring Option Realized Variances Empirically

In this section, we first present our datasets. Next, we analyze nonparametrically the information contained in option realized variances. This is done by constructing surfaces of these variations. We complement this investigation with a principal component analysis (PCA) of option realized variances. Finally, we provide evidence of the

²⁵We run the following regression: $\Delta OQV_i = \beta_0 + \beta_1 ORV_i + \varepsilon_i$ and compute the R -squared associated with it. The higher the R -squared, the better the quality of the approximation.

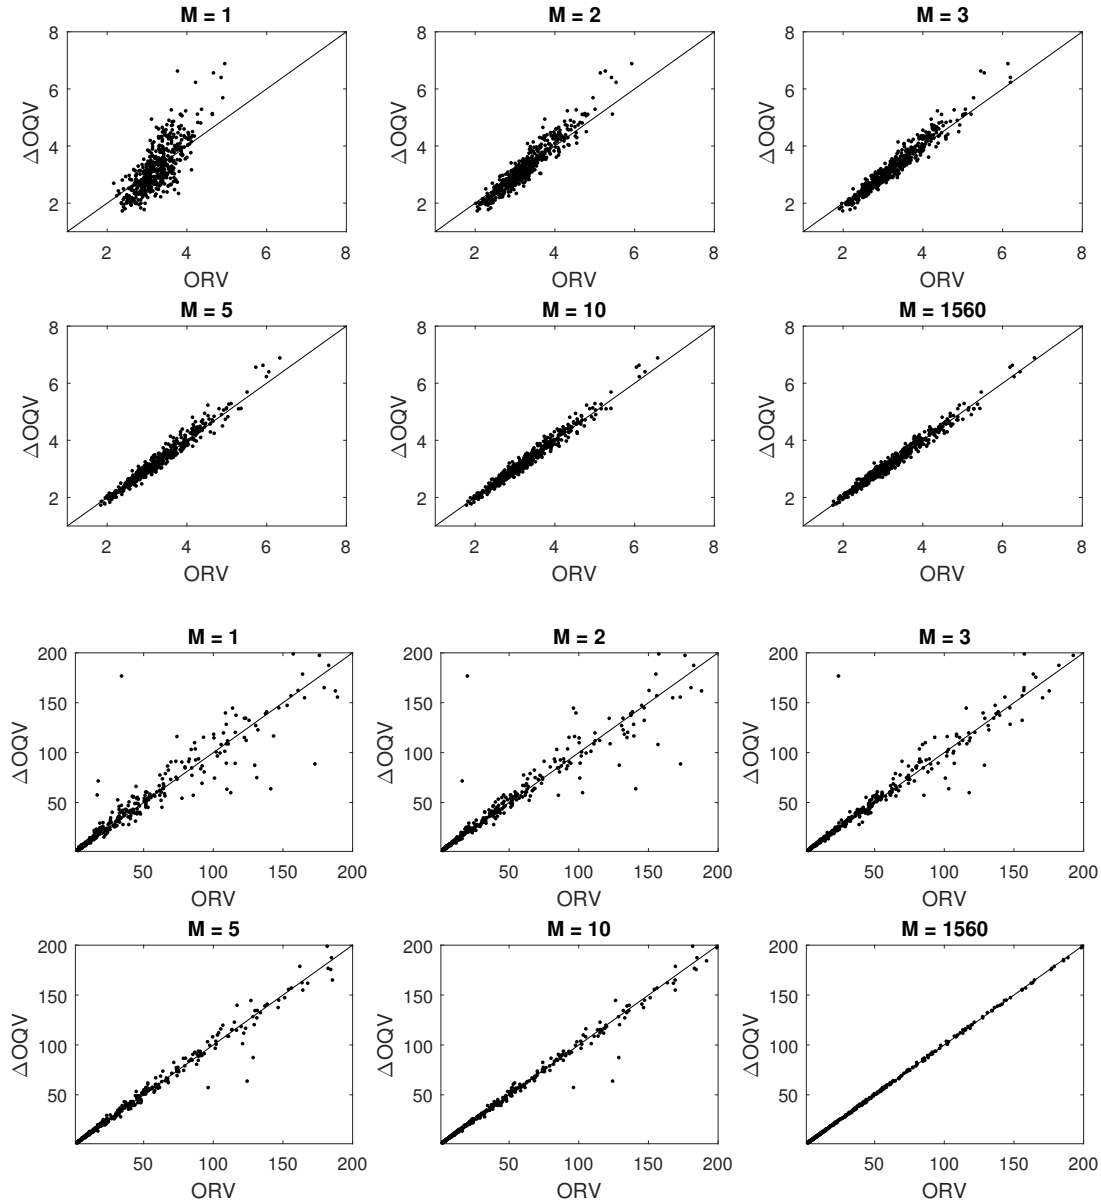


Figure 5.2: Option realized variance against option quadratic variance increments for a call-equivalent delta of 0.50 and a maturity of 30 business days, on days without (top panels) and without jumps (bottom panels).

This figure is generated using 500 simulated paths of one day and a European call option that has a call-equivalent delta of 0.50 and a maturity of 30 business days. These paths contain no jump. *ORVs* are compared to six different values for which various M are used (i.e 1, 2, 3, 5, 10 and 1,560). Parameters of Table 5.1 are used.

Table 5.6: **RMSE, relative RMSE and regression R -squared for ΔOQV approximation across 500 days.**

Panel A: RMSE										
Δ^e	Days without jumps					Days with jumps				
	0.2	0.35	0.5	0.65	0.8	0.2	0.35	0.5	0.65	0.8
1	0.867	1.851	2.794	1.295	0.563	21.925	35.893	55.659	62.941	30.808
2	0.458	1.006	1.599	0.713	0.295	16.286	28.140	45.653	58.119	28.371
3	0.311	0.711	1.191	0.560	0.220	14.131	22.775	32.613	37.994	20.383
5	0.232	0.559	0.987	0.449	0.169	11.956	15.303	20.103	20.851	12.772
10	0.172	0.462	0.855	0.381	0.137	8.960	13.258	18.675	19.392	11.282
1560	0.146	0.433	0.806	0.348	0.126	0.725	1.714	2.924	1.890	0.787
Panel B: Relative RMSE										
Δ^e	Days without jumps					Days with jumps				
	0.2	0.35	0.5	0.65	0.8	0.2	0.35	0.5	0.65	0.8
1	0.205	0.151	0.121	0.141	0.175	0.289	0.240	0.205	0.176	0.199
2	0.109	0.081	0.068	0.076	0.089	0.209	0.179	0.158	0.144	0.159
3	0.077	0.058	0.052	0.059	0.065	0.141	0.118	0.103	0.108	0.120
5	0.057	0.046	0.043	0.046	0.049	0.111	0.091	0.077	0.077	0.089
10	0.043	0.039	0.038	0.040	0.040	0.091	0.079	0.069	0.070	0.078
1560	0.038	0.038	0.036	0.037	0.038	0.026	0.027	0.027	0.025	0.024
Panel C: Regression R-squareds										
Δ^e	Days without jumps					Days with jumps				
	0.2	0.35	0.5	0.65	0.8	0.2	0.35	0.5	0.65	0.8
1	0.439	0.475	0.532	0.585	0.573	0.915	0.892	0.864	0.916	0.918
2	0.867	0.861	0.855	0.889	0.899	0.953	0.935	0.910	0.926	0.928
3	0.944	0.933	0.922	0.930	0.941	0.964	0.957	0.952	0.960	0.956
5	0.967	0.957	0.946	0.954	0.963	0.977	0.981	0.981	0.988	0.981
10	0.978	0.968	0.958	0.965	0.973	0.987	0.985	0.984	0.989	0.985
1560	0.981	0.970	0.960	0.969	0.976	1.000	1.000	1.000	1.000	1.000

The real ORV value is computed using option prices at a frequency of $1/1,560$. The error corresponds to the difference between ΔOQV and ORV . 0.20, 0.35, 0.50, 0.65 and 0.80 correspond to the different call-equivalent deltas considered. The OTM option maturity is 30 business days. $ORVs$ are compared to six different values for which various M are used (i.e. 1, 2, 3, 5, 10 and 1,560). To compute the R -squared, we run the following regression: $\Delta OQV_i = \beta_0 + \beta_1 ORV_i + \varepsilon_i$.

economic relationship between option realized variances and index return variability using predictive regressions.

5.5.1 Data

The sample period employed in this study is from July 2004 to December 2012. This period offers several empirical features such as periods of high economic uncertainty, which makes it appealing for studying different empirical features of equity index returns and volatility.

We start by constructing a time series measure of the option realized variance at the daily level. We employ tick-by-tick Level I quote data provided by Tick Data for European options written on the S&P 500 index. Tick Data prices come from the Options Price Reporting Authority (OPRA), the national market system that provides information about last sale reports and quotation information. We employ midquote prices instead of trade prices to mitigate the effect of bid-ask spread bounces in the total variation of the option price. For each trading day, we start with 390 one-minute prices from 9:30 AM to 4:00 PM. Next, we construct five different grids with five-minute prices and compute the daily variation according to Equation (5.11) over each grid. The average of these five values provides an estimate of the option daily variation. This procedure helps to mitigate the presence of microstructure noise, as suggested by Zhang et al. (2005) for the case of realized volatility estimation. For a quote to be included in the dataset, we require its bid price to be higher than zero and the quote not to have any condition code or be eligible for automatic execution.^{26,27} Each day, we restrict our attention to a representative sample set composed of OTM and ATM options with maturities closest to 30 and 90 business days. As was argued in Section 5.2.2, *ORVs* for ITM options yield equivalent information to that of *RV*, so we exclude these options from our sample. Only options with positive volume and bid prices are included in the sample, as well as those satisfying the no-arbitrage conditions of Bakshi et al. (1997).

Additional to the option's price variation, we compute several variables that capture different aspects of market activity for each day in the sample. The first is the daily Black-Scholes implied volatility (*IV*), which is computed for the option with a forward-to-strike ratio closest to 1 (ATM option) and a maturity closest to 30 business days. The second is the realized variance *RV* of index returns, constructed from one-minute returns of the E-mini S&P futures contract prices. Finally, the third variable is the bipower variation *BV*, which accounts for the variability of the diffusive component governing the return process. These last two variables are constructed with five sub-grids, following Zhang et al. (2005) sub-sampling methodology. Descriptive statistics about these variables and the panel of option realized variances are provided in the Appendix B.4.

Figure 5.3 presents the time series associated with some of these variables. It shows in the first three panels the daily time series of the previous three variables and in the

²⁶Our final *ORV* dataset contains 682,380 data points, for an average of 286 option realized variance per day.

²⁷Errors in the option quote dataset could artificially increase the computed *ORV*. To discard such outliers, the last permille (0.1%) of the *ORV* sample (i.e. largest values) is removed.

last panel the option realized variance time series for the ATM option with a maturity of 30 business days. We observe that all series increase during the financial crisis period and exhibit sporadic spikes across the sample following the flash crash episode (May 6, 2010) and the downgrade of U.S. debt (August 5, 2011). This commonality shows the degree of interdependence across different markets in response to important events. We now proceed to explore in more detail the full panel of option realized variances.

5.5.2 Understanding Option Realized Variances

The surfaces (indexed by time) constitute a parsimonious way to study and understand the behaviour of option prices, as shown in Andersen et al. (2015b). In opposition to previous authors who studied implied volatilities, we analyze the behaviour of option realized variances. To interpret the large cross section of option data that spans distinct maturities and moneynesses, we construct a sequence of option realized variance surfaces across these two dimensions. For a given day, we collect the *ORVs* for all the ATM and OTM options in our dataset and perform a locally weighted scatter plot smoothing across moneynesses and maturities.²⁸ Figure 5.4 shows the surface induced by option realized variances on July 6, 2004. As it is clear from the figure, the surface has an inverted U-shape, in which the highest values are observed for options that are at-the-money and the lowest for those that are out-of-the-money.

Following Andersen et al. (2015b), Figure 5.5 presents specific *ORV* surface characteristics such as its level, term structure, skew, and skew term structure across time. The *ORV* level comes from the ATM option ($\Delta^e = 0.5$) with 30 days to maturity. The *ORV* term structure (TS) is defined as the difference between the *ORV* of the ATM with 90 days to maturity minus the *ORV* of the ATM option with 30 days to maturity. The *ORV* skew represents the difference between shorter dated OTM put options (i.e. $\Delta^e = 0.9$) and OTM call options (i.e. $\Delta^e = 0.1$), both with 30 days to maturity. Finally, the *ORV* skew term structure (Skew TS) is the difference between longer and shorter dated skew, with the longer dated (i.e. 90 business days) skew defined analogously to the shorter one.

The surface level presents sporadic spikes that generate some persistence after their occurrence (top-left panel of Figure 5.5). ATM options exhibit realized variations that

²⁸The locally smoothing quadratic regression is performed using Matlab procedure Lowess on $\log(ORV)$, and then transformed back to *ORV*. The Matlab procedure performs a local regression using weighted linear least squares with a second order polynomial model.




Figure 5.3: Realized volatility, square root of the bipower variation, ATM option implied volatility for options with maturity closest to 30 business days and ATM option realized volatility (square root of ORV) for options with maturity closest to 30 business days.

Annualized realized variance and bipower variation are computed from intraday S&P futures prices. Since futures are quarterly contracts, we build a time series from these data by rolling contracts two weeks before expiration. We follow Zhang et al. (2005) and compute a microstructure-noise robust estimate of the daily realized variance as the average of RV and BV estimates based on different subsets of prices. To compute the daily realized variance of an option, we employ tick-by-tick Level I quote data from options provided by Tick Data. From the data, we construct one-minute midquote series and compute the daily variation according to Equation (5.11). We again follow Zhang et al. (2005) and compute a microstructure-noise robust estimate of the daily option realized variance. Each time series is displayed from July 2004 to December 2012.




Figure 5.4: **Option realized variances and fitted surface as a function of both the option Black and Scholes (1973) call-equivalent delta and business days to maturity.**

Option realized variances are computed using the daily variation according to Equation (5.11). The example is given for July 6, 2004. The fitted surface uses the locally smoothing quadratic regression method.

mimic the behaviour of the underlying realized variance (top-left and bottom-middle panels of Figure 5.5, as well as Table 5.7 showing that the correlation coefficient between the level and RV is 92%). Consequently, ATM option realized variance brings little new information when RV or BV are part of the sample. The sporadic spikes are also observed at the same times in the other characteristics of the surface (top-middle, top-right and bottom-left panels of Figure 5.5), showing a level of commonality associated with large shocks. Since the other surface characteristics present no sign of persistence and fluctuate around zero, this behaviour reveals that these characteristics are particularly elevated during turbulent market periods. More precisely, ORV term structure (top-middle panel of Figure 5.5) shows that shocks are more important for short-dated options. Its low correlation with RV and $JV = \max(RV - BV, 0)$ (see Table 5.7) indicates that the ORV term structure is a non-redundant information source. Regarding the ORV skew (top-right panel of Figure 5.5), this characteristic generally takes positive values, which means that OTM puts are more responsive to shocks than OTM calls. This type of responsiveness is observed more for short dated options than for longer dated ones, as evidenced from the negative sign associated with values of the skew term structure (bottom-left panel of Figure 5.5).

Figure 5.5 also presents realized volatility and realized jump variation, defined as $JV = \max(RV - BV, 0)$, of index returns. Realized jump variation captures the variability induced by discontinuous activity in the index return, which explains its erratic

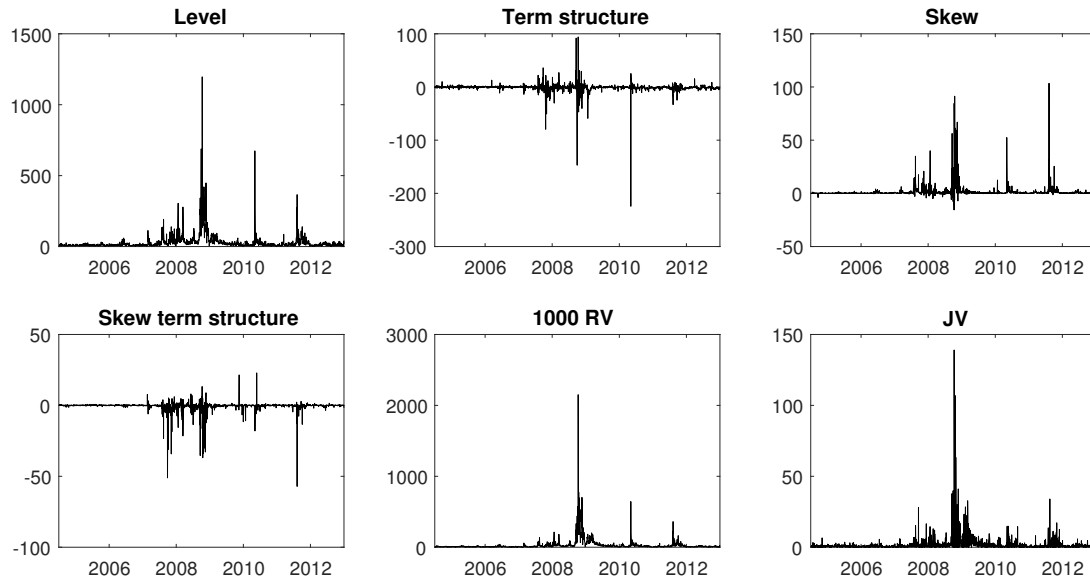


Figure 5.5: Option realized variance surface characteristics and realized measures.

The option realized variance level is the option realized volatility for ATM options (i.e. $\Delta^e = 0.5$ and 30 business days). The option realized volatility term structure is the difference between long- and short-dated ATM options (i.e. 90 and 30 business days in our case). The option realized variance skew is the difference between short-dated OTM put options (i.e. $\Delta^e = 0.9$, 30 business days) and OTM call options (i.e. $\Delta^e = 0.1$, 30 business days). The option realized volatility skew term structure is the difference between long- and short-dated skew, with the long-dated (i.e. 90 business days) skew defined analogously to the short one. The realized volatility and the bipower variation are computed using Zhang et al.'s (2005) microstructure-noise robust estimate. *RV* stands for realized variance and *JV* for realized jump variation.

spikes and less persistent behaviour. Increases in different characteristics of the *ORV* surface are associated with shocks to realized variance and jump activity of the option's underlying asset (Figure 5.5 and Table 5.7).

We now employ the principal component analysis (PCA) of the *ORV* surface to look more closely at different aspects of the commonality among *ORV* surface characteristics, realized variance, and jump activity of the S&P index. We extract from the *ORV* surface 18 values per day and conduct a PCA over these values for the full sample. More precisely, we take nine equally spaced points over the call-equivalent delta dimension between 0.1 and 0.9 for maturities of 30 and 90 business days.

Figure 5.6 shows the first six in-sample principal components (PCs) of the option realized variance surface. Similar to what is observed for the implied volatility surface in Andersen et al. (2015b), the *ORV* surface displays a dominant level type effect, as the first PC accounts for 94.44% of the total variation and displays a high degree of commonality with the surface level (top-left panel of Figure 5.6). The second PC

Table 5.7: **Correlation matrix for option realized variance characteristics and realized measures.**

	Level	TS	Skew	Skew TS	RV	JV
Level	1.000					
TS	-0.260	1.000				
Skew	0.630	0.049	1.000			
Skew TS	-0.384	-0.098	-0.763	1.000		
RV	0.920	-0.099	0.675	-0.390	1.000	
JV	0.626	-0.008	0.423	-0.217	0.703	1.000

The option realized variance level is the option realized variance for ATM options (i.e. $\Delta^e = 0.5$). The option realized variance term structure (TS) is the difference between longer and shorter dated ATM options (i.e. 90 and 30 business days in our case). The option realized variance skew is the difference between shorter dated OTM put options (i.e. $\Delta^e = 0.9$, 30 business days) and OTM call options (i.e. $\Delta^e = 0.1$, 30 business days). The option realized variance skew term structure (Skew TS) is the difference between longer and shorter dated skew, with the longer dated (i.e. 90 business days) skew defined analogously to the shorter one. The realized variance and the bipower variation are computed using Zhang et al.'s (2005) microstructure-noise robust estimate and are multiplied by 1000. *JV* is the jump variation and is computed as the positive difference between the realized variance and the bipower variation, i.e. $\max(RV - BV, 0)$.

captures 4.73% of the total variation, while the following ones account for 0.32%, 0.18%, 0.11%, and 0.05%, respectively (all the panels of Figure 5.6, except for the top-left panel).

We now turn to the question of how much information is shared between *ORV* surface characteristics (or variations of the index returns) and the principal components. The aim of this exercise is to determine whether the *ORV* surface can be summarized in one simple specification (e.g. the first component of a PCA) or whether different *ORV* measures are needed to appropriately capture the nonlinear changes in the surface. Indeed, we have no guarantee that extracted PCs would represent the whole surface adequately as nonlinear behaviour cannot be captured efficiently by a few summaries of the said surface.

This analysis is performed by computing in-sample regressions of these variables on PCs as follows:

$$\text{Char}_t = \beta_0 + \beta_1 \text{PC}_{1,t} + \beta_2 \text{PC}_{2,t} + \beta_3 \text{PC}_{3,t} + \beta_4 \text{PC}_{4,t} + \beta_5 \text{PC}_{5,t} + \beta_6 \text{PC}_{6,t} + \varepsilon_t,$$

where Char_t is the characteristic of interest at time t (i.e. Level, TS, Skew, Skew TS, *RV* and *JV*) and $\text{PC}_{n,t}$ represents the n^{th} principal component at time t . Newey and West (1987) regressions are run to account for heteroskedasticity and autocorrelation.

Table 5.8 displays regression coefficients, standard errors, *R*-squareds, and autocorrelations in the residuals. Observe that PCs are associated with surface characteristics




Figure 5.6: Principal components of the option realized volatility surface.

The figure shows the first six principal components extracted from the S&P 500 option realized variance surface from July 2004 to December 2012. On each day, option realized variances are interpolated from a locally smoothing quadratic regression surface estimated on every available option realized volatility. Over the call-equivalent delta dimension, we use the grid of values 0.1, 0.2, ..., 0.9; over the tenor dimension, we employ maturities of 30 and 90 business days. These partitions require a total of 18 option realized volatilities per day.

and index return variations, as well, providing further evidence of the commonality existing between these two sets of variables. This commonality effect is largely driven by the first PC, as other PCs are not always statistically significant at explaining index return variations. Notice also that PCs are successful at capturing the behaviour of index return variations, as evidenced from R -squared values of 90% for realized variance and 45% for realized jump variation. However, PCs impact these variables differently, suggesting that the information contained in the surface is impounded differently in these variations. This difference is also evidenced by looking at the persistence of regression residuals, which have different autocorrelation patterns. Persistent residuals in the realized variance regression suggest that either the relation is nonlinear, or that there are missing factors in the model (e.g. lag values of RV). On the other hand, the jump realized variation regression exhibits low persistent residuals.

What can we learn from these results? First, option realized variations are intertwined with measures of index return variation, revealing an important degree of commonality that is useful for analyzing index return dynamics such as variances and sporadic shocks. Notwithstanding, option realized variations also exhibit information that is not shared with index return variations, suggesting their pertinence as an alternative

Table 5.8: Option realized volatility characteristics, realized measures, and principal component regressions.

	Level	TS	Skew	Skew TS	RV	JV
PC ₁	0.395 (0.001)	-0.009 (0.002)	0.029 (0.002)	-0.011 (0.002)	0.515 (0.036)	0.025 (0.003)
PC ₂	0.268 (0.005)	-0.072 (0.006)	-0.107 (0.007)	0.051 (0.007)	-0.060 (0.074)	0.004 (0.010)
PC ₃	-0.434 (0.016)	0.938 (0.034)	-0.014 (0.046)	-0.005 (0.043)	0.697 (0.358)	0.070 (0.041)
PC ₄	-0.360 (0.016)	0.230 (0.028)	0.048 (0.031)	-0.066 (0.046)	-0.536 (0.445)	-0.088 (0.057)
PC ₅	0.027 (0.029)	-0.126 (0.045)	-0.028 (0.032)	0.170 (0.044)	1.203 (0.575)	0.079 (0.086)
PC ₆	-0.427 (0.043)	0.375 (0.057)	0.101 (0.050)	-0.050 (0.056)	0.988 (0.642)	0.207 (0.111)
R^2	1.000	0.950	0.928	0.537	0.901	0.445
AC(1)	0.065	0.234	0.129	0.051	0.623	0.025
AC(2:10)	0.053	0.109	0.070	0.038	0.479	0.074
AC(11:20)	-0.009	0.000	0.036	0.015	0.320	0.051

The following linear regressions are performed:

$$\text{Char}_t = \beta_0 + \beta_1 \text{PC}_{1,t} + \beta_2 \text{PC}_{2,t} + \beta_3 \text{PC}_{3,t} + \beta_4 \text{PC}_{4,t} + \beta_5 \text{PC}_{5,t} + \beta_6 \text{PC}_{6,t} + \varepsilon_t,$$

where Char_t is the characteristic of interest at time t (i.e. Level, TS, Skew, Skew TS, RV and JV) and $\text{PC}_{n,t}$ represents the n^{th} principal component at time t . R^2 and Newey and West (1987) standard errors (in parentheses) are reported. Values in bold are statistically significant at a confidence level of 95%. The first six principal components extracted from the S&P 500 option realized volatility surface from July 2004 to December 2012. The first sample autocorrelation coefficient and the average sample autocorrelation over two to ten and eleven to twenty lags of the regression residuals are exhibited. The values of RV and JV are multiplied by 1000 in the regressions. JV is the jump variation and is computed as the positive difference between the realized variance and the bipower variation, i.e. $\max(RV - BV, 0)$.

source of information about the dynamics of the underlying generating process. Second, the fact that different characteristics of the ORV surface cannot be summarized in one simple specification of PCs suggests that various option realized variances should be used in parametric studies; given the large number of options available at a given point in time, a well-selected subset of option realized variance should be employed for empirical analyses.

5.5.3 Predicting Index Return Variations

The economic relationship between option realized variations and index return variations is now studied with predictive regressions. The motivation behind this exercise is to look at the role of option realized variations as economic variables that predict different types of index return variability.

The first model consists of explaining one-day ahead index return realized variances. We employ the logarithm of realized variance as a dependent variable and estimate the

Table 5.9: Correlation matrix of realized values used in the predictive regressions.

	$\log(RV)$	JV	$\log(IV^{\text{ATM}})$	$\log(ORV^{\Delta^e=0.1})$	$\log(ORV^{\Delta^e=0.9})$
$\log(RV)$	1.0000				
JV	0.4876	1.0000			
$\log(IV^{\text{ATM}})$	0.8511	0.4121	1.0000		
$\log(ORV^{\Delta^e=0.1})$	0.8470	0.3961	0.7116	1.0000	
$\log(ORV^{\Delta^e=0.9})$	0.8691	0.3993	0.7202	0.8194	1.0000

JV is the jump variation and is computed as the positive difference between the realized variance and the bipower variation, i.e. $\max(RV - BV, 0)$.

following model:

$$\begin{aligned} \log(RV_{t+1}) = & \beta_0 + \beta_1 \log(RV_t) + \beta_2 JV_t + \beta_3 IV_{\log, \perp, t}^{\text{ATM}} \\ & + \beta_4 ORV_{\log, \perp, t}^{\Delta^e=0.1} + \beta_5 ORV_{\log, \perp, t}^{\Delta^e=0.9} + \varepsilon_t \end{aligned} \quad (5.18)$$

The variables of interest in this model are lagged values of realized jump variation (JV_t), ATM implied volatility (IV^{ATM}), and option realized variances of OTM options with call-equivalent deltas (Δ^e) of 0.1 and 0.9. The maturity of all options is 30 days. To remove potential collinearity issues across the regressors (see Table 5.9), we employ the orthogonal component that results from the projection of a given measure on RV in our regressions.

Panel A of Table 5.10 presents estimates and t -statistics of six specifications of the regression in Equation (5.18). The first four specifications show that only the orthogonal components of IV^{ATM} and $ORV^{\Delta^e=0.1}$ are statistically significant explaining the one-day ahead realized variation of index returns after controlling for the lagged value of the dependent variable. They enter in the regression model with positive signs, which confirms the forward-looking nature of IV^{ATM} . Interestingly, it is the residual component of ORV associated with OTM calls that predicts the subsequent realized variance, even after controlling for IV^{ATM} . As discussed in Subsection 5.2.2, the ORV for an OTM option is related to jump activity in the log-equity price and variance processes. Thus, the explanation power of this variable might come from two sources. The first one is the high persistence of the variance process, so that volatility jumps lead to higher activity in subsequent periods. The second possibility is that positive jumps induce future variability by arriving in clusters or by increasing the volatility directly. Fulop et al. (2014) provide evidence of self-exciting jump clustering during turbulent market periods.

Table 5.10: Predictive regressions.

Panel A: Logarithm of realized variance							
Regression	β_0	β_1	β_2	β_3	β_4	β_5	R^2
(1)	-0.5493 (0.0981)	0.8733 (0.0199)	-3.2276 (3.6529)				0.7490
(2)	-0.2279 (0.0932)	0.9425 (0.0193)	-2.9616 (3.2157)	0.9432 (0.0705)			0.7839
(3)	-0.8316 (0.1101)	0.8128 (0.0228)	-2.8130 (3.6017)		0.1197 (0.0257)		0.7515
(4)	-0.3872 (0.1391)	0.9081 (0.0289)	-3.4894 (3.6881)			-0.0438 (0.0238)	0.7494
(5)	-0.5424 (0.1003)	0.8750 (0.0216)	-2.4922 (3.1548)	0.9529 (0.0723)	0.1347 (0.0280)		0.7871
(6)	-0.1785 (0.1270)	0.9531 (0.0269)	-3.0435 (3.2375)	0.9408 (0.0705)		-0.0136 (0.0238)	0.7839
Panel B: Jump variation (Positive difference between realized variance and bipower variation)							
Regression	β_0	β_1	β_2	β_3	β_4	β_5	R^2
(1)	0.0102 (0.0019)	0.0019 (0.0004)	0.0963 (0.0460)				0.1885
(2)	0.0110 (0.0019)	0.0021 (0.0004)	0.0969 (0.0457)	0.0023 (0.0004)			0.1963
(3)	0.0100 (0.0020)	0.0019 (0.0004)	0.0966 (0.0462)		0.0001 (0.0003)		0.1885
(4)	0.0126 (0.0024)	0.0024 (0.0005)	0.0924 (0.0455)			-0.0007 (0.0002)	0.1921
(5)	0.0107 (0.0020)	0.0020 (0.0004)	0.0974 (0.0460)	0.0023 (0.0004)	0.0001 (0.0003)		0.1964
(6)	0.0131 (0.0024)	0.0026 (0.0005)	0.0934 (0.0453)	0.0022 (0.0004)		-0.0006 (0.0002)	0.1992

Variations of following linear regression are performed:

$$X_{t+1} = \beta_0 + \beta_1 \log(RV_t) + \beta_2 JV_t + \beta_3 IV_{\log, \perp, t}^{\text{ATM}} + \beta_4 ORV_{\log, \perp, t}^{\Delta^c=0.1} + \beta_5 ORV_{\log, \perp, t}^{\Delta^c=0.9} + \varepsilon_t,$$

where $X_{t+1} \in \{RV_{t+1}, JV_t + 1\}$, RV_t is the realized variance at time t , JV_t is the jump variation at time t , $IV_{\log, \perp, t}^{\text{ATM}}$ is the residual of the following regression

$$\log(IV_t^{\text{ATM}}) = \alpha_1 \log(RV_t) + IV_{\log, \perp, t}^{\text{ATM}},$$

where IV_t^{ATM} is the ATM implied volatility (with maturity of 30 days). The variables $ORV_{\log, \perp, t}^{\Delta^c=0.1}$ and $ORV_{\log, \perp, t}^{\Delta^c=0.9}$ are computed similarly from the option realized variance for OTM calls with call-equivalent delta of 0.1 and the option realized variance for OTM calls with call-equivalent delta of 0.9. We compute Newey-West standard errors. These are in parentheses in the above table. Values in bold are statistically significant at a confidence level of 95%. Results for bipower variation are both qualitatively and quantitatively similar to those of realized variance.

The second model we consider consists of explaining one-day ahead jump activity on the index return realized variances. This time we run the regression:

$$JV_{t+1} = \beta_0 + \beta_1 \log(RV_t) + \beta_2 JV_t + \beta_3 IV_{\log, \perp, t}^{\text{ATM}} + \beta_4 ORV_{\log, \perp, t}^{\Delta^c=0.1} + \beta_5 ORV_{\log, \perp, t}^{\Delta^c=0.9} + \varepsilon_t. \quad (5.19)$$

Panel B of Table 5.10 shows that lagged values of RV and jump realized variance have explanatory power over future jump activity. This relationship provides evidence in favour of the parametric model we are considering in Section 5.2, since the intensity of the counting process governing jump returns is a function of volatility. Regarding IV ,

its residual component also has a positive effect on jump realized variance. Contrary to what was observed in the *RV* regression, it is the *ORV* of the OTM put that now has a significant relation with the dependent variable. However, this time, the coefficient of the relationship is negative. Given that jump activity is measured as the difference between total and diffusive variance, the negative sign might indicate that the jump activity driving the *ORV* of OTM puts exerts more influence over future diffusive variance than return jumps. Notwithstanding, if we compare the magnitude of this coefficient with the one of OTM calls in Equation (5.18), we note that the significance of the relationship is more important in the latter case.

Lastly, we emphasize that the regression models (5.18) and (5.19) are of predictive nature, so our analysis focuses on the identification of variables that are important in the determination of future realized variances. The fact that *R*-squareds in Panel A are higher than those reported in Panel B just shows how much easier it is to predict future total variance than to predict variance induced by jump activity.

5.6 Option Pricing Implications

This section provides empirical results using the model introduced in Section 5.2. We start by analyzing parameter estimates of the model with different sets of information. We then use these parameters to disentangle latent variables and analyze the information content of option realized variances. Finally, we look at the fitting performance using in- and out-of-sample analyses.

5.6.1 Parameter Estimates

We obtain model parameters using the estimation procedure described in Subsection 5.3.2. As explained in Subsection 5.2.2, this procedure combines several observables into a likelihood function that is computed using a Monte Carlo based approximation. The first set of observable variables are related to S&P 500 returns. We use daily log returns, realized variance, and bipower variation.²⁹ The second set is composed of daily option prices and their realized variances (*ORVs*). Option prices come from OptionMetrics and correspond to European S&P 500 index option contracts. As argued in Bates (2000), the daily overabundance of option data becomes a hurdle for estimation routines, so we use a representative sample of options by restricting our

²⁹The computation of these last two variables is explained in Subsection 5.5.1.

Table 5.11: S&P 500 parameter estimates.

Panel A: With <i>ORV</i>								
Log-price process			Variance process			Standard deviations of error terms		
η_Y	0.7635	(0.000230)	η_V	-0.2199	(0.000013)	η_1	0.2160	(0.000023)
γ_Y	0.0058	(0.000010)	Γ_V	0.5811	(0.000019)	η_2	0.2296	(0.000023)
ρ	-0.4336	(0.000011)	κ	5.7808	(0.000061)	η_3	0.2471	(0.000018)
$\lambda_{Y,0}$	2.1530	(0.000498)	θ	0.0085	(0.000177)	η_4	0.4600	(0.000028)
$\lambda_{Y,1}$	29.0744	(0.000682)	σ	0.8935	(0.000007)			
μ_Y	-0.0050	(0.000018)	$\lambda_{V,0}$	7.6498	(0.000090)			
σ_Y	0.0163	(0.000129)	μ_V	0.0214	(0.000101)			
			V_0	0.0118	(0.000174)			
Panel B: Without <i>ORV</i>								
Log-price process			Variance process			Standard deviations of error terms		
η_Y	1.8734	(0.000311)	η_V	-1.5870	(0.000033)	η_1	0.2407	(0.000034)
γ_Y	0.0009	(0.000004)	Γ_V	0.7061	(0.000016)	η_2	0.2288	(0.000022)
ρ	-0.3912	(0.000015)	κ	4.7614	(0.000107)	η_3	0.2453	(0.000018)
$\lambda_{Y,0}$	0.0045	(0.025987)	θ	0.0030	(0.000825)	η_4	-	-
$\lambda_{Y,1}$	30.0106	(0.000992)	σ	0.9121	(0.000010)			
μ_Y	-0.0062	(0.000049)	$\lambda_{V,0}$	11.4438	(0.000065)			
σ_Y	0.0189	(0.000464)	μ_V	0.0174	(0.000041)			
			V_0	0.0114	(0.000242)			

The index parameters are estimated using daily index returns, realized variances, bipower variations, daily option prices, and option realized variances (in the first case), from July 2004 to December 2012. Parameters are estimated using multiple simplex search method optimizations (fminsearch in Matlab). Robust standard errors are computed from the outer product of the gradient at optimal parameter values. Each day, we restrict our attention to OTM or ATM options with maturities closest to 30 and 90 days, and call-equivalent deltas closest to 0.2, 0.35, 0.5, 0.65 and 0.8.

attention to OTM or ATM options with maturities closest to 30 and 90 days, and call-equivalent deltas closest to 0.2, 0.35, 0.5, 0.65 and 0.8. Options with positive volume and bid price are included in the sample, as well as those satisfying the no-arbitrage conditions of Bakshi et al. (1997).³⁰ Thus, the final sample of options is composed of a panel of 21,400 contracts.

To assess the information contained in *ORVs*, we first estimate the model excluding these variables from the set of observables and then re-estimate it using the complete set. Table 5.11 reports parameter estimates and standard errors (in parentheses) obtained with and without information from option realized variances. We first consider parameters governing the jump process intensities. We observe a decrease in the intensity of the jump process governing volatility jumps, $\lambda_{V,0}$, when *ORVs* are introduced, but an increase in the parameter governing the size of the jumps, μ_V . This suggests that *ORV* is detecting less frequent jumps with higher magnitudes in the volatility process. Regarding jumps in the log-equity price process, we observe an increase in the base intensity, $\lambda_{Y,0}$, as well as in the average size of jumps, μ_Y , in absolute value. These

³⁰Option realized variances are computed following the procedure explained in Subsection 5.5.1.

values show that *ORV* is favouring more frequent negative jumps with lower magnitudes, which is also observed from the average intensity -0.0112 for the estimation with *ORV* and 0.0028 without.

We now look at compensation for bearing different types of risk. Jump return risk, or crash risk, is captured in this model by the parameter γ_Y times λ_{Y,t^-} , as shown in Equation (5.3). Since *ORV* is favouring more frequent negative jumps, compensation for bearing this type of risk increases with the addition of this new source of information (i.e. 0.06% without *ORV* vs. 1.63% with *ORV*, on average). The diffusive risk premium parameter η_Y captures the risk premium associated with diffusive shocks when multiplied by V_{t^-} . To find the compensation for this type of risk, we multiply the filtered instantaneous variance by η_Y and observe that this value is higher when *ORV* is not included in the estimation, on average (4.41% when *ORV* is excluded and 1.75% when *ORV* is used). Therefore, in total, the average risk premiums calculated using both models are similar, although the breakdown between diffusive and jump risks is different. These findings suggest that the risk premiums of discontinuous risk is underestimated when information about intraday evolution of option prices is not accounted for during the estimation. Estimates for jump risk are close but somewhat lower than those reported in Broadie et al. (2007), Christoffersen et al. (2012), and Ornathanalai (2014).³¹

5.6.2 Informativeness of Data Sources

Next, we focus on the informativeness of *ORV* about different latent variables of the model. Using the set of parameters estimated in Table 5.11, we filter different latent quantities using all data sources and compare the average standard deviations of these quantities with those obtained when filtering without *ORV*. Table 5.12 shows the posterior standard deviation (PSD) of these latent quantities under the two filtering procedures.³² The PSD of variance jumps decreases approximately five times when *ORVs* are added. This dramatic decrease confirms our simulated results that show how *ORV* embeds vital information about the presence of variance jumps. Note also, as it is the case in the simulation results, the PSDs of the instantaneous variances, quadratic variations, and integrated variances also fall when *ORVs* are included. Regarding the slight increase in the PSD for log-return jumps, this could be associated with the

³¹Variations in the sampling period or the datasets could also explain these small differences, to some extent.

³²The posterior standard deviation is the standard deviation of the posterior density (as given by the particle filter). It allows us to assess the uncertainty around estimated latent variables.

Table 5.12: **Posterior standard deviation of filtered values with and without ORV.**

	With <i>ORV</i>	Without <i>ORV</i>
Instantaneous variance	4.821×10^{-03}	5.961×10^{-03}
Quadratic variation	1.252×10^{-05}	1.410×10^{-05}
Integrated variance	1.088×10^{-05}	1.255×10^{-05}
Log-equity price jumps	2.886×10^{-04}	1.804×10^{-04}
Variance jumps	7.038×10^{-04}	3.297×10^{-03}

The posterior standard deviation is calculated from the particles generated in the particle filtering scheme. Parameters from Table 5.11 are used to infer the filtered values and their posterior standard deviation.

fact that these jumps are less frequent than jumps in volatility, which increases the estimation uncertainty about this value.

We next provide visual evidence of the informativeness of *ORVs* for disentangling jumps. Figure 5.7 displays filtered price jumps and variance jumps excluding *ORVs* (first row) and with all data sources (second row). The dynamics of both jump processes consistently differ across sources, showing that jumps for log-prices are sporadic and largely negative on average and that jumps in volatility are more frequent and tend to be small on average. *ORV* plays an important role identifying jumps in volatility. Note the striking decrease in the number and sizes of volatility jumps when *ORV* is employed in the filtering process. This result, coupled with the fact that jump risk premiums vary depending on the data source employed, show the economic gains from including intraday option prices in the analysis of option pricing models.

5.6.3 In- and Out-of-Sample Assessment

We now investigate the goodness of fit of the two parameter sets under analysis. We start with an in-sample analysis that employs data over the same period the parameters were estimated. Then, we employ the parameters obtained over the period from July 2004 to December 2012 to perform out-of-sample comparisons using 2013 option data.

5.6.3.1 In-Sample

We first assess the ability of both parameter sets to fit historical implied volatilities. To perform this exercise, we use the relative implied volatility root mean square error




Figure 5.7: Filtered log-equity price jumps (left panels) and variance jumps (right panels).

The figure shows the log-equity price jumps (left panels) and variance jumps (right panels). The top panels show the filtered jumps based on returns, option implied volatilities, realized variances and bipower variations. The bottom panels also include option realized variances as a new source of information.

(RIVRMSE), defined as follows:

$$\text{RIVRMSE} = \sqrt{\left(\frac{1}{\sum_k O_k}\right) \sum_k \sum_{i=1}^{O_k} \left(\frac{IV_{k\Delta,i}(Y_{\Delta k}, \hat{V}_{\Delta k}) - \sigma_{k\Delta,i}^{\text{BS}}}{\sigma_{k\Delta,i}^{\text{BS}}} \right)^2}, \quad (5.20)$$

where IV is the model implied volatility, $\hat{V}_{\Delta k}$ is the filtered instantaneous variance on day k , and O_k represents the number of options in a subset of all the options available on day k . Here, σ^{BS} denotes the Black-Scholes implied volatility associated with the observed option price.

The leftmost columns of Panel A (Table 5.13) show the RIVRMSE for the panel of 21,400 options employed in the estimation. We observe that the fit provided by the parameter set without $ORVs$ is more accurate on average than the one that includes these variances – an RIVRMSE of 24.56 with ORV vs. 23.45 without. This result stems from the fact that maximization of the likelihood in the former case backs parameters that better fit IVs . However, the average fit is not that different, so the addition of $ORVs$ does not substantially deteriorate the overall fit of option prices.

In order to identify which contracts would benefit more from the inclusion of $ORVs$, we assess the goodness of fit of both parameter sets for these variances. To this end, we

use a criterion similar to the RIVRMSE in Equation (5.20) and compute the relative option realized variance root mean square error (RORVMSE), defined as:

$$\text{RORVMSE} = \sqrt{\left(\frac{1}{\sum_k O_k}\right) \sum_k \sum_{i=1}^{O_k} \left(\frac{\Delta OQV_{(k-1)\Delta, k\Delta, i} - ORV_{(k-1)\Delta, k\Delta, i}}{ORV_{(k-1)\Delta, k\Delta, i}}\right)^2},$$

where ΔOQV is the option quadratic variation increment computed from the model and ORV is the observed option realized variance.

Rightmost columns of Panel A (Table 5.13) exhibit the RORVMSE by moneyness and maturity. As expected, the average fit of these quantities is lower when $ORVs$ are included in the information set. Nonetheless, it is remarkable to observe that the overall RORVMSE is about 3 times lower when ORV is included and that there are significant differences across contracts. Whereas error differences for call options are almost fourfold, those differences for put options are about twofold. The differences across contracts suggest that $ORVs$ bring more information from OTM calls than from OTM puts, which goes in line with the importance of the coefficient associated with call contracts in explaining future index variation as discussed in Subsection 5.5.3.

We continue with the in-sample analysis by looking at one-day ahead predicted option prices from both sets of parameters. To analyse the fit of implied volatilities, we enlarge our estimation sample of 21,400 options to include all OTM options available for the S&P 500 index in OptionMetrics between July 2004 and December 2012.³³ We restrict our analysis to maturities of at least one week and at most one year. As before, observations violating no-arbitrage restrictions are excluded. Our new sample is composed of a total of 401,081 contracts. Option prices are converted to implied volatilities with the Black and Scholes's (1973) pricing formula. To compute model-predicted implied volatilities on day t , we calculate one-day ahead expectations of model variables for day t using the filter's predictive distribution resulting from day $t - 1$. Using observed and predicted implied volatilities, we compute RIVRMSEs according to maturity, moneyness, and year, which provides a better picture of the fit associated with each parameter set.

The overall RIVRMSEs (leftmost columns of Panel B in Table 5.13) are very close on average for both parameter sets, with lower values of RIVRMSE observed when $ORVs$ are included (32.79 against 32.87). We use the Diebold and Mariano (1995,

³³We continue using OTM options to keep a comparable sample with that employed in our previous analyses.

Table 5.13: In-sample performance (2004–2012).

Panel A: In-sample option pricing and option realized variance performances				
	RIVRMSE		RORVMSE	
	With <i>ORV</i>	Without <i>ORV</i>	With <i>ORV</i>	Without <i>ORV</i>
DTM = 30, $\Delta^e = 0.20$	19.13	17.62	50.17	152.95
DTM = 30, $\Delta^e = 0.35$	25.05	23.04	36.29	103.59
DTM = 30, $\Delta^e = 0.50$	28.51	27.93	41.81	90.90
DTM = 30, $\Delta^e = 0.65$	27.27	25.94	55.26	106.11
DTM = 30, $\Delta^e = 0.80$	25.03	23.16	61.43	95.19
DTM = 90, $\Delta^e = 0.20$	21.97	26.01	61.02	272.48
DTM = 90, $\Delta^e = 0.35$	23.37	22.90	40.78	145.63
DTM = 90, $\Delta^e = 0.50$	25.23	23.16	40.77	101.02
DTM = 90, $\Delta^e = 0.65$	25.17	22.40	44.92	122.24
DTM = 90, $\Delta^e = 0.80$	23.57	20.69	53.56	137.20
All	24.56	23.45	49.34	142.17
Panel B: One-day ahead in-sample performances				
	RIVRMSE		RORVMSE	
	With <i>ORV</i>	Without <i>ORV</i>	With <i>ORV</i>	Without <i>ORV</i>
DTM < 60	26.78	25.39	233.08	286.97
60 ≤ DTM < 120	26.20	24.24	230.46	337.40
120 ≤ DTM < 180	26.07	25.08	203.42	335.05
180 ≤ DTM	51.76	55.28	180.71	270.46
$\Delta^e < 0.20$	23.25	28.56	358.69	512.51
$0.20 \leq \Delta^e < 0.35$	29.55	31.11	216.98	302.68
$0.35 \leq \Delta^e < 0.50$	34.52	35.51	152.92	198.51
$0.50 \leq \Delta^e < 0.65$	41.42	41.24	188.62	239.69
$0.65 \leq \Delta^e < 0.80$	37.41	36.83	200.04	267.02
$0.80 \leq \Delta^e$	31.74	29.62	210.69	285.87
2004	25.73	28.52	204.64	268.17
2005	28.44	33.71	270.87	348.78
2006	26.58	32.05	316.43	407.00
2007	29.29	31.87	220.31	290.46
2008	31.17	30.56	77.09	92.51
2009	38.47	35.74	79.21	93.69
2010	34.99	33.49	188.20	284.60
2011	35.33	34.38	193.82	274.96
2012	31.68	31.39	314.55	437.40
All	32.79	32.87	221.94	303.31

In Panel A, only options used for estimation are employed. In Panel B, the sample of S&P 500 options is acquired via OptionMetrics. Options violating arbitrage conditions were discarded. In Panel C, the sample of S&P 500 options is acquired via Tick Data. Options violating arbitrage conditions were discarded. Model prices and *ORVs* are calculated by using the parameters of Table 5.11. The relative implied volatility root mean square error (RIVRMSE) is computed as follows:

$$\text{RIVRMSE} = \sqrt{\left(\frac{1}{\sum_k O_k}\right) \sum_k \sum_{i=1}^{O_k} \left(\frac{IV_{k\Delta,i}(Y_{\Delta k}, \hat{V}_{\Delta k}) - \sigma_{k\Delta,i}^{\text{BS}}}{\sigma_{k\Delta,i}^{\text{BS}}} \right)^2},$$

where IV is the model implied volatility, $\hat{V}_{\Delta k}$ is the filtered instantaneous variance on day k , and O_k represents the number of options in a subset of all the options available on day k . The relative option realized variance root mean square error (RORVMSE) is computed similarly: $IV_{k\Delta,i}(Y_{\Delta k}, \hat{V}_{\Delta k})$ is replaced by $\Delta OQV_{(k-1)\Delta,k\Delta,i}$ and $\sigma_{k\Delta,i}^{\text{BS}}$ by $ORV_{(k-1)\Delta,k\Delta,i}$. RIVRMSEs and RORVMSEs are given in percentage.

henceforth DM) test to see if the apparent predictive superiority of *ORV*-based forecasts is not particular to this sample. Using both RIVRMSE time series, we compare their forecasting accuracy and test for:³⁴

$$H_0 : \mathbb{E}[d_t] = 0, \forall t \quad H_1 : \mathbb{E}[d_t] > 0, \forall t$$

where $d_t = \text{RIVRMSE}_{\text{without},t} - \text{RIVRMSE}_{\text{with},t}$ is the time- t loss differential between the forecast produced without *ORV* and the one including it. The DM test statistic is 5.32 and is significant at a 1% level, confirming that there exists a differential between the two forecasts and that the one based on *ORV* information produces more accurate results on average. A closer scrutiny of the data reveals that including *ORVs* is especially helpful in the pricing of options with long maturities, OTM calls, and options for the years before 2008.

Finally, we carry out a similar exercise with option realized variances and analyse one-day ahead forecasts for both parameter sets. For this exercise, we restrict our enlarged sample to those options for which *ORV* data is available in the TickData database. Out of 401,081 options, 282,534 contracts were included in our analysis. As reported in the rightmost columns of Panel B, Table 5.13, we observe again that, across different option characteristics and years, RORVRMSE are lower for the parameter set that was obtained with *ORVs*. Similar to the RIVRMSE case, we apply the DM test to both RORVRMSE series and obtain a value of 13.63, confirming statistically the important differences between the two sets of parameters.³⁵

5.6.3.2 Out-of-Sample

We provide further evidence on the forecast differences of both parameter sets by constructing out-of-sample errors. We employ the parameters obtained over the sample period between July 2004 and December 2012, and use daily information between January 2013 and December 2013 to compute one-day ahead forecast errors. As in the previous section, we employ the filter's predictive distribution resulting from day $t - 1$ to compute expectations for day t . Regarding the sample, we employ all OTM options available in OptionMetrics for 2013 to compute RIVRMSE, yielding a total of

³⁴The lag in the Diebold and Mariano (1995) is selected as the first partial autocorrelation that is within confidence bounds. The estimated lag in this exercise is 2.

³⁵We repeated our analysis of RIVRMSE with the sample of 282,534 contracts and found similar results.

77,310 contracts. To compute RORVRMSEs, we employed a dataset of 68,857 OTM options for which *ORVs* were available.

Table 5.14 shows RIVRMSE and RORVRMSE by moneyness and maturity. The parameter set based on *ORVs* produces the lowest forecast errors, as shown by the average means for the whole sample. This is also confirmed with the DM test, which gives a value of 8.09 for RIVRMSE and 16.17 for RORVRMSE, both significant at a 1% level. When we look at the RIVRMSE across maturities, we observe that *ORV*-based forecasts consistently outperform those forecasts that do not employ the information of *ORVs*. Regarding performance by type of contract, it is still the case that the largest gains are obtained for OTM calls. Nonetheless, OTM puts that are close to the money also exhibit lower RIVRMSE.

The above evidence is consistent with the view of *ORV* as a new source of information to disentangle jumps in volatility and prices. The parameter set obtained with the addition of *ORV* supports a variance process that has less frequent jumps with higher magnitudes and a price process that has more frequent negative jumps with lower magnitudes. This set also has a different attribution of risk premiums between diffusive and discontinuous innovations. It seems most likely that these features are important in the pricing of options, as they consistently produce lower forecasting errors of implied volatilities in-sample and out-of-sample.

Table 5.14: One-day ahead out-of-sample performances, in terms of RIVRMSE and RORVRMSE (2013).

	RIVRMSE		RORVRMSE	
	With <i>ORV</i>	Without <i>ORV</i>	With <i>ORV</i>	Without <i>ORV</i>
DTM < 60	18.85	19.09	512.96	635.93
60 ≤ DTM < 120	18.62	20.75	559.93	810.20
120 ≤ DTM < 180	22.50	27.59	431.52	691.71
180 ≤ DTM	50.14	55.54	390.10	585.19
$\Delta^e < 0.20$	34.97	45.40	871.92	1214.42
$0.20 \leq \Delta^e < 0.35$	33.16	39.70	465.52	617.28
$0.35 \leq \Delta^e < 0.50$	36.38	41.39	275.85	361.76
$0.50 \leq \Delta^e < 0.65$	37.38	38.37	357.39	464.95
$0.65 \leq \Delta^e < 0.80$	30.85	31.08	526.85	701.61
$0.80 \leq \Delta^e$	21.18	19.61	433.91	573.80
All	28.94	31.95	504.30	680.05

The sample of S&P 500 options is acquired via OptionMetrics. Options violating arbitrage conditions were discarded. The *IV* sample size is 77,310. The sample of S&P 500 options is acquired via Tick Data. Options violating arbitrage conditions were discarded. The *ORV* sample size is 68,857. Model prices and *ORVs* are calculated by using the parameters of Table 5.11. RIVRMSEs and RORVRMSEs are given in percentages.

5.7 Concluding Remarks

The observable variables used in this paper provide a portrait of the incremental information of these variables for the identification of latent features governing dynamics of asset prices. The paper proposes the option realized variance as a new observable quantity. This variable provides evidence of the usefulness of high frequency option prices to infer characteristics such as instantaneous variance, as well as jumps in returns and the variance process.

One of the reasons behind the information gains of this measure is its ability to capture different properties of the DGP depending on the moneyness of the option employed. Whereas ITM and OTM options provide *ORVs* that are responsive to discontinuous and diffusive innovations, the *ORV* of OTM puts are sensitive to changes in the discontinuous part of the variance and the return process.

The paper empirically documents how option realized variance behaves for options on the S&P 500 and studies its relationship with the realized variance of the index. For instance, using a principal component analysis for the surface of option realized variances, we find that *ORV* contains additional information when compared to the one embedded in realized variance. The paper further explores the economic implications of not using high frequency option prices for model estimation. Model parameters estimated without this information are not able to correctly disentangle diffusive and discontinuous innovations, which produces an incorrect attribution of risk premiums among these two competitive sources of risk.

5.A Pricing

5.A.1 Radon-Nikodym Derivative

Let $\mathcal{F}_t = \sigma\{W_{V,u}, W_{\perp,u}, J_{Y,u}, J_{V,u}\}_{0 \leq u \leq t}$ be the σ -field generated by the past and actual noise terms. The market being incomplete, there are infinitely many equivalent martingale measures. We restrict the choice among those that have a Radon-Nikodym derivative of the form

$$\left. \frac{d\mathbb{Q}}{d\mathbb{P}} \right|_{\mathcal{F}_t} = \frac{\exp\left(-\int_0^t \Lambda_{V,u} dW_{V,u} - \int_0^t \Lambda_{\perp,u} dW_{\perp,u} + \Gamma_Y J_{Y,t} + \Gamma_V J_{V,t}\right)}{\mathbb{E}^{\mathbb{P}}\left[\exp\left(-\int_0^t \Lambda_{V,u} dW_{V,u} - \int_0^t \Lambda_{\perp,u} dW_{\perp,u} + \Gamma_Y J_{Y,t} + \Gamma_V J_{V,t}\right)\right]}, \quad (5.21)$$

the latter being an extended version of the Girsanov theorem. The predictable processes $\{\Lambda_{V,t^-}\}_{t \geq 0}$ and $\{\Lambda_{\perp,t^-}\}_{t \geq 0}$ and the constants Γ_Y and Γ_V characterize the risk premiums embedded in this framework by linking the \mathbb{P} - and \mathbb{Q} -parameters together.

This approach is different from the one proposed in Duffie et al. (2000), Pan (2002), and Broadie et al. (2007) in which the \mathbb{P} - and \mathbb{Q} - parameters are allow to vary independently. However, it shares similarities to Christoffersen et al. (2012) and Ornathanalai (2014) that consider GARCH models with jumps. It is also related to Bates (2000): in the latter, the author restricts the value of certain parameters to be consistent with the time series behaviour of returns.

Building on the properties of exponential martingales,³⁶

$$\begin{aligned} & \mathbb{E}^{\mathbb{P}} \left[\exp \left(- \int_0^t \Lambda_{V,u^-} dW_{V,u} - \int_0^t \Lambda_{\perp,u^-} dW_{\perp,u} + \Gamma_Y J_{Y,t} + \Gamma_V J_{V,t} \right) \right] \\ &= \exp \left(\frac{1}{2} \int_0^t (\Lambda_{V,u^-}^2 + \Lambda_{\perp,u^-}^2) du + (\varphi_{Z_Y}^{\mathbb{P}}(\Gamma_Y) - 1) \int_0^t \lambda_{Y,u^-} du \right. \\ & \quad \left. + (\varphi_{Z_V}^{\mathbb{P}}(\Gamma_V) - 1) \int_0^t \lambda_{V,u^-} du \right) \end{aligned}$$

where $\varphi_{Z_Y}^{\mathbb{P}}(\Gamma_Y)$ and $\varphi_{Z_V}^{\mathbb{P}}(\Gamma_V)$ represent the moment generating functions of the log-equity price and variance jump size,³⁷

$$\varphi_{Z_Y}^{\mathbb{P}}(\Gamma_Y) = \exp \left(\mu_Y \Gamma_Y + \frac{1}{2} \sigma_Y^2 \Gamma_Y^2 \right) \quad \text{and} \quad \varphi_{Z_V}^{\mathbb{P}}(\Gamma_V) = (1 - \Gamma_V \mu_V)^{-1}. \quad (5.22)$$

5.A.2 Model Under the Risk-Neutral Measure \mathbb{Q}

For the diffusion components of the model, the risk-neutral Brownian motion are constructed in the usual way:

$$W_{V,t}^{\mathbb{Q}} = W_{V,t} + \int_0^t \Lambda_{V,u^-} du, \quad W_{\perp,t}^{\mathbb{Q}} = W_{\perp,t} + \int_0^t \Lambda_{\perp,u^-} du \quad (5.23)$$

and $W_{Y,t}^{\mathbb{Q}} = \rho W_{V,t} + \sqrt{1 - \rho^2} W_{\perp,t} = W_{Y,t} + \int_0^t \Lambda_{Y,u^-} du$ where $\Lambda_{Y,u^-} = \rho \Lambda_{V,u^-} + \sqrt{1 - \rho^2} \Lambda_{\perp,u^-}$.

The risk-neutral jump components are obtained from a direct comparison of the \mathbb{P} - and \mathbb{Q} -versions of the moment generating functions of the jump increments $(J_{Y,t} - J_{Y,s})$

³⁶It is required that $\mathbb{E}^{\mathbb{P}} \left[\int_0^t \Lambda_{V,u^-}^2 du \right] < \infty$ and $\mathbb{E}^{\mathbb{P}} \left[\int_0^t \Lambda_{\perp,u^-}^2 du \right] < \infty$. Details are provided in Lemmas B.4 and B.6 of Appendix B.

³⁷Provided that $\Gamma_V < \frac{1}{\mu_V}$.

and $J_{V,t} - J_{V,s}$.³⁸ Indeed, the change of measure affects the parameters:

$$J_{Y,t}^{\mathbb{Q}} = \sum_{n=1}^{N_{Y,t}^{\mathbb{Q}}} Z_{Y,n}^{\mathbb{Q}}, \text{ and } J_{V,t}^{\mathbb{Q}} = \sum_{n=1}^{N_{V,t}^{\mathbb{Q}}} Z_{V,n}^{\mathbb{Q}},$$

where $(N_{Y,t}^{\mathbb{Q}})_{t \geq 0}$ is a Cox process with predictable intensity $(\lambda_{Y,t})_{t \geq 0}$, $(N_{V,t}^{\mathbb{Q}})_{t \geq 0}$ is a Poisson process of intensity $\lambda_{V,0}^{\mathbb{Q}}$ and

$$\begin{aligned} \lambda_{Y,t}^{\mathbb{Q}} &= \varphi_{Z_Y}^{\mathbb{P}}(\Gamma_Y) \lambda_{Y,t}, & \mu_Y^{\mathbb{Q}} &= \mu_Y + \Gamma_Y \sigma_Y^2, & \sigma_Y^{\mathbb{Q}} &= \sigma_Y, \\ \lambda_{V,0}^{\mathbb{Q}} &= \varphi_{Z_V}^{\mathbb{P}}(\Gamma_V) \lambda_{V,0}, & \mu_V^{\mathbb{Q}} &= \varphi_{Z_V}^{\mathbb{P}}(\Gamma_V) \mu_V. \end{aligned}$$

The risk-neutral dynamics of the log-price process is established by imposing the discounted price process $\{\exp((q-r)t) \exp(Y_t)\}_{t \geq 0}$ to be a \mathbb{Q} -martingale where r is the risk-free rate and q is the dividend rate.³⁹ To get a semi-closed form for option prices, the risk-neutral stochastic differential equation (SDE) of the variance process is assumed to have a mean reverting behaviour, as in Heston (1993) among others. Implicitly, it constrains⁴⁰

$$\Lambda_{V,t} = \eta_V \sqrt{V_{t-}} \quad (5.24)$$

and the model under the risk-neutral measure is thus

$$dY_t = \alpha_t^{\mathbb{Q}} dt + \sqrt{V_{t-}} dW_{Y,t}^{\mathbb{Q}} + dJ_{Y,t}^{\mathbb{Q}}, \quad (5.25)$$

$$dV_t = \kappa^{\mathbb{Q}} (\theta^{\mathbb{Q}} - V_{t-}) dt + \sigma \sqrt{V_{t-}} dW_{V,t}^{\mathbb{Q}} + dJ_{V,t}^{\mathbb{Q}}, \quad (5.26)$$

³⁸Lemma B.8 of Appendix B shows that

$$\begin{aligned} \varphi_{J_{Y,t}-J_{Y,s}}^{\mathbb{P}}(a) &= \mathbb{E}_{\mathcal{F}_s}^{\mathbb{P}} \left[\left(\exp \left(\mu_Y a + \frac{1}{2} \sigma_Y^2 a^2 \right) - 1 \right) \exp \left(\int_s^t \lambda_{Y,u} du \right) \right], \\ \varphi_{J_{Y,t}-J_{Y,s}}^{\mathbb{Q}}(a) &= \mathbb{E}_{\mathcal{F}_s}^{\mathbb{P}} \left[\left(\exp \left((\mu_Y + \Gamma_Y \sigma_Y^2) a + \frac{1}{2} a^2 \sigma_Y^2 \right) - 1 \right) \exp \left(\varphi_{Z_Y}^{\mathbb{P}}(\Gamma_Y) \int_s^t \lambda_{Y,u} du \right) \right], \\ \varphi_{J_{V,t}-J_{V,s}}^{\mathbb{P}}(a) &= \mathbb{E}_{\mathcal{F}_s}^{\mathbb{P}} \left[\exp \left(((1 - a\mu_V)^{-1} - 1) \int_s^t \lambda_{V,u} du \right) \right], \\ \varphi_{J_{V,t}-J_{V,s}}^{\mathbb{Q}}(a) &= \mathbb{E}_{\mathcal{F}_s}^{\mathbb{P}} \left[\exp \left(\left((1 - a\varphi_{Z_V}^{\mathbb{P}}(\Gamma_V) \mu_V)^{-1} - 1 \right) \varphi_{Z_V}^{\mathbb{P}}(\Gamma_V) \int_s^t \lambda_{Y,u} du \right) \right]. \end{aligned}$$

³⁹Details are provided in Lemma B.9 of Appendix B.

⁴⁰Indeed,

$$\begin{aligned} dV_t &= \kappa(\theta - V_{t-}) dt + \sigma \sqrt{V_{t-}} dW_{V,t} + dJ_{V,t} \\ &= \kappa(\theta - V_{t-}) - \sigma \sqrt{V_{t-}} \Lambda_{V,t} dt + \sigma \sqrt{V_{t-}} dW_{V,t}^{\mathbb{Q}} + dJ_{V,t}^{\mathbb{Q}} \\ &= (\kappa + \sigma \eta_V) \left(\frac{\kappa \theta}{\kappa + \sigma \eta_V} - V_{t-} \right) dt + \sigma \sqrt{V_{t-}} dW_{V,t}^{\mathbb{Q}} + dJ_{V,t}^{\mathbb{Q}}. \end{aligned}$$

where the correspondence between the \mathbb{P} - and \mathbb{Q} - parameters is

$$\begin{aligned}\alpha_{r^-}^{\mathbb{Q}} &= r - q - \frac{1}{2}V_{r^-} - \left(\varphi_{Z_Y}^{\mathbb{P}}(1) - 1\right)\lambda_{Y_{r^-}}^{\mathbb{Q}}, & \kappa^{\mathbb{Q}} &= \kappa + \sigma\eta_V, \\ \lambda_{Y_{r^-}}^{\mathbb{Q}} &= \varphi_{Z_Y}^{\mathbb{P}}(\Gamma_Y)\lambda_{Y_{r^-}}, & \theta^{\mathbb{Q}} &= \frac{\kappa\theta}{\kappa + \sigma\eta_V}.\end{aligned}$$

5.A.3 Pricing

The price of a European call option with strike price K and maturity T is

$$C_t(Y_t, V_t) = \exp(Y_t) \exp(-q(T-t)) P_1(Y_t, V_t) - K \exp(-r(T-t)) P_2(Y_t, V_t) \quad (5.27)$$

where

$$\begin{aligned}P_1(y, v) &= \frac{1}{2} + \frac{1}{\pi} \int_0^\infty \operatorname{Re} \left(\frac{\exp(-iuk - y) \varphi_{Y_T|Y_t, V_t}^{\mathbb{Q}}(ui + 1, y, v)}{ui} \right) du \\ P_2(y, v) &= \frac{1}{2} + \frac{1}{\pi} \int_0^\infty \operatorname{Re} \left(\frac{\exp(-iuk) \varphi_{Y_T|Y_t, V_t}^{\mathbb{Q}}(ui, y, v)}{ui} \right) du\end{aligned}$$

and $\varphi_{Y_T|Y_t, V_t}^{\mathbb{Q}}(u, y, v) = \mathbb{E}_t^{\mathbb{Q}}[\exp(uY_T) | Y_t = y, V_t = v]$ is the moment generating function of Y_T conditional on time t information. More precisely,

$$\varphi_{Y_T|Y_t, V_t}^{\mathbb{Q}}(u, Y_t, V_t) = \exp(\mathcal{A}(u, t, T) + uY_t + C(u, t, T)V_t)$$

where

$$\begin{aligned}C(u, t, T) &= \frac{2C_0(\exp(-C_2(T-t)) - 1)}{C_2(\exp(-C_2(T-t)) + 1) - C_1(\exp(-C_2(T-t)) - 1)}, \quad (5.28) \\ C_0 &= \lambda_{Y,1}^{\mathbb{Q}}(\varphi_{Z_Y}^{\mathbb{Q}}(1) - 1)u - \lambda_{Y,1}^{\mathbb{Q}}(\varphi_{Z_Y}^{\mathbb{Q}}(u) - 1) + \frac{u - u^2}{2}, \\ C_1 &= \kappa^{\mathbb{Q}} - \rho\sigma u, \\ C_2 &= \sqrt{C_1^2 + 2\sigma^2 C_0},\end{aligned}$$

and

$$\begin{aligned}\mathcal{A}(u; t, T) &= D_0(T-t) + \theta^{\mathbb{Q}}\kappa^{\mathbb{Q}}g_1(t, T) + \lambda_{V,0}^{\mathbb{Q}}g_2(t, T), \quad (5.29) \\ D_0 &= -r + (r-q)u + \lambda_{Y,0}^{\mathbb{Q}}(\varphi_{Z_Y}^{\mathbb{Q}}(u) - 1) - \lambda_{Y,0}^{\mathbb{Q}}(\varphi_{Z_Y}^{\mathbb{Q}}(1) - 1)u, \\ g_1(t, T) &= -\frac{1}{\sigma^2} \left(2\log(2C_2) - 2\log\left(C_1(e^{C_2(T-t)} - 1) + C_2(e^{C_2(T-t)} + 1)\right) \right)\end{aligned}$$

$$\begin{aligned}
 & +(C_1 + C_2)(T - t)), \\
 g_2(t, T) = & \frac{\mu_V^Q}{2C_0(\mu_V^Q)^2 + 2C_1\mu_V^Q - \sigma^2} \left(2\log(2C_2) \right. \\
 & - 2\log\left(C_1\left(e^{C_2(T-t)} - 1\right) + C_2\left(e^{C_2(T-t)} + 1\right) + 2C_0\mu_V^Q\left(e^{C_2(T-t)} - 1\right)\right) \\
 & \left. + 2C_0\mu_V^Q + C_1 - C_2(T - t)\right).
 \end{aligned}$$

The approach is inspired from Heston (1993) that relies on an inversion similar to the one of Gil-Pelaez (1951). The explicit form of the moment generating function is similar to the ones found by Filipovic and Mayerhofer (2009) and Duffie et al. (2000).⁴¹

5.B Estimation

5.B.1 Drift Term

The drift term process of Equation (5.1) is

$$\begin{aligned}
 \alpha_{u^-}^{\mathbb{P}} &= r - q - \frac{1}{2}V_{u^-} + \sqrt{V_{u^-}}\Lambda_{Y,u^-} + \left(\varphi_{Z_Y}^{\mathbb{P}}(\Gamma_Y) - \varphi_{Z_Y}^{\mathbb{P}}(1 + \Gamma_Y)\right)\lambda_{Y,u^-} \\
 &= r - q - \frac{1}{2}V_{u^-} + \sqrt{V_{u^-}}\Lambda_{Y,u^-} + \left(\gamma_Y - \left(\varphi_{Z_Y}^{\mathbb{P}}(1) - 1\right)\right)\lambda_{Y,u^-}
 \end{aligned}$$

where the moment generating functions, $\varphi_{Z_Y}^{\mathbb{P}}(\Gamma_Y)$ and $\varphi_{Z_V}^{\mathbb{P}}(\Gamma_V)$, of the log-equity price and variance jump size are defined at Equation (5.22). To obtain a drift term that depends on the instantaneous variance only, it is assumed that

$$\Lambda_{\perp,u^-} = \eta_{\perp} \sqrt{V_{t^-}}. \quad (5.30)$$

Consequently, letting $\eta_Y = \rho\eta_V + \sqrt{1 - \rho^2}\eta_{\perp}$ yields $\Lambda_{Y,u^-} = (\rho\eta_V + \sqrt{1 - \rho^2}\eta_{\perp}) \sqrt{V_{t^-}}$ and

$$\alpha_{u^-}^{\mathbb{P}} = r - q + \left(\eta_Y - \frac{1}{2}\right)V_{u^-} + \left(\gamma_Y - \left(\varphi_{Z_Y}^{\mathbb{P}}(1) - 1\right)\right)\lambda_{Y,u^-}. \quad (5.31)$$

⁴¹Details are available in Appendix B.2.

Sketch of the proof. Let

$$Z_t = \left(\begin{array}{l} - \int_0^t \Lambda_{V,u^-} dW_{V,u} - \int_0^t \Lambda_{\perp,u^-} dW_{\perp,u} - \frac{1}{2} \int_0^t (\Lambda_{V,u^-}^2 + \Lambda_{\perp,u^-}^2) du \\ + \Gamma_Y J_{Y,t} - (\varphi_{Z_Y}^{\mathbb{P}}(\Gamma_Y) - 1) \int_0^t \lambda_{Y,u^-} du + \Gamma_V J_{V,t} - (\varphi_{Z_V}^{\mathbb{P}}(\Gamma_V) - 1) \int_0^t \lambda_{V,u^-} du \end{array} \right)$$

be associated to the Radon-Nikodym derivative of (5.21) and $D_t = \exp(-(r-q)t)$ be the combination of the discount factor and the dividend yield. Since the discounted price has to be a \mathbb{Q} -martingale, then for all $0 < s < t$,

$$1 = \mathbb{E}_s^{\mathbb{Q}} \left[\frac{D_t \exp(Y_t)}{D_s \exp(Y_s)} \right] = \mathbb{E}_s^{\mathbb{P}} \left[\frac{D_t \exp(Y_t + Z_t)}{D_s \exp(Y_s + Z_s)} \right],$$

which means that $\{D_t \exp(Y_t + Z_t)\}_{t \geq 0}$ is a \mathbb{P} -martingale. The computation of this last expectation leads to the final result.⁴² \square

5.B.2 Filtering Procedure

5.B.2.1 Intraday Simulation

In what follows, h is set to Δ/M implying that the path simulation is performed with M steps per day. Assuming that Y_t and V_t are known, V_{t+h} , $\Delta I_{t,t+h}$ and Y_{t+h} are generated as follows:

1. The jump indicator functions $\mathbf{1}_{\{N_{Y,t+h}-N_{Y,t}=1\}}$ and $\mathbf{1}_{\{N_{V,t+h}-N_{V,t}=1\}}$ are generated from independent Bernoulli random variables with a success probabilities of $\lambda_{Y,t}h$ and $\lambda_{V,t}h$ respectively.⁴³
2. If needed, $Z_{Y,t+h} \sim \mathcal{N}(\mu_Y; \sigma_Y^2)$ and $Z_{V,t+h} \sim \text{Exp}(\mu_V)$ are simulated.
3. Time $t+h$ variance V_{t+h} is simulated according to

$$\begin{aligned} V_{t+h} &\cong V_t + \kappa(\theta - V_t)h + \sigma \sqrt{V_t} (W_{V,t+h} - W_{V,t}) + Z_{V,t+h} \mathbf{1}_{\{N_{V,t+h}-N_{V,t}=1\}}. \\ &= V_{(t+h)^-} + Z_{V,t+h} \mathbf{1}_{\{N_{V,t+h}-N_{V,t}=1\}} \end{aligned} \quad (5.32)$$

where $\mathbf{1}_{\{N_{V,t+h}-N_{V,t}=1\}}$ is an indicator function that is worth 1 if there is a jump during the interval and 0 otherwise. Indeed, it corresponds to the Euler approximation of the variance process (5.2).

⁴²Lemma B.10 of Appendix B provides all details.

⁴³Because of the Poisson process properties, the probability of observing a single jump is approximately $\lambda_{\cdot,t}$ when h is small.

4. The integrated variance increment $\Delta I_{t,t+h} = \int_t^{t+h} V_u^- du$ is simulated conditionally on V_t and $V_{(t+h)^-}$. More precisely, the cumulant generating function of $\Delta I_{t,t+h}$ is

$$\begin{aligned} & \log \left(\mathbb{E}^{\mathbb{P}} [\exp(a \Delta I_{t,t+h}) | V_t, V_{(t+h)^-}] \right) \\ &= \log \left(1 + \sum_{m=1}^{\infty} \mathbb{E}^{\mathbb{P}} [(\Delta I_{t,t+h})^m | V_t, V_{(t+h)^-}] \frac{a^m}{m!} \right) \\ &\cong \mathbb{E}^{\mathbb{P}} [\Delta I_{t,t+h} | V_t, V_{(t+h)^-}] a + \text{Var}^{\mathbb{P}} [\Delta I_{t,t+h} | V_t, V_{(t+h)^-}] \frac{a^2}{2}, \end{aligned}$$

which is the cumulant generating function of a Gaussian distribution. The conditional density function of $\Delta I_{t,t+h}$ is approximated by

$$\frac{1}{\sqrt{2\pi \text{Var}^{\mathbb{P}} [\Delta I_{t,t+h} | V_t, V_{(t+h)^-}]}} \exp \left(-\frac{1}{2} \frac{(x - \mathbb{E}^{\mathbb{P}} [\Delta I_{t,t+h} | V_t, V_{(t+h)^-}])^2}{\text{Var}^{\mathbb{P}} [\Delta I_{t,t+h} | V_t, V_{(t+h)^-}]} \right).$$

The first two centred moments of $\Delta I_{t,t+h}$ are available in closed form in Tse and Wan (2013).

5. The log-price is simulated according to ^{44,45,46}

$$\begin{aligned} Y_{t+h} &= Y_t + c_1 h + c_2 \Delta I_{t,t+h} + \rho \sigma^{-1} (V_{(t+h)^-} - V_t) \\ &\quad + \sqrt{1 - \rho^2} \int_t^{t+h} \sqrt{V_u^-} dW_{\perp,u} - \frac{\rho}{\sigma} \left(\sum_{n=N_{V,t}+1}^{N_{V,t+h}} Z_{V,n} \right) + \sum_{n=N_{Y,t}+1}^{N_{Y,t+h}} Z_{Y,n}. \end{aligned} \quad (5.33)$$

where

$$\begin{aligned} c_1 &= r - q + \left(\varphi_{Z_Y}^{\mathbb{P}}(\Gamma_Y) - \varphi_{Z_Y}^{\mathbb{P}}((1 + \Gamma_Y)) \right) \lambda_{Y,0} - \frac{\rho \kappa \theta}{\sigma}, \\ c_2 &= \eta_Y - \frac{1}{2} + \left(\varphi_{Z_Y}^{\mathbb{P}}(\Gamma_Y) - \varphi_{Z_Y}^{\mathbb{P}}((1 + \Gamma_Y)) \right) \lambda_{Y,1} + \frac{\rho}{\sigma} \kappa. \end{aligned}$$

⁴⁴Again, for small $h > 0$, the two last terms may be approximated with

$$-\frac{\rho}{\sigma} Z_{V,t+h} \mathbf{1}_{\{N_{V,t+h} - N_{V,t} = 1\}} + Z_{Y,t+h} \mathbf{1}_{\{N_{Y,t+h} - N_{Y,t} = 1\}}.$$

⁴⁵As in Broadie and Kaya (2006), $\int_t^{t+h} \sqrt{V_u^-} dW_{\perp,u}$ is approximated with a Gaussian distribution of mean zero and variance $\Delta I_{t,t+h}$. Indeed, since the instantaneous variance is stochastic, the distribution of $\int_t^{t+h} \sqrt{V_u^-} dW_{\perp,u}$ is not truly Gaussian. Though, it would have been the case if the variance was a deterministic function of time. The variance of $\int_t^{t+h} \sqrt{V_u^-} dW_{\perp,u}$ is $\mathbb{E}_t^{\mathbb{P}} \left[\int_t^{t+h} V_u^- du \right] = \mathbb{E}_t^{\mathbb{P}} [\Delta I_{t,t+h}]$.

⁴⁶Note that the log-prices are simulated intraday but the observed log-price is used whenever it is available, that is at the end of each day. These intermediary log-prices are required in the computation of ΔOQV_{t-h} .

In fact, from the integral version of Equation (5.2),

$$\int_t^{t+h} \sqrt{V_{u^-}} dW_{V,u} = \sigma^{-1} \left(V_{(t+h)^-} - V_t - \kappa \theta h + \kappa \int_t^{t+h} V_{u^-} du - \sum_{n=N_{V,t}+1}^{N_{V,t+h}} Z_{V,n} \right).$$

Replacing back in the integral form of the log-equity price (5.1), substituting the drift term of Equation (5.31) and noting that $\lambda_{Y,t^-} = \lambda_{Y,0} + \lambda_{Y,1} V_{t^-}$ completes the construction.

5.B.2.2 Time Aggregation

Let $t = (k-1)\Delta$. Once Steps 1 to 5 of Appendix 5.B.2.1 have been executed for each intraday periods $\{t + jh\}_{j=1}^M$, the integrated variance and quadratic variation increments are approximated by the aggregation of the simulated variables:

$$\Delta I_{t,t+\Delta} \cong \sum_{j=1}^M \Delta I_{t+(j-1)h,t+jh} \text{ and } \Delta QV_{t,t+\Delta} \cong \Delta I_{t,t+\Delta} + \sum_{j=1}^M Z_{Y,t+jh} \mathbf{1}_{\{N_{Y,t+jh} - N_{Y,t+(j-1)h} = 1\}}.$$

Note that a time aggregation of Equation (5.33) suggests the approximation

$$Y_{t+\Delta} - Y_t = \sum_{j=1}^M (Y_{t+jh} - Y_{t+(j-1)h}) \cong \mu_{t+\Delta} + \sigma_{t+\Delta}^2 \varepsilon_{t+h}$$

where ε_{t+h} is a standard normal random variable,

$$\begin{aligned} \mu_{t+\Delta} = & c_1 \Delta + c_2 \Delta I_{t,t+\Delta} + \frac{\rho}{\sigma} (V_{t+\Delta} - V_t) \\ & - \frac{\rho}{\sigma} \sum_{j=1}^M Z_{V,t+jh} \mathbf{1}_{\{N_{V,t+jh} - N_{V,t+(j-1)h} = 1\}} + \sum_{j=1}^M Z_{Y,t+jh} \mathbf{1}_{\{N_{Y,t+jh} - N_{Y,t+(j-1)h} = 1\}} \end{aligned} \quad (5.34)$$

and $\sigma_{t+\Delta}^2 = (1 - \rho^2) \Delta I_{t,t+\Delta}$. The model option implied volatility $IV_{t+\Delta}(Y_{t+\Delta}, V_{t+\Delta})$ is calculated based on Equation (5.27).⁴⁷ Finally, the Euler discretization of Equation

⁴⁷ $IV_{t+\Delta}(Y_{t+\Delta}, V_{t+\Delta})$ is the Black-Scholes implied volatility obtained from the model option price $O_{t,t+\Delta}^{(i)}(Y_{t,t+\Delta}, V_{t,t+\Delta})$.

(5.10) provides an approximation of the option quadratic variation:⁴⁸

$$\begin{aligned} \Delta OQV_{t,t+\Delta} = & h \sum_{j=0}^{M-1} \left(\left(\frac{\partial O_{t+jh}}{\partial y} (Y_{t+jh}, V_{t+jh}) \right)^2 + \sigma^2 \left(\frac{\partial O_{t+jh}}{\partial v} (Y_{t+jh}, V_{t+jh}) \right)^2 \right. \\ & + 2\sigma\rho \frac{\partial O_{t+jh}}{\partial y} (Y_{t+jh}, V_{t+jh}) \frac{\partial O_{t+jh}}{\partial v} (Y_{t+jh}, V_{t+jh}) \left. \right) \Delta I_{t+(j-1)h,t+jh} \\ & + \sum_{j=1}^M \left\{ O_{t+jh} (Y_{t+jh}, V_{t+jh}) - O_{t+jh} (Y_{(t+jh)^-}, V_{(t+jh)^-}) \right\}^2. \end{aligned} \quad (5.35)$$

Appendix B.2.2.1 provides details about how these derivatives are computed. At the end of this stage, a simulated vector

$$x_{t+\Delta} = (Y_{t+\Delta}, V_{t+\Delta}, \Delta I_{t,t+\Delta}, QV_{t,t+\Delta}, O_{t+\Delta}^{(i)}(Y_{t+\Delta}, V_{t+\Delta}), \Delta OQV_{t,t+\Delta}^{(i)} : i \in \{1, 2, \dots, n_t\})$$

is obtained at the end of the day.

References

- Aït-Sahalia, Y., and A. W. Lo. 1998. Nonparametric estimation of state-price densities implicit in financial asset prices. *Journal of Finance* 53:499–547.
- Alizadeh, S., M. W. Brandt, and F. X. Diebold. 2002. Range-based estimation of stochastic volatility models. *Journal of Finance* 57:1047–1091.
- Andersen, T. G., T. Bollerslev, F. X. Diebold, and H. Ebens. 2001. The distribution of realized stock return volatility. *Journal of Financial Economics* 61:43–76.
- Andersen, T. G., O. Bondarenko, V. Todorov, and G. Tauchen. 2015a. The fine structure of equity-index option dynamics. *Journal of Econometrics* 187:532–546.
- Andersen, T. G., N. Fusari, and V. Todorov. 2015b. The risk premia embedded in index options. *Journal of Financial Economics* 117:558–584.
- Audrino, F., and M. R. Fengler. 2015. Are classical option pricing models consistent with observed option second-order moments? Evidence from high-frequency data. *Journal of Banking & Finance* 61:46–63.

⁴⁸The log-equity process is only observable on a daily basis. However, $\Delta OQV_{t+(j-1)h,t+jh}^{(i)}$ requires intraday log-equity values. Therefore, it is convenient to treat intraday log-equity values as latent variables.

- Bakshi, G., C. Cao, and Z. Chen. 1997. Empirical performance of alternative option pricing models. *Journal of Finance* 52:2003–2049.
- Bandi, F. M., and R. Reno. 2016. Price and volatility co-jumps. *Journal of Financial Economics* 119:107–146.
- Bardgett, C., E. Gourier, and M. Leippold. 2015. Inferring volatility dynamics and risk premia from the S&P 500 and VIX markets. *Working paper*.
- Barndorff-Nielsen, O. E., and N. Shephard. 2002. Estimating quadratic variation using realized variance. *Journal of Applied Econometrics* 17:457–477.
- Barndorff-Nielsen, O. E., and N. Shephard. 2004. Power and bipower variation with stochastic volatility and jumps. *Journal of Financial Econometrics* 2:1–37.
- Bates, D. S. 1996. Jumps and stochastic volatility: exchange rate processes implicit in Deutsche Mark options. *Review of Financial Studies* 9:69–107.
- Bates, D. S. 2000. Post-'87 crash fears in the S&P 500 futures option market. *Journal of Econometrics* 94:181–238.
- Bates, D. S. 2006. Maximum likelihood estimation of latent affine processes. *Review of Financial Studies* 19:909–965.
- Black, F., and M. Scholes. 1973. The pricing of options and corporate liabilities. *Journal of Political Economy* 81:637–654.
- Bollerslev, T. 1986. Generalized autoregressive conditional heteroskedasticity. *Journal of Econometrics* 31:307–327.
- Broadie, M., M. Chernov, and M. Johannes. 2007. Model specification and risk premia: Evidence from futures options. *Journal of Finance* 62:1453–1490.
- Broadie, M., and O. Kaya. 2006. Exact simulation of stochastic volatility and other affine jump diffusion processes. *Operations Research* 54:217–231.
- Chernov, M., A. R. Gallant, E. Ghysels, and G. Tauchen. 2003. Alternative models for stock price dynamics. *Journal of Econometrics* 116:225–257.
- Chernov, M., and E. Ghysels. 2000. A study towards a unified approach to the joint estimation of objective and risk neutral measures for the purpose of options valuation. *Journal of Financial Economics* 56:407–458.

- Christoffersen, P., B. Feunou, K. Jacobs, and N. Meddahi. 2014. The economic value of realized volatility: Using high-frequency returns for option valuation. *Journal of Financial and Quantitative Analysis* 49:663–697.
- Christoffersen, P., B. Feunou, and Y. Jeon. 2015. Option valuation with observable volatility and jump dynamics. *Journal of Banking & Finance* 61:S101–S120.
- Christoffersen, P., K. Jacobs, and C. Ornathanalai. 2012. Dynamic jump intensities and risk premiums: Evidence from S&P 500 returns and options. *Journal of Financial Economics* 106:447–472.
- Cont, R., and T. Kokholm. 2013. A consistent pricing model for index options and volatility derivatives. *Mathematical Finance* 23:248–274.
- Diebold, F. X., and R. S. Mariano. 1995. Comparing predictive accuracy. *Journal of Business & Economic Statistics* 13:253–263.
- Duffie, D., J. Pan, and K. Singleton. 2000. Transform analysis and asset pricing for affine jump-diffusions. *Econometrica* 68:1343–1376.
- Duffie, D., and K. Singleton. 1993. Simulated moments estimation of Markov models of asset prices. *Econometrica* 61:929–952.
- Dumas, B., J. Fleming, and R. E. Whaley. 1998. Implied volatility functions: Empirical tests. *Journal of Finance* 53:2059–2106.
- Engle, R. F. 1982. Autoregressive Conditional Heteroscedasticity with Estimates of the Variance of United Kingdom Inflation. *Econometrica* 50:987–1007.
- Eraker, B. 2001. MCMC analysis of diffusion models with application to finance. *Journal of Business & Economic Statistics* 19:177–191.
- Eraker, B. 2004. Do stock prices and volatility jump? Reconciling evidence from spot and option prices. *Journal of Finance* 59:1367–1404.
- Eraker, B., M. Johannes, and N. Polson. 2003. The impact of jumps in volatility and returns. *Journal of Finance* 58:1269–1300.
- Filipovic, D., and E. Mayerhofer. 2009. Affine diffusion processes: Theory and applications. *Advanced Financial Modelling* 8:125.
- Fulop, A., J. Li, and J. Yu. 2014. Self-exciting jumps, learning, and asset pricing implications. *Review of Financial Studies* 28:876–912.

- Gallant, A. R., and G. Tauchen. 2010. Simulated score methods and indirect inference for continuous-time models. *Handbook of Financial Econometrics* 1:427–477.
- Gil-Pelaez, J. 1951. Note on the inversion theorem. *Biometrika* 38:481–482.
- Gordon, N. J., D. J. Salmond, and A. F. Smith. 1993. Novel approach to nonlinear/non-Gaussian Bayesian state estimation. In *IEE Proceedings F (Radar and Signal Processing)*, vol. 140, pp. 107–113.
- Heston, S. 1993. A closed-form solution for options with stochastic volatility with applications to bond and currency options. *Review of Financial Studies* 6:327.
- Hu, F., and J. V. Zidek. 2002. The weighted likelihood. *Canadian Journal of Statistics* 30:347–371.
- Hürzeler, M., and H. R. Künsch. 2001. Approximating and maximising the likelihood for a general state-space model. In *Sequential Monte Carlo Methods in Practice*, pp. 159–175. Springer.
- Jackwerth, J. C., and M. Rubinstein. 1996. Recovering probability distributions from option prices. *Journal of Finance* 51:1611–1631.
- Johannes, M. S., N. G. Polson, and J. R. Stroud. 2009. Optimal filtering of jump diffusions: Extracting latent states from asset prices. *Review of Financial Studies* 22:2759–2799.
- Jones, C. S. 1998. Bayesian estimation of continuous-time finance models. *Working paper*.
- Malik, S., and M. K. Pitt. 2011. Particle filters for continuous likelihood evaluation and maximisation. *Journal of Econometrics* 165:190–209.
- Merton, R. C. 1980. On estimating the expected return on the market: An exploratory investigation. *Journal of Financial Economics* 8:323–361.
- Newey, W., and K. West. 1987. A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix. *Econometrica* 55:703–708.
- Ornthanalai, C. 2014. Levy jump risk: Evidence from options and returns. *Journal of Financial Economics* 112:69–90.
- Pan, J. 2002. The jump-risk premia implicit in options: Evidence from an integrated time-series study. *Journal of Financial Economics* 63:3–50.

- Protter, P. 2004. *Stochastic integration and differential equations*. Springer.
- Renault, E. 1997. Econometric models of option pricing errors. *Econometric Society Monographs* 28:223–278.
- Santa-Clara, P., and S. Yan. 2010. Crashes, volatility, and the equity premium: lessons from S&P 500 options. *Review of Economics and Statistics* 92:435–451.
- Stephan, J. A., and R. E. Whaley. 1990. Intraday price change and trading volume relations in the stock and stock option markets. *Journal of Finance* 45:191–220.
- Todorov, V., and G. Tauchen. 2011. Volatility jumps. *Journal of Business & Economic Statistics* 29:356–371.
- Tse, S. T., and J. W. Wan. 2013. Low-bias simulation scheme for the Heston model by inverse Gaussian approximation. *Quantitative Finance* 13:919–937.
- van Haastrecht, A., and A. Pelsser. 2010. Efficient, almost exact simulation of the Heston stochastic volatility model. *International Journal of Theoretical and Applied Finance* 13:1–43.
- Zhang, L., P. A. Mykland, and Y. Aït-Sahalia. 2005. A tale of two time scales. *Journal of the American Statistical Association* 100:1394–1411.

Chapter 6

Concluding Remarks

Filtering methods are extremely powerful as they allow us to understand the behaviour of latent variables which cannot be directly observed. This thesis presents a few applications of filtering methods in finance. The first two essays focus on credit risk. In the first essay, credit risk is investigated before, during and after the last financial crisis with a flexible credit risk model. A filtering method based on the work of Tugnait (1982) is tailored for the issue at hand. The filter is extended in the second essay to account for co-movements of firm leverages. Systemic risk in both the banking and the insurance subsectors is also analyzed.

In the third essay, a particle filter is used to estimate a discrete time GARCH-type jump-diffusion model. This filter is an adaptation of Gordon et al.'s (1993) sequential importance resampling (SIR) methodology. Based on the estimation results of 260 firms, we find that systematic risk factors explain close to 60% of the risk premium on average, while idiosyncratic factors explain more than 40%. We also show that tail risk plays a central role in the pricing of idiosyncratic risk.

In the fourth essay, the SIR is generalized to incorporate high frequency option information in the estimation of posterior distribution of latent variables. This is done through the introduction of a novel quantity: the option realized variance. Our results show that the information contained in these variances improves the inference of the latent variables such as the instantaneous variance and jumps.

As shown in the present thesis, filtering methods are of paramount importance in finance as most modelling framework embed these latent variables. These techniques are getting more efficient, reliable and flexible with time, making their potential applications virtually limitless.

Yet, a major obstacle to getting better and faster filtering methods is related to their numerical burdens. Fortunately for us, computers are getting faster every year. Break-throughs such as parallel computing and GPU computing already allow us for more efficient schemes. However, it is hard to foresee the technology of the future; for instance, twenty years ago, today's technology was beyond most people's wildest dreams. Nonetheless, one thing is certain: we will see more of these methods in the years to come.

Appendix A

Appendices of *Idiosyncratic Jump Risk Matters: Evidence from Equity Returns and Options*

A.1 Normal-Inverse Gaussian Distribution

The jumps $y_{u,t+1}$ have a NIG distribution with location parameter 0, a scale parameter $h_{u,y,t+1}$, an asymmetry parameter δ_u and a tail heaviness parameter α_u . The first standardized moments are

$$\begin{aligned}\mathbb{E}_t^{\mathbb{P}} [y_{u,t+1}] &= \frac{\delta_u}{\sqrt{\alpha_u^2 - \delta_u^2}} h_{u,y,t+1}, \\ \text{Var}_t^{\mathbb{P}} [y_{u,t+1}] &= \frac{\alpha_u^2}{\left(\sqrt{\alpha_u^2 - \delta_u^2}\right)^3} h_{u,y,t+1}, \\ \text{Skew}_t^{\mathbb{P}} [y_{u,t+1}] &= \frac{3\delta_u}{\alpha_u (\alpha_u^2 - \delta_u^2)^{\frac{1}{4}}} \frac{1}{\sqrt{h_{u,y,t+1}}}\end{aligned}$$

and the excess kurtosis is

$$\text{ExKurt}_t^{\mathbb{P}} [y_{u,t+1}] = 3 \left(1 + \frac{4\delta_u^2}{\alpha_u^2} \right) \frac{1}{\sqrt{\alpha_u^2 - \delta_u^2}} \frac{1}{h_{u,y,t+1}}.$$

The moment generating function is

$$\varphi_{y_{u,t+1}}(\phi) = \exp \left(\left(\sqrt{\alpha_u^2 - \delta_u^2} - \sqrt{\alpha_u^2 - (\delta_u + \phi)^2} \right) h_{u,y,t+1} \right)$$

A.1.1 Interpretation of the Jump Intensity Parameter

For comparison, let N_{t+1} be a Poisson random variable of intensity λ_{t+1} , and consider the compound Poisson random variable $\sum_{j=0}^{N_t} J_j$ where the jumps J_j are independent $\text{NIG}(0, h', \delta', \alpha')$ random variables. The moment generating function of $\sum_{j=0}^{N_t} J_j$ is

$$\begin{aligned}
 & \varphi_{\sum_{j=0}^{N_t} J_j}(\phi) \\
 = & \exp(-\lambda_t) \\
 & + \sum_{j=1}^{\infty} \exp\left(j\left(\sqrt{(\alpha')^2 - (\delta')^2} - \sqrt{(\alpha')^2 - (\delta' + \phi)^2}\right)h'\right) \exp(-\lambda_t) \frac{\lambda_t^j}{j!} \\
 = & \exp(-\lambda_t) \sum_{j=0}^{\infty} \left[\lambda_t \exp\left(\left(\sqrt{(\alpha')^2 - (\delta')^2} - \sqrt{(\alpha')^2 - (\delta' + \phi)^2}\right)h'\right) \right]^j \frac{1}{j!} \\
 = & \exp\left(\lambda_t \left[\exp\left(\left(\sqrt{(\alpha')^2 - (\delta')^2} - \sqrt{(\alpha')^2 - (\delta' + \phi)^2}\right)h'\right) - 1 \right]\right) \\
 \cong & \exp\left(\lambda_t \left(\sqrt{(\alpha')^2 - (\delta')^2} - \sqrt{(\alpha')^2 - (\delta' + \phi)^2} \right)h'\right)
 \end{aligned}$$

where the last approximation holds from a first order Taylor expansion, provided that h' is close to zero. Letting $\alpha' = \alpha_u^2$, and $\delta' = \delta_u$, a direct comparison between $\varphi_{\sum_{j=0}^{N_t} J_j}(\phi)$ and $\varphi_{y_{u,t+1}}(\phi)$ implies that

$$h_{u,y,t+1} \cong \lambda_t h',$$

that is $h_{u,y,t+1}$ may be interpreted as a scaled version of the jump intensity.

A.1.2 Returns' Conditional Moments

The conditional moment generation function of $az_{M,t+1} + by_{M,t+1} + c(z_{S,t+1} + y_{S,t+1})$ is

$$\begin{aligned}
 & \mathbb{E}_t^{\mathbb{P}} [\exp(\phi (az_{M,t+1} + by_{M,t+1} + c(z_{S,t+1} + y_{S,t+1})))] \\
 = & \mathbb{E}_t^{\mathbb{P}} [\exp(\phi az_{M,t+1})] \mathbb{E}_t^{\mathbb{P}} [\exp(\phi by_{M,t+1})] \mathbb{E}_t^{\mathbb{P}} [\exp(\phi cz_{S,t+1})] \mathbb{E}_t^{\mathbb{P}} [\exp(\phi cy_{S,t+1})] \\
 = & \exp\left(\frac{a^2 \phi^2}{2} h_{M,z,t} + \Pi_M(b\phi) h_{M,y,t} + \frac{c^2 \phi^2}{2} h_{S,z,t} + \Pi_S(c\phi) h_{S,y,t}\right)
 \end{aligned}$$

where

$$\Pi_u(\phi) = \sqrt{\alpha_u^2 - \delta_u^2} - \sqrt{\alpha_u^2 - (\delta_u + \phi)^2}.$$

Note that

$$\begin{aligned}\frac{\partial \Pi_u}{\partial \phi}(\phi) &= \frac{(\delta_u + \phi)}{\sqrt{\alpha_u^2 - (\delta_u + \phi)^2}} & \frac{\partial^2 \Pi_u}{\partial \phi^2}(\phi) &= \frac{\alpha_u^2}{(\alpha_u^2 - (\delta_u + \phi)^2)^{\frac{3}{2}}} \\ \frac{\partial^3 \Pi_u}{\partial \phi^3}(\phi) &= 3 \frac{\alpha_u^2(\delta_u + \phi)}{(\alpha_u^2 - (\delta_u + \phi)^2)^{\frac{5}{2}}} & \frac{\partial^4 \Pi_u}{\partial \phi^4}(\phi) &= 3\alpha_u^2 \frac{\alpha_u^2 + 4(\delta_u + \phi)^2}{(\alpha_u^2 - (\delta_u + \phi)^2)^{\frac{7}{2}}}\end{aligned}$$

The cumulant generating function is therefore

$$\xi(\phi; a, b, c) = \frac{a^2 \phi^2}{2} h_{M,z,t} + \Pi_M(b\phi) h_{M,y,t} + \frac{c^2 \phi^2}{2} h_{S,z,t} + \Pi_S(c\phi) h_{S,y,t}$$

Note that

$$\begin{aligned}\frac{\partial \xi}{\partial \phi}(\phi; a, b, c) &= a^2 \phi h_{M,z,t} + b \frac{\partial \Pi_M}{\partial \phi}(b\phi) h_{M,y,t} + c^2 \phi h_{S,z,t} + c \frac{\partial \Pi_S}{\partial \phi}(c\phi) h_{S,y,t}, \\ \frac{\partial^2 \xi}{\partial \phi^2}(\phi; a, b, c) &= a^2 h_{M,z,t} + b^2 \frac{\partial^2 \Pi_M}{\partial \phi^2}(b\phi) h_{M,y,t} + c^2 h_{S,z,t} + c^2 \frac{\partial^2 \Pi_S}{\partial \phi^2}(c\phi) h_{S,y,t}, \\ \frac{\partial^3 \xi}{\partial \phi^3}(\phi; a, b, c) &= b^3 \frac{\partial^3 \Pi_M}{\partial \phi^3}(b\phi) h_{M,y,t} + c^3 \frac{\partial^3 \Pi_S}{\partial \phi^3}(c\phi) h_{S,y,t}, \\ \frac{\partial^4 \xi}{\partial \phi^4}(\phi; a, b, c) &= b^4 \frac{\partial^4 \Pi_M}{\partial \phi^4}(b\phi) h_{M,y,t} + c^4 \frac{\partial^4 \Pi_S}{\partial \phi^4}(c\phi) h_{S,y,t}.\end{aligned}$$

The first moment of the market and stock returns are

$$\begin{aligned}\mathbb{E}_t^{\mathbb{P}}[R_{M,t+1}] &= \mu_{M,t+1}^{\mathbb{P}} - \xi_{M,t+1}^{\mathbb{P}} + \frac{\partial \xi}{\partial \phi}(0; 1, 1, 0), \\ \mathbb{E}_t^{\mathbb{P}}[R_{S,t+1}] &= \mu_{S,t+1}^{\mathbb{P}} - \xi_{S,t+1}^{\mathbb{P}} + \frac{\partial \xi}{\partial \phi}(0; \beta_{S,z}, \beta_{S,y}, 1).\end{aligned}$$

Their variances correspond to

$$\begin{aligned}\text{Var}_t^{\mathbb{P}}[R_{M,t+1}] &= \text{Var}_t^{\mathbb{P}}[z_{M,t+1} + y_{M,t+1}] = \frac{\partial^2 \xi}{\partial \phi^2}(0; 1, 1, 0) \\ \text{Var}_t^{\mathbb{P}}[R_{S,t+1}] &= \text{Var}_t^{\mathbb{P}}[\beta_{S,z} z_{M,t+1} + \beta_{S,y} y_{M,t+1} + z_{S,t+1} + y_{S,t+1}] = \frac{\partial^2 \xi}{\partial \phi^2}(0; \beta_{S,z}, \beta_{S,y}, 1).\end{aligned}$$

Similarly, since the third cumulant corresponds to the third centered moment, the third standardized moment are respectively

$$\text{Skew}_t^{\mathbb{P}}[R_{M,t+1}] = \frac{\frac{\partial^3 \xi}{\partial \phi^3}(0; 1, 1, 0)}{\left(\frac{\partial^2 \xi}{\partial \phi^2}(0; 1, 1, 0)\right)^{\frac{3}{2}}} \text{ and } \text{Skew}_t^{\mathbb{P}}[R_{S,t+1}] = \frac{\frac{\partial^3 \xi}{\partial \phi^3}(0; \beta_{S,z}, \beta_{S,y}, 1)}{\left(\frac{\partial^2 \xi}{\partial \phi^2}(0; \beta_{S,z}, \beta_{S,y}, 1)\right)^{\frac{3}{2}}}.$$

Finally, the excess kurtosis are

$$\begin{aligned} \mathbb{E}_t^{\mathbb{P}} \left[\left(\frac{R_{M,t+1} - \mathbb{E}_t^{\mathbb{P}} [R_{M,t+1}]}{\sqrt{\text{Var}_t^{\mathbb{P}} [R_{M,t+1}]}} \right)^4 \right] - 3 &= \frac{\frac{\partial^4 \xi}{\partial \phi^4} (0; 1, 1, 0)}{\left(\frac{\partial^2 \xi}{\partial \phi^2} (0; 1, 1, 0) \right)^2}, \\ \mathbb{E}_t^{\mathbb{P}} \left[\left(\frac{R_{M,t+1} - \mathbb{E}_t^{\mathbb{P}} [R_{M,t+1}]}{\sqrt{\text{Var}_t^{\mathbb{P}} [R_{M,t+1}]}} \right)^4 \right] - 3 &= \frac{\frac{\partial^4 \xi}{\partial \phi^4} (0; \beta_{S,z}, \beta_{S,y}, 1)}{\left(\frac{\partial^2 \xi}{\partial \phi^2} (0; \beta_{S,z}, \beta_{S,y}, 1) \right)^2}. \end{aligned}$$

A.1.3 Conditional Variance and Jump Intensity Variable Moments

Lemma A.1.

$$\begin{aligned} \text{Var}_{t-1}^{\mathbb{P}} [h_{M,z,t+1}] &= 2a_{M,z}^2 (1 + 2c_{M,z}^2 h_{M,z,t}), \\ \text{Var}_{t-1}^{\mathbb{P}} [h_{M,y,t+1}] &= 2a_{M,y}^2 (1 + 2c_{M,y}^2 h_{M,z,t}), \\ \text{Var}_{t-1}^{\mathbb{P}} [h_{S,z,t+1}] &= \kappa_{S,z}^2 \text{Var}_{t-1}^{\mathbb{P}} [h_{M,z,t+1}] + 2a_{S,z}^2 (1 + 2c_{S,z}^2 h_{S,z,t}), \\ \text{Var}_{t-1}^{\mathbb{P}} [h_{S,y,t+1}] &= \kappa_{S,y}^2 \text{Var}_{t-1}^{\mathbb{P}} [h_{M,y,t+1}] + 2a_{S,y}^2 (1 + 2c_{S,y}^2 h_{S,z,t}). \end{aligned}$$

Proof. Recall that the market conditional variance is

$$h_{M,z,t+1} = w_{M,z} + b_{M,z} h_{M,z,t} + a_{M,z} (\varepsilon_{M,t} - c_{M,z} \sqrt{h_{M,z,t}})^2.$$

Therefore, $\mathbb{E}_{t-1}^{\mathbb{P}} [h_{M,z,t+1}] = w_{M,z} + b_{M,z} h_{M,z,t} + a_{M,z} (1 + c_{M,z}^2 h_{M,z,t})$ and

$$\begin{aligned} &\text{Var}_{t-1}^{\mathbb{P}} [h_{M,z,t+1}] \\ &= a_{M,z}^2 \mathbb{E}_{t-1}^{\mathbb{P}} \left[\left((\varepsilon_{M,t} - c_{M,z} \sqrt{h_{M,z,t}})^2 - (1 + c_{M,z}^2 h_{M,z,t}) \right)^2 \right] \\ &= a_{M,z}^2 \mathbb{E}_{t-1}^{\mathbb{P}} \left[\varepsilon_{M,t}^4 - 4c_{M,z} \sqrt{h_{M,z,t}} \varepsilon_{M,t}^3 + 2(2c_{M,z}^2 h_{M,z,t} - 1) \varepsilon_{M,t}^2 + 4c_{M,z} \sqrt{h_{M,z,t}} \varepsilon_{M,t} + 1 \right] \\ &= a_{M,z}^2 (3 + 2(2c_{M,z}^2 h_{M,z,t} - 1) + 1) \\ &= 2a_{M,z}^2 (1 + 2c_{M,z}^2 h_{M,z,t}). \end{aligned}$$

The market jump scale parameter is

$$h_{M,y,t+1} = w_{M,y} + b_{M,y} h_{M,y,t} + a_{M,y} (\varepsilon_{M,t} - c_{M,y} \sqrt{h_{M,z,t}})^2.$$

Hence, $\mathbb{E}_{t-1}^{\mathbb{P}} [h_{M,y,t+1}] = w_{M,y} + b_{M,y}h_{M,y,t} + a_{M,y}(1 + c_{M,y}^2 h_{M,z,t})$ and

$$\begin{aligned}\text{Var}_{t-1}^{\mathbb{P}} [h_{M,y,t+1}] &= a_{M,y}^2 \mathbb{E}_{t-1}^{\mathbb{P}} \left[\left((\varepsilon_{M,t} - c_{M,y} \sqrt{h_{M,z,t}})^2 - (1 + c_{M,y}^2 h_{M,z,t}) \right)^2 \right] \\ &= 2a_{M,y}^2 (1 + 2c_{M,y}^2 h_{M,z,t}).\end{aligned}$$

The stock conditional variance satisfies

$$h_{S,z,t+1} = \kappa_{S,z} h_{M,z,t+1} + b_{S,z} (h_{S,z,t} - \kappa_{S,z} h_{M,z,t}) + a_{S,z} (\varepsilon_{S,t}^2 - 1 - 2c_{S,z} \sqrt{h_{S,z,t}} \varepsilon_{S,t}).$$

Therefore, $\mathbb{E}_{t-1}^{\mathbb{P}} [h_{S,z,t+1}] = \kappa_{S,z} \mathbb{E}_{t-1}^{\mathbb{P}} [h_{M,z,t+1}] + b_{S,z} (h_{S,z,t} - \kappa_{S,z} h_{M,z,t})$ and

$$\begin{aligned}\text{Var}_{t-1}^{\mathbb{P}} [h_{S,z,t+1}] &= \mathbb{E}_{t-1}^{\mathbb{P}} \left[\left(\kappa_{S,z} (h_{M,z,t+1} - \mathbb{E}_{t-1}^{\mathbb{P}} [h_{M,z,t+1}]) + a_{S,z} (\varepsilon_{S,t}^2 - 1 - 2c_{S,z} \sqrt{h_{S,z,t}} \varepsilon_{S,t}) \right)^2 \right] \\ &= \kappa_{S,z}^2 \mathbb{E}_{t-1}^{\mathbb{P}} \left[(h_{M,z,t+1} - \mathbb{E}_{t-1}^{\mathbb{P}} [h_{M,z,t+1}])^2 \right] + a_{S,z}^2 \mathbb{E}_{t-1}^{\mathbb{P}} \left[(\varepsilon_{S,t}^2 - 1 - 2c_{S,z} \sqrt{h_{S,z,t}} \varepsilon_{S,t})^2 \right] \\ &\quad + 2\kappa_{S,z} a_{S,z} \mathbb{E}_{t-1}^{\mathbb{P}} \left[(h_{M,z,t+1} - \mathbb{E}_{t-1}^{\mathbb{P}} [h_{M,z,t+1}]) (\varepsilon_{S,t}^2 - 1 - 2c_{S,z} \sqrt{h_{S,z,t}} \varepsilon_{S,t}) \right] \\ &= \kappa_{S,z}^2 \text{Var}_{t-1}^{\mathbb{P}} [h_{M,z,t+1}] + 2a_{S,z}^2 (1 + 2c_{S,z}^2 h_{S,z,t}) \\ &\quad + 2\kappa_{S,z} a_{S,z} \mathbb{E}_{t-1}^{\mathbb{P}} \left[(\varepsilon_{M,t}^2 - 1 - 2c_{M,z} \sqrt{h_{M,z,t}} \varepsilon_{M,t}) (\varepsilon_{S,t}^2 - 1 - 2c_{S,z} \sqrt{h_{S,z,t}} \varepsilon_{S,t}) \right] \\ &= \kappa_{S,z}^2 \text{Var}_{t-1}^{\mathbb{P}} [h_{M,z,t+1}] + 2a_{S,z}^2 (1 + 2c_{S,z}^2 h_{S,z,t}) \\ &\quad + 2\kappa_{S,z} a_{S,z} \left(\begin{aligned} &\mathbb{E}_{t-1}^{\mathbb{P}} [(\varepsilon_{M,t}^2 - 1)] \mathbb{E}_{t-1}^{\mathbb{P}} [(\varepsilon_{S,t}^2 - 1)] \\ &- 2c_{S,z} \sqrt{h_{S,z,t}} \mathbb{E}_{t-1}^{\mathbb{P}} [(\varepsilon_{M,t}^2 - 1)] \mathbb{E}_{t-1}^{\mathbb{P}} [\varepsilon_{S,t}] \\ &- 2c_{M,z} \sqrt{h_{M,z,t}} \mathbb{E}_{t-1}^{\mathbb{P}} [(\varepsilon_{M,t}) (\varepsilon_{S,t}^2 - 1)] \mathbb{E}_{t-1}^{\mathbb{P}} [(\varepsilon_{M,t}) (\varepsilon_{S,t}^2 - 1)] \\ &+ 4c_{M,z} c_{S,z} \sqrt{h_{M,z,t}} \sqrt{h_{S,z,t}} \mathbb{E}_{t-1}^{\mathbb{P}} [\varepsilon_{M,t}] \mathbb{E}_{t-1}^{\mathbb{P}} [\varepsilon_{S,t}] \end{aligned} \right) \\ &= \kappa_{S,z}^2 \text{Var}_{t-1}^{\mathbb{P}} [h_{M,z,t+1}] + 2a_{S,z}^2 (1 + 2c_{S,z}^2 h_{S,z,t}).\end{aligned}$$

Finally,

$$h_{S,y,t+1} = \kappa_{S,y} h_{M,y,t+1} + b_{S,y} (h_{S,y,t} - \kappa_{S,y} h_{M,y,t}) + a_{S,y} (\varepsilon_{S,t}^2 - 1 - 2c_{S,y} \sqrt{h_{S,y,t}} \varepsilon_{S,t})$$

implies that $\mathbb{E}_{t-1}^{\mathbb{P}} [h_{S,y,t+1}] = \kappa_{S,y} \mathbb{E}_{t-1}^{\mathbb{P}} [h_{M,y,t+1}] + b_{S,y} (h_{S,y,t} - \kappa_{S,y} h_{M,y,t})$ and

$$\begin{aligned}\text{Var}_{t-1}^{\mathbb{P}} [h_{S,y,t+1}] &= \mathbb{E}_{t-1}^{\mathbb{P}} \left[\left(\kappa_{S,y} (h_{M,y,t+1} - \mathbb{E}_{t-1}^{\mathbb{P}} [h_{M,y,t+1}]) + a_{S,y} (\varepsilon_{S,t}^2 - 1 - 2c_{S,y} \sqrt{h_{S,y,t}} \varepsilon_{S,t}) \right)^2 \right] \\ &= \kappa_{S,y}^2 \text{Var}_{t-1}^{\mathbb{P}} [h_{M,y,t+1}] + 2a_{S,y}^2 (1 + 2c_{S,y}^2 h_{S,y,t}).\end{aligned}$$

A.2 Detailed Proofs of Section 4.B's Results

Proof of Lemma 4.1. The conditional cumulant generating function of $\varepsilon_{u,t}$ under \mathbb{Q} is

$$\begin{aligned}
 & \xi_{\varepsilon_{u,t}}^{\mathbb{Q}}(\phi) \\
 &= \log \mathbb{E}_{t-1}^{\mathbb{Q}} [\exp(\phi \varepsilon_{u,t})] \\
 &= \log \mathbb{E}_{t-1}^{\mathbb{P}} \left[\frac{\exp(-\Lambda_M z_{M,t} - \Gamma_M y_{M,t} - \sum_{S \in \mathbb{S}} \Lambda_S z_{S,t} - \sum_{S \in \mathbb{S}} \Gamma_S y_{S,t})}{\mathbb{E}_{t-1}^{\mathbb{P}} [\exp(-\Lambda_M z_{M,t} - \Gamma_M y_{M,t} - \sum_{S \in \mathbb{S}} \Lambda_S z_{S,t} - \sum_{S \in \mathbb{S}} \Gamma_S y_{S,t})]} \exp(\phi \varepsilon_{u,t}) \right] \\
 &= \log \mathbb{E}_{t-1}^{\mathbb{P}} \left[\frac{\exp(-\Lambda_u z_{u,t})}{\mathbb{E}_{t-1}^{\mathbb{P}} [\exp(-\Lambda_u z_{u,t})]} \exp(\phi \varepsilon_{u,t}) \right] \\
 &= \xi_{\varepsilon_{u,t}}^{\mathbb{P}}(\phi - \Lambda_u \sqrt{h_{u,z,t}}) - \xi_{\varepsilon_{u,t}}^{\mathbb{P}}(-\Lambda_u \sqrt{h_{u,z,t}}) \\
 &= \left(\frac{1}{2} (\phi - \Lambda_u \sqrt{h_{u,z,t}})^2 - \frac{1}{2} (-\Lambda_u \sqrt{h_{u,z,t}})^2 \right) \\
 &= \left(\frac{1}{2} \phi^2 - \Lambda_u \sqrt{h_{u,z,t}} \phi \right).
 \end{aligned}$$

□

Proof of Lemma 4.2. For any $y_{u,t} \in \{y_{M,t}, y_{S,t} : S \in \mathbb{S}\}$, the conditional cumulant generating function of $y_{u,t}$ under \mathbb{Q} is

$$\begin{aligned}
 & \xi_{y_{u,t}}^{\mathbb{Q}}(\phi) \\
 &= \log \mathbb{E}_{t-1}^{\mathbb{Q}} [\exp(\phi y_{u,t})] \\
 &= \log \mathbb{E}_{t-1}^{\mathbb{P}} \left[\frac{\exp(-\Lambda_M z_{M,t} - \Gamma_M y_{M,t} - \sum_{S \in \mathbb{S}} \Lambda_S z_{S,t} - \sum_{S \in \mathbb{S}} \Gamma_S y_{S,t})}{\mathbb{E}_{t-1}^{\mathbb{P}} [\exp(-\Lambda_M z_{M,t} - \Gamma_M y_{M,t} - \sum_{S \in \mathbb{S}} \Lambda_S z_{S,t} - \sum_{S \in \mathbb{S}} \Gamma_S y_{S,t})]} \exp(\phi y_{u,t}) \right] \\
 &= \log \mathbb{E}_{t-1}^{\mathbb{P}} \left[\frac{\exp(-\Gamma_u y_{u,t})}{\mathbb{E}_{t-1}^{\mathbb{P}} [\exp(-\Gamma_u y_{u,t})]} \exp(\phi y_{u,t}) \right] \\
 &= \xi_{y_{u,t}}^{\mathbb{P}}(\phi - \Gamma_u) - \xi_{y_{u,t}}^{\mathbb{P}}(-\Gamma_u) \\
 &= \Pi_u(\phi - \Gamma_u) h_{u,y,t} - \Pi_u(-\Gamma_u) h_{u,y,t} \\
 &= \left(\sqrt{\alpha_u^2 - (\delta_u - \Gamma_u)^2} - \sqrt{\alpha_u^2 - (\delta_u + \phi - \Gamma_u)^2} \right) h_{u,y,t}
 \end{aligned}$$

which is the cumulant generating function of a NIG of parameter $\mu_u^* = \mu_u = 0$, $\alpha_u^* = \alpha_u$, $\delta_u^* = \delta_u - \Gamma_u$ and $h_{u,y,t}^* = h_{u,y,t}$. □

A.3 Detailed Proofs of Section 4.C's Results

A.3.1 Market Drift Under \mathbb{P}

Recall that

$$\log\left(\frac{M_t}{M_{t-1}}\right) = R_{M,t} = \mu_{M,t}^{\mathbb{P}} - \xi_{M,t}^{\mathbb{P}} + z_{M,t} + y_{M,t}$$

Since the discounted stock price should behave as a \mathbb{Q} -martingale,

$$\begin{aligned} 1 &= \mathbb{E}_{t-1}^{\mathbb{Q}} \left[\frac{\exp(-r_t) M_t}{M_{t-1}} \right] \\ &= \mathbb{E}_{t-1}^{\mathbb{Q}} [\exp(R_{M,t} - r_t)] \\ &= \mathbb{E}_{t-1}^{\mathbb{P}} \left[\frac{\exp(-\Lambda_M z_{M,t} - \Gamma_M y_{M,t} - \sum_{S \in \mathbb{S}} \Lambda_S z_{S,t} - \sum_{S \in \mathbb{S}} \Gamma_S y_{S,t})}{\mathbb{E}_{t-1}^{\mathbb{P}} [\exp(-\Lambda_M z_{M,t} - \Gamma_M y_{M,t} - \sum_{S \in \mathbb{S}} \Lambda_S z_{S,t} - \sum_{S \in \mathbb{S}} \Gamma_S y_{S,t})]} \exp(R_{M,t} - r_t) \right] \\ &= \mathbb{E}_{t-1}^{\mathbb{P}} \left[\frac{\exp(-\Lambda_M z_{M,t} - \Gamma_M y_{M,t})}{\mathbb{E}_{t-1}^{\mathbb{P}} [\exp(-\Lambda_M z_{M,t} - \Gamma_M y_{M,t})]} \exp(\mu_{M,t}^{\mathbb{P}} - \xi_{M,t}^{\mathbb{P}} + z_{M,t} + y_{M,t} - r_t) \right] \\ &= \mathbb{E}_{t-1}^{\mathbb{P}} \left[\frac{\exp(\mu_{M,t}^{\mathbb{P}} - r_t - \xi_{M,t}^{\mathbb{P}} + (1 - \Lambda_M) z_{M,t} + (1 - \Gamma_M) y_{M,t})}{\mathbb{E}_{t-1}^{\mathbb{P}} [\exp(-\Lambda_M z_{M,t} - \Gamma_M y_{M,t})]} \right] \\ &= \exp \left(\begin{aligned} &\mu_{M,t}^{\mathbb{P}} - r_t - \xi_{M,t}^{\mathbb{P}} (1) - \xi_{y_{M,t}}^{\mathbb{P}} (1) \\ &+ \xi_{z_{M,t}}^{\mathbb{P}} (1 - \Lambda_M) + \xi_{y_{M,t}}^{\mathbb{P}} (1 - \Gamma_M) - \xi_{z_{M,t}}^{\mathbb{P}} (-\Lambda_M) - \xi_{y_{M,t}}^{\mathbb{P}} (-\Gamma_M) \end{aligned} \right). \end{aligned}$$

Because,

$$\begin{aligned} & -\xi_{z_{M,t}}^{\mathbb{P}} (1) + \xi_{z_{M,t}}^{\mathbb{P}} (1 - \Lambda_M) - \xi_{z_{M,t}}^{\mathbb{P}} (-\Lambda_M) \\ &= -\frac{1}{2} h_{M,z,t} + \frac{1}{2} h_{M,z,t} (1 - \Lambda_M)^2 - \frac{1}{2} h_{M,z,t} \Lambda_M^2 = -\Lambda_M h_{M,z,t} \end{aligned}$$

and

$$-\xi_{y_{M,t}}^{\mathbb{P}} (1) + \xi_{y_{M,t}}^{\mathbb{P}} (1 - \Gamma_M) - \xi_{y_{M,t}}^{\mathbb{P}} (-\Gamma_M) = -h_{M,y,t} \gamma_M$$

where

$$\begin{aligned} \gamma_M &= \sqrt{\alpha_M^2 - \delta_M^2} - \sqrt{\alpha_M^2 - (\delta_M + 1)^2} + \sqrt{\alpha_M^2 - (\delta_M + 1 - \Gamma_M)^2} \\ &\quad - \sqrt{\alpha_M^2 - (\delta_M - \Gamma_M)^2} \\ &= \Pi_M(1) - \Pi_M^*(1), \end{aligned}$$

we conclude that

$$1 = \exp(\mu_{M,t}^{\mathbb{P}} - r_t - \Lambda_M h_{M,z,t} - h_{M,y,t} \gamma_M).$$

Therefore,

$$\mu_{M,t}^{\mathbb{P}} = r_t + \Lambda_M h_{M,z,t} + \gamma_M h_{M,y,t}.$$

A.3.2 Stock Drift Under \mathbb{P}

Recall that

$$\log \left(\frac{S_t}{S_{t-1}} \right) = R_{S,t} = \mu_{S,t}^{\mathbb{P}} - \xi_{S,t}^{\mathbb{P}} + \beta_{S,z} z_{M,t} + \beta_{S,y} y_{M,t} + z_{S,t} + y_{S,t}$$

Since the discounted stock price should behave as a \mathbb{Q} -martingale,

$$\begin{aligned} & 1 \\ &= \mathbb{E}_{t-1}^{\mathbb{Q}} \left[\frac{\exp(-r_t) S_t}{S_{t-1}} \right] \\ &= \mathbb{E}_{t-1}^{\mathbb{Q}} [\exp(R_{S,t} - r_t)] \\ &= \mathbb{E}_{t-1}^{\mathbb{P}} \left[\frac{\exp(-\Lambda_M z_{M,t} - \Gamma_M y_{M,t} - \sum_{S \in \mathbb{S}} \Lambda_S z_{S,t} - \sum_{S \in \mathbb{S}} \Gamma_S y_{S,t})}{\mathbb{E}_{t-1}^{\mathbb{P}} [\exp(-\Lambda_M z_{M,t} - \Gamma_M y_{M,t} - \sum_{S \in \mathbb{S}} \Lambda_S z_{S,t} - \sum_{S \in \mathbb{S}} \Gamma_S y_{S,t})]} \exp(R_{S,t} - r_t) \right] \\ &= \mathbb{E}_{t-1}^{\mathbb{P}} \left[\frac{\exp \left(-\Lambda_M z_{M,t} - \Gamma_M y_{M,t} - \sum_{S \in \mathbb{S}} \Lambda_S z_{S,t} - \sum_{S \in \mathbb{S}} \Gamma_S y_{S,t} \right)}{\mathbb{E}_{t-1}^{\mathbb{P}} [\exp(-\Lambda_M z_{M,t} - \Gamma_M y_{M,t} - \sum_{S \in \mathbb{S}} \Lambda_S z_{S,t} - \sum_{S \in \mathbb{S}} \Gamma_S y_{S,t})]} \right. \\ &\quad \left. + \mu_{S,t}^{\mathbb{P}} - \xi_{S,t}^{\mathbb{P}} + \beta_{S,z} z_{M,t} + \beta_{S,y} y_{M,t} + z_{S,t} + y_{S,t} - r_t \right] \\ &= \mathbb{E}_{t-1}^{\mathbb{P}} \left[\frac{\exp \left(\begin{aligned} & \mu_{S,t}^{\mathbb{P}} - r_t - \xi_{S,t}^{\mathbb{P}} + (\beta_{S,z} - \Lambda_M) z_{M,t} \\ & + (\beta_{S,y} - \Gamma_M) y_{M,t} + (1 - \Lambda_S) z_{S,t} + (1 - \Gamma_S) y_{S,t} \end{aligned} \right)}{\mathbb{E}_{t-1}^{\mathbb{P}} [\exp(-\Lambda_M z_{M,t} - \Gamma_M y_{M,t} - \Lambda_S z_{S,t} - \Gamma_S y_{S,t})]} \right] \\ &= \mathbb{E}_{t-1}^{\mathbb{P}} \left[\frac{\exp \left(\begin{aligned} & \mu_{S,t}^{\mathbb{P}} - r_t - \xi_{z_{M,t}}^{\mathbb{P}} (\beta_{S,z}) - \xi_{y_{M,t}}^{\mathbb{P}} (\beta_{S,y}) - \xi_{z_{S,t}}^{\mathbb{P}} (1) - \xi_{y_{S,t}}^{\mathbb{P}} (1) \\ & + \xi_{z_{M,t}}^{\mathbb{P}} (\beta_{S,z} - \Lambda_M) + \xi_{y_{M,t}}^{\mathbb{P}} (\beta_{S,y} - \Gamma_M) + \xi_{z_{S,t}}^{\mathbb{P}} (1 - \Lambda_S) + \xi_{y_{S,t}}^{\mathbb{P}} (1 - \Gamma_S) \\ & - \xi_{z_{M,t}}^{\mathbb{P}} (-\Lambda_M) - \xi_{y_{M,t}}^{\mathbb{P}} (-\Gamma_M) - \xi_{z_{S,t}}^{\mathbb{P}} (-\Lambda_S) - \xi_{y_{S,t}}^{\mathbb{P}} (-\Gamma_S) \end{aligned} \right)}{\mathbb{E}_{t-1}^{\mathbb{P}} [\exp(-\Lambda_M z_{M,t} - \Gamma_M y_{M,t} - \Lambda_S z_{S,t} - \Gamma_S y_{S,t})]} \right]. \end{aligned}$$

Because,

$$\begin{aligned} -\xi_{z_{M,t}}^{\mathbb{P}} (\beta_{S,z}) + \xi_{z_{M,t}}^{\mathbb{P}} (\beta_{S,z} - \Lambda_M) - \xi_{z_{M,t}}^{\mathbb{P}} (-\Lambda_M) &= -\frac{1}{2} h_{M,z,t} \beta_{S,z}^2 + \frac{1}{2} h_{M,z,t} (\beta_{S,z} - \Lambda_M)^2 \\ &\quad - \frac{1}{2} h_{M,z,t} \Lambda_M^2 \\ &= -\Lambda_M \beta_{S,z} h_{M,z,t}, \end{aligned}$$

$$\begin{aligned}
 -\xi_{zS,t}^{\mathbb{P}}(1) + \xi_{zS,t}^{\mathbb{P}}(1 - \Lambda_S) - \xi_{zS,t}^{\mathbb{P}}(-\Lambda_S) &= -\frac{1}{2}h_{S,z,t} + \frac{1}{2}h_{S,z,t}(1 - \Lambda_S)^2 - \frac{1}{2}h_{S,z,t}\Lambda_S^2 \\
 &= -\Lambda_S h_{S,z,t}, \\
 -\xi_{yM,t}^{\mathbb{P}}(\beta_{S,y}) + \xi_{yM,t}^{\mathbb{P}}(\beta_{S,y} - \Gamma_M) - \xi_{yM,t}^{\mathbb{P}}(-\Gamma_M) &= \left(-\Pi_M(\beta_{S,y}) + \Pi_M(\beta_{S,y} - \Gamma_M) \right. \\
 &\quad \left. - \Pi_M(-\Gamma_M) \right) h_{M,y,t} \\
 &= \Pi_M^*(\beta_{S,y}) - \Pi_M(\beta_{S,y}) \\
 &= -\gamma_{M,S}(\beta_{S,y}) h_{M,y,t} \\
 -\xi_{yS,t}^{\mathbb{P}}(1) + \xi_{yS,t}^{\mathbb{P}}(1 - \Gamma_S) - \xi_{yS,t}^{\mathbb{P}}(-\Gamma_S) &= \Pi_S^*(1) - \Pi_S(1) \\
 &= -\gamma_S h_{S,y,t},
 \end{aligned}$$

we conclude that

$$1 = \exp\left(\mu_{S,t}^{\mathbb{P}} - r_t - \Lambda_M \beta_{S,z} h_{M,z,t} - h_{M,y,t} \gamma_{M,S}(\beta_{S,y}) - \Lambda_S h_{S,z,t} - h_{S,y,t} \gamma_S\right).$$

where

$$\gamma_{M,S}(\beta_{S,y}) = \Pi_M(\beta_{S,y}) - \Pi_M^*(\beta_{S,y}) \text{ and } \gamma_S = \Pi_S(1) - \Pi_S^*(1).$$

Therefore,

$$\mu_{S,t}^{\mathbb{P}} = r_t + \Lambda_M \beta_{S,z} h_{M,z,t} + h_{M,y,t} \gamma_{M,S}(\beta_{S,y}) + \Lambda_S h_{S,z,t} + h_{S,y,t} \gamma_S.$$

A.4 Calculation Associated with Section 4.D

Replacing (4.27) in the market conditional variance leads to

$$\begin{aligned}
 h_{M,z,t+1} &= w_{M,z} + b_{M,z} h_{M,z,t} + a_{M,z} \left(\varepsilon_{M,t} - c_{M,z} \sqrt{h_{M,z,t}} \right)^2 \\
 &= w_{M,z} + b_{M,z} h_{M,z,t} + a_{M,z} \left(\left(\varepsilon_{M,t}^* - \Lambda_M \sqrt{h_{M,z,t}} \right) - c_{M,z} \sqrt{h_{M,z,t}} \right)^2 \\
 &= w_{M,z} + b_{M,z} h_{M,z,t} + a_{M,z} \left(\varepsilon_{M,t}^* - (c_{M,z} + \Lambda_M) \sqrt{h_{M,z,t}} \right)^2 \\
 &= w_{M,z} + b_{M,z} h_{M,z,t} + a_{M,z} \left(\left(\varepsilon_{M,t}^* \right)^2 - 2(c_{M,z} + \Lambda_M) \sqrt{h_{M,z,t}} \varepsilon_{M,t}^* + (c_{M,z} + \Lambda_M)^2 h_{M,z,t} \right) \\
 &= w_{M,z} + \left(b_{M,z} + a_{M,z} (c_{M,z} + \Lambda_M)^2 \right) h_{M,z,t} + a_{M,z} \left(\varepsilon_{M,t}^* \right)^2 - 2a_{M,z} (c_{M,z} + \Lambda_M) \sqrt{h_{M,z,t}} \varepsilon_{M,t}^*.
 \end{aligned}$$

A similar argument leads to the stock conditional variance:

$$h_{S,z,t+1} = \kappa_{S,z} h_{M,z,t+1} + b_{S,z} (h_{S,z,t} - \kappa_{S,z} h_{M,z,t}) + a_{S,z} \left(\varepsilon_{S,t}^2 - 1 - 2c_{S,z} z_{S,t} \right)$$

$$\begin{aligned}
 &= \kappa_{S,z} h_{M,z,t+1}^* + b_{S,z} \left(h_{S,z,t}^* - \kappa_{S,z} h_{M,z,t}^* \right) \\
 &\quad + a_{S,z} \left((2c_{S,z} + \Lambda_S) \Lambda_S h_{S,z,t}^* + \left((\varepsilon_{S,t}^*)^2 - 1 - 2(c_{S,z} + \Lambda_S) \sqrt{h_{S,z,t}^*} \varepsilon_{S,t}^* \right) \right) \\
 &= \kappa_{S,z} h_{M,z,t+1}^* + b_{S,z} \left(h_{S,z,t}^* - \kappa_{S,z} h_{M,z,t}^* \right) + a_{S,z} (2c_{S,z} + \Lambda_S) \Lambda_S h_{S,z,t}^* \\
 &\quad + a_{S,z} \left((\varepsilon_{S,t}^*)^2 - a_{S,z} - 2a_{S,z} (c_{S,z} + \Lambda_S) \sqrt{h_{S,z,t}^*} \varepsilon_{S,t}^* \right) \\
 &= \pi_{S,z,1} + \pi_{S,z,2} h_{M,z,t}^* + 0h_{S,y,t}^* + \pi_{S,z,4} h_{S,z,t}^* + 0h_{S,y,t}^* \\
 &\quad + \pi_{S,z,6} \left(\varepsilon_{M,t}^* \right)^2 + \pi_{S,z,7} \sqrt{h_{M,z,t}^*} \varepsilon_{M,t}^* + \pi_{S,z,8} \left(\varepsilon_{S,t}^* \right)^2 + \pi_{S,z,9} \sqrt{h_{S,z,t}^*} \varepsilon_{S,t}^*.
 \end{aligned}$$

The risk-neutral market jump scale parameter is

$$\begin{aligned}
 h_{M,y,t} &= w_{M,y} + b_{M,y} h_{M,y,t} + a_{M,y} \left(\varepsilon_{M,t} - c_{M,y} \sqrt{h_{M,z,t}} \right)^2 \\
 &= w_{M,y} + a_{M,y} c_{M,y}^2 h_{M,z,t} + b_{M,y} h_{M,y,t} + a_{M,y} \left(\varepsilon_{M,t}^2 - 2c_{M,y} \sqrt{h_{M,z,t}} \varepsilon_{M,t} \right) \\
 &= \left(\begin{aligned} &w_{M,y} + a_{M,y} c_{M,y}^2 h_{M,z,t} + b_{M,y} h_{M,y,t} \\ &+ a_{M,y} \left(\left(\varepsilon_{M,t}^* - \Lambda_M \sqrt{h_{M,z,t}} \right)^2 - 2c_{M,y} \sqrt{h_{M,z,t}} \left(\varepsilon_{M,t}^* - \Lambda_M \sqrt{h_{M,z,t}} \right) \right) \end{aligned} \right) \\
 &= w_{M,y} + a_{M,y} \left(c_{M,y} + \Lambda_M \right)^2 h_{M,z,t} + b_{M,y} h_{M,y,t} \\
 &\quad + a_{M,y} \left(\left(\varepsilon_{M,t}^* \right)^2 - 2(c_{M,y} + \Lambda_M) \sqrt{h_{M,z,t}} \varepsilon_{M,t}^* \right) \\
 &= w_{M,y} + a_{M,y} \left(c_{M,y} + \Lambda_M \right)^2 h_{M,z,t}^* + b_{M,y} h_{M,y,t}^* \\
 &\quad + a_{M,y} \left(\left(\varepsilon_{M,t}^* \right)^2 - 2(c_{M,y} + \Lambda_M) \sqrt{h_{M,z,t}^*} \varepsilon_{M,t}^* \right) \\
 &= \pi_{M,y,1} + \pi_{M,y,2} h_{M,z,t}^* + \pi_{M,y,3} h_{M,y,t}^* + 0h_{S,z,t}^* + 0h_{S,y,t}^* \\
 &\quad + \pi_{M,y,6} \left(\varepsilon_{M,t}^* \right)^2 + \pi_{M,y,7} \sqrt{h_{M,z,t}^*} \varepsilon_{M,t}^* + 0 \left(\varepsilon_{S,t}^* \right)^2 + 0 \sqrt{h_{S,z,t}^*} \varepsilon_{S,t}^*.
 \end{aligned}$$

Similarly, the risk-neutral stock jump scale parameter is

$$\begin{aligned}
 h_{S,y,t} &= \kappa_{S,y} h_{M,y,t+1} + b_{S,y} \left(h_{S,y,t} - \kappa_{S,y} h_{M,y,t} \right) + a_{S,y} \left(\varepsilon_{S,t}^2 - 1 - 2c_{S,y} \sqrt{h_{S,z,t}} \varepsilon_{S,t} \right) \\
 &= \kappa_{S,y} h_{M,y,t+1}^* + b_{S,y} \left(h_{S,y,t}^* - \kappa_{S,y} h_{M,y,t}^* \right) \\
 &\quad + a_{S,y} \left((2c_{S,y} + \Lambda_S) \Lambda_S h_{S,z,t}^* + \left((\varepsilon_{S,t}^*)^2 - 1 - 2(c_{S,y} + \Lambda_S) \sqrt{h_{S,z,t}^*} \varepsilon_{S,t}^* \right) \right) \\
 &= \kappa_{S,y} h_{M,y,t+1}^* - a_{S,y} - b_{S,y} \kappa_{S,y} h_{M,y,t}^* + b_{S,y} h_{S,y,t}^* + a_{S,y} (2c_{S,y} + \Lambda_S) \Lambda_S h_{S,z,t}^* \\
 &\quad + a_{S,y} \left((\varepsilon_{S,t}^*)^2 - 2a_{S,y} (c_{S,y} + \Lambda_S) \sqrt{h_{S,z,t}^*} \varepsilon_{S,t}^* \right) \\
 &= \pi_{S,y,1} + \pi_{S,y,2} h_{M,z,t}^* + \pi_{S,y,3} h_{M,y,t}^* + \pi_{S,y,4} h_{S,z,t}^* + \pi_{S,y,5} h_{S,y,t}^* \\
 &\quad + \pi_{S,y,6} \left(\varepsilon_{M,t}^* \right)^2 + \pi_{S,y,7} \sqrt{h_{M,z,t}^*} \varepsilon_{M,t}^* + \pi_{S,y,8} \left(\varepsilon_{S,t}^* \right)^2 + \pi_{S,y,9} \sqrt{h_{S,z,t}^*} \varepsilon_{S,t}^*.
 \end{aligned}$$

A.5 Detailed Proofs of Section 4.E's Results

Lemma A.2. *if ε represents a standard normal random variable, then*

$$\mathbb{E} \left[\exp(a\varepsilon^2 + b\varepsilon) \right] = \exp \left(-\frac{1}{2} \log(1 - 2a) + \frac{1}{2} \frac{b^2}{(1 - 2a)} \right)$$

provided that $a < \frac{1}{2}$.

Proof of Lemma A.2.

$$\begin{aligned} & \mathbb{E} \left[\exp(a\varepsilon^2 + b\varepsilon) \right] \\ &= \int \exp(a\varepsilon^2 + b\varepsilon) \frac{1}{\sqrt{2\pi}} \exp\left(-\frac{1}{2}\varepsilon^2\right) d\varepsilon \\ &= \int \frac{1}{\sqrt{2\pi}} \exp\left(-\frac{(1-2a)}{2} \left(\varepsilon^2 - 2\frac{b}{(1-2a)}\varepsilon\right)\right) d\varepsilon \\ &= \exp\left(\frac{1}{2} \frac{b^2}{(1-2a)}\right) \sqrt{\frac{1}{(1-2a)}} \int \frac{1}{\sqrt{2\pi}} \frac{1}{\sqrt{\frac{1}{(1-2a)}}} \exp\left(-\frac{1}{2} \left(\varepsilon - \frac{b}{(1-2a)}\right)^2\right) d\varepsilon. \end{aligned}$$

If $1 - 2a > 0$, then the integral is one since it corresponds to the area under the density function of a $\mathcal{N}\left(\frac{b}{(1-2a)}, \frac{1}{(1-2a)}\right)$ random variable. Hence,

$$\mathbb{E} \left[\exp(a\varepsilon^2 + b\varepsilon) \right] = \exp \left(\log \sqrt{\frac{1}{(1-2a)}} + \frac{1}{2} \frac{b^2}{(1-2a)} \right).$$

A.5.1 Moment Generating Function

For $u \in \{M, S\}$, The risk neutral returns process is

$$\begin{aligned} \log \left(\frac{u_{t+1}}{u_t} \right) &= R_{u,t+1} = r_{t+1} - \xi_{u,t+1}^Q + \beta_{u,z} z_{M,t+1}^* + \beta_{u,y} y_{M,t+1}^* + \beta'_{u,z} z_{S,t+1}^* + \beta'_{u,y} y_{S,t+1}^* \\ z_{u,t+1}^* &= \sqrt{h_{u,z,t+1}^*} \varepsilon_{u,t+1}^*, \quad \varepsilon_{u,t+1}^* \sim \mathcal{N}(0, 1) \\ y_{u,t+1}^* &\sim \text{NIG}(0, \alpha_u^*, \delta_u^2, h_{u,z,t+1}^*) \end{aligned}$$

where the convexity correction¹ is

$$\xi_{u,t}^{\mathbb{Q}} = \xi_{z_{M,t}^*}^{\mathbb{Q}}(\beta_{u,z}) + \xi_{y_{M,t}^*}^{\mathbb{Q}}(\beta_{u,y}) + \xi_{z_{S,t}^*}^{\mathbb{Q}}(\beta'_{u,z}) + \xi_{y_{S,t}^*}^{\mathbb{Q}}(\beta'_{u,y}).$$

For the market case, $\beta_{M,z} = \beta_{M,y} = 1$ and $\beta'_{M,z} = \beta'_{M,y} = 0$. For the stock, $\beta'_{S,z} = \beta'_{S,y} = 1$.

Proof of Lemma 4.5. For $u \in \{M, S\}$, let $\tilde{R}_{u,t+j}$ denotes the excess return. Its risk neutral dynamics is

$$\begin{aligned} \tilde{R}_{u,t+j} &= R_{u,t+j} - r_{t+j} \\ &= -\frac{1}{2}\beta_{u,z}^2 h_{M,z,t+j}^* - \Pi_M^*(\beta_{u,y}) h_{M,y,t+j}^* - \frac{1}{2}(\beta'_{u,z})^2 h_{S,z,t+j}^* - \Pi_S^*(\beta'_{u,y}) h_{S,y,t+j}^* \\ &\quad + \beta_{u,z} z_{M,t+j}^* + \beta_{u,y} y_{M,t+j}^* + \beta'_{u,z} z_{S,t+j}^* + \beta'_{u,y} y_{S,t+j}^*. \end{aligned}$$

For the market case, $\beta_{M,z} = \beta_{M,y} = 1$ and $\beta'_{M,z} = \beta'_{M,y} = 0$. For the stock, $\beta'_{S,z} = \beta'_{S,y} = 1$. The proof is based on a backward recursion over time. Indeed, the moment generating function of $\sum_{j=1}^{T-t} \tilde{R}_{u,t+j}$ given \mathcal{F}_t^S is

$$\begin{aligned} &\varphi_{\tilde{R},t,T}^{\mathbb{Q}}(\phi) \\ &= \mathbb{E}_t^{\mathbb{Q}} \left[\exp \left(\phi \sum_{j=1}^{T-t} \tilde{R}_{u,t+j} \right) \right] \\ &= \mathbb{E}_t^{\mathbb{Q}} \left[\exp(\phi \tilde{R}_{u,t+1}) \mathbb{E}_{t+1}^{\mathbb{Q}} \left[\exp \left(\phi \sum_{j=1}^{T-t-1} \tilde{R}_{u,t+1+j} \right) \right] \right] \\ &= \mathbb{E}_t^{\mathbb{Q}} \left[\exp \left(\phi \left(-\frac{1}{2}\beta_{u,z}^2 h_{M,z,t+1}^* - \Pi_M^*(\beta_{u,y}) h_{M,y,t+1}^* - \frac{1}{2}(\beta'_{u,z})^2 h_{S,z,t+1}^* \right. \right. \right. \\ &\quad \left. \left. \left. - \Pi_S^*(\beta'_{u,y}) h_{S,y,t+1}^* + \beta_{u,z} z_{M,t+1}^* + \beta_{u,y} y_{M,t+1}^* + \beta'_{u,z} z_{S,t+1}^* + \beta'_{u,y} y_{S,t+1}^* \right) \right) \right. \\ &\quad \left. \times \exp \left(\mathcal{A}_{u,T-t-1}(\phi) + \mathcal{B}_{u,T-t-1}(\phi) h_{M,z,t+2}^* + \mathcal{C}_{u,T-t-1}(\phi) h_{M,y,t+2}^* \right. \right. \\ &\quad \left. \left. + \mathcal{D}_{u,T-t-1}(\phi) h_{S,z,t+2}^* + \mathcal{E}_{u,T-t-1}(\phi) h_{S,y,t+2}^* \right) \right] \end{aligned}$$

from the induction hypothesis. Therefore,

$$\varphi_{\tilde{R},t,T}^{\mathbb{Q}}(\phi)$$

¹Section 4.B shows that

$$\xi_{z_{u,t}^*}^{\mathbb{Q}}(\phi) = \frac{\phi^2}{2} h_{u,z,t}^* \text{ and } \xi_{y_{u,t}^*}^{\mathbb{Q}}(\phi) = \Pi_u^*(\phi) h_{u,y,t}^*$$

where $\Pi_u^*(\phi) = \exp\left(\frac{1}{2}\delta_u^2 \phi^2 + \alpha_u^* \phi\right)$ and $\alpha_u^* = \alpha_u - \delta_u^2 \Gamma_u$.

$$\begin{aligned}
 &= \mathbb{E}_t^Q \left[\exp \left(\phi \left(\begin{aligned} &-\frac{1}{2} \beta_{u,z}^2 h_{M,z,t+1}^* - \Pi_M^* (\beta_{u,y}) h_{M,y,t+1}^* - \frac{1}{2} (\beta'_{u,z})^2 h_{S,z,t+1}^* \\ &-\Pi_S^* (\beta'_{u,y}) h_{S,y,t+1}^* + \beta_{u,z} z_{M,t+1}^* + \beta_{u,y} y_{M,t+1}^* + \beta'_{u,z} z_{S,t+1}^* \\ &+ \beta'_{u,y} y_{S,t+1}^* \end{aligned} \right) \right) \right] \\
 &= \exp \left(\begin{aligned} &\mathcal{A}_{u,T-t-1}(\phi) + \zeta_{u,T-t-1,1}(\phi) + \zeta_{u,T-t-1,2}(\phi) h_{M,z,t+1}^* \\ &+ \zeta_{u,T-t-1,3}(\phi) h_{M,y,t+1}^* + \zeta_{u,T-t-1,4}(\phi) h_{S,z,t+1}^* + \zeta_{u,T-t-1,5}(\phi) h_{S,y,t+1}^* \\ &+ \zeta_{u,T-t-1,6}(\phi) (\varepsilon_{M,t+1}^*)^2 + \zeta_{u,T-t-1,7}(\phi) \sqrt{h_{M,z,t+1}^*} \varepsilon_{M,t+1}^* \\ &+ \zeta_{u,T-t-1,8}(\phi) (\varepsilon_{S,t+1}^*)^2 + \zeta_{u,T-t-1,9}(\phi) \sqrt{h_{S,z,t+1}^*} \varepsilon_{S,t+1}^* \end{aligned} \right) \\
 &= \exp \left(\begin{aligned} &\mathcal{A}_{u,T-t-1}(\phi) + \zeta_{u,T-t-1,1}(\phi) \\ &+ (\zeta_{u,T-t-1,2}(\phi) - \frac{1}{2} \beta_{u,z}^2 \phi) h_{M,z,t+1}^* + (\zeta_{u,T-t-1,3}(\phi) - \Pi_M^* (\beta_{u,y}) \phi) h_{M,y,t+1}^* \\ &+ (\zeta_{u,T-t-1,4}(\phi) - \frac{1}{2} (\beta'_{u,z})^2 \phi) h_{S,z,t+1}^* + (\zeta_{u,T-t-1,5}(\phi) - \Pi_S^* (\beta'_{u,y}) \phi) h_{S,y,t+1}^* \end{aligned} \right) \\
 &\quad \mathbb{E}_t^Q \left[\exp \left(\begin{aligned} &\zeta_{u,T-t-1,6}(\phi) (\varepsilon_{M,t+1}^*)^2 + (\zeta_{u,T-t-1,7}(\phi) + \beta_{u,z} \phi) \sqrt{h_{M,z,t+1}^*} \varepsilon_{M,t+1}^* \\ &+ \beta_{u,y} \phi y_{M,t+1}^* + \zeta_{u,T-t-1,8}(\phi) (\varepsilon_{S,t+1}^*)^2 \\ &+ (\zeta_{u,T-t-1,9}(\phi) + \beta'_{u,z} \phi) \sqrt{h_{S,z,t+1}^*} \varepsilon_{S,t+1}^* + \beta'_{u,y} \phi y_{S,t+1}^* \end{aligned} \right) \right].
 \end{aligned}$$

But the moment generating function of the risk-neutral jump component Equation (4.24) gives $\mathbb{E}_{\mathcal{F}_t^S}^Q [\exp(\beta \phi y_{u,t+1}^*)] = \exp(\Pi_u^* (\beta \phi) h_{u,y,t+1}^*)$. Therefore,

$$\mathbb{E}_{\mathcal{F}_t^S}^Q [\exp(\beta_{u,y} \phi y_{M,t+1}^*)] = \exp(\Pi_M^* (\beta_{u,y} \phi) h_{M,y,t+1}^*)$$

and

$$\mathbb{E}_{\mathcal{F}_t^S}^Q [\exp(\beta'_{u,y} \phi y_{S,t+1}^*)] = \exp(\Pi_S^* (\beta'_{u,y} \phi) h_{S,y,t+1}^*).$$

Moreover, Lemma A.2 implies that

$$\begin{aligned}
 &\mathbb{E}_t^Q \left[\exp \left(\zeta_{u,T-t-1,6}(\phi) (\varepsilon_{M,t+1}^*)^2 + (\zeta_{u,T-t-1,7}(\phi) + \beta_{u,z} \phi) \sqrt{h_{M,z,t+1}^*} \varepsilon_{M,t+1}^* \right) \right] \\
 &= \exp \left(-\frac{1}{2} \ln(1 - 2\zeta_{u,T-t-1,6}(\phi)) + \frac{1}{2} \frac{(\zeta_{u,T-t-1,7}(\phi) + \beta_{u,z} \phi)^2}{(1 - 2\zeta_{u,T-t-1,6}(\phi))} h_{M,z,t+1}^* \right)
 \end{aligned}$$

if $\zeta_{u,T-t-1,6}(\phi) < \frac{1}{2}$ and

$$\begin{aligned}
 &\mathbb{E}_t^Q \left[\exp \left(\zeta_{u,T-t-1,8}(\phi) (\varepsilon_{S,t+1}^*)^2 + (\zeta_{u,T-t-1,9}(\phi) + \beta'_{u,z} \phi) \sqrt{h_{S,z,t+1}^*} \varepsilon_{S,t+1}^* \right) \right] \\
 &= \exp \left(-\frac{1}{2} \ln(1 - 2\zeta_{u,T-t-1,8}(\phi)) + \frac{1}{2} \frac{(\zeta_{u,T-t-1,9}(\phi) + \beta'_{u,z} \phi)^2}{(1 - 2\zeta_{u,T-t-1,8}(\phi))} h_{S,z,t+1}^* \right)
 \end{aligned}$$

provided that $\zeta_{u,T-t-1,8}(\phi) < \frac{1}{2}$. Therefore,

$$\varphi_{\bar{R},t,T}^Q(\phi)$$

$$\begin{aligned}
 &= \exp \left(\begin{aligned} &\mathcal{A}_{u,T-t-1,T}(\phi) + \zeta_{u,T-t-1,1}(\phi) \\ &+ \left(\zeta_{u,T-t-1,2}(\phi) - \frac{1}{2} \beta_{u,z}^2 \phi \right) h_{M,z,t+1}^* + \left(\zeta_{u,T-t-1,3}(\phi) - \Pi_M^*(\beta_{u,y}) \phi \right) h_{M,y,t+1}^* \\ &+ \left(\zeta_{u,T-t-1,4}(\phi) - \frac{1}{2} (\beta'_{u,z})^2 \phi \right) h_{S,z,t+1}^* + \left(\zeta_{u,T-t-1,5}(\phi) - \Pi_S^*(\beta'_{u,y}) \phi \right) h_{S,y,t+1}^* \end{aligned} \right) \\
 &\times \exp \left(-\frac{1}{2} \ln(1 - 2\zeta_{u,T-t-1,6}(\phi)) + \frac{1}{2} \frac{(\zeta_{u,T-t-1,7}(\phi) + \beta_{u,z}\phi)^2}{(1 - 2\zeta_{u,T-t-1,6}(\phi))} h_{M,z,t+1}^* \right) \\
 &\times \exp(\Pi_M^*(\beta_{u,y}\phi) h_{M,y,t+1}^*) \\
 &\times \exp \left(-\frac{1}{2} \ln(1 - 2\zeta_{u,T-t-1,8}(\phi)) + \frac{1}{2} \frac{(\zeta_{u,T-t-1,9}(\phi) + \beta'_{u,z}\phi)^2}{(1 - 2\zeta_{u,T-t-1,8}(\phi))} h_{S,z,t+1}^* \right) \\
 &\times \exp(\Pi_S^*(\beta'_{u,y}\phi) h_{S,y,t+1}^*).
 \end{aligned}$$

A comparison with

$$\varphi_{\tilde{R},t,T}^{\mathbb{Q}}(\phi) = \exp \left(\begin{aligned} &\mathcal{A}_{u,T-t}(\phi) + \mathcal{B}_{u,T-t}(\phi) h_{M,z,t+1}^* + \mathcal{C}_{u,T-t}(\phi) h_{M,y,t+1}^* \\ &+ \mathcal{D}_{u,T-t}(\phi) h_{S,z,t+1}^* + \mathcal{E}_{u,T-t}(\phi) h_{S,y,t+1}^* \end{aligned} \right)$$

leads to the recursive formulation of the coefficients.

A.5.2 European Call Option Price

Given the moment generating function (mgf) $\varphi_{\tilde{R},t,T}^{\mathbb{Q}}(\phi)$ of the risk-neutral excess returns, we build on the work of Heston and Nandi (2000) and obtain a closed-form solution for the price of European index and stock options. More precisely, for $u_t \in \{M_t, S_t\}$, the price of an European call option is

$$\begin{aligned}
 C_t(u_t, K, T) &= e^{-r_{t,T}(T-t)} \mathbb{E}_t^{\mathbb{Q}}[\max(u_T - K, 0)] \\
 &= e^{-r_{t,T}(T-t)} \left(\mathbb{E}_t^{\mathbb{Q}}[u_T \mathcal{I}\{u_T > K\}] - K \mathbb{E}_t^{\mathbb{Q}}[\mathcal{I}\{u_T > K\}] \right),
 \end{aligned}$$

where $r_{t,T} = \frac{1}{T-t} \sum_{j=1}^{T-t} r_{t+j}$, in which r_{t+j} is the deterministic risk-free rate at time $t+j$ and $\mathcal{I}\{A\}$ is the indicator function that worth 1 if the event A is realized and 0 otherwise. Since

$$u_T = u_t \exp \left(\sum_{j=1}^{T-t} R_{u,t+j} \right) = u_t \exp(r_{t,T}(T-t)) \exp \left(\sum_{j=1}^{T-t} \tilde{R}_{u,t+j} \right),$$

then

$$P_{2,t,T} = \mathbb{E}_t^{\mathbb{Q}}[\mathcal{I}\{u_T > K\}]$$

$$\begin{aligned}
 &= \mathbb{Q} \left[u_t e^{r_{t,T}(T-t)} \exp \left(\sum_{j=1}^{T-t} \tilde{R}_{u,t+j} \right) > K \right] \\
 &= \mathbb{Q} \left[\sum_{j=1}^{T-t} \tilde{R}_{u,t+j} > \ln \tilde{K}_{t,T} \right] \\
 &= 1 - F_{\tilde{R},t,T}^{\mathbb{Q}} (\log(\tilde{K}_{t,T}))
 \end{aligned}$$

where $\tilde{K}_{t,T} = \frac{K e^{-r_{t,T}(T-t)}}{u_t}$, and $F_{\tilde{R},t,T}^{\mathbb{Q}}$ is the cumulative distribution function associated with the MGF $\varphi_{\tilde{R},t,T}^{\mathbb{Q}}(\phi)$. Using results from Feller (1971) and Kendall and Stuart (1977),

$$P_{2,t,T} = \frac{1}{2} + \frac{1}{\pi} \int_0^\infty \operatorname{Re} \left[\frac{1}{\phi i} e^{-i\phi \log \tilde{K}_t} \varphi_{\tilde{R},t,T}^{\mathbb{Q}}(\phi i) \right] d\phi.$$

Moreover,

$$\mathbb{E}_t^{\mathbb{Q}} [u_T \mathcal{I}\{u_T > K\}] = u_t e^{r_{t,T}(T-t)} \mathbb{E}_t^{\mathbb{Q}} \left[\exp \left(\sum_{j=1}^{T-t} \tilde{R}_{u,t+j} \right) \mathcal{I}\{u_T > K\} \right] = u_t e^{r_{t,T}(T-t)} P_{1,t,T}$$

where

$$P_{1,t,T} = \mathbb{E}_t^{\mathbb{Q}} \left[\exp \left(\sum_{j=1}^{T-t} \tilde{R}_{u,t+j} \right) \mathcal{I}\{u_T > K\} \right] = \int_{\log(\tilde{K}_{t,T})}^\infty \exp(x) f_{\tilde{R},t,T}^{\mathbb{Q}}(x) dx = \int_{\log(\tilde{K})}^\infty \tilde{p}(x) dx$$

where $f_{\tilde{R},t,T}^{\mathbb{Q}}$ is the density function of the excess returns and the last equality stands by letting $\tilde{p}(x) = \exp(x) f_{\tilde{R},t,T}^{\mathbb{Q}}(x)$. Note that \tilde{p} is a well-defined distribution since $\exp(x)$ is always positive and

$$\int_{-\infty}^\infty \tilde{p}(x) dx = \int_{-\infty}^\infty \exp(x) p(x) dx = \varphi_{\tilde{R},t,T}^{\mathbb{Q}}(1) = 1,$$

given that $f_{\tilde{R},t,T}^{\mathbb{Q}}(1)$ is the gross expected excess return under \mathbb{Q} , that is

$$\int_{-\infty}^\infty \tilde{p}(x) dx = \mathbb{E}_t^{\mathbb{Q}} \left[\exp \left(\sum_{j=1}^{T-t} \tilde{R}_{u,t+j} \right) \right] = \mathbb{E}_t^{\mathbb{Q}} \left[\exp(-r_{t,T}(T-t)) u_T | u_t = 1 \right] = 1.$$

Hence, following Heston and Nandi (2000), we note that mgf corresponding to \tilde{p} is simply

$$\tilde{\varphi}_{t,T}^{\mathbb{Q}}(\phi) = \int_{-\infty}^\infty \exp(\phi x) \tilde{p}(x) dx = \int_{-\infty}^\infty \exp(\phi x) \exp(x) \varphi_{\tilde{R},t,T}^{\mathbb{Q}}(x) dx = \varphi_{t,T}^{\mathbb{Q}}(\phi + 1)$$

and thus

$$\begin{aligned} P_{1,t,T} &= \frac{1}{2} + \frac{1}{\pi} \int_0^\infty \operatorname{Re} \left[\frac{1}{\phi i} e^{-i\phi \log \tilde{K}} \tilde{\varphi}_{t,T}^{\mathbb{Q}}(\phi i) \right] d\phi. \\ &= \frac{1}{2} + \frac{1}{\pi} \int_0^\infty \operatorname{Re} \left[\frac{1}{\phi i} e^{-i\phi \log \tilde{K}} \varphi_{t,T}^{\mathbb{Q}}(\phi i + 1) \right] d\phi. \end{aligned}$$

Finally, we have that

$$C_t(u_t, K, T) = u_T P_{1,t,T} - K e^{-r_{t,T}(T-t)} P_{2,t,T}.$$

Appendix B

Appendices of *Extracting Latent States with High Frequency Option Prices*

B.1 Model Construction

B.1.1 General Properties of Moment Generating Functions

The jumps component are model using Cox processes. Moment generating function for such processes are involved in many proofs.

Lemma B.1. *Let $\varphi_X^{\mathbb{P}}(a) = \mathbb{E}^{\mathbb{P}}[\exp(aX)]$ be the moment generating function of the random variable X . Assume that $\{X_k\}_{k=1}^{\infty}$ is a sequence of independent and identically distributed random variables. Then*

$$\varphi_{\sum_{i=1}^n X_i}^{\mathbb{P}}(a) = \prod_{i=1}^n \mathbb{E}^{\mathbb{P}}[\exp(aX_i)] = \left(\varphi_{X_1}^{\mathbb{P}}(a)\right)^n. \quad (\text{B.1})$$

If N is a Poisson distributed random variable of expectation λ , then the law of iterated expectation implies that

$$\begin{aligned} \varphi_{\sum_{i=1}^N X_i}^{\mathbb{P}}(a) &= \mathbb{E}^{\mathbb{P}} \left[\mathbb{E}^{\mathbb{P}} \left[\exp \left(a \sum_{k=1}^N X_k \right) \middle| N \right] \right] \\ &= \sum_{n=0}^{\infty} \left(\varphi_{X_1}^{\mathbb{P}}(a) \right)^n \exp(-\lambda) \frac{\lambda^n}{n!} = \exp \left(\lambda \left(\varphi_{X_1}^{\mathbb{P}}(a) - 1 \right) \right). \end{aligned} \quad (\text{B.2})$$

If $\{N_t\}_{t \geq 0}$ is a Cox process with predictable intensity $\{\lambda_t\}_{t \geq 0}$, then

$$\begin{aligned} \varphi_{\sum_{i=1}^{N_t} X_i}^{\mathbb{P}}(a) &= \mathbb{E}^{\mathbb{P}} \left[\mathbb{E}^{\mathbb{P}} \left[\exp \left(a \sum_{i=1}^{N_t} X_i \right) \middle| \int_0^t \lambda_u du \right] \right] \\ &= \mathbb{E}^{\mathbb{P}} \left[\exp \left(\left(\int_0^t \lambda_u du \right) (\varphi_{X_1}^{\mathbb{P}}(a) - 1) \right) \right]. \end{aligned} \quad (\text{B.3})$$

Lemma B.2. If $X \sim \mathcal{N}(\mu, \sigma^2)$, then $\varphi_X^{\mathbb{P}}(a) = \exp(\mu a + \frac{1}{2} a^2 \sigma^2)$.

Lemma B.3. If X is exponentially distributed of expectation μ , then $\varphi_X^{\mathbb{P}}(a) = (1 - a\mu)^{-1}$ provided that $a < \frac{1}{\mu}$.

B.1.2 Exponential Martingales

Many proofs are based on properties of exponential martingales. The following three lemmas are the key ingredients of many proofs.

Lemma B.4. If $\mathbb{E}^{\mathbb{P}} \left[\int_0^t \Lambda_{V,u}^2 du \right] < \infty$ and $\mathbb{E}^{\mathbb{P}} \left[\int_0^t \Lambda_{\perp,u}^2 du \right] < \infty$, then

$$\left\{ \exp \left(- \int_0^t \Lambda_{V,u} dW_{V,u} - \int_0^t \Lambda_{\perp,u} dW_{\perp,u} - \frac{1}{2} \int_0^t (\Lambda_{V,u}^2 + \Lambda_{\perp,u}^2) du \right) \right\}_{t \geq 0}$$

is a \mathbb{P} -martingale of expectation 1.

Lemma B.5. For $Z \in \{Y, V\}$, $\left\{ J_{Z,t} - \mu_Z \int_0^t \lambda_{Z,s} ds \right\}_{t \geq 0}$ is a \mathbb{P} -martingale.

Lemma B.6. For $X \in \{Y, V\}$, $\left\{ \exp \left(\Gamma_X J_{X,t} - (\varphi_{Z_X}^{\mathbb{P}}(\Gamma_X) - 1) \int_0^t \lambda_{X,u} du \right) \right\}_{t \geq 0}$ is a \mathbb{P} -martingale where $\varphi_{Z_X}^{\mathbb{P}}$ is the moment generating function of the jump size Z_X .

B.1.2.1 Proofs

Proof of Lemma B.4. The continuous process $\{X_t\}_{t \geq 0}$ where

$$X_t = - \int_0^t \Lambda_{V,u} dW_{V,u} - \int_0^t \Lambda_{\perp,u} dW_{\perp,u}$$

is a \mathbb{P} -martingale. Its quadratic variation is $[X, X]_t = \int_0^t (\Lambda_{V,u}^2 + \Lambda_{\perp,u}^2) du$. Using Itô's lemma,

$$d \exp \left(X_t - \frac{1}{2} [X, X]_t \right)$$

$$\begin{aligned}
 &= \exp\left(X_t - \frac{1}{2}[X, X]_t\right) \left(dX_t - \frac{1}{2}d[X, X]_t\right) + \frac{1}{2} \exp\left(X_t - \frac{1}{2}[X, X]_t\right) d[X, X]_t \\
 &= \exp\left(X_t - \frac{1}{2}[X, X]_t\right) dX_t.
 \end{aligned}$$

Because $\exp\left(X_t - \frac{1}{2}[X, X]_t\right)$ is a stochastic integral with respect to a martingale,

$$\left\{ \exp\left(X_t - \frac{1}{2}[X, X]_t\right) \right\}_{t \geq 0}$$

is a \mathbb{P} -martingale. Since $X_0 = 0$, then

$$\mathbb{E}^{\mathbb{P}} \left[\exp\left(X_t - \frac{1}{2}[X, X]_t\right) \right] = \mathbb{E}^{\mathbb{P}} \left[\exp\left(X_0 - \frac{1}{2}[X, X]_0\right) \right] = 1.$$

□

Proof of Lemma B.5.

$$\begin{aligned}
 &\mathbb{E}_{\mathcal{F}_s}^{\mathbb{P}} \left[J_{Z,t} - J_{Z,s} - \mu_Z \int_0^t \lambda_{Z,u^-} du \right] \\
 &= \mathbb{E}_{\mathcal{F}_s}^{\mathbb{P}} \left[\mathbb{E}_{\mathcal{F}_s}^{\mathbb{P}} \left[J_{Z,t} - J_{Z,s} - \mu_Z \int_0^t \lambda_{Z,u^-} du \middle| \int_0^t \lambda_{Z,u^-} du \right] \right] \\
 &= \mathbb{E}_{\mathcal{F}_s}^{\mathbb{P}} \left[\mathbb{E}_{\mathcal{F}_s}^{\mathbb{P}} \left[\sum_{n=N_{Z,s}+1}^{N_{Z,t}} Z_{Y,n} - \mu_Z \int_0^t \lambda_{Z,u^-} du \middle| \int_0^t \lambda_{Z,u^-} du \right] \right] \\
 &= \mathbb{E}_{\mathcal{F}_s}^{\mathbb{P}} \left[\mathbb{E}_{\mathcal{F}_s}^{\mathbb{P}} \left[\mathbb{E}_{\mathcal{F}_s}^{\mathbb{P}} \left[\sum_{n=N_{Z,s}+1}^{N_{Z,t}} Z_{Y,n} \middle| \int_0^t \lambda_{Z,u^-} du, N_{Z,t} - N_{Z,s} \right] - \mu_Z \int_0^t \lambda_{Z,u^-} du \middle| \int_0^t \lambda_{Z,u^-} du \right] \right] \\
 &= \mathbb{E}_{\mathcal{F}_s}^{\mathbb{P}} \left[\mathbb{E}_{\mathcal{F}_s}^{\mathbb{P}} \left[\sum_{n=0}^{\infty} n \mu_Z \exp\left(-\int_0^t \lambda_{Z,u^-} du\right) \frac{\left(\int_0^t \lambda_{Z,u^-} du\right)^n}{n!} - \mu_Z \int_0^t \lambda_{Z,u^-} du \middle| \int_0^t \lambda_{Z,u^-} du \right] \right] \\
 &= \mathbb{E}_{\mathcal{F}_s}^{\mathbb{P}} \left[\mathbb{E}_{\mathcal{F}_s}^{\mathbb{P}} \left[\mu_Z \int_0^t \lambda_{Z,u^-} du - \mu_Z \int_0^t \lambda_{Z,u^-} du \middle| \int_0^t \lambda_{Z,u^-} du \right] \right] \\
 &= 0.
 \end{aligned}$$

□

Proof of Lemma B.6. From the moment generating function of a compound Poisson process with stochastic jump intensity of Equation (B.3),

$$\begin{aligned}
 &\mathbb{E}_{\mathcal{F}_s}^{\mathbb{P}} \left[\exp\left(\Gamma_X(J_{X,t} - J_{X,s}) - \left(\int_s^t \lambda_{X,u^-} du\right) (\varphi_{Z_X}^{\mathbb{P}}(\Gamma_X) - 1)\right) \right] \\
 &= \mathbb{E}_{\mathcal{F}_s}^{\mathbb{P}} \left[\exp\left(\Gamma_X\left(\sum_{n=N_{X,s}+1}^{N_{X,t}} Z_{X,n}\right) - \left(\int_s^t \lambda_{X,u^-} du\right) (\varphi_{Z_X}^{\mathbb{P}}(\Gamma_X) - 1)\right) \right]
 \end{aligned}$$

$$= \mathbb{E}_{\mathcal{F}_s}^{\mathbb{P}} \left[\exp \left(\left(\int_s^t \lambda_{X,u} du \right) (\varphi_{Z_X}^{\mathbb{P}}(\Gamma_X) - 1) - \left(\int_s^t \lambda_{X,u} du \right) (\varphi_{Z_X}^{\mathbb{P}}(\Gamma_X) - 1) \right) \right] = 0.$$

For the log-equity price jumps,

$$\varphi_{Z_Y}^{\mathbb{P}}(\Gamma_Y) - 1 = \exp \left(\mu_Y \Gamma_Y + \frac{1}{2} \sigma_Y^2 \Gamma_Y^2 \right) - 1,$$

while for the variance jumps,

$$\varphi_{Z_V}^{\mathbb{P}}(\Gamma_V) - 1 = (1 - \Gamma_V \mu_V)^{-1} - 1$$

provided that $\Gamma_V < \mu_V^{-1}$. □

B.1.3 Risk-Neutral Innovations

Lemma B.7 justifies Equation (5.23). Lemma B.8 determines the jump dynamics under \mathbb{Q} . These are components of the risk-neutral version of the model presented in Section 5.A.1 and 5.A.2.

Lemma B.7. *For $X \in \{Y, \perp\}$, the risk-neutral moment generating function of the Brownian increments $W_{X,t} - W_{X,s}$ satisfies*

$$\varphi_{W_{X,t} - W_{X,s}}^{\mathbb{Q}}(a) = \mathbb{E}_{\mathcal{F}_s}^{\mathbb{P}} \left[\exp \left(-a \int_s^t \Lambda_{X,u} du + \frac{a^2 (t-s)}{2} \right) \right]$$

which is the moment generating function of a Gaussian distribution of expectation

$$E_{\mathcal{F}_s}^{\mathbb{P}} \left[-a \int_s^t \Lambda_{X,u} du \right]$$

and variance $t - s$.

Lemma B.8. *For $X \in \{Y, V\}$, the risk-neutral moment generating function of the jump increments $J_{X,t} - J_{X,s}$ is given by*

$$\mathbb{E}_{\mathcal{F}_s}^{\mathbb{Q}} [\exp(a(J_{X,t} - J_{X,s}))] = \mathbb{E}_{\mathcal{F}_s}^{\mathbb{P}} \left[\exp \left(\left(\varphi_{Z_X}^{\mathbb{P}}(a + \Gamma_X) - \varphi_{Z_X}^{\mathbb{P}}(\Gamma_X) \right) \int_s^t \lambda_{X,u} du \right) \right].$$

In particular

$$\begin{aligned} & \varphi_{J_{Y,t} - J_{Y,s}}^{\mathbb{Q}}(a) \\ &= \mathbb{E}_{\mathcal{F}_s}^{\mathbb{Q}} [\exp(a(J_{Y,t} - J_{Y,s}))] \end{aligned}$$

$$= \mathbb{E}_{\mathcal{F}_s}^{\mathbb{P}} \left[\exp \left(\left(\exp \left(\mu_Y (a + \Gamma_Y) + \frac{1}{2} (a + \Gamma_Y)^2 \sigma_Y^2 \right) - \exp \left(\mu_Y \Gamma_Y + \frac{1}{2} \Gamma_Y^2 \sigma_Y^2 \right) \right) \int_s^t \lambda_{Y,u} du \right) \right]$$

and

$$\begin{aligned} \varphi_{J_{V,t}-J_{V,s}}^{\mathbb{Q}}(a) &= \mathbb{E}_{\mathcal{F}_s}^{\mathbb{Q}} [\exp(a(J_{V,t} - J_{V,s}))] \\ &= \mathbb{E}_{\mathcal{F}_s}^{\mathbb{P}} \left[\exp \left(\left((1 - (a + \Gamma_V) \mu_{\mu_V})^{-1} - (1 - \Gamma_V \mu_{\mu_V})^{-1} \right) \int_s^t \lambda_{Y,u} du \right) \right] \end{aligned}$$

provided that $\Gamma_V < \mu_{\mu_V}^{-1}$ and $a + \Gamma_V < \mu_{\mu_V}^{-1}$.

Here is a comparison of the \mathbb{P} - and the \mathbb{Q} -versions of the model generating functions of $J_{Y,t} - J_{Y,s}$ and $J_{V,t} - J_{V,s}$:

$$\begin{aligned} \varphi_{J_{Y,t}-J_{Y,s}}^{\mathbb{P}}(a) &= \mathbb{E}_{\mathcal{F}_s}^{\mathbb{P}} \left[\exp \left(\int_s^t \lambda_{Y,u} du \left(\exp \left(\mu_Y a + \frac{1}{2} \sigma_Y^2 a^2 \right) - 1 \right) \right) \right], \\ \varphi_{J_{Y,t}-J_{Y,s}}^{\mathbb{Q}}(a) &= \mathbb{E}_{\mathcal{F}_s}^{\mathbb{P}} \left[\exp \left(\exp \left(\mu_Y \Gamma_Y + \frac{1}{2} \Gamma_Y^2 \sigma_Y^2 \right) \int_s^t \lambda_{Y,u} du \left(\frac{\exp(\mu_Y(a+\Gamma_Y) + \frac{1}{2}(a+\Gamma_Y)^2 \sigma_Y^2)}{\exp(\mu_Y \Gamma_Y + \frac{1}{2} \Gamma_Y^2 \sigma_Y^2)} - 1 \right) \right) \right] \\ &= \mathbb{E}_{\mathcal{F}_s}^{\mathbb{P}} \left[\exp \left(\exp \left(\mu_Y \Gamma_Y + \frac{1}{2} \Gamma_Y^2 \sigma_Y^2 \right) \right. \right. \\ &\quad \left. \left. \times \int_s^t \lambda_{Y,u} du \left(\exp \left((\mu_Y + \Gamma_Y \sigma_Y^2) a + \frac{1}{2} a^2 \sigma_Y^2 \right) - 1 \right) \right) \right], \\ \varphi_{J_{V,t}-J_{V,s}}^{\mathbb{P}}(a) &= \mathbb{E}_{\mathcal{F}_s}^{\mathbb{P}} \left[\exp \left(\int_s^t \lambda_{V,u} du \left((1 - a \mu_V)^{-1} - 1 \right) \right) \right], \\ \varphi_{J_{V,t}-J_{V,s}}^{\mathbb{Q}}(a) &= \mathbb{E}_{\mathcal{F}_s}^{\mathbb{P}} \left[\exp \left((1 - \Gamma_V \mu_V)^{-1} \int_s^t \lambda_{Y,u} du \left(\frac{1 - \Gamma_V \mu_V}{1 - (a + \Gamma_V) \mu_V} - 1 \right) \right) \right] \\ &= \mathbb{E}_{\mathcal{F}_s}^{\mathbb{P}} \left[\exp \left((1 - \Gamma_V \mu_V)^{-1} \int_s^t \lambda_{Y,u} du \left(\left(1 - a \frac{\mu_V}{1 - \Gamma_V \mu_V} \right)^{-1} - 1 \right) \right) \right]. \end{aligned}$$

B.1.3.1 Proofs

Proof of Lemma B.7. The moment generating function of $W_{X,t} - W_{X,s}$ under \mathbb{Q} is

$$\begin{aligned} \varphi_{W_{X,t}-W_{X,s}}^{\mathbb{Q}}(a) \\ = \mathbb{E}_{\mathcal{F}_s}^{\mathbb{Q}} [\exp(a(W_{X,t} - W_{X,s}))] \end{aligned}$$

$$= \mathbb{E}_{\mathcal{F}_s}^{\mathbb{P}} \left[\exp(a(W_{X,t} - W_{X,s})) \exp \left(\begin{aligned} & - \int_s^t \Lambda_{V,u^-} dW_{V,u} - \int_s^t \Lambda_{\perp,u^-} dW_{\perp,u} \\ & - \frac{1}{2} \int_s^t (\Lambda_{V,u^-}^2 + \Lambda_{\perp,u^-}^2) du \\ & + \Gamma_Y (J_{Y,t} - J_{Y,s}) \\ & - \left(\exp(\mu_Y \Gamma_Y + \frac{1}{2} \sigma_Y^2 \Gamma_Y^2) - 1 \right) \int_s^t \lambda_{Y,u^-} du \\ & + \Gamma_V (J_{V,t} - J_{V,s}) - \left((1 - \mu_V \Gamma_V)^{-1} - 1 \right) \int_s^t \lambda_{V,u^-} du \end{aligned} \right) \right].$$

Because the variance jump intensity is independent of the other components inside the exponential function, the last expression is equal to

$$\begin{aligned} & \mathbb{E}_{\mathcal{F}_s}^{\mathbb{P}} \left[\exp(a(W_{X,t} - W_{X,s})) \exp \left(\begin{aligned} & - \int_s^t \Lambda_{V,u^-} dW_{V,u} - \int_s^t \Lambda_{\perp,u^-} dW_{\perp,u} \\ & - \frac{1}{2} \int_s^t (\Lambda_{V,u^-}^2 + \Lambda_{\perp,u^-}^2) du \\ & + \Gamma_Y (J_{Y,t} - J_{Y,s}) \\ & - \left(\exp(\mu_Y \Gamma_Y + \frac{1}{2} \sigma_Y^2 \Gamma_Y^2) - 1 \right) \int_s^t \lambda_{Y,u^-} du \end{aligned} \right) \right] \\ & \times \mathbb{E}_{\mathcal{F}_s}^{\mathbb{P}} \left[\exp \left(\Gamma_V (J_{V,t} - J_{V,s}) - \left((1 - \mu_V \Gamma_V)^{-1} - 1 \right) \int_s^t \lambda_{V,u^-} du \right) \right]. \end{aligned}$$

Lemma B.6 implies that the last expectation is 1. Using the tower property of conditional expectations, conditioning on the log-equity price intensity, the previous expression becomes

$$\begin{aligned} & \mathbb{E}_{\mathcal{F}_s}^{\mathbb{P}} \left[\begin{aligned} & \mathbb{E}_{\mathcal{F}_s}^{\mathbb{P}} \left[\begin{aligned} & \exp(a(W_{X,t} - W_{X,s})) \\ & \times \exp \left(- \int_s^t \Lambda_{V,u^-} dW_{V,u} - \int_s^t \Lambda_{\perp,u^-} dW_{\perp,u} \right) \left| \int_s^t \lambda_{Y,u^-} du \right| \times \\ & \times \exp \left(- \frac{1}{2} \int_s^t (\Lambda_{V,u^-}^2 + \Lambda_{\perp,u^-}^2) du \right) \end{aligned} \right] \\ & \mathbb{E}_{\mathcal{F}_s}^{\mathbb{P}} \left[\exp \left(\Gamma_Y (J_{Y,t} - J_{Y,s}) - \left(\exp(\mu_Y \Gamma_Y + \frac{1}{2} \sigma_Y^2 \Gamma_Y^2) - 1 \right) \int_s^t \lambda_{Y,u^-} du \right) \left| \int_s^t \lambda_{Y,u^-} du \right| \right] \end{aligned} \right] \\ & = \mathbb{E}_{\mathcal{F}_s}^{\mathbb{P}} \left[\mathbb{E}_{\mathcal{F}_s}^{\mathbb{P}} \left[\begin{aligned} & \exp(a(W_{X,t} - W_{X,s})) \\ & \exp \left(- \int_s^t \Lambda_{V,u^-} dW_{V,u} - \int_s^t \Lambda_{\perp,u^-} dW_{\perp,u} \right) \left| \int_s^t \lambda_{Y,u^-} du \right| \\ & \exp \left(- \frac{1}{2} \int_s^t (\Lambda_{V,u^-}^2 + \Lambda_{\perp,u^-}^2) du \right) \end{aligned} \right] \right]. \end{aligned}$$

Finally, from Lemma B.4,

$$\begin{aligned} & \varphi_{W_{X,t} - W_{X,s}}^{\mathbb{Q}}(a) \\ & = \mathbb{E}_{\mathcal{F}_s}^{\mathbb{P}} \left[\exp(a(W_{X,t} - W_{X,s})) \exp \left(- \int_s^t \Lambda_{X,u^-} dW_{X,u} - \frac{1}{2} \int_s^t \Lambda_{X,u^-}^2 du \right) \right] \\ & = \mathbb{E}_{\mathcal{F}_s}^{\mathbb{P}} \left[\exp \left(\int_s^t (a - \Lambda_{X,u^-}) dW_{V,u} - \frac{1}{2} \int_s^t \Lambda_{X,u^-}^2 du \right) \right] \\ & = \mathbb{E}_{\mathcal{F}_s}^{\mathbb{P}} \left[\exp \left(\int_s^t (a - \Lambda_{X,u^-}) dW_{V,u} - \frac{1}{2} \int_s^t (a - \Lambda_{X,u^-})^2 du \right. \right. \\ & \quad \left. \left. - \frac{1}{2} \int_s^t \Lambda_{X,u^-}^2 du + \frac{1}{2} \int_s^t (a - \Lambda_{X,u^-})^2 du \right) \right] \end{aligned}$$

$$= \mathbb{E}_{\mathcal{F}_s}^{\mathbb{P}} \left[\exp \left(-a \int_s^t \Lambda_{X,u} du + \frac{a^2 (t-s)}{2} \right) \right].$$

□

Proof of Lemma B.8.

$$\begin{aligned} & \mathbb{E}_{\mathcal{F}_s}^{\mathbb{Q}} [\exp(a(J_{X,t} - J_{X,s}))] \\ &= \mathbb{E}_{\mathcal{F}_s}^{\mathbb{P}} \left[\exp \left(\begin{aligned} & a(J_{X,t} - J_{X,s}) \\ & - \int_s^t \Lambda_{V,u} dW_{V,u} - \int_s^t \Lambda_{\perp,u} dW_{\perp,u} - \frac{1}{2} \int_s^t (\Lambda_{V,u}^2 + \Lambda_{\perp,u}^2) du \\ & + \Gamma_X (J_{Y,t} - J_{Y,s}) - (\varphi_{Z_Y}^{\mathbb{P}}(\Gamma_Y) - 1) \int_s^t \lambda_{Y,u} du \\ & + \Gamma_V (J_{V,t} - J_{V,t}) - (\varphi_{Z_Y}^{\mathbb{P}}(\Gamma_V) - 1) \int_s^t \lambda_{V,u} du \end{aligned} \right) \right]. \end{aligned}$$

Using the tower property of conditional expectations, the last expression becomes

$$\begin{aligned} & \mathbb{E}_{\mathcal{F}_s}^{\mathbb{P}} \left[\mathbb{E}_{\mathcal{F}_s}^{\mathbb{P}} \left[\exp \left(- \int_s^t \Lambda_{V,u} dW_{V,u} - \int_s^t \Lambda_{\perp,u} dW_{\perp,u} \right. \right. \right. \\ & \quad \left. \left. - \frac{1}{2} \int_s^t (\Lambda_{V,u}^2 + \Lambda_{\perp,u}^2) du \right) \middle| \int_s^t \lambda_{Y,u} du \right] \right] \\ & \times \mathbb{E}_{\mathcal{F}_s}^{\mathbb{P}} \left[\mathbb{E}_{\mathcal{F}_s}^{\mathbb{P}} \left[\exp \left(\begin{aligned} & \exp(a(J_{X,t} - J_{X,s})) \\ & + \Gamma_X (J_{Y,t} - J_{Y,s}) - (\varphi_{Z_Y}^{\mathbb{P}}(\Gamma_Y) - 1) \int_s^t \lambda_{Y,u} du \\ & + \Gamma_V (J_{V,t} - J_{V,t}) - (\varphi_{Z_Y}^{\mathbb{P}}(\Gamma_V) - 1) \int_s^t \lambda_{V,u} du \end{aligned} \right) \middle| \int_s^t \lambda_{Y,u} du \right] \right]. \end{aligned}$$

Lemma B.4 implies that the first term vanishes. Therefore,

$$\begin{aligned} & \mathbb{E}_{\mathcal{F}_s}^{\mathbb{Q}} [\exp(a(J_{X,t} - J_{X,s}))] \\ &= \mathbb{E}_{\mathcal{F}_s}^{\mathbb{P}} \left[\exp \left((a + \Gamma_X)(J_{X,t} - J_{X,s}) - (\varphi_{Z_{X,1}}^{\mathbb{P}}(\Gamma_X) - 1) \int_s^t \lambda_{X,u} du \right) \right] \\ &= \mathbb{E}_{\mathcal{F}_s}^{\mathbb{P}} \left[\exp \left(\begin{aligned} & (a + \Gamma_X)(J_{X,t} - J_{X,s}) - (\varphi_{Z_X}^{\mathbb{P}}(a + \Gamma_X) - 1) \int_s^t \lambda_{X,u} du \\ & + (\varphi_{Z_X}^{\mathbb{P}}(a + \Gamma_X) - 1) \int_s^t \lambda_{X,u} du - (\varphi_{Z_X}^{\mathbb{P}}(\Gamma_X) - 1) \int_s^t \lambda_{X,u} du \end{aligned} \right) \right] \\ &= \mathbb{E}_{\mathcal{F}_s}^{\mathbb{P}} \left[\mathbb{E}_{\mathcal{F}_s}^{\mathbb{P}} \left[\exp \left(\begin{aligned} & (a + \Gamma_X)(J_{X,t} - J_{X,s}) \\ & - (\varphi_{Z_X}^{\mathbb{P}}(a + \Gamma_X) - 1) \int_s^t \lambda_{X,u} du \\ & + (\varphi_{Z_X}^{\mathbb{P}}(a + \Gamma_X) - 1) \int_s^t \lambda_{X,u} du \\ & - (\varphi_{Z_X}^{\mathbb{P}}(\Gamma_X) - 1) \int_s^t \lambda_{X,u} du \end{aligned} \right) \middle| \int_s^t \lambda_{X,u} du \right] \right] \\ &= \mathbb{E}_{\mathcal{F}_s}^{\mathbb{P}} \left[\exp \left((\varphi_{Z_X}^{\mathbb{P}}(a + \Gamma_X) - \varphi_{Z_X}^{\mathbb{P}}(\Gamma_X)) \int_s^t \lambda_{X,u} du \right) \right]. \end{aligned}$$

□

B.1.4 Log-Equity Price Drift

This section provides a detailed proof of the log-equity price drift term described in Section 5.B.1.

Lemma B.9. *Under the risk-neutral measure, the log-equity price is characterized by*

$$\begin{aligned} dY_t &= \alpha_{t^-}^{\mathbb{Q}} dt + \sqrt{V_{t^-}} dW_{Y,t}^{\mathbb{Q}} + dJ_{Y,t}^{\mathbb{Q}}, \\ \alpha_{t^-}^{\mathbb{Q}} &= r - q - \frac{1}{2} V_{t^-} - \xi_Y^{\mathbb{Q}} \lambda_{Y,t^-}^{\mathbb{Q}}, \\ \xi_Y^{\mathbb{Q}} &= \varphi_{Z_Y}^{\mathbb{Q}}(1) - 1 = \exp\left(\mu_Y^{\mathbb{Q}} + \frac{1}{2} \sigma_Y^2\right) - 1, \\ \lambda_{Y,t^-}^{\mathbb{Q}} &= \varphi_{Z_Y}^{\mathbb{Q}}(\Gamma_Y) \lambda_{Y,t^-}^{\mathbb{P}} = \exp\left(\mu_Y \Gamma_Y + \frac{1}{2} \Gamma_Y^2 \sigma_Y^2\right) \lambda_{Y,t^-}^{\mathbb{P}}. \end{aligned}$$

Proof of Lemma B.9. Using Itô's lemma,

$$\begin{aligned} \exp(Y_t) - \exp(Y_0) &= \int_0^t \exp(Y_{u^-}) dY_u + \frac{1}{2} \int_0^t \exp(Y_{u^-}) d[Y, Y]_u^c \\ &\quad + \sum_{0 < u \leq t} \{\exp(Y_u) - \exp(Y_{u^-}) - \exp(Y_{u^-}) \Delta Y_u\}. \end{aligned}$$

The last term is equal to

$$\begin{aligned} &\sum_{0 < u \leq t} \exp(Y_{u^-}) \{\exp(Y_u - Y_{u^-}) - 1\} - \sum_{0 < u \leq t} \{\exp(Y_{u^-}) (J_{Y,u}^{\mathbb{Q}} - J_{Y,u^-}^{\mathbb{Q}})\} \\ &= \sum_{0 < u \leq t} \exp(Y_{u^-}) \{\exp(J_{Y,u}^{\mathbb{Q}} - J_{Y,u^-}^{\mathbb{Q}}) - 1\} - \sum_{0 < u \leq t} \{\exp(Y_{u^-}) (J_{Y,u}^{\mathbb{Q}} - J_{Y,u^-}^{\mathbb{Q}})\} \\ &= \sum_{0 < u \leq t} \exp(Y_{u^-} - J_{Y,u^-}^{\mathbb{Q}}) \{\exp(J_{Y,u}^{\mathbb{Q}}) - \exp(J_{Y,u^-}^{\mathbb{Q}})\} \\ &\quad - \sum_{0 < u \leq t} \{\exp(Y_{u^-}) (J_{Y,u}^{\mathbb{Q}} - J_{Y,u^-}^{\mathbb{Q}})\}. \end{aligned}$$

Therefore, by substituting Equation (5.25) in the first term of $\exp(Y_t) - \exp(Y_0)$, we obtain

$$\begin{aligned} &\exp(Y_t) - \exp(Y_0) \\ &= \int_0^t \exp(Y_{u^-}) \alpha_{u^-}^{\mathbb{Q}} du + \int_0^t \exp(Y_{u^-}) \sqrt{V_{u^-}} dW_{Y,u}^{\mathbb{Q}} + \int_0^t \exp(Y_{u^-}) dJ_{Y,u}^{\mathbb{Q}} \\ &\quad + \frac{1}{2} \int_0^t \exp(Y_{u^-}) V_{u^-} du + \int_0^t \exp(Y_{u^-} - J_{Y,u^-}^{\mathbb{Q}}) d\exp(J_{Y,u}^{\mathbb{Q}}) - \int_0^t \exp(Y_{u^-}) dJ_{Y,u}^{\mathbb{Q}} \\ &= \int_0^t \exp(Y_{u^-}) \left(\alpha_{u^-}^{\mathbb{Q}} + \frac{1}{2} V_{u^-}\right) du + \int_0^t \exp(Y_{u^-}) \sqrt{V_{u^-}} dW_{Y,u}^{\mathbb{Q}} \end{aligned}$$

$$+ \int_0^t \exp(Y_{u^-} - J_{Y,u}^Q) d \exp(J_{Y,u}^Q).$$

Lemma B.6 states that $\left\{ \exp(J_{Y,t}^Q - (\varphi_{Z_Y}^Q(1) - 1) \int_0^t \lambda_{Y,u}^Q du) \right\}_{t \geq 0}$ is a \mathbb{Q} -martingale, where $\varphi_{Z_Y}^Q$ is the moment generating function of the log-equity price jump size. Therefore, if

$$\xi_Y^Q = \varphi_{Z_Y}^Q(1) - 1 = \exp\left(\mu_Y^Q + \frac{1}{2}\sigma_Y^2\right) - 1,$$

then

$$\begin{aligned} & \exp(Y_t) - \exp(Y_0) \\ &= \int_0^t \exp(Y_{u^-}) \left(\alpha_{u^-}^Q + \frac{1}{2} V_{u^-} \right) du + \int_0^t \exp(Y_{u^-}) \sqrt{V_{u^-}} dW_{Y,u}^Q \\ & \quad + \int_0^t \exp(Y_{u^-} - J_{Y,u}^Q) d \left(\exp(J_{Y,u}^Q - \xi_Y^Q \int_0^u \lambda_{Y,s}^Q ds) \exp\left(\xi_Y^Q \int_0^u \lambda_{Y,s}^Q ds\right) \right) \\ &= \int_0^t \exp(Y_{u^-}) \left(\alpha_{u^-}^Q + \frac{1}{2} V_{u^-} \right) du + \int_0^t \exp(Y_{u^-}) \sqrt{V_{u^-}} dW_{Y,u}^Q \\ & \quad + \int_0^t \exp(Y_{u^-} - J_{Y,u}^Q) \exp\left(\xi_Y^Q \int_0^u \lambda_{Y,s}^Q ds\right) d \exp\left(J_{Y,u}^Q - \xi_Y^Q \int_0^u \lambda_{Y,s}^Q ds\right) \\ & \quad + \int_0^t \exp(Y_{u^-} - J_{Y,u}^Q) \exp\left(J_{Y,u}^Q - \xi_Y^Q \int_0^u \lambda_{Y,s}^Q ds\right) \xi_Y^Q \lambda_{Y,u}^Q \exp\left(\xi_Y^Q \int_0^u \lambda_{Y,s}^Q ds\right) du \\ &= \int_0^t \exp(Y_{u^-}) \left(\alpha_{u^-}^Q + \frac{1}{2} V_{u^-} + \xi_Y^Q \lambda_{Y,u}^Q \right) du + \int_0^t \exp(Y_{u^-}) \sqrt{V_{u^-}} dW_{Y,u}^Q \\ & \quad + \int_0^t \exp(Y_{u^-} - J_{Y,u}^Q) \exp\left(\xi_Y^Q \int_0^u \lambda_{Y,s}^Q ds\right) d \exp\left(J_{Y,u}^Q - \xi_Y^Q \int_0^u \lambda_{Y,s}^Q ds\right). \end{aligned}$$

Note that the last two terms are martingales. Because the discount factor and the dividend correction are both continuous, i.e.

$$D_t = \exp(-(r - q)t), \quad (\text{B.4})$$

integration by part implies that

$$\begin{aligned} & D_t \exp(Y_t) - D_0 \exp(Y_0) \\ &= \int_0^t D_u d \exp(Y_u) - \int_0^t (r - q) D_u \exp(Y_{u^-}) du \\ &= \int_0^t D_u \exp(Y_{u^-}) \left(\alpha_{u^-}^Q + \frac{1}{2} V_{u^-} + \xi_Y^Q \lambda_{Y,u}^Q - r + q \right) du \\ & \quad + \int_0^t D_u \exp(Y_{u^-}) \sqrt{V_{u^-}} dW_{Y,u}^Q \end{aligned}$$

$$+ \int_0^t D_u \exp(Y_{u^-} - J_{Y,u^-}^{\mathbb{Q}}) \exp\left(\xi_Y^{\mathbb{Q}} \int_0^u \lambda_{Y,s^-}^{\mathbb{Q}} du\right) d \exp\left(J_{Y,u}^{\mathbb{Q}} - \xi_Y^{\mathbb{Q}} \int_0^t \lambda_{Y,u^-}^{\mathbb{Q}} du\right).$$

The pricing theory implied that $\{D_t \exp(Y_t)\}_{t \geq 0}$ must be a \mathbb{Q} -martingale. Therefore, the drift component must be nil, that is

$$\alpha_{u^-}^{\mathbb{Q}} = r - q - \frac{1}{2} V_{u^-} - \xi_Y^{\mathbb{Q}} \lambda_{Y,u^-}^{\mathbb{Q}}.$$

□

Lemma B.10. *Under the \mathbb{P} -measure, the returns' are characterized by*

$$\begin{aligned} dY_t &= \alpha_t^{\mathbb{P}} dt + \sqrt{V_t} dW_{Y,t} + dJ_{Y,t}, \\ \alpha_{u^-}^{\mathbb{P}} &= r - q - \frac{1}{2} V_{u^-} + \sqrt{V_{u^-}} \Lambda_{Y,u^-} + (\xi_Y^{\mathbb{P}} - \zeta_Y^{\mathbb{P}}) \lambda_{Y,u^-} + (\xi_V^{\mathbb{P}} - \zeta_V^{\mathbb{P}}) \lambda_{V,u^-} \\ &= r - q - \frac{1}{2} V_{u^-} + \sqrt{V_{u^-}} \Lambda_{Y,u^-} + (\varphi_{Z_Y}^{\mathbb{P}}(\Gamma_Y) - \varphi_{Z_Y}^{\mathbb{P}}(1 + \Gamma_Y)) \lambda_{Y,u^-}, \\ \xi_Y^{\mathbb{P}} &= \varphi_{Z_Y}^{\mathbb{P}}(\Gamma_Y) - 1, \\ \xi_V^{\mathbb{P}} &= \varphi_{Z_V}^{\mathbb{P}}(\Gamma_V) - 1, \\ \zeta_Y^{\mathbb{P}} &= \varphi_{Z_Y}^{\mathbb{P}}(1 + \Gamma_Y) - 1, \\ \zeta_V^{\mathbb{P}} &= \varphi_{Z_V}^{\mathbb{P}}(\Gamma_V) - 1, \end{aligned}$$

Proof of Lemma B.10. Let

$$Z_t = \left(\begin{aligned} & - \int_0^t \Lambda_{V,u^-} dW_{V,u} - \int_0^t \Lambda_{\perp,u^-} dW_{\perp,u} - \frac{1}{2} \int_0^t (\Lambda_{V,u^-}^2 + \Lambda_{\perp,u^-}^2) du \\ & + \Gamma_Y J_{Y,t} - \xi_Y^{\mathbb{P}} \int_0^t \lambda_{Y,u^-} du + \Gamma_V J_{V,t} - \xi_V^{\mathbb{P}} \int_0^t \lambda_{V,u^-} du \end{aligned} \right)$$

be associated to the Radon-Nikodym derivative (5.21) and $D_t \exp(-(r - q)t)$ be the combination of the discount factor and the dividend yield. Since the discounted price must be a \mathbb{Q} -martingale,

$$1 = \mathbb{E}_s^{\mathbb{Q}} \left[\frac{D_t \exp(Y_t)}{D_s \exp(Y_s)} \right] = \mathbb{E}_s^{\mathbb{P}} \left[\frac{D_t \exp(Y_t + Z_t)}{D_s \exp(Y_s + Z_s)} \right]$$

which means that $\{D_t \exp(Y_t + Z_t)\}_{t \geq 0}$ is a \mathbb{P} -martingale. Note that

$$\begin{aligned} Y_t + Z_t &= Y_0 + Z_0 + \int_0^t \left(\alpha_{u^-}^{\mathbb{P}} - \frac{1}{2} (\Lambda_{V,u^-}^2 + \Lambda_{\perp,u^-}^2) - \xi_Y^{\mathbb{P}} \lambda_{Y,u^-} - \xi_V^{\mathbb{P}} \lambda_{V,u^-} \right) du \\ &\quad + \int_0^t (\rho \sqrt{V_{u^-}} - \Lambda_{V,u^-}) dW_{V,u} + \int_0^t (\sqrt{1 - \rho^2} \sqrt{V_{u^-}} - \Lambda_{\perp,u^-}) dW_{\perp,u} \\ &\quad + (1 + \Gamma_Y) \int_0^t dJ_{Y,u} + \Gamma_V \int_0^t dJ_{V,u}. \end{aligned}$$

Itô's lemma implies that

$$\begin{aligned} \exp(Y_t + Z_t) - \exp(Y_0 + Z_0) &= \int_0^t \exp(Y_{u^-} + Z_{u^-}) d(Y_u + Z_u) \\ &+ \frac{1}{2} \int_0^t \exp(Y_{u^-} + Z_{u^-}) \left(V_{u^-} - 2\sqrt{V_{u^-}} \Lambda_{Y,u^-} + \Lambda_{V,u^-}^2 + \Lambda_{\perp,u^-}^2 \right) du \\ &+ \sum_{0 < u \leq t} \{ \exp(Y_u + Z_u) - \exp(Y_{u^-} + Z_{u^-}) - \exp(Y_{u^-} + Z_{u^-}) \Delta(Y_u + Z_u) \}. \end{aligned}$$

Replacing the first expression in the second one leads to

$$\begin{aligned} &\exp(Y_t + Z_t) - \exp(Y_0 + Z_0) \\ &= \int_0^t \exp(Y_{u^-} + Z_{u^-}) \left(\alpha_{u^-}^{\mathbb{P}} - \frac{1}{2} (\Lambda_{V,u^-}^2 + \Lambda_{\perp,u^-}^2) - \xi_Y^{\mathbb{P}} \lambda_{Y,u^-} - \xi_V^{\mathbb{P}} \lambda_{V,u^-} \right) du \\ &+ \int_0^t \exp(Y_{u^-} + Z_{u^-}) (\rho \sqrt{V_{u^-}} - \Lambda_{V,u^-}) dW_{V,u} \\ &+ \int_0^t \exp(Y_{u^-} + Z_{u^-}) (\sqrt{1 - \rho^2} \sqrt{V_{u^-}} - \Lambda_{\perp,u^-}) dW_{\perp,u} \\ &+ \int_0^t \exp(Y_{u^-} + Z_{u^-}) (1 + \Gamma_Y) dJ_{Y,u} + \int_0^t \exp(Y_{u^-} + Z_{u^-}) \Gamma_V dJ_{V,u} \\ &+ \frac{1}{2} \int_0^t \exp(Y_{u^-} + Z_{u^-}) \left(V_{u^-} - 2\sqrt{V_{u^-}} \Lambda_{Y,u^-} + \Lambda_{V,u^-}^2 + \Lambda_{\perp,u^-}^2 \right) du \\ &+ \sum_{0 < u \leq t} \{ \exp(Y_u + Z_u) - \exp(Y_{u^-} + Z_{u^-}) - \exp(Y_{u^-} + Z_{u^-}) \Delta(Y_u + Z_u) \}. \end{aligned}$$

The last term can be rewritten as

$$\begin{aligned} &\sum_{0 < u \leq t} \exp(Y_{u^-} + Z_{u^-}) \{ \exp(\Delta Y_u + \Delta Z_u) - 1 - \Delta(Y_u + Z_u) \} \\ &= \sum_{0 < u \leq t} \exp(Y_{u^-} + Z_{u^-}) \{ \exp((1 + \Gamma_Y) \Delta J_{Y,u} + \Gamma_V \Delta J_{V,u}) - 1 - (1 + \Gamma_Y) \Delta J_{Y,u} - \Gamma_V \Delta J_{V,u} \} \\ &= \sum_{0 < u \leq t} \left[\exp(Y_{u^-} + Z_{u^-} - (1 + \Gamma_Y) J_{Y,u^-} - \Gamma_V J_{V,u^-}) \right. \\ &\quad \times \{ \exp((1 + \Gamma_Y) J_{Y,u} + \Gamma_V J_{V,u}) - \exp((1 + \Gamma_Y) J_{Y,u^-} + \Gamma_V J_{V,u^-}) \} \left. \right] \\ &\quad - \sum_{0 < u \leq t} \exp(Y_{u^-} + Z_{u^-}) \{ (1 + \Gamma_Y) \Delta J_{Y,u} + \Gamma_V \Delta J_{V,u} \} \\ &= \int_0^t \exp(Y_{u^-} + Z_{u^-} - (1 + \Gamma_Y) J_{Y,u^-} - \Gamma_V J_{V,u^-}) d \exp((1 + \Gamma_Y) J_{Y,u} + \Gamma_V J_{V,u}) \\ &\quad - \int_0^t \exp(Y_{u^-} + Z_{u^-}) (1 + \Gamma_Y) dJ_{Y,u} - \int_0^t \exp(Y_{u^-} + Z_{u^-}) \Gamma_V dJ_{V,u}. \end{aligned}$$

Therefore,

$$\begin{aligned}
 & \exp(Y_t + Z_t) - \exp(Y_0 + Z_0) \\
 = & \int_0^t \exp(Y_{u^-} + Z_{u^-}) \left(\alpha_{u^-}^{\mathbb{P}} - \xi_Y^{\mathbb{P}} \lambda_{Y,u^-} - \xi_V^{\mathbb{P}} \lambda_{V,u^-} + \frac{1}{2} V_{u^-} - \sqrt{V_{u^-}} \Lambda_{Y,u^-} \right) du \\
 & + \int_0^t \exp(Y_{u^-} + Z_{u^-}) \left(\rho \sqrt{V_{u^-}} - \Lambda_{V,u^-} \right) dW_{V,u} \\
 & + \int_0^t \exp(Y_{u^-} + Z_{u^-}) \left(\sqrt{1 - \rho^2} \sqrt{V_{u^-}} - \Lambda_{\perp,u^-} \right) dW_{\perp,u} \\
 & + \int_0^t \exp(Y_{u^-} + Z_{u^-} - (1 + \Gamma_Y) J_{Y,u^-} - \Gamma_V J_{V,u^-}) d \exp((1 + \Gamma_Y) J_{Y,u} + \Gamma_V J_{V,u}).
 \end{aligned}$$

We need to construct a martingale out of the last term. From Lemma B.6, $\{M_t\}_{t \geq 0}$ is a \mathbb{P} -martingale where

$$M_t = \exp \left((1 + \Gamma_Y) J_{Y,t} + \Gamma_V J_{V,t} - \zeta_Y^{\mathbb{P}} \int_0^t \lambda_{Y,u^-} du - \zeta_V^{\mathbb{P}} \int_0^t \lambda_{V,u^-} du \right).$$

Indeed,

$$\begin{aligned}
 & \mathbb{E}_{\mathcal{F}_s}^{\mathbb{P}} \left[\exp \left((1 + \Gamma_Y) (J_{Y,t} - J_{Y,s}) + \Gamma_V (J_{V,t} - J_{V,s}) - \zeta_Y^{\mathbb{P}} \left(\int_s^t \lambda_{Y,u^-} du \right) - \zeta_V^{\mathbb{P}} \left(\int_s^t \lambda_{V,u^-} du \right) \right) \right] \\
 = & \mathbb{E}_{\mathcal{F}_s}^{\mathbb{P}} \left[\mathbb{E}_{\mathcal{F}_s}^{\mathbb{P}} \left[\exp \left((1 + \Gamma_Y) (J_{Y,t} - J_{Y,s}) + \Gamma_V (J_{V,t} - J_{V,s}) - \zeta_Y^{\mathbb{P}} \left(\int_s^t \lambda_{Y,u^-} du \right) - \zeta_V^{\mathbb{P}} \left(\int_s^t \lambda_{V,u^-} du \right) \right) \right] \right] \\
 = & 1.
 \end{aligned}$$

Moreover,

$$\begin{aligned}
 & d \exp((1 + \Gamma_Y) J_{Y,t} + \Gamma_V J_{V,t}) \\
 = & dM_t \exp \left(\zeta_Y^{\mathbb{P}} \int_0^t \lambda_{Y,u^-} du + \zeta_V^{\mathbb{P}} \int_0^t \lambda_{V,u^-} du \right) \\
 = & M_{t^-} d \exp \left(\zeta_Y^{\mathbb{P}} \int_0^t \lambda_{Y,u^-} du + \zeta_V^{\mathbb{P}} \int_0^t \lambda_{V,u^-} du \right) + \exp \left(\zeta_Y^{\mathbb{P}} \int_0^t \lambda_{Y,u^-} du + \zeta_V^{\mathbb{P}} \int_0^t \lambda_{V,u^-} du \right) dM_t \\
 = & (\zeta_Y^{\mathbb{P}} \lambda_{Y,t^-} + \zeta_V^{\mathbb{P}} \lambda_{V,t^-}) M_{t^-} \exp \left(\zeta_Y^{\mathbb{P}} \int_0^t \lambda_{Y,u^-} du + \zeta_V^{\mathbb{P}} \int_0^t \lambda_{V,u^-} du \right) dt \\
 & + \exp \left(\zeta_Y^{\mathbb{P}} \int_0^t \lambda_{Y,u^-} du + \zeta_V^{\mathbb{P}} \int_0^t \lambda_{V,u^-} du \right) dM_t \\
 = & (\zeta_Y^{\mathbb{P}} \lambda_{Y,t^-} + \zeta_V^{\mathbb{P}} \lambda_{V,t^-}) \exp((1 + \Gamma_Y) J_{Y,t^-} + \Gamma_V J_{V,t^-}) dt \\
 & + \exp \left(\zeta_Y^{\mathbb{P}} \int_0^t \lambda_{Y,u^-} du + \zeta_V^{\mathbb{P}} \int_0^t \lambda_{V,u^-} du \right) dM_t.
 \end{aligned}$$

Therefore,

$$\begin{aligned}
 & \exp(Y_t + Z_t) - \exp(Y_0 + Z_0) \\
 &= \int_0^t \exp(Y_{u^-} + Z_{u^-}) \left(\alpha_{u^-}^{\mathbb{P}} - \xi_Y^{\mathbb{P}} \lambda_{Y,u^-} - \xi_V^{\mathbb{P}} \lambda_{V,u^-} + \frac{1}{2} V_{u^-} - \sqrt{V_{u^-}} \Lambda_{Y,u^-} \right) du \\
 & \quad + \int_0^t \exp(Y_{u^-} + Z_{u^-}) \left(\rho \sqrt{V_{u^-}} - \Lambda_{V,u^-} \right) dW_{V,u} \\
 & \quad + \int_0^t \exp(Y_{u^-} + Z_{u^-}) \left(\sqrt{1 - \rho^2} \sqrt{V_{u^-}} - \Lambda_{\perp,u^-} \right) dW_{\perp,u} \\
 & \quad + \int_0^t \left(\xi_Y^{\mathbb{P}} \lambda_{Y,u^-} + \xi_V^{\mathbb{P}} \lambda_{V,u^-} \right) \exp(Y_{u^-} + Z_{u^-}) du \\
 & \quad + \int_0^t \exp \left(Y_{u^-} + Z_{u^-} - (1 + \Gamma_Y) J_{Y,u^-} - \Gamma_V J_{V,u^-} + \zeta_Y^{\mathbb{P}} \int_0^u \lambda_{Y,s^-} ds + \zeta_V^{\mathbb{P}} \int_0^u \lambda_{V,s^-} ds \right) dM_u \\
 &= \int_0^t \exp(Y_{u^-} + Z_{u^-}) \left(\alpha_{u^-}^{\mathbb{P}} - \xi_Y^{\mathbb{P}} \lambda_{Y,u^-} - \xi_V^{\mathbb{P}} \lambda_{V,u^-} + \frac{1}{2} V_{u^-} - \sqrt{V_{u^-}} \Lambda_{Y,u^-} + \zeta_Y^{\mathbb{P}} \lambda_{Y,u^-} + \zeta_V^{\mathbb{P}} \lambda_{V,u^-} \right) du \\
 & \quad + \int_0^t \exp(Y_{u^-} + Z_{u^-}) \left(\rho \sqrt{V_{u^-}} - \Lambda_{V,u^-} \right) dW_{V,u} \\
 & \quad + \int_0^t \exp(Y_{u^-} + Z_{u^-}) \left(\sqrt{1 - \rho^2} \sqrt{V_{u^-}} - \Lambda_{\perp,u^-} \right) dW_{\perp,u} \\
 & \quad + \int_0^t \exp \left(Y_{u^-} + Z_{u^-} - (1 + \Gamma_Y) J_{Y,u^-} - \Gamma_V J_{V,u^-} + \zeta_Y^{\mathbb{P}} \int_0^u \lambda_{Y,s^-} ds + \zeta_V^{\mathbb{P}} \int_0^u \lambda_{V,s^-} ds \right) dM_u.
 \end{aligned}$$

Since

$$\begin{aligned}
 & D_t \exp(Y_t + Z_t) - D_0 \exp(Y_0 + Z_0) \\
 &= \int_0^t D_u d \exp(Y_u + Z_u) + \int_0^t (q - r) D_u \exp(Y_{u^-} + Z_{u^-}) du,
 \end{aligned}$$

then

$$\begin{aligned}
 & D_t \exp(Y_t + Z_t) - D_0 \exp(Y_0 + Z_0) \\
 &= \int_0^t D_u \exp(Y_{u^-} + Z_{u^-}) \left(\alpha_{u^-}^{\mathbb{P}} - \xi_Y^{\mathbb{P}} \lambda_{Y,u^-} - \xi_V^{\mathbb{P}} \lambda_{V,u^-} + \frac{1}{2} V_{u^-} - \sqrt{V_{u^-}} \Lambda_{Y,u^-} \right. \\
 & \quad \left. + \zeta_Y^{\mathbb{P}} \lambda_{Y,u^-} + \zeta_V^{\mathbb{P}} \lambda_{V,u^-} + q - r \right) du \\
 & \quad + \int_0^t D_u \exp(Y_{u^-} + Z_{u^-}) \left(\rho \sqrt{V_{u^-}} - \Lambda_{V,u^-} \right) dW_{V,u} \\
 & \quad + \int_0^t D_u \exp(Y_{u^-} + Z_{u^-}) \left(\sqrt{1 - \rho^2} \sqrt{V_{u^-}} - \Lambda_{\perp,u^-} \right) dW_{\perp,u} \\
 & \quad + \int_0^t D_u \exp \left(Y_{u^-} + Z_{u^-} - (1 + \Gamma_Y) J_{Y,u^-} - \Gamma_V J_{V,u^-} + \zeta_Y^{\mathbb{P}} \int_0^u \lambda_{Y,s^-} ds \right. \\
 & \quad \left. + \zeta_V^{\mathbb{P}} \int_0^u \lambda_{V,s^-} ds \right) dM_u.
 \end{aligned}$$

Finally, as $\{D_t \exp(Y_t + Z_t)\}_{t \geq 0}$ needs to be a \mathbb{P} -martingale, the drift term must be nil, implying that

$$\alpha_{u^-}^{\mathbb{P}} = r - q - \frac{1}{2}V_{u^-} + \sqrt{V_{u^-}}\Lambda_{Y,u^-} + (\xi_Y^{\mathbb{P}} - \zeta_Y^{\mathbb{P}})\lambda_{Y,u^-} + (\xi_V^{\mathbb{P}} - \zeta_V^{\mathbb{P}})\lambda_{V,u^-}.$$

□

Corollary B.11.

$$\mathbb{E}_t^{\mathbb{P}}[\exp(Y_T)] = \exp(Y_t) \mathbb{E}_t^{\mathbb{P}}\left[\exp\left(\int_t^T m_{u^-}^{\mathbb{P}} du\right)\right]$$

where

$$\begin{aligned} m_{u^-}^{\mathbb{P}} &= \alpha_{u^-}^{\mathbb{P}} + \frac{1}{2}V_{u^-} du + (\varphi_{Z_Y}^{\mathbb{P}}(1) - 1)\lambda_{Y,u^-} \\ &= r - q + \sqrt{V_{u^-}}\Lambda_{Y,u^-} + (\xi_Y^{\mathbb{P}} - \zeta_Y^{\mathbb{P}})\lambda_{Y,u^-} + (\varphi_{Z_Y}^{\mathbb{P}}(1) - 1)\lambda_{Y,u^-} \\ &= r - q + \sqrt{V_{u^-}}\Lambda_{Y,u^-} + (\varphi_{Z_Y}^{\mathbb{P}}(\Gamma_Y) - \varphi_{Z_Y}^{\mathbb{P}}(1 + \Gamma_Y) + \varphi_{Z_Y}^{\mathbb{P}}(1) - 1)\lambda_{Y,u^-}. \end{aligned}$$

Proof of Corollary B.11. Starting from Equation (5.1),

$$\begin{aligned} &\mathbb{E}_t^{\mathbb{P}}[\exp(Y_T)] \\ &= \exp(Y_t) \mathbb{E}_t^{\mathbb{P}}\left[\exp\left(\int_t^T \alpha_{u^-}^{\mathbb{P}} du + \int_t^T \sqrt{V_{u^-}} dW_{Y,u} + \int_t^T dJ_{Y,u}\right)\right] \\ &= \exp(Y_t) \mathbb{E}_t^{\mathbb{P}}\left[\exp\left(\begin{aligned} &\int_t^T \alpha_{u^-}^{\mathbb{P}} du + \frac{1}{2}\int_t^T V_{u^-} du + (\varphi_{Z_Y}^{\mathbb{P}}(1) - 1)\int_t^T \lambda_{Y,u^-} du \\ &+ \int_t^T \sqrt{V_{u^-}} dW_{Y,u} - \frac{1}{2}\int_t^T V_{u^-} du \\ &+ \int_t^T dJ_{Y,u} - (\varphi_{Z_Y}^{\mathbb{P}}(1) - 1)\int_t^T \lambda_{Y,u^-} du \end{aligned}\right)\right] \\ &= \exp(Y_t) \\ &\quad \times \mathbb{E}_t^{\mathbb{P}}\left[\mathbb{E}_t^{\mathbb{P}}\left[\exp\left(\begin{aligned} &\int_t^T \alpha_{u^-}^{\mathbb{P}} du + \frac{1}{2}\int_t^T V_{u^-} du + (\varphi_{Z_Y}^{\mathbb{P}}(1) - 1)\int_t^T \lambda_{Y,u^-} du \\ &+ \int_t^T \sqrt{V_{u^-}} dW_{Y,u} - \frac{1}{2}\int_t^T V_{u^-} du \\ &+ \int_t^T dJ_{Y,u} - (\varphi_{Z_Y}^{\mathbb{P}}(1) - 1)\int_t^T \lambda_{Y,u^-} du \end{aligned}\right)\right]\int_t^T \lambda_{Y,u^-} du\right] \\ &= \exp(Y_t) \mathbb{E}_t^{\mathbb{P}}\left[\exp\left(\int_t^T \left(\alpha_{u^-}^{\mathbb{P}} + \frac{1}{2}V_{u^-} du + (\varphi_{Z_Y}^{\mathbb{P}}(1) - 1)\lambda_{Y,u^-}\right) du\right)\right] \\ &= \exp(Y_t) \mathbb{E}_t^{\mathbb{P}}\left[\exp\left(\int_t^T m_{u^-}^{\mathbb{P}} du\right)\right]. \end{aligned}$$

Therefore,

$$m_{u^-}^{\mathbb{P}} = \alpha_{u^-}^{\mathbb{P}} + \frac{1}{2}V_{u^-} du + (\varphi_{Z_Y}^{\mathbb{P}}(1) - 1)\lambda_{Y,u^-}$$

$$\begin{aligned}
 &= r - q - \frac{1}{2}V_{u^-} + \sqrt{V_{u^-}}\Lambda_{Y,u^-} + (\xi_Y^{\mathbb{P}} - \zeta_Y^{\mathbb{P}})\lambda_{Y,u^-} + (\xi_V^{\mathbb{P}} - \zeta_V^{\mathbb{P}})\lambda_{V,u^-} \\
 &\quad + \frac{1}{2}V_{u^-}du + (\varphi_{Z_Y}^{\mathbb{P}}(1) - 1)\lambda_{Y,u^-} \\
 &= r - q + \sqrt{V_{u^-}}\Lambda_{Y,u^-} + (\xi_Y^{\mathbb{P}} - \zeta_Y^{\mathbb{P}} + \varphi_{Z_Y}^{\mathbb{P}}(1) - 1)\lambda_{Y,u^-} + (\xi_V^{\mathbb{P}} - \zeta_V^{\mathbb{P}})\lambda_{V,u^-} \\
 &= r - q + \sqrt{V_{u^-}}\Lambda_{Y,u^-} \\
 &\quad + (\varphi_{Z_Y}^{\mathbb{P}}(1) - \varphi_{Z_Y}^{\mathbb{P}}(1 + \Gamma_Y) + \varphi_{Z_Y}^{\mathbb{P}}(1) - 1)\lambda_{Y,u^-} + (\varphi_{Z_V}^{\mathbb{P}}(1) - \varphi_{Z_V}^{\mathbb{P}}(\Gamma_V))\lambda_{V,u^-}.
 \end{aligned}$$

□

B.2 Option Pricing

B.2.1 Moment Generating Function of Log-Equity Price Variation

Lemma B.12. *The risk-neutral moment generating function of Y_T satisfies*

$$\varphi_{Y_T|Y_t,V_t}^{\mathbb{Q}}(u, Y_t, V_t) = \mathbb{E}^{\mathbb{Q}}[\exp(uY_T) | Y_t, V_t] = \exp(\mathcal{A}(u, t, T) + uY_t + C(u, t, T)V_t)$$

where $C(u, t, T)$ and $\mathcal{A}(u; t, T)$ are provided at Equations (5.28) and (5.29) respectively.

Proof of Lemma B.12. The proof is based on Filipovic and Mayerhofer (2009) and Duffie et al. (2000).

According to Duffie et al. (2000), $\mathcal{A}(u; t, T)$, $\mathcal{B}(u; t, T)$ and $C(u; t, T)$ satisfy the complex valued ordinary differential equations (ODEs)

$$\begin{aligned}
 \mathcal{A}'(u; t, T) &= r - \left[r - q - \lambda_{Y,0}^{\mathbb{Q}}(\varphi_{Z_Y}^{\mathbb{Q}}(1) - 1) \right]^{\top} \begin{bmatrix} \mathcal{B}(u; t, T) \\ C(u; t, T) \end{bmatrix} - \lambda_{Y,0}^{\mathbb{Q}}(\varphi_{Z_Y}^{\mathbb{Q}}(\mathcal{B}(u; t, T)) - 1) \\
 &\quad - \lambda_{V,0}^{\mathbb{Q}}(\varphi_{Z_V}^{\mathbb{Q}}(C(u; t, T)) - 1), \tag{B.5}
 \end{aligned}$$

$$\mathcal{B}'(u; t, T) = 0,$$

$$\begin{aligned}
 C'(u; t, T) &= \left[\frac{1}{2} + \lambda_{Y,1}^{\mathbb{Q}}(\varphi_{Z_Y}^{\mathbb{Q}}(1) - 1) \right]^{\top} \begin{bmatrix} \mathcal{B}(u; t, T) \\ C(u; t, T) \end{bmatrix} \\
 &\quad - \frac{1}{2} \begin{bmatrix} \mathcal{B}(u; t, T) \\ C(u; t, T) \end{bmatrix}^{\top} \begin{bmatrix} 1 & \sigma\rho \\ \sigma\rho & \sigma^2 \end{bmatrix} \begin{bmatrix} \mathcal{B}(u; t, T) \\ C(u; t, T) \end{bmatrix} - \lambda_{Y,1}^{\mathbb{Q}}(\varphi_{Z_Y}^{\mathbb{Q}}(u) - 1),
 \end{aligned}$$

with initial conditions $\mathcal{A}(u; T, T) = 0$, $\mathcal{B}(u; T, T) = u$ and $C(u; T, T) = 0$. The second ODE is obvious: since $\mathcal{B}(u; T, T) = u$ and $\mathcal{B}'(u; t, T) = 0$, it means that $\mathcal{B}(u; t, T) = u$. Then, from this equation, we can try to find the solution to the third ODE (i.e. $C(u; t, T)$). As in Filipovic and Mayerhofer (2009), the solution to this ODE is given by Equation (5.28). Finally, the expression for $\mathcal{A}(u; t, T)$ provided at Equation (5.29) is obtained from a simple integration since there are no \mathcal{A} in the right hand side of Equation (B.5). \square

B.2.2 Quadratic Variation of Option Prices

The option price is a function of time, returns and variance: $O_t = O_t(Y_t, V_t)$. Assume that O is twice continuously differentiable. Itô's formula implies that

$$\begin{aligned}
 & O_t(Y_t, V_t) - O_0(Y_0, V_0) \\
 &= \int_0^t \frac{\partial O_u}{\partial u}(Y_{u-}, V_{u-}) du + \int_0^t \frac{\partial O_u}{\partial y}(Y_{u-}, V_{u-}) dY_u + \int_0^t \frac{\partial O_u}{\partial v}(Y_{u-}, V_{u-}) dV_u \\
 &+ \frac{1}{2} \int_0^t \frac{\partial^2 O_u}{\partial y^2}(Y_{u-}, V_{u-}) d[Y, Y]_u^c + \frac{1}{2} \int_0^t \frac{\partial^2 O_u}{\partial v^2}(Y_{u-}, V_{u-}) d[V, V]_u^c \\
 &+ \int_0^t \frac{\partial^2 O_u}{\partial v \partial y}(Y_{u-}, V_{u-}) d[Y, V]_u^c \\
 &+ \sum_{0 < u \leq t} \{O_u(Y_u, V_u) - O_u(Y_{u-}, V_{u-})\} - \sum_{0 < u \leq t} \left\{ \frac{\partial O_u}{\partial y}(Y_{u-}, V_{u-}) (Y_u - Y_{u-}) \right\} \\
 &- \sum_{0 < u \leq t} \left\{ \frac{\partial O_u}{\partial v}(Y_{u-}, V_{u-}) (V_u - V_{u-}) \right\}.
 \end{aligned}$$

Replacing Equation (5.1) in the latter leads to

$$\begin{aligned}
 & O_t(Y_t, V_t) - O_0(Y_0, V_0) \\
 &= \int_0^t \frac{\partial O_u}{\partial u}(Y_{u-}, V_{u-}) du + \int_0^t \frac{\partial O_u}{\partial y}(Y_{u-}, V_{u-}) \alpha_{u-} du + \int_0^t \frac{\partial O_u}{\partial y}(Y_{u-}, V_{u-}) \rho \sqrt{V_{u-}} dW_{V,u} \\
 &+ \int_0^t \frac{\partial O_u}{\partial y}(Y_{u-}, V_{u-}) \sqrt{1 - \rho^2} \sqrt{V_{u-}} dW_{\perp,u} + \int_0^t \frac{\partial O_u}{\partial y}(Y_{u-}, V_{u-}) dJ_{Y,u} \\
 &+ \int_0^t \frac{\partial O_u}{\partial v}(Y_{u-}, V_{u-}) \kappa(\theta - V_{u-}) du + \int_0^t \frac{\partial O_u}{\partial v}(Y_{u-}, V_{u-}) \sigma \sqrt{V_{u-}} dW_{V,u} \\
 &+ \int_0^t \frac{\partial O_u}{\partial v}(Y_{u-}, V_{u-}) dJ_{V,u} + \frac{1}{2} \int_0^t \frac{\partial^2 O_u}{\partial y^2}(Y_{u-}, V_{u-}) V_{u-} du \\
 &+ \frac{1}{2} \int_0^t \frac{\partial^2 O_u}{\partial v^2}(Y_{u-}, V_{u-}) \sigma^2 V_{u-} du + \int_0^t \frac{\partial^2 O_u}{\partial v \partial y}(Y_{u-}, V_{u-}) \sigma \rho V_{u-} du
 \end{aligned}$$

$$\begin{aligned}
 & + \sum_{0 < u \leq t} \{O_u(Y_u, V_u) - O_u(Y_{u^-}, V_{u^-})\} - \int_0^t \frac{\partial O_u}{\partial y}(Y_{u^-}, V_{u^-}) dJ_{Y,u} \\
 & - \int_0^t \frac{\partial O_u}{\partial v}(Y_{u^-}, V_{u^-}) dJ_{V,u}.
 \end{aligned}$$

Then,

$$\begin{aligned}
 & O_t(Y_t, V_t) - O_0(Y_0, V_0) \\
 & = \int_0^t \left\{ \frac{\partial O_u}{\partial u}(Y_{u^-}, V_{u^-}) + \alpha_{u^-} \frac{\partial O_u}{\partial y}(Y_{u^-}, V_{u^-}) + \frac{\partial O_u}{\partial v}(Y_{u^-}, V_{u^-}) \kappa(\theta - V_{u^-}) \right. \\
 & \quad \left. + \left(\frac{1}{2} \frac{\partial^2 O_u}{\partial y^2}(Y_{u^-}, V_{u^-}) + \frac{1}{2} \frac{\partial^2 O_u}{\partial v^2}(Y_{u^-}, V_{u^-}) \sigma^2 + \frac{\partial^2 O_u}{\partial v \partial y}(Y_{u^-}, V_{u^-}) \sigma \rho \right) V_{u^-} \right\} du \\
 & \quad + \int_0^t \left(\rho \frac{\partial O_u}{\partial y}(Y_{u^-}, V_{u^-}) + \sigma \frac{\partial O_u}{\partial v}(Y_{u^-}, V_{u^-}) \right) \sqrt{V_{u^-}} dW_{V,u} \\
 & \quad + \int_0^t \frac{\partial O_u}{\partial y}(Y_{u^-}, V_{u^-}) \sqrt{1 - \rho^2} \sqrt{V_{u^-}} dW_{\perp,u} + \sum_{0 < u \leq t} \{O_u(Y_u, V_u) - O_u(Y_{u^-}, V_{u^-})\}.
 \end{aligned}$$

Finally, the quadratic variation is

$$\begin{aligned}
 & [O, O]_t \\
 & = \int_0^t \left(\rho \frac{\partial O_u}{\partial y}(Y_{u^-}, V_{u^-}) + \sigma \frac{\partial O_u}{\partial v}(Y_{u^-}, V_{u^-}) \right)^2 V_{u^-} du \\
 & \quad + \int_0^t \left(\frac{\partial O_u}{\partial y}(Y_{u^-}, V_{u^-}) \right)^2 (1 - \rho^2) V_{u^-} du + \sum_{0 < u \leq t} \{O_u(Y_u, V_u) - O_u(Y_{u^-}, V_{u^-})\}^2 \\
 & = \int_0^t \left(\left(\frac{\partial O_u}{\partial y}(Y_{u^-}, V_{u^-}) \right)^2 + 2\sigma\rho \frac{\partial O_u}{\partial y}(Y_{u^-}, V_{u^-}) \frac{\partial O_u}{\partial v}(Y_{u^-}, V_{u^-}) \right. \\
 & \quad \left. + \sigma^2 \left(\frac{\partial O_u}{\partial v}(Y_{u^-}, V_{u^-}) \right)^2 \right) V_{u^-} du + \sum_{0 < u \leq t} \{O_u(Y_u, V_u) - O_u(Y_{u^-}, V_{u^-})\}^2.
 \end{aligned}$$

B.2.2.1 Derivative Calculation for ΔOQV

Starting from Equation (5.27),

$$\frac{\partial}{\partial y} C_t(y, v) = \exp(y - q(T - t)) \left(P_1(y, v) + \frac{\partial P_1}{\partial x} P_1(y, v) \right) - K \exp(-r(T - t)) \frac{\partial P_2}{\partial y}(y, v)$$

and

$$\frac{\partial}{\partial v} C_t(y, v) = \exp(y - q(T - t)) \frac{\partial P_1}{\partial v} P_1(y, v) - K \exp(-r(T - t)) \frac{\partial P_2}{\partial v}(y, v).$$

The derivatives $\frac{\partial}{\partial y} P_1(y, v, k; t, T)$ and $\frac{\partial}{\partial y} P_2(y, v, k; t, T)$ can be calculated by computing the derivative of the integrand.¹ For $x \in \{y, v\}$,

$$\begin{aligned} \frac{\partial P_1}{\partial x}(y, v) &= \frac{\partial}{\partial x} \left[\frac{1}{2} + \frac{1}{\pi} \int_0^\infty \operatorname{Re} \left(\frac{1}{ui} \exp(-iuk - y) \varphi_{Y_T|Y_t, V_t}^{\mathbb{Q}}(ui + 1, y, v) \right) du \right] \\ &= \frac{1}{\pi} \int_0^\infty \frac{\partial}{\partial x} \operatorname{Re} \left(\frac{1}{ui} \exp(-iuk - y) \varphi_{Y_T|Y_t, V_t}^{\mathbb{Q}}(ui + 1, y, v) \right) du \\ &= \frac{1}{\pi} \int_0^\infty \frac{\partial}{\partial x} \operatorname{Re} \left(\frac{1}{ui} \exp(-iuk - y) \exp \left(\frac{\mathcal{A}(ui + 1, t, T) + (ui + 1)y}{+C(ui + 1, t, T)v} \right) \right) du \end{aligned}$$

where the last expression comes from the specific shape of the moment generating function. Because $\exp(a + ib) = \exp(a) \times (\cos b + i \sin b)$, the last expression becomes

$$\begin{aligned} &\frac{1}{\pi} \int_0^\infty \frac{\partial}{\partial x} \operatorname{Re} \left(\frac{1}{ui} \exp \left(\frac{\operatorname{Re}[\mathcal{A}(ui + 1, t, T)] + \operatorname{Re}[C(ui + 1, t, T)]v}{+i \operatorname{Im}[\mathcal{A}(ui + 1, t, T)] + i \operatorname{Im}[C(ui + 1, t, T)]v} + iu(y - k) \right) \right) du \\ &= \frac{1}{\pi} \int_0^\infty \frac{\partial}{\partial x} \left(\frac{1}{u} \left(\frac{\exp(\operatorname{Re}[\mathcal{A}(ui + 1, t, T)] + \operatorname{Re}[C(ui + 1, t, T)]v)}{\times \sin(\operatorname{Im}[\mathcal{A}(ui + 1, t, T)] + \operatorname{Im}[C(ui + 1, t, T)]v + u(y - k))} \right) \right) du. \end{aligned}$$

If $x = y$, then

$$\begin{aligned} \frac{\partial P_1}{\partial y}(y, v) &= \frac{1}{\pi} \int_0^\infty \left(\frac{\exp(\operatorname{Re}[\mathcal{A}(ui + 1, t, T)] + \operatorname{Re}[\gamma(ui + 1, t, T)v])}{\times \cos(\operatorname{Im}[\mathcal{A}(ui + 1, t, T)] + \operatorname{Im}[\gamma(ui + 1, t, T)v] + u(y - k))} \right) du \end{aligned}$$

For $x = v$,

$$\begin{aligned} \frac{\partial P_1}{\partial v}(y, v) &= \frac{1}{\pi} \int_0^\infty \left(\frac{\frac{1}{u} \left(\frac{\operatorname{Re}[C(ui + 1, t, T)]}{\times \exp(\operatorname{Re}[\mathcal{A}(ui + 1, t, T)] + \operatorname{Re}[C(ui + 1, t, T)]v)} \right)}{\times \sin(\operatorname{Im}[\mathcal{A}(ui + 1, t, T)] + \operatorname{Im}[C(ui + 1, t, T)]v + u(y - k))} \right) du. \\ &\quad + \frac{1}{u} \left(\frac{\operatorname{Im}[C(ui + 1, t, T)]}{\times \exp(\operatorname{Re}[\mathcal{A}(ui + 1, t, T)] + \operatorname{Re}[C(ui + 1, t, T)]v)} \right) \times \cos(\operatorname{Im}[\mathcal{A}(ui + 1, t, T)] + \operatorname{Im}[C(ui + 1, t, T)]v + u(y - k)) \end{aligned}$$

¹Note that Leibniz integral rule is used here. One should verify that the integrand is a Lebesgue-integrable function of x for each u , that for almost all x , the derivative of the integrand (w.r.t. x) exists for all x , and that there is an integrable function θ such that $\left| \frac{\partial}{\partial y} \operatorname{Re} \left(\frac{e^{-iuk-y} f(ui+1, v; t, T)}{iu} \right) \right| \leq \theta(u)$ for all x and almost every u .

B.3 Simulation-Based Results: Additional Results

B.3.1 Filtering Tests

Figures B.1 to B.5 show an example of the filtered values based on different data sources. For this experiment, M is set to 5. In general, using the five different data sources yields accurate filtered values (with narrower confidence intervals). Instantaneous variance and log-equity jumps are more precisely filtered when using the option realized variance. When only log-equity values are used, the various filtered values are very imprecise. This observation is consistent with Table 5.2: the error measures are higher when we only consider the log-equity values in the filtering step.

B.3.2 Option Data: How Much Is Enough?

We use the simulation experiment of Section 5.4.1, but this time, fifteen implied volatilities are observed on each day. These European OTM options have a maturity of 30, 90 and 150 business days, and call-equivalent deltas of 0.20, 0.35, 0.50, 0.65 and 0.80 respectively. The filter is run using $M = 5$.

The idea behind this test is to only use part of the option data available and select a limited number of maturities or call-equivalent deltas. From this subsample, we apply the filter and compare filtered values to real simulated values and to filtered values computed using the whole dataset.

Table B.1 shows RMSE for various filtered values. These error measures are computed with respect to true values and filtered means again.

Regarding the option maturity, it has a marginal impact on the results. However, according to Panel A, it seems that the use of two maturities is better than one. Therefore, adding more options helps the filtering of latent quantities, regardless of the maturity of the option. According to Panel B of Table B.1, OTM put options filter the jump components more adequately.

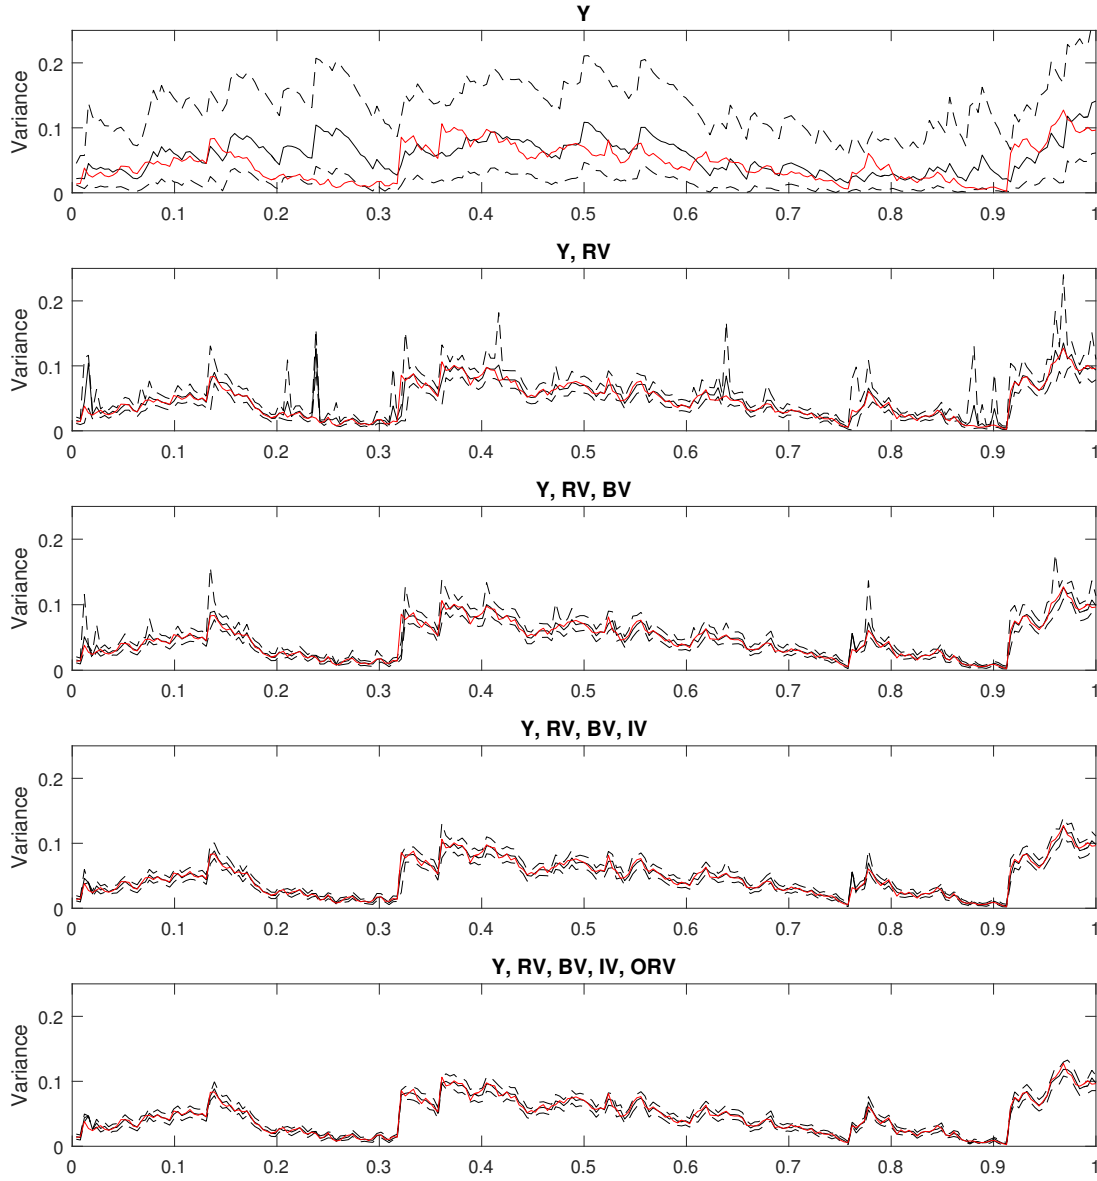


Figure B.1: Filtered instantaneous variance using various data sources.

Variance path's means obtained by the particle filters, true spot variance path, and 95% and confidence intervals computed using empirical quantiles (from particle filters) are shown in this figure using different data sources. The mean of the filtered density (obtained using particle filters) is taken as the filtered instantaneous variance. Y means daily log-equity value, RV means realized variance, BV means bipower variation, IV means implied volatility here, and ORV means option realized variance.

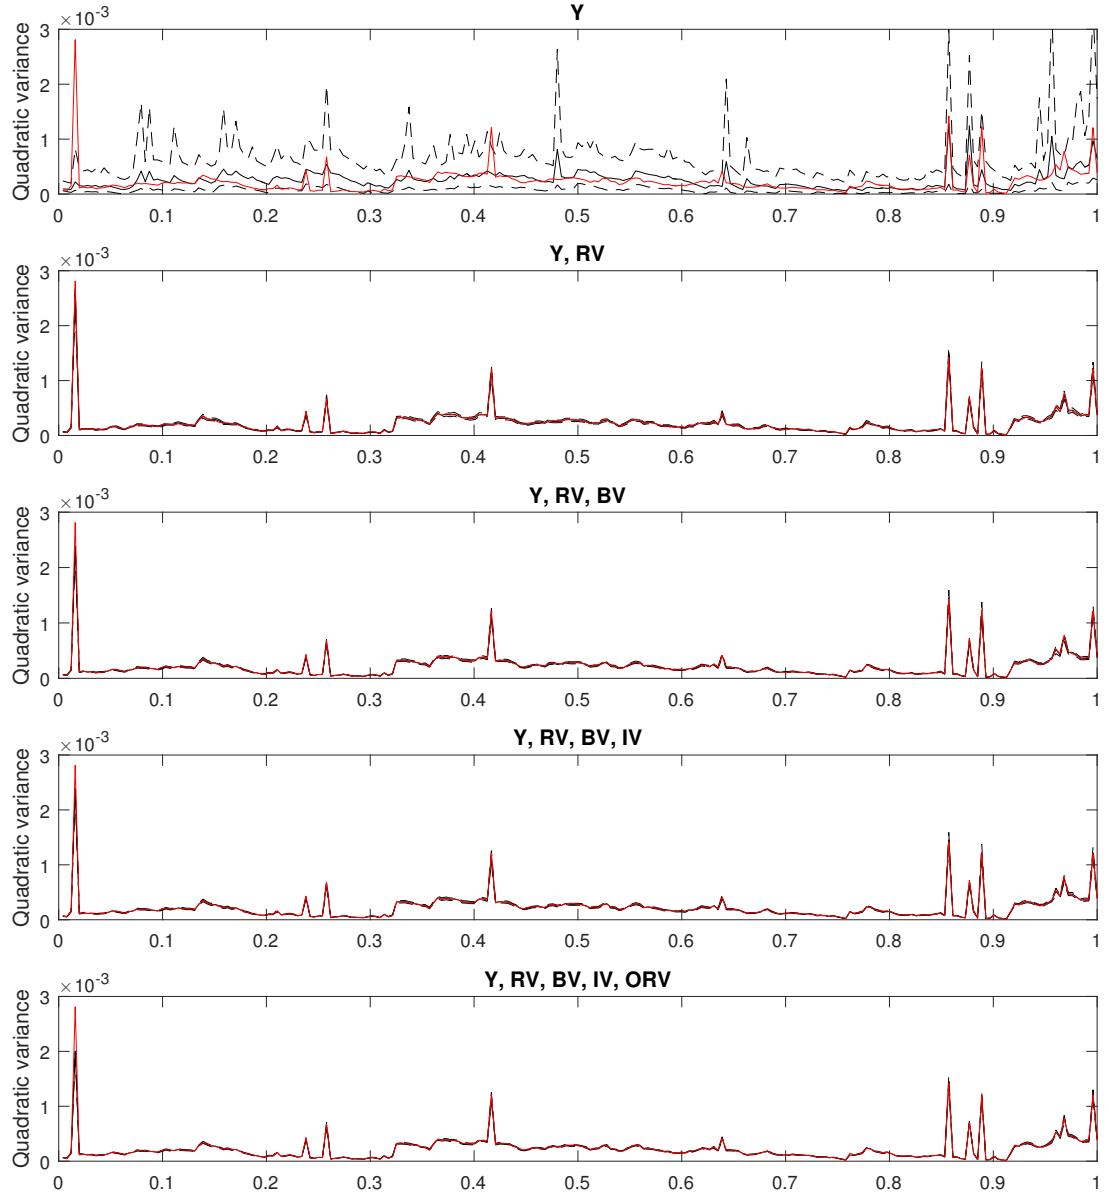


Figure B.2: Filtered quadratic variation using various data sources.

Quadratic variation path's means obtained by the particle filters, true quadratic variation path, and 95% and confidence intervals computed using empirical quantiles (from particle filters) are shown in this figure using different data sources. Y means daily log-equity value, RV means realized variance, BV means bipower variation, IV means implied volatility here, and ORV means option realized variance.

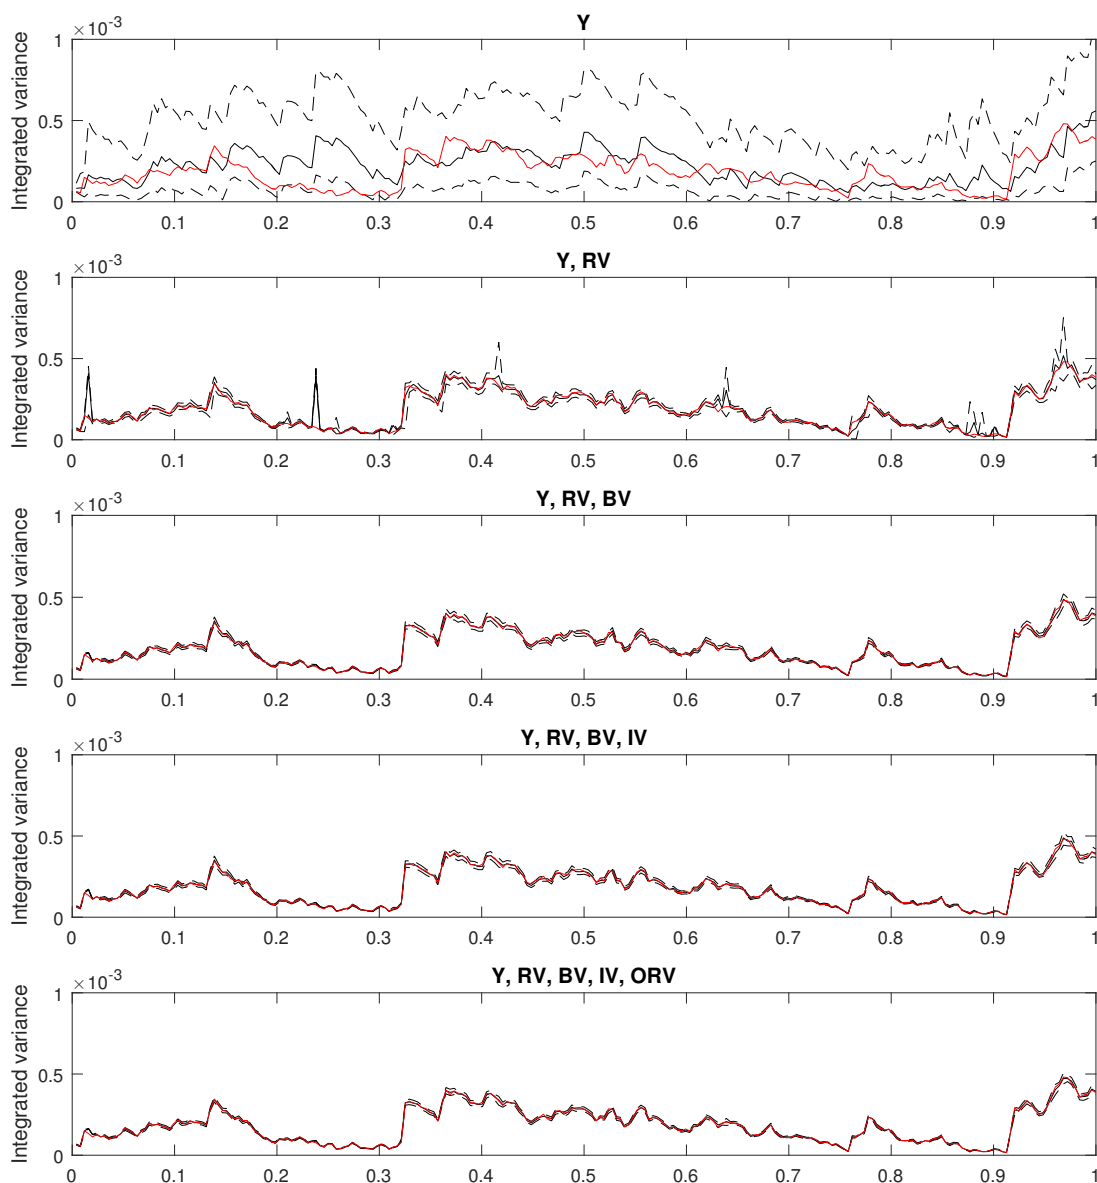


Figure B.3: Filtered integrated variance using various data sources.

Integrated variance path's means obtained by the particle filters, true integrated variance path, and 95% and confidence intervals computed using empirical quantiles (from particle filters) are shown in this figure using different data sources. Y means daily log-equity value, RV means realized variance, BV means bipower variation, IV means implied volatility here, and ORV means option realized variance.

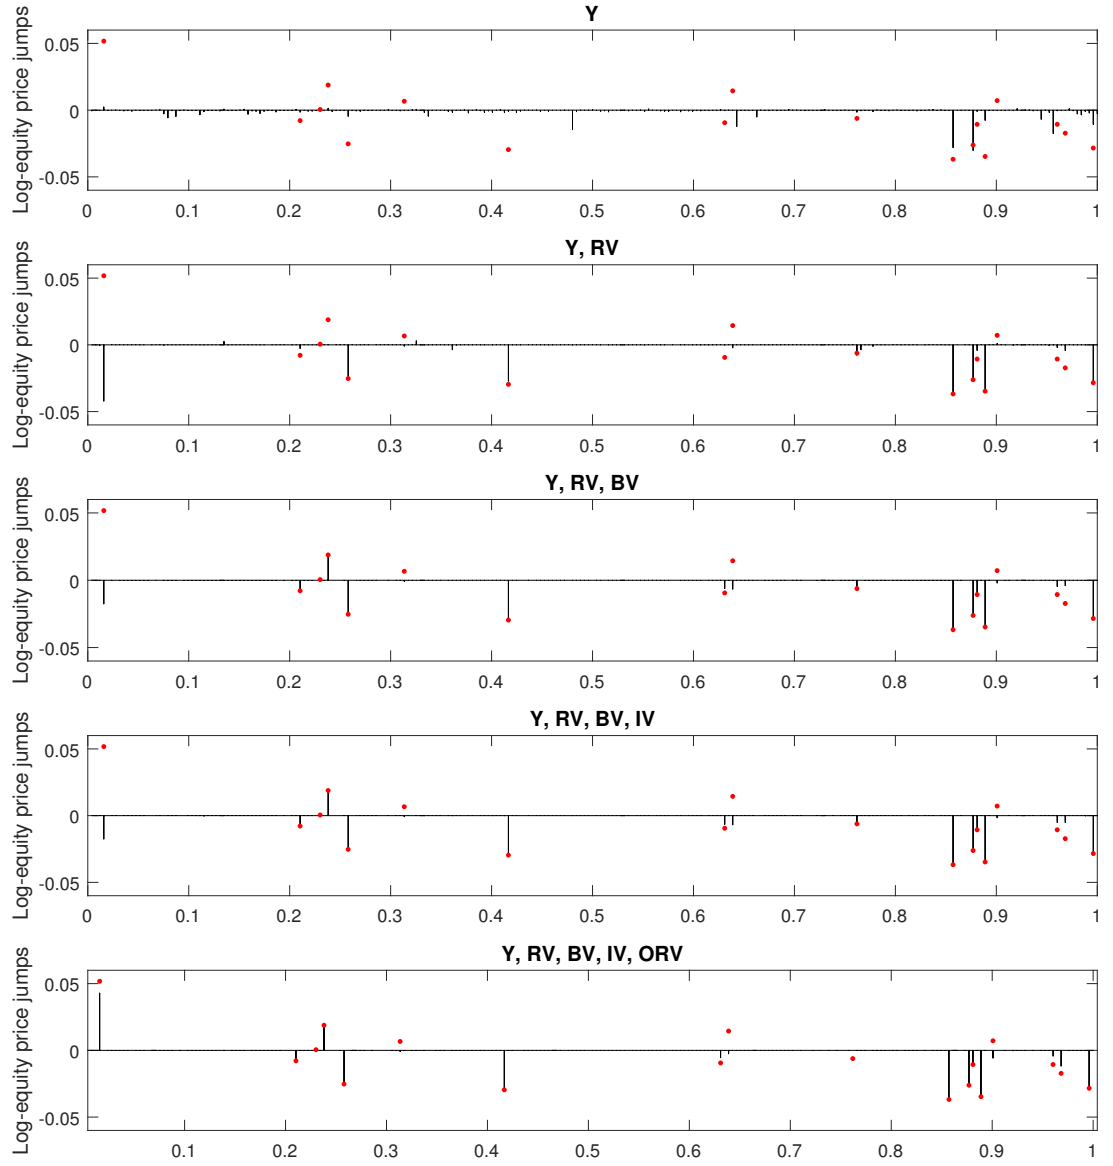


Figure B.4: Filtered log-equity jumps using various data sources.

Return jump path's means obtained by the particle filters and true return jump path are shown in this figure using different data sources. Y means daily log-equity value, RV means realized variance, BV means bipower variation, IV means implied volatility here, and ORV means option realized variance.

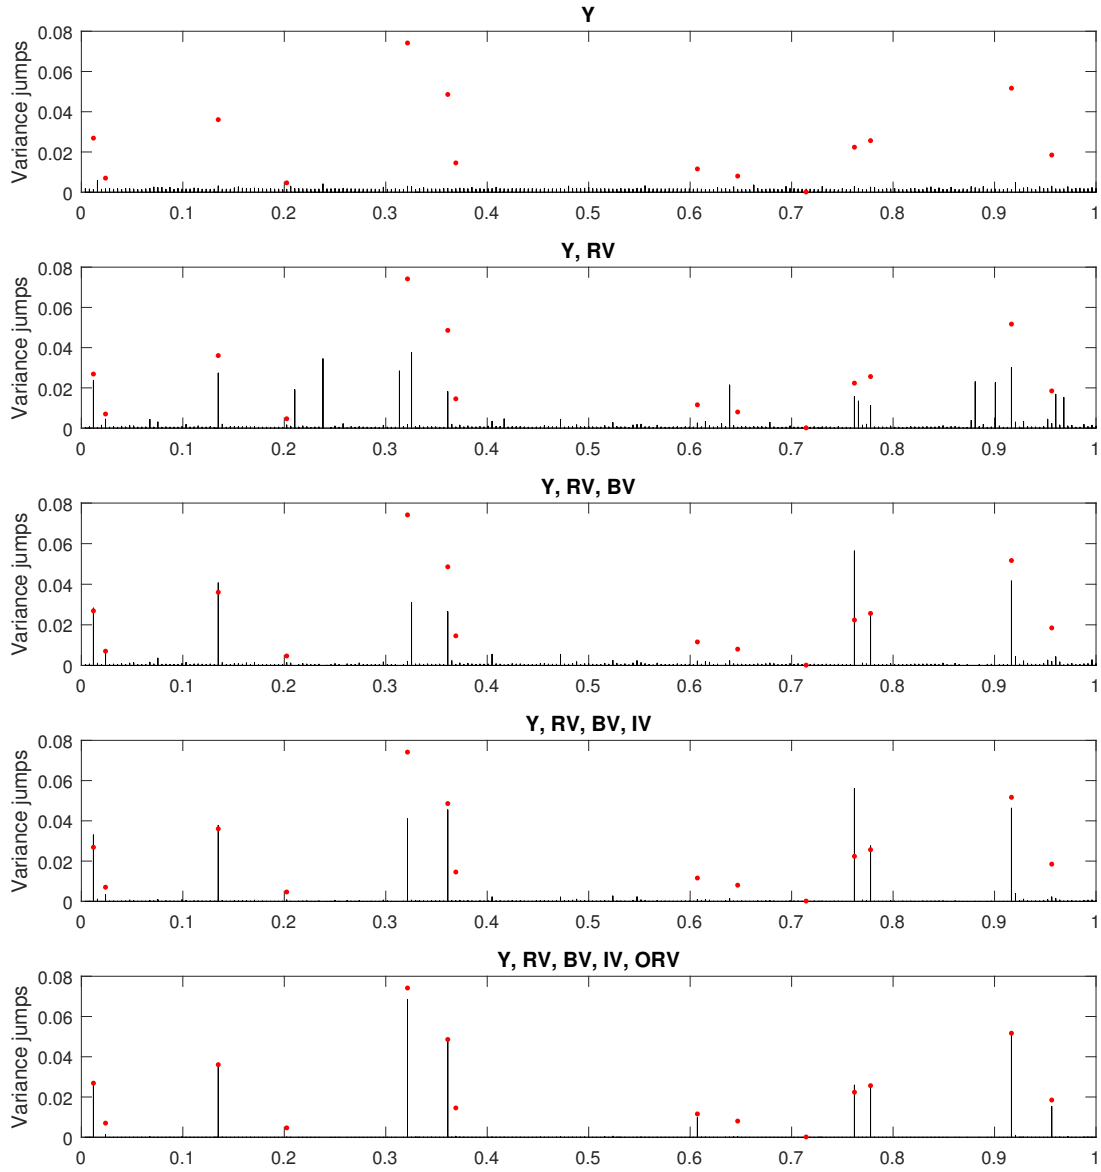


Figure B.5: Filtered variance jumps using various data sources.

Variance jump path's means obtained by the particle filters and true variance jump path are shown in this figure using different data sources. Y means daily returns, RV means realized variance, BV means bipower variation, IV means implied volatility here, and ORV means option realized variance.

Table B.1: RMSE for various quantities across 100 simulated paths when using only a limited number of options.

Panel A: Maturities										
	RMSE (true)					RMSE (filtered)				
	V	ΔQV	ΔI	Z_Y	Z_V	V	ΔQV	ΔI	Z_Y	Z_V
DTM = 30	5.9196	0.0212	0.0116	1.8850	2.2890	3.9730	0.0238	0.0135	0.7718	1.0889
DTM = 90	6.3944	0.0202	0.0122	1.8452	1.8596	4.0821	0.0226	0.0140	0.7283	0.9411
DTM = 150	6.6394	0.0198	0.0126	1.9275	1.9132	4.1860	0.0229	0.0143	0.7591	0.8998
DTM = 30, 90	6.0780	0.0201	0.0116	1.8852	1.9806	3.8558	0.0226	0.0133	0.7273	0.8593
DTM = 30, 150	6.1739	0.0188	0.0118	1.8551	2.0669	3.8661	0.0232	0.0135	0.7240	0.8453
DTM = 120, 150	6.4748	0.0194	0.0122	1.8500	1.8894	3.9892	0.0224	0.0137	0.7676	0.7847
Panel B: Moneynesses										
	RMSE (true)					RMSE (filtered)				
	V	ΔQV	ΔI	Z_Y	Z_V	V	ΔQV	ΔI	Z_Y	Z_V
$\Delta^e = 0.20$	6.7777	0.0271	0.0147	1.8839	2.3313	4.6669	0.0290	0.0158	0.8560	1.2977
$\Delta^e = 0.35$	6.8674	0.0265	0.0149	1.9021	2.6694	4.7140	0.0286	0.0158	0.7991	1.7129
$\Delta^e = 0.50$	6.9247	0.0246	0.0150	1.9514	3.1289	4.8637	0.0272	0.0162	0.8819	2.1768
$\Delta^e = 0.65$	6.8310	0.0244	0.0149	1.8926	2.4371	4.9126	0.0298	0.0167	0.9845	1.5027
$\Delta^e = 0.80$	6.7287	0.0263	0.0145	1.9438	2.0052	4.7973	0.0307	0.0163	0.9010	1.2448
$\Delta^e = 0.20, 0.35$	6.6490	0.0250	0.0144	1.8983	2.3443	4.4762	0.0266	0.0154	0.8107	1.3127
$\Delta^e = 0.65, 0.80$	6.6995	0.0244	0.0146	1.8956	2.1555	4.7616	0.0292	0.0164	0.8497	1.2778

Quantities were multiplied by 1,000. Filtered values are computed as the mean of resampled particles obtained via a SIR particle filter with $M = 5$ using Y , RV , BV , I and ORV and every maturities and moneynesses. A maximum of fifteen implied volatilities are observed on each day: these European OTM options have a maturity of 30, 90 and 150 business days, and call-equivalent deltas of 0.20, 0.35, 0.50, 0.65 and 0.80 respectively. In the table, DTM means day to maturity and $\Delta^e = 0.20$ represents call-equivalent delta.

Table B.2: Descriptive statistics of realized variance, bipower variation and option realized variance.

	Mean	S.D.	Skewness	Kurtosis	10%	90%
$RV \times 1000$	27.523	75.247	13.846	320.527	3.990	53.757
$BV \times 1000$	25.783	72.283	13.371	294.644	3.591	49.767
ORV , DTM = 30, $\Delta^e = 0.20$	24.001	83.350	8.827	104.327	0.792	55.478
ORV , DTM = 30, $\Delta^e = 0.35$	33.794	87.637	8.314	92.257	2.832	68.863
ORV , DTM = 30, $\Delta^e = 0.50$	33.373	68.679	8.518	104.435	6.387	65.347
ORV , DTM = 30, $\Delta^e = 0.65$	29.858	67.238	9.305	130.376	4.279	59.223
ORV , DTM = 30, $\Delta^e = 0.80$	19.484	46.732	6.753	66.278	1.307	44.229
ORV , DTM = 90, $\Delta^e = 0.20$	17.434	78.232	9.779	119.597	0.695	22.679
ORV , DTM = 90, $\Delta^e = 0.35$	26.705	77.032	10.151	141.495	2.510	53.110
ORV , DTM = 90, $\Delta^e = 0.50$	34.453	74.512	8.202	97.525	5.806	66.649
ORV , DTM = 90, $\Delta^e = 0.65$	25.184	53.298	8.336	109.830	3.671	47.970
ORV , DTM = 90, $\Delta^e = 0.80$	15.027	55.173	16.408	363.312	1.272	30.695

S.D. stands for standard deviation. 10% and 90% represent the 10% and 90% quantiles of empirical samples respectively. RV and BV are given on an annualized basis.

B.4 Descriptive Statistics of Realized Variations

Table B.2 exhibits descriptive statistics of realized variance, bipower variation and option realized variance.

Bibliography

- Acharya, V., R. Engle, and M. Richardson. 2012. Capital shortfall: A new approach to ranking and regulating systemic risks. *American Economic Review* 102:59–64.
- Acharya, V., L. H. Pedersen, T. Philippon, and M. P. Richardson. 2010. Measuring systemic risk. *Working paper*.
- Acharya, V. V., S. T. Bharath, and A. Srinivasan. 2007. Does industry-wide distress affect defaulted firms? Evidence from creditor recoveries. *Journal of Financial Economics* 85:787–821.
- Adrian, T., and M. K. Brunnermeier. 2009. CoVaR. *Working paper*.
- Aït-Sahalia, Y., and A. Lo. 1998a. Nonparametric estimation of state-price densities implicit in financial asset prices. *Journal of Finance* 53:499–547.
- Aït-Sahalia, Y., and A. W. Lo. 1998b. Nonparametric estimation of state-price densities implicit in financial asset prices. *Journal of Finance* 53:499–547.
- Albuquerque, R. 2012. Skewness in stock returns: Reconciling the evidence on firm versus aggregate returns. *Review of Financial Studies* 25:1630–1673.
- Alexander, C., and A. Kaeck. 2008. Regime dependent determinants of credit default swap spreads. *Journal of Banking & Finance* 32:1008–1021.
- Alizadeh, S., M. W. Brandt, and F. X. Diebold. 2002. Range-based estimation of stochastic volatility models. *Journal of Finance* 57:1047–1091.
- Allen, M., C. B. Rosenberg, C. Keller, B. Setser, and N. Roubini. 2002. A balance sheet approach to financial crisis. *Working paper*.
- Altman, E. I., B. Brady, A. Resti, and A. Sironi. 2005. The link between default and recovery rates: Theory, empirical evidence, and implications. *Journal of Business* 78:2203–2228.

- Amihud, Y. 2002. Illiquidity and stock returns: Cross-section and time-series effects. *Journal of Financial Markets* 5:31–56.
- Andersen, T. G., T. Bollerslev, F. X. Diebold, and H. Ebens. 2001. The distribution of realized stock return volatility. *Journal of Financial Economics* 61:43–76.
- Andersen, T. G., O. Bondarenko, V. Todorov, and G. Tauchen. 2015a. The fine structure of equity-index option dynamics. *Journal of Econometrics* 187:532–546.
- Andersen, T. G., N. Fusari, and V. Todorov. 2015b. The risk premia embedded in index options. *Journal of Financial Economics* 117:558–584.
- Ang, A., R. J. Hodrick, Y. Xing, and X. Zhang. 2006. The cross-section of volatility and expected returns. *Journal of Finance* 61:259–299.
- Arora, N., P. Gandhi, and F. A. Longstaff. 2012. Counterparty Credit Risk and the Credit Default Swap Market. *Journal of Financial Economics* 103:280–293.
- Audrino, F., and M. R. Fengler. 2015. Are classical option pricing models consistent with observed option second-order moments? Evidence from high-frequency data. *Journal of Banking & Finance* 61:46–63.
- Augustin, P. 2012. The term structure of CDS spreads and sovereign credit risk. *Working paper*.
- Babaoğlu, K. 2015. The pricing of market jumps in the cross-section of stocks and options. *Working paper*.
- Bade, B., D. Rösch, and H. Scheule. 2011. Default and recovery risk dependencies in a simple credit risk model. *European Financial Management* 17:120–144.
- Baek, E., and W. Brock. 1992. A general test for nonlinear Granger causality: Bivariate model. *Working paper*.
- Bai, J., and P. Collin-Dufresne. 2011. The determinants of the CDS-bond basis during the financial crisis of 2007–2009. *Working paper*.
- Bakshi, G., C. Cao, and Z. Chen. 1997. Empirical performance of alternative option pricing models. *Journal of Finance* 52:2003–2049.
- Bakshi, G., P. Carr, and L. Wu. 2008. Stochastic risk premiums, stochastic skewness in currency options, and stochastic discount factors in international economies. *Journal of Financial Economics* 87:132–156.

- Bakshi, G., N. Kapadia, and D. Madan. 2003. Stock returns characteristics, skew laws, and the differential pricing of individual equity options. *Review of Financial Studies* 16:101–143.
- Bakshi, G., D. Madan, and G. Panayotov. 2010. Returns of claims on the upside and the viability of U-shaped pricing kernels. *Journal of Financial Economics* 97:130–154.
- Bakshi, G., D. Madan, and F. X. Zhang. 2006a. Investigating the role of systematic and firm-specific factors in default risk: Lessons from empirically evaluating credit risk models. *Journal of Business* 79:1955–1987.
- Bakshi, G., D. Madan, and F. X. Zhang. 2006b. Understanding the role of recovery in default risk models: Empirical comparisons and implied recovery rates. *Working paper*.
- Bandi, F. M., and R. Reno. 2016. Price and volatility co-jumps. *Journal of Financial Economics* 119:107–146.
- Bardgett, C., E. Gourier, and M. Leippold. 2015. Inferring volatility dynamics and risk premia from the S&P 500 and VIX markets. *Working paper*.
- Barndorff-Nielsen, O. E., and N. Shephard. 2002. Estimating quadratic variation using realized variance. *Journal of Applied Econometrics* 17:457–477.
- Barndorff-Nielsen, O. E., and N. Shephard. 2004. Power and bipower variation with stochastic volatility and jumps. *Journal of Financial Econometrics* 2:1–37.
- Bates, D. S. 1991. The Crash of '87: Was it expected? The evidence from options markets. *Journal of Finance* 46:1009–1044.
- Bates, D. S. 1996. Jumps and stochastic volatility: exchange rate processes implicit in Deutsche Mark options. *Review of Financial Studies* 9:69–107.
- Bates, D. S. 2000. Post-'87 crash fears in the S&P 500 futures option market. *Journal of Econometrics* 94:181–238.
- Bates, D. S. 2006. Maximum likelihood estimation of latent affine processes. *Review of Financial Studies* 19:909–965.
- Bates, D. S. 2008. The market for crash risk. *Journal of Economic Dynamics and Control* 32:2291–2321.

- Bégin, J.-F., M. Boudreault, and G. Gauthier. 2016. Firm-specific credit risk modelling in the presence of statistical regimes and noisy prices. *Working paper*.
- Bharath, S. T., and T. Shumway. 2008. Forecasting default with the Merton distance to default model. *Review of Financial Studies* 21:1339–1369.
- Bielecki, T., D. Brigo, and F. Patras. 2011. *Credit risk frontiers: Subprime crisis, pricing and hedging, CVA, MBS, ratings, and liquidity*. John Wiley and Sons.
- Billio, M., M. Getmansky, A. W. Lo, and L. Pelizzon. 2012. Econometric measures of connectedness and systemic risk in the finance and insurance sectors. *Journal of Financial Economics* 104:535–559.
- Black, F., and J. C. Cox. 1976. Valuing corporate securities: Some effects of bond indenture provisions. *Journal of Finance* 31:351–367.
- Black, F., and M. Scholes. 1973. The pricing of options and corporate liabilities. *Journal of Political Economy* 81:637–654.
- Bollen, N. P. 1998. Valuing options in regime-switching models. *Journal of Derivatives* 6:38–49.
- Bollerslev, T. 1986. Generalized autoregressive conditional heteroskedasticity. *Journal of Econometrics* 31:307–327.
- Bollerslev, T., and V. Todorov. 2011. Tails, fears, and risk premia. *Journal of Finance* 66:2165–2211.
- Boloorforoosh, A. 2014. Is idiosyncratic volatility risk priced? Evidence from the physical and risk-neutral distributions. *Working paper*.
- Boudreault, M., G. Gauthier, and T. Thomassin. 2013. Recovery rate risk and credit spreads in a hybrid credit risk model. *Journal of Credit Risk* 9:3–39.
- Boudreault, M., G. Gauthier, and T. Thomassin. 2014. Contagion effect on bond portfolio risk measures in a hybrid credit risk model. *Finance Research Letters* 11:131–139.
- Boudreault, M., G. Gauthier, and T. Thomassin. 2015. Estimation of correlations in portfolio credit risk models based on noisy security prices. *Journal of Economic Dynamics and Control* 61:334–349.
- Brennan, M. 1979. The pricing of contingent claims in discrete time models. *Journal of Finance* 34:53–68.

- Brigo, D., M. Predescu, A. Capponi, T. R. Bielecki, and F. Patras. 2011. Liquidity modeling for credit default swaps: An overview. In *Credit risk frontiers: Subprime crisis, pricing and hedging, CVA, MBS, ratings, and liquidity*, pp. 585–617. John Wiley and Sons.
- Broadie, M., M. Chernov, and M. Johannes. 2007. Model specification and risk premia: Evidence from futures options. *Journal of Finance* 62:1453–1490.
- Broadie, M., and O. Kaya. 2006. Exact simulation of stochastic volatility and other affine jump diffusion processes. *Operations Research* 54:217–231.
- Bruche, M., and C. González-Aguado. 2010. Recovery rates, default probabilities, and the credit cycle. *Journal of Banking & Finance* 34:754–764.
- Bühler, W., and M. Trapp. 2009a. Explaining the bond-CDS basis: The role of credit risk and liquidity. *Working paper* pp. 1–35.
- Bühler, W., and M. Trapp. 2009b. Time-varying credit risk and liquidity premia in bond and CDS markets. *Working paper*.
- Carhart, M. 1997. On persistence in mutual fund performance. *Journal of Finance* 52:57–82.
- Carr, P., H. Geman, D. Madan, and M. Yor. 2002. The fine structure of asset returns: An empirical investigation. *Journal of Business* 75:305–332.
- Carr, P., and L. Wu. 2007. Stochastic skew in currency options. *Journal of Financial Economics* 86:213–247.
- Carr, P., and L. Wu. 2010. Stock options and credit default swaps: A joint framework for valuation and estimation. *Journal of Financial Econometrics* 8:409–449.
- Chen, H., J. D. Cummins, K. S. Viswanathan, and M. A. Weiss. 2014. Systemic risk and the interconnectedness between banks and insurers: An econometric analysis. *Journal of Risk and Insurance* 81:623–652.
- Chernov, M. 2003. Empirical reverse engineering of the pricing kernel. *Journal of Econometrics* 116:329–364.
- Chernov, M., A. R. Gallant, E. Ghysels, and G. Tauchen. 2003. Alternative models for stock price dynamics. *Journal of Econometrics* 116:225–257.

- Chernov, M., and E. Ghysels. 2000. A study towards a unified approach to the joint estimation of objective and risk neutral measures for the purpose of options valuation. *Journal of Financial Economics* 56:407–458.
- Christoffersen, P., C. Dorion, K. Jacobs, and L. Karoui. 2014a. Nonlinear Kalman filtering in affine term structure models. *Management Science* 60:2248–2268.
- Christoffersen, P., R. Elkamhi, B. Feunou, and K. Jacobs. 2010. Option valuation with conditional heteroskedasticity and nonnormality. *Review of Financial Studies* 23:2139–2183.
- Christoffersen, P., B. Feunou, K. Jacobs, and N. Meddahi. 2014b. The economic value of realized volatility: Using high-frequency returns for option valuation. *Journal of Financial and Quantitative Analysis* 49:663–697.
- Christoffersen, P., B. Feunou, and Y. Jeon. 2015. Option valuation with observable volatility and jump dynamics. *Journal of Banking & Finance* 61:S101–S120.
- Christoffersen, P., M. Fournier, and K. Jacobs. 2013. The factor structure in equity options. *Working paper*.
- Christoffersen, P., K. Jacobs, and C. Ornathanalai. 2012. Dynamic jump intensities and risk premiums: Evidence from S&P 500 returns and options. *Journal of Financial Economics* 106:447–472.
- Christoffersen, P. F., K. Jacobs, C. Ornathanalai, and Y. Wang. 2008. Option valuation with long-run and short-run volatility components. *Journal of Financial Economics* 90:272–297.
- Cont, R., and T. Kokholm. 2013. A consistent pricing model for index options and volatility derivatives. *Mathematical Finance* 23:248–274.
- Cox, J. C., S. A. Ross, and M. Rubinstein. 1979. Option pricing: A simplified approach. *Journal of Financial Economics* 7:229–263.
- Creal, D. 2012. A survey of sequential Monte Carlo methods for economics and finance. *Econometric Reviews* 31:245–296.
- Cummins, J. D., and M. A. Weiss. 2014. Systemic risk and the U.S. insurance sector. *Journal of Risk and Insurance* 81:489–528.
- Davis, M., and V. Lo. 2001. Infectious defaults. *Quantitative Finance* 1:382–387.

- Dick-Nielsen, J. 2009. Liquidity biases in TRACE. *Journal of Fixed Income* 19:43–55.
- Dick-Nielsen, J., P. Feldhütter, and D. Lando. 2012. Corporate bond liquidity before and after the onset of the subprime crisis. *Journal of Financial Economics* 103:471–492.
- Diebold, F. X., and R. S. Mariano. 1995. Comparing predictive accuracy. *Journal of Business & Economic Statistics* 13:253–263.
- Diks, C., and V. Panchenko. 2006. A new statistic and practical guidelines for non-parametric Granger causality testing. *Journal of Economic Dynamics and Control* 30:1647–1669.
- Dionne, G., G. Gauthier, K. Hammami, M. Maurice, and J.-G. Simonato. 2010. Default risk in corporate yield spreads. *Financial Management* 39:707–731.
- Dionne, G., and O. Maalaoui Chun. 2013. Default and liquidity regimes in the bond market during the 2002–2012 period. *Canadian Journal of Economics* 46:1160–1195.
- Doshi, H., J. Ericsson, K. Jacobs, and S. M. Turnbull. 2013. Pricing credit default swaps with observable covariates. *Review of Financial Studies* 26:2049–2094.
- Douglas, G. W. 1969. Risk in the equity markets: An empirical appraisal of market efficiency. *Yale Economic Essays* 9:3–45.
- Duan, J.-C., and A. Fulop. 2009. Estimating the structural credit risk model when equity prices are contaminated by trading noises. *Journal of Econometrics* 150:288–296.
- Duan, J.-C., and J. Wei. 2009. Systematic risk and the price structure of individual equity options. *Review of Financial Studies* 22:1981–2206.
- Duffie, D., and N. Gârleanu. 2001. Risk and valuation of collateralized debt obligations. *Financial Analysis Journal* 57:41–59.
- Duffie, D., and D. Lando. 2001. Term structures of credit spreads with incomplete accounting information. *Econometrica* 69:633–664.
- Duffie, D., J. Pan, and K. Singleton. 2000. Transform analysis and asset pricing for affine jump-diffusions. *Econometrica* 68:1343–1376.

- Duffie, D., and K. Singleton. 1993. Simulated moments estimation of Markov models of asset prices. *Econometrica* 61:929–952.
- Duffie, D., and K. J. Singleton. 1999. Modeling term structures of defaultable bonds. *Review of Financial Studies* 12:687–720.
- Dumas, B., J. Fleming, and R. E. Whaley. 1998. Implied volatility functions: Empirical tests. *Journal of Finance* 53:2059–2106.
- Durham, G., J. Geweke, and P. Ghosh. 2015. A comment on Christoffersen, Jacobs, and Ornathanalai (2012), “Dynamic jump intensities and risk premiums: Evidence from S&P 500 returns and options”. *Journal of Financial Economics* 115:210–214.
- Elizalde, A., S. Doctor, and Y. Saltuk. 2009. Bond-CDS basis handbook. Tech. rep., JP Morgan European Credit Derivatives Research.
- Elkamhi, R., and C. Ornathanalai. 2010. Market jump risk and the price structure of individual equity options. *Working paper* pp. 1–55.
- Elliott, R. J., L. Aggoun, and J. B. Moore. 1995. *Hidden Markov models: Estimation and control*. Springer Science & Business Media.
- Engle, R. 2002. Dynamic conditional correlation: A simple class of multivariate generalized autoregressive conditional heteroskedasticity models. *Journal of Business & Economic Statistics* 20:339–350.
- Engle, R., E. Ghysels, and B. Sohn. 2013. Stock market volatility and macroeconomic fundamentals. *Review of Economics and Statistics* 95:776–797.
- Engle, R. F. 1982. Autoregressive Conditional Heteroscedasticity with Estimates of the Variance of United Kingdom Inflation. *Econometrica* 50:987–1007.
- Engle, R. F., and G. Lee. 1999. A permanent and transitory component model of stock return volatility. In *Cointegration, causality, and forecasting: A festschrift in honor of Clive W.J. Granger*, pp. 475–497. Oxford University Press.
- Engle, R. F., and J. G. Rangel. 2008. The spline-GARCH model for low-frequency volatility and its global macroeconomic causes. *Review of Financial Studies* 21:1187–1222.
- Eraker, B. 2001. MCMC analysis of diffusion models with application to finance. *Journal of Business & Economic Statistics* 19:177–191.

- Eraker, B. 2004. Do stock prices and volatility jump? Reconciling evidence from spot and option prices. *Journal of Finance* 59:1367–1404.
- Eraker, B. 2008. Affine general equilibrium models. *Management Science* 54:2068–2080.
- Eraker, B., M. Johannes, and N. Polson. 2003. The impact of jumps in volatility and returns. *Journal of Finance* 58:1269–1300.
- Ericsson, J., K. Jacobs, and R. Oviedo. 2009. The determinants of credit default swap premia. *Journal of Financial and Quantitative Analysis* 44:109–132.
- Çetin, U., R. Jarrow, P. Protter, and Y. Yildirim. 2004. Modeling credit risk with partial information. *The Annals of Applied Probability* 14:1167–1178.
- Fabozzi, F. J., and S. V. Mann. 2012. *The handbook of fixed income securities*. McGraw Hill Professional.
- Fama, E. F., and K. French. 1992. The cross-section of expected stock returns. *Journal of Finance* 47:427–465.
- Fama, E. F., and K. French. 1993. Common risk factors in the returns on stocks and bonds. *Journal of Financial Economics* 33:3–56.
- Fama, E. F., and K. French. 2015. A five-factor asset pricing model. *Journal of Financial Economics* 116:1–22.
- Fearnhead, P., and P. Clifford. 2003. On-line inference for hidden Markov models via particle filters. *Journal of the Royal Statistical Society* 65:887–899.
- Feldhütter, P. 2012. The same bond at different prices: Identifying search frictions and selling pressures. *Review of Financial Studies* 25:1155–1206.
- Feller, W. 1971. *An introduction to probability and its applications*. John Wiley & Sons.
- Filipovic, D., and E. Mayerhofer. 2009. Affine diffusion processes: Theory and applications. *Advanced Financial Modelling* 8:125.
- Fontana, A. 2010. The negative CDS-bond basis and convergence trading during the 2007/09 financial crisis. *Working paper*.
- Frey, R., and A. J. McNeil. 2003. Dependent defaults in models of portfolio credit risk. *Journal of Risk* 6:59–92.

- Friewald, N., R. Jankowitsch, and M. G. Subrahmanyam. 2012. Illiquidity or credit deterioration: A study of liquidity in the US corporate bond market during financial crises. *Journal of Financial Economics* 105:18–36.
- Fulop, A., J. Li, and J. Yu. 2014. Self-exciting jumps, learning, and asset pricing implications. *Review of Financial Studies* 28:876–912.
- Gallant, A. R., and G. Tauchen. 2010. Simulated score methods and indirect inference for continuous-time models. *Handbook of Financial Econometrics* 1:427–477.
- Garzarelli, F. 2009. The 2007-09 credit crisis and its aftermath. *Working paper*.
- Giesecke, K. 2006. Default and information. *Journal of Economic Dynamics and Control* 30:2281–2303.
- Giesecke, K., and L. Goldberg. 2003. Forecasting default in the face of uncertainty. *Journal of Derivatives* 12:11–25.
- Gil-Pelaez, J. 1951. Note on the inversion theorem. *Biometrika* 38:481–482.
- Gordon, N. J., D. J. Salmond, and A. F. Smith. 1993. Novel approach to nonlinear/non-Gaussian Bayesian state estimation. In *IEE Proceedings F (Radar and Signal Processing)*, vol. 140, pp. 107–113.
- Gourier, E. 2014. Pricing of idiosyncratic equity and variance risks. *Working paper*.
- Goyal, A., and P. Santa-Clara. 2003. Idiosyncratic risk matters! *Journal of Finance* 58:975–1008.
- Granger, C. W. 1969. Investigating causal relations by econometric models and cross-spectral methods. *Econometrica* 37:424–438.
- Guarin, A., X. Liu, and W. L. Ng. 2014. Recovering default risk from CDS spreads with a nonlinear filter. *Journal of Economic Dynamics and Control* 38:87–104.
- Gupton, G. M., C. C. Finger, and M. Bhatia. 2007. *Creditmetrics: Technical document*. CreditMetrics.
- Hamilton, J. D. 1994. *Time series analysis*. Princeton University Press.
- Han, S., and H. Zhou. 2008. Effects of liquidity on the nondefault component of corporate yield spreads: evidence from intraday transactions data. *Working paper*.

- Harrington, S. E. 2009. The financial crisis, systemic risk, and the future of insurance regulation. *Journal of Risk and Insurance* 76:785–819.
- Herskovic, B., B. Kelly, H. Lustig, and S. Van Nieuwerburgh. 2016. The common factor in idiosyncratic volatility: Quantitative asset pricing implications. *Journal of Financial Economics* 119:249–283.
- Heston, S. 1993. A closed-form solution for options with stochastic volatility with applications to bond and currency options. *Review of Financial Studies* 6:327.
- Heston, S., and S. Nandi. 2000. A closed-form GARCH option valuation model. *Review of Financial Studies* 13:585–625.
- Hiemstra, C., and J. D. Jones. 1994. Testing for linear and nonlinear Granger causality in the stock price-volume relation. *Journal of Finance* 49:1639–1664.
- Hu, F., and J. V. Zidek. 2002. The weighted likelihood. *Canadian Journal of Statistics* 30:347–371.
- Huang, A. Y., and W.-C. Hu. 2012. Regime switching dynamics in credit default swaps: Evidence from smooth transition autoregressive model. *Physica A: Statistical Mechanics and its Applications* 391:1497–1508.
- Huang, S. J., and J. Yu. 2010. Bayesian analysis of structural credit risk models with microstructure noises. *Journal of Economic Dynamics and Control* 34:2259–2272.
- Huang, X., H. Zhou, and H. Zhu. 2009. A framework for assessing the systemic risk of major financial institutions. *Journal of Banking and Finance* 33:2036–2049.
- Huang, X., H. Zhou, and H. Zhu. 2012. Assessing the systemic risk of a heterogeneous portfolio of banks during the recent financial crisis. *Journal of Financial Stability* 8:193–205.
- Hull, J., M. Predescu, and A. White. 2010. The valuation of correlation-dependent credit derivatives using a structural model. *Journal of Credit Risk* 6:99–132.
- Hürzeler, M., and H. R. Künsch. 2001. Approximating and maximising the likelihood for a general state-space model. In *Sequential Monte Carlo Methods in Practice*, pp. 159–175. Springer.
- Jackwerth, J. 2000. Recovering risk aversion from option prices and realized returns. *Review of Financial Studies* 13:433–451.

- Jackwerth, J. C., and M. Rubinstein. 1996. Recovering probability distributions from option prices. *Journal of Finance* 51:1611–1631.
- Jacobs, K., and K. Q. Wang. 2004. Idiosyncratic consumption risk and the cross section of asset returns. *Journal of Finance* 59:2211–2252.
- Jarrow, R., and P. Protter. 2004. Structural versus reduced form models: a new information based perspective. *Journal of Investment Management* 2:1–10.
- Jensen, C. S., D. Lando, and L. H. Pedersen. 2016. Generalized recovery. *Working paper*.
- Joe, H. 2014. *Dependence modeling with copulas*. CRC Press.
- Johannes, M. S., N. G. Polson, and J. R. Stroud. 2009. Optimal filtering of jump diffusions: Extracting latent states from asset prices. *Review of Financial Studies* 22:2759–2799.
- Jones, C. S. 1998. Bayesian estimation of continuous-time finance models. *Working paper*.
- Julier, S., and J. Uhlmann. 1997. A new extension of the Kalman Filter to nonlinear systems. In *SPIE proceedings series*, pp. 182–193. Society of Photo-Optical Instrumentation Engineers.
- Kalman, R. E. 1960. A new approach to linear filtering and prediction problems. *Journal of Basic Engineering* 82:35–45.
- Kendall, M. G., and A. Stuart. 1977. *The advanced theory of statistics*. Charles Griffin & Co.
- Kim, C.-J. 1994. Dynamic linear models with Markov-switching. *Journal of Econometrics* 60:1–22.
- Kraus, A., and R. H. Litzenberger. 1976. Skewness preference and the valuation of risk assets. *Journal of Finance* 31:1085–1100.
- Kwon, T. Y., and Y. Lee. 2015. Estimating structural credit risk models when market prices are contaminated with noise. *Applied Stochastic Models in Business and Industry* 32:18–32.
- Kyle, A. S. 1985. Continuous auctions and insider trading. *Econometrica* 53:1315–1336.

- Li, D. X. 2000. On default correlation: A copula function approach. *Journal of Fixed Income* 9:43–54.
- Lintner, J. 1965a. Security prices and risk: The theory and comparative analysis of A.T.&T. and leading industrials. *Presented at the conference on "The Economics of Regulated Public Utilities" at the University of Chicago Business School.*
- Lintner, J. 1965b. The valuation of risk assets and the selection of risky investments in stock portfolios and capital budgets. *Review of Economics and Statistics* 47:13–37.
- Liu, J., J. Pan, and T. Wang. 2005. An equilibrium model of rare-event premia and its implication for option smirks. *Review of Financial Studies* 18:131–164.
- Longstaff, F. A., S. Mithal, and E. Neis. 2005. Corporate yield spreads: default risk or liquidity? New evidence from the credit default swap market. *Journal of Finance* 60:2213–2253.
- Maalaoui Chun, O., G. Dionne, and P. Francois. 2014. Detecting regime shifts in credit spreads. *Journal of Financial and Quantitative Analysis* 49:1339–1364.
- Madan, D., and H. Unal. 2000. A two-factor hazard rate model for pricing risky debt and the term structure of credit spreads. *Journal of Financial and Quantitative Analysis* 35:43–66.
- Mahanti, S., A. Nashikkar, M. Subrahmanyam, G. Chacko, and G. Mallik. 2008. Latent liquidity: a new measure of liquidity, with an application to corporate bonds. *Journal of Financial Economics* 88:272–298.
- Malik, S., and M. K. Pitt. 2011. Particle filters for continuous likelihood evaluation and maximisation. *Journal of Econometrics* 165:190–209.
- Markose, S., S. Giansante, and A. R. Shaghaghi. 2012. Too interconnected to fail financial network of US CDS market: Topological fragility and systemic risk. *Journal of Economic Behavior and Organization* 83:627–646.
- Markowitz, H. 1952. Portfolio selection. *Journal of Finance* 7:77–91.
- Martin, I. 2016. What is the expected return on the market? *Working paper.*
- Martin, I., and C. Wagner. 2016. What is the expected return on a stock? *Working paper.*
- Meine, C., H. Supper, and G. N. Weiß. 2016. Is tail risk priced in credit default swap premia? *Review of Finance* 20:287–336.

- Merton, R. 1973. An intertemporal capital asset pricing model. *Econometrica* 41:867–887.
- Merton, R. C. 1974. On the pricing of corporate debt: the risk structure of interest rates. *Journal of Finance* 29:449–470.
- Merton, R. C. 1980. On estimating the expected return on the market: An exploratory investigation. *Journal of Financial Economics* 8:323–361.
- Milne, A. 2014. Distance to default and the financial crisis. *Journal of Financial Stability* 12:26–36.
- Moody's. 2009. Corporate Default and Recovery Rates, 1920-2008. Tech. rep., Moody's. URL <http://v2.moodys.com/cust/content/content.ashx?source=StaticContent/Free%20Pages/Credit%20Policy%20Research/documents/current/20074000000578875.pdf>.
- Mueller, P. 2008. Credit spreads and real activity. *Working paper*.
- NAIC. 2015. Update on the insurance industry's use of derivatives and exposure trends. Tech. rep., National Association of Insurance Commissioners. URL http://www.naic.org/capital_markets_archive/150807.htm.
- Newey, W., and K. West. 1987. A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix. *Econometrica* 55:703–708.
- Nier, E., J. Yang, T. Yorulmazer, and A. Alentorn. 2007. Network models and financial stability. *Journal of Economic Dynamics and Control* 31:2033–2060.
- Ornthanalai, C. 2014. Levy jump risk: Evidence from options and returns. *Journal of Financial Economics* 112:69–90.
- Pan, J. 2002. The jump-risk premia implicit in options: Evidence from an integrated time-series study. *Journal of Financial Economics* 63:3–50.
- Protter, P. 2004. *Stochastic integration and differential equations*. Springer.
- Qi, H., and D. Sun. 2006. A quadratically convergent Newton method for computing the nearest correlation matrix. *SIAM Journal on Matrix Analysis and Applications* 28:360–385.
- Qiu, J., and F. Yu. 2012. Endogenous liquidity in credit derivatives. *Journal of Financial Economics* 103:611–631.

- Renault, E. 1997. Econometric models of option pricing errors. *Econometric Society Monographs* 28:223–278.
- Roll, R. 1984. A simple implicit measure of the effective bid-ask spread in an efficient market. *Journal of Finance* 39:1127–1139.
- Ross, S. 1976. The arbitrage theory of capital asset pricing. *Journal of Economic Theory* 13:341–360.
- Ross, S. 2015. The recovery theorem. *Journal of Finance* 70:615–648.
- Rubinstein, M. 1975. The strong case for the generalized logarithmic utility model of financial markets. *Journal of Finance* 31:551–571.
- Saldías, M. 2013. Systemic risk analysis using forward-looking distance-to-default series. *Journal of Financial Stability* 9:498–517.
- Santa-Clara, P., and S. Yan. 2010. Crashes, volatility, and the equity premium: lessons from S&P 500 options. *Review of Economics and Statistics* 92:435–451.
- Saunders, A., and L. Allen. 2010. *Credit risk management in and out of the financial crisis: New approaches to value at risk and other paradigms*. John Wiley & Sons.
- Schönbucher, P. J. 2002. A tree implementation of a credit spread model for credit derivatives. *Journal of Computational Finance* 6:1–38.
- Sharpe, W. F. 1964. Capital asset prices: A theory of market equilibrium under conditions of risk. *Journal of Finance* 19:425–442.
- Stambaugh, R. F., J. Yu, and Y. Yuan. 2015. Arbitrage asymmetry and the idiosyncratic volatility puzzle. *Journal of Finance* 70:1903–1948.
- Stephan, J. A., and R. E. Whaley. 1990. Intraday price change and trading volume relations in the stock and stock option markets. *Journal of Finance* 45:191–220.
- Tang, D. Y., and H. Yan. 2007. Liquidity and credit default swap spreads. *Working paper*.
- Todorov, V., and G. Tauchen. 2011. Volatility jumps. *Journal of Business & Economic Statistics* 29:356–371.
- Trolle, A., and E. Schwartz. 2009. Unspanned stochastic volatility and the pricing of commodity derivatives. *Review of Financial Studies* 22:4423–4461.

- Tse, S. T., and J. W. Wan. 2013. Low-bias simulation scheme for the Heston model by inverse Gaussian approximation. *Quantitative Finance* 13:919–937.
- Tugnait, J. K. 1982. Detection and estimation for abruptly changing systems. *Automatica* 18:607–615.
- van Haastrecht, A., and A. Pelsser. 2010. Efficient, almost exact simulation of the Heston stochastic volatility model. *International Journal of Theoretical and Applied Finance* 13:1–43.
- van Trees, H. L. 1968. *Detection, estimation and linear modulation theory*. John Wiley & Sons.
- Vasicek, O. 1977. An equilibrium characterization of the term structure. *Journal of Financial Economics* 5:177–188.
- Vazza, D., and E. Gunter. 2012. Recovery study (U.S.): Recoveries come into focus as the speculative-grade cycle turns negative. Tech. rep., Standard & Poor's Rating Services.
- Viterbi, A. J. 1967. Error bounds for convolutional codes and an asymptotically optimum decoding algorithm. *IEEE Transactions on Information Theory* 13:260–269.
- Wei, G. N., and J. Mhlnickel. 2014. Why do some insurers become systemically relevant? *Journal of Financial Stability* 13:95–117.
- Xiao, X., and C. Zhou. 2014. Systematic and idiosyncratic jump risks in the expected stock returns. *Working paper*.
- Yuen, F. L., and H. Yang. 2010. Option pricing with regime switching by trinomial tree method. *Journal of Computational and Applied Mathematics* 233:1821–1833.
- Zhang, B. Y., H. Zhou, and H. Zhu. 2009. Explaining credit default swap spreads with the equity volatility and jump risks of individual firms. *Review of Financial Studies* 22:5099–5131.
- Zhang, L., P. A. Mykland, and Y. Ait-Sahalia. 2005. A tale of two time scales. *Journal of the American Statistical Association* 100:1394–1411.

