



A University of Sussex PhD thesis

Available online via Sussex Research Online:

<http://sro.sussex.ac.uk/>

This thesis is protected by copyright which belongs to the author.

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

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

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

Please visit Sussex Research Online for more information and further details

**Model specification and news
announcements: theoretical and
empirical aspects of option pricing**

Nathaniel Wiesendanger Shaw

Submitted for the degree of Doctor of Philosophy

University of Sussex

February 2020

Declaration

I hereby declare that this thesis has not been and will not be submitted in whole or in part to another University for the award of any other degree.

The work which forms chapter two has been published in the Journal of Futures Markets (DOI: fut.22061) and is a direct copy of the paper which is joint work with my supervisor Prof. Xavier Calmet.

Furthermore, the ideas which makes up chapters three and four are a result from the joint discussions between myself and Prof. Andreas Kaeck. However, the undertaking of the analysis and writing are entirely my own.

Signature:

Nathaniel Wiesendanger Shaw

Acknowledgements

First and foremost I would like to thank Andreas Kaeck for his dedication and support all the way through this journey. It has been a pleasure to work alongside such a brilliant academic. I am further indebted to Xavier Calmet for his constant support and teachings throughout. I have enjoyed enormously working with such a distinguished academic as Xavier. I do not hesitate for a moment to dedicate this thesis to those who are most important to me, my partner Isha, my father and mother Peter and Sarah. I would also like to dedicate in part this thesis to my closest friends William and Yusufu for their endless encouragement and the strength which they have afforded me throughout.

UNIVERSITY OF SUSSEX

NATHANIEL MAURICE WIESENDANGER SHAW

DOCTOR OF PHILOSOPHY

MODEL SPECIFICATION AND NEWS ANNOUNCEMENTS:
THEORETICAL AND EMPIRICAL ASPECTS OF OPTION PRICING

SUMMARY

This thesis covers the effects of macroeconomic announcements and model specification on option pricing. In the first main chapter a novel design for an approximative solution to a non-affine stochastic volatility model for option pricing is derived, inspired by methods frequently used in physics. This methodology is then compared against numerical solutions and is found to perform very well. In the second main chapter we conduct an empirical study of high-frequency option pricing on S&P 500 options. The chapter has two objectives: to examine the equity price uncertainty surrounding information released in FOMC announcements. Additionally, we quantify this impact on formal option pricing models. Secondly, to investigate option market maker biases to volatility accrual and this biases ability to explain why option returns, on average, are negative overnight and positive intraday. We also investigate this effect on formal option pricing models. In the third main chapter we investigate equity option returns and their behaviour around earnings announcements. We document a robust novel finding that delta-hedged equity option returns increase with increasing earnings announcement variance of the underlying stock. This result is separate from existing anomalies in stock and options markets and cannot be explained by standard risk factors. The thoughts demonstrated in this thesis lay the foundations for further interesting research and practically relevant applications. Please note that the first main chapter has been publish in the journal of Futures Markets, DOI: fut.22061.

Contents

List of Tables	ix
List of Figures	x
1 Introduction	1
2 An Analytical Perturbative Solution to the Merton-Garman Model Using Symmetries	6
2.1 Introduction	6
2.2 The Merton-Garman Model	8
2.3 Reduction to the Symmetrical Model and Perturbative Solution to the Merton-Garman Model	10
2.3.1 Symmetrical Model	10
2.3.2 Solution of the Symmetrical Model	12
2.3.3 Symmetry Breaking Terms and Solution to the Merton-Garman Model	13
2.4 Comparison with Numerical Simulations	15
2.4.1 Static Cross-Section Comparison	15
2.4.2 Simulated Time Series Calibration	19
2.5 Model building in finance, symmetries and group theory	24
2.6 Conclusion	27
3 Macroeconomic Announcements and Volatility Seasonalities, a High-Frequency Approach to Option Pricing	29
3.1 Literature Review and Discussion	31
3.1.1 Volatility Seasonalities	33

3.1.2	Macroeconomic and Political Announcements	36
3.2	Incorporating FOMC Announcements in Index Price Models	42
3.2.1	Anticipated Uncertainty Estimators	44
3.2.2	Calibration Methodology	46
3.2.3	VIX Construction	49
3.2.4	Preliminary Data Analysis	51
3.2.5	Summary Statistics	59
3.2.6	Non-Parametric Tests	61
3.2.7	Characterising Anticipated FOMC Volatility	62
3.2.8	Option Pricing Implications	67
3.2.9	Summary	80
3.3	Volatility Seasonality Adjustments	80
3.3.1	Time Notation	82
3.3.2	Model Framework	83
3.3.3	Overnight Effect	87
3.3.4	Weekend Effect	91
3.3.5	Conclusion	95
4	Earnings Announcements and Expected Option Returns	96
4.1	Introduction	96
4.1.1	Literature Review	99
4.2	Data and Delta-Hedged Option Returns	103
4.2.1	Data	103
4.2.2	Delta-Hedged Option Returns	105
4.2.3	Earnings Surprises	106
4.3	Empirical Results	107
4.3.1	Summary Statistics	108
4.3.2	Fama-MacBeth Regressions	109
4.4	Earnings Announcement Option Trading Strategy	118
4.4.1	Time Series Regressions	118
4.4.2	Average Portfolio Returns	120

4.4.3	Controlling for Trading Costs	125
4.5	Robustness Checks	126
4.6	Discussion	132
4.7	Conclusion	134
5	Conclusion	136
	Bibliography	138
6	Appendix	153
6.1	Heat Equation	153
6.2	Solution of the Symmetrical Model	156
6.3	Symmetry Breaking Terms and Solution to the Merton-Garman Model . . .	158
6.4	Time Series Simulation Merton-Garman Model	163
6.5	Characteristic Function	164
6.6	Impact of Stochastic Volatility on FOMC Estimators	166
6.7	Pricing Method	170
6.8	Extracting Latent Volatility	175
6.9	Time Series Simulation Multi-Factor Models	178
6.10	Short Maturity Data Set Calibrations	181
6.11	Maturity Related Effects	183

List of Tables

2.1	Static Cross-Sectional Calibration IVRMSE	19
2.2	Simulated Time Series Parameter Estimates	23
2.3	Simulated Time Series Error Estimates	24
3.1	Number of Options	54
3.2	Summary Statistics of Returns on S&P 500 Index	57
3.3	Summary Statistics on VIX Estimates	59
3.4	Summary Statistics for FOMC and Non-FOMC Implied Volatilities	60
3.5	Wilcoxon Test	62
3.6	Individual and Average Uncertainty Estimates	66
3.7	Anticipated Uncertainty (Term Structure Estimator)	67
3.8	Anticipated Uncertainty (Time Series Estimator)	67
3.9	Average 1-hr Variability of Front Maturity ATM Implied Variance	68
3.10	Parameter Estimates (Full Maturity Data Set)	73
3.11	Individual FOMC Volatility Estimates	77
3.12	Average Pricing Errors Around FOMC Announcements	78
3.13	Overnight Effect Time Adjusted Parameter Estimates	90
3.14	Weekend Effect Time Adjusted Parameter Estimates	94
4.1	Summary Statistics of Individual Equity Options Data	108
4.2	Fama-Macbeth Regressions, Stock-Related Characteristics	111
4.3	Fama-Macbeth Regressions, Volatility-Related Characteristics	114
4.4	Fama-MacBeth Regressions, Limits to Arbitrage	116
4.5	Fama-Macbeth Regressions, Alternate Measures of Delta-Hedged Returns	117
4.6	Time Series Regressions	120

4.7	Average Monthly Returns of Option Trading Strategy	121
4.8	Average Monthly Returns, Double Portfolio Sorts	125
4.9	Controlling for Trading Costs	127
4.10	Controlling for Trading Costs, Alternate Weighting Schemes	129
4.11	Time Series Regressions, Held Until Maturity	131
4.12	Controlling for Trading Costs, Held Until Maturity	132
4.13	Investor Demand Sorted By earnings announcement variance	134
6.1	Term Structure Slope Estimates	170
6.2	Static VWRMSE of Monte Carlo Against Numerical ODE Solution	180
6.3	Simulated Calibration Exercise	181
6.4	Parameter Estimates (Reduced Maturity Data Sets)	182

List of Figures

2.1	Static Price and Implied Volatility Fits	18
2.2	Pricing Corrections of First Order Approximation	20
3.1	Returns of S&P 500 index	55
3.2	VIX Time Series (30 and 360 Day Horizon)	58
3.3	Time Series of Pricing Errors Around FOMC Announcements	79
4.1	Time Series of Returns	122
4.2	Time Series of Returns With Effective Spread	128
6.1	Trading Ratio VWRMSE	184

Acronyms

ATM at-the-money.

CBOE Chicago Board of Options Exchange.

DTM Days to maturity.

EAD earnings announcement date.

FOMC federal open market committee.

ITM in-the-money.

IV implied variance.

IVRMSE Implied Volatility Root Mean Squared Error.

MC Monte Carlo.

ODE ordinary differential equation.

OMM option market maker.

OTM out-of-the-money.

RMSE Root Mean Squared Error.

S&P 500 Standard & Poor's 500 Stock Market Index.

SDE stochastic differential equation.

SV stochastic volatility.

SVD stochastic volatility with FOMC jumps.

SVJ stochastic volatility with price jumps.

SVJD stochastic volatility with price jumps and FOMC jumps.

SV2 two-factor stochastic volatility.

SV2D two-factor stochastic volatility with FOMC jumps.

SVCJ2 two-factor stochastic volatility with contemporaneous jumps.

SVCJ2D two-factor stochastic volatility with contemporaneous jumps and FOMC jumps.

VIX CBOE Volatility Index.

VWRMSE Vega Weighted Root Mean Squared Error.

Chapter 1

Introduction

In the three main chapters of this thesis we discuss a range of aspects relevant to option pricing. Since the seminal work of Black-Scholes (1973), empirical option pricing and model specification has received a substantial amount of attention in the literature. The model of Black-Scholes (1973) was the first to give mathematical rigour to option pricing.¹ The model gives a theoretical estimate of the price of European-style options, which is unique regardless of the expected return of the underlying asset. While the Black-Scholes (1973) model has been incredibly successful, it is not without its drawbacks. For example, the main problem of the Black-Scholes model is the assumption of constant volatility regardless of moneyness (defined as the ratio of strike price to underlying asset price) or time to maturity (defined as the amount of time left until the option expires, measured in years) of the option. This assumption is strongly at odds with empirical findings. For instance Macbeth and Merville (1979) find implied volatilities are not flat and often exhibit a "volatility smile" or "smirk" pattern. Where implied volatility is defined as the volatility value of the underlying asset which when supplied to an option pricing model returns a theoretical value equal to the current market value of the option contract. For the purposes of this thesis the model in question is always the Black-Scholes (1973) model, however, in theory other models would be valid. The empirical shortcomings of the Black-Scholes model prompted research along

¹However, Bronzin was a precursor to Black-Scholes, publishing his "Theory of Premium Contracts" work in 1908, containing many elements of modern option pricing. However, his work was quickly forgotten and less well-known.

dimensions relaxing the restrictive assumptions. Consequently, a very substantial body of literature exists utilising stochastic volatility models, which allow the instantaneous volatility of asset returns to evolve stochastically, with respect to time. An example of such a stochastic volatility model is the Heston (1993) model:

$$dS_t = \mu S_t dt + \sqrt{V_t} S_t dW_t^S, \quad (1.1)$$

$$dV_t = \kappa(\theta - V_t)dt + \xi\sqrt{V_t}dW_t^V, \quad (1.2)$$

where S_t denotes the price of the underlying asset at time t , with rate of return μ . V_t is the instantaneous latent variance process of S_t . The variance process possesses the parameters κ, θ, ξ with the following definitions: θ is the long run variance level, where κ represents the rate at which V_t tends to θ . ξ represents the volatility of volatility and determines the variance of V_t . While these stochastic volatility models are very popular and certainly provide much empirical improvement against the Black-Scholes (1973) model there are many different potential model specifications, making this still an actively researched area. A popular choice is to model the volatility using a diffusion process, such as in Merton (1976). This leads to the popular affine square-root volatility model of Heston (1993). However, there is a substantial strand of literature which finds that the popular square-root stochastic volatility model is misspecified (see e.g., Eraker et al. (2003), Duan and Yeh (2010) to name a few). Model misspecification is defined as models which are specified in error, as any model is only an approximation to reality. Misspecified models are known to manifest biased parameter estimates. While using a non-affine framework is more realistic to model the volatility of returns, there is a drawback of these models which is a general lack of closed form characteristic function, which makes pricing much more challenging. It is this aspect of option pricing model specification which we are concerned with in the first main chapter of this thesis. Specifically we are concerned with attempting to overcome the problems of pricing such models in a timely and accurate fashion.

Continuing the theme of affine volatility models additional research has shown there to be substantial empirical evidence to support large magnitude jump-like movements in returns, for example the 1987 crash. Bates (1991) (initially) uses random Poisson jumps in

the price process to model such events. This leads to the jump-diffusion class of models, for example Merton (1976). Further research in jump-diffusion models suggests the volatility of returns can increase suddenly. Without jumps, the volatility of returns, is driven by a Brownian motion, thus it can only increase gradually via a sequence of normally distributed increments. Evidence from Bates (2000) and Duffie et al. (2000) has shown empirically that the conditional volatility of returns can increase rapidly. Eraker et al. (2003) demonstrate that jumps in volatility allow the volatility of returns to increase suddenly.

While single factor jump-diffusion models are very popular in the literature there is strong empirical evidence to suggest a multi-factor volatility framework is needed, with one highly persistent and one quickly mean-reverting factor. For example, Christoffersen et al. (2009) find using a two-factor framework improves the ability to model the slope and level of the volatility smile, which move largely independently. As such their findings demonstrate significant improvement over a benchmark Heston model. It is this affine two-factor jump-diffusion model type that we primarily investigate for the second main chapter of this thesis. We conduct extensive empirical exercises to determine which model processes and features are relevant to model high-frequency options data on the S&P 500. This links to the first main chapter in part by analysing model specification for option pricing, this time from an affine volatility framework.

There is also a further branch of the literature which investigates systematic patterns in stock returns and or return volatilities. For example, French and Roll (1986) find stock volatility is higher during trading hours mainly due to private information accrued during periods of non-trading. More recently, Kaplanski and Levy (2015) investigate a number of volatility seasonalities, showing weekend, holiday and overnight trading breaks generate excessive perceived risk in options markets. A part of the second main chapter deals with capturing such volatility seasonalities and incorporating them into traditional option pricing models, further linking to the first main chapter by building upon improving option pricing model specifications.

The second main chapter further links to the first by investigating option pricing model specification in the face of macroeconomic announcements. The notion of uncertainty as a

fundamental driving force for asset prices is intuitive and has been extensively documented. For example, Kelly et al. (2016) study international equity option quotes and demonstrate that political uncertainty from national elections and global summits is priced by investors with the result that options spanning important events are on average more expensive.

The third main chapter of the thesis links to the second by investigating the effect of news announcements on option pricing. The third main chapter investigates the effects of earnings announcements on individual stock options and constructs a robust trading strategy on earnings announcement variance. The following paragraphs provide a more detailed overview of each chapter.

The second chapter investigates option pricing model specifications under a non-affine volatility framework. To quote Chourdakis and Dotsis (2011): "Does analytical tractability come at the cost of empirical misspecification?", we propose to use methods common in physics to find an analytical approximation to a general one-factor non-affine stochastic volatility model. We find our method performs well in simulated time series calibration exercises and provides very fast computation of option prices.

The third chapter investigates novel empirical option pricing model specifications under high-frequency, intradaily options data on the S&P 500. We wish to determine which factors and processes are relevant in this new age of data. We decompose this research question into three sub-questions, each building on a respective area of the literature. Firstly, we wish to determine if standard factors and processes, which have been shown in prior studies using lower-frequency daily or weekly data (such as Christoffersen et al. (2009), Eraker et al. (2003), Duffie et al. (2000)), are still relevant. We find empirical evidence to confirm, at high-frequencies, multiple variance factors and jumps in the price and variance process are important. Secondly, we are interested in incorporating model processes to capture option market maker biases towards volatility accrual during different periods of the week. Our work builds on that of the overnight effect by Muravyev and Ni (2019) and also on the weekend effect by Jones and Shemesh (2018). We support the conclusion of Muravyev and Ni (2019) that market makers ignore the day-night volatility seasonality in options and find our proposed model specification to capture this effect reduces out-of-sample errors by approxi-

mately 7%. Thirdly, we wish to understand the effect of macroeconomic announcements on option prices. We investigate the Federal Open Market Committee (FOMC) announcement, utilising the high-frequency data it is possible to observe the effect of the announcement on options in minute detail. We find evidence to suggest that FOMC announcements contain important information for option pricing.

The fourth chapter follows empirical studies by Cao and Han (2013) and Cao et al. (2005) who analyse delta-hedged option returns on individual stock options, specifically we are interested in the effect of the idiosyncratic volatility component which earnings announcements provide. We define idiosyncratic volatility as the component of total volatility of the assets returns that cannot be explained by market returns. While Cao and Han (2013) consider total idiosyncratic volatility, finding a negative relationship between delta-hedged returns and idiosyncratic volatility we construct an earnings announcement variance measure which is positively related to delta-hedged option returns. We form an option trading strategy based on sorting firms by their earnings announcement variance. Our trading strategy is found to be robust against standard risk factors and the idiosyncratic volatility trading strategy of Cao and Han (2013). We also document our strategy is robust to trading costs.

Chapter five concludes by summarising the main results. The Appendix contains a range of tools that are used throughout this thesis: Appendices 6.1 - 6.4 provides mathematical details needed for the derivation of our solution to the non-affine model considered in Chapter two and details of the simulated time series calibration exercise. Appendices 6.5 - 6.11 discuss the pricing of the model specification for the third Chapter and also provide checks as to our methodology at various points.

Chapter 2

An Analytical Perturbative Solution to the Merton-Garman Model Using Symmetries

2.1 Introduction

Calculating the price of an option is an important challenge in mathematical finance. The first attempts in that direction are attributed to Louis Bachelier who in his Doctoral thesis, *Théorie de la spéculation*, published in 1900, considered a mathematical model of Brownian motion and its use for valuing options. This work provided the foundations for the Black-Scholes model (Black, Scholes (1973)). However, while the Black-Scholes model was a breakthrough in the field, it is widely accepted that it has limitations. In particular, the volatility is treated as a constant which is not very realistic.

Since the seminal works of Black, Scholes (1973) and Merton (1973), more sophisticated models with a time dependent volatility have been proposed. For example, the affine Heston model (See Heston (1993)), which assumes a time-dependent volatility, with a stochastic process involving the square-root of the stochastic volatility, and a leverage effect, has been implemented in a large number of empirical studies (Andersen et al. (2002), Bakshi et al.

(1997), Bates (2000), Bates (2006), Chernov et al. (2003), Huang and Wu (2004), Pan (2002), Eraker (2004) to name a few). Such models have however limitations and are often modified artificially by combining them with models of jumps in returns and/or in volatility (such as in Jones (2003) and Benzoni (2002)). As a consequence, there is a substantial strand of literature devoted to non-affine volatility models, which note that the popular square-root stochastic volatility model is not very realistic (see e.g., (Eraker et al. (2003) Duan and Yeh (2010), Aït-Sahalia, Kimmel (2007) Christoffersen et al. (2010), Chourdakis and Dotsis (2011) and Kaeck and Alexander (2012)) to name a few). However, the issue with such models is a general lack of closed form characteristic function, which makes pricing much more challenging. As stated in Chourdakis and Dotsis (2011) when regarding the place of non-affine models and the debate of their tractability against affine models: “does analytically tractability come at the cost of empirical misspecification?”. It is a useful endeavor to study non-affine model as we propose in this paper, if an analytical solution for the option pricing formula can be found.

A well-known example of such models is the Merton-Garman model (Garman (1976) and Merton (1973)) which is indeed a more realistic model as it allows for a time-dependent volatility and it is not restricted to an affine model for the volatility. However, solving non-affine models is time consuming, as it involves numerical methods. Thus, many practitioners are still using the Black-Scholes formula to obtain a fast, albeit not necessarily very reliable, price quote for an option.

The aim of our work is twofold. We will derive an analytical approximative solution to the partial differential equation describing the Merton-Garman model which enables fast calculations of option prices. This requires us to identify a “symmetric” version of the model which can easily be solved analytically. One can then reintroduce the symmetry breaking terms of the original Merton-Garman differential equation and do perturbation theory around the symmetric solution thereby obtaining an approximative but analytical solution to the original Merton-Garman differential equation. We then propose a new approach to model building in option pricing based on the concept of symmetry groups and representation theory. This concept has been extremely successful in modern physics. It is at the origin

of all successful models in physics, e.g., in particle physics, cosmology or solid state physics. We note that perturbation theory has been used in option pricing models (Baaquie (1997), Baaquie et al. (2003), Blazhyevskiy and Yanishevsky (2011), Aguilar (2017), Kleinert and Korbel (2016), Utama and Purqon (2016)) but here we organise perturbation theory around a very specific solution, namely that of the symmetrical model which we will introduce in this paper.

This paper is organised as follows. In section 2, we derive the partial differential equation which describes the Merton-Garman model. In section 3, we explain how to reduce the original Merton-Garman to a simple, symmetrical, model. We present an exact analytical solution to the symmetrical model. We then restore the original Merton-Garman by re-introducing the symmetry breaking terms and provide an analytical perturbative solution to the Merton-Garman model. In section 4, we compare our solutions to different numerical solutions found in the literature. In section 5, we propose a new approach to model building in mathematical finance. Finally, we conclude in section 6.

2.2 The Merton-Garman Model

In the Merton-Garman model, the price of an option is dependent on the time t , the price of the underlying S and the volatility V . Both S and V are taken to be time-dependent functions and thus the Merton-Garman model has the potential to provide a more accurate calculation of an option price than e.g., the Black-Scholes model.

We start from the stochastic differential equations for the price of the underlying S and for the volatility V , under the risk-neutral measure

$$dS = rSdt + \sqrt{V}SdW^S, \tag{2.1}$$

$$dV = \kappa(\theta - V)dt + \xi V^\alpha dW^V, \tag{2.2}$$

which is a stochastic, mean reverting, volatility regime. Here, ξ is the standard deviation of

the volatility and κ is the speed of mean reversion to the long run variance θ . The interest rate r is assumed to be constant. The model described by Equations (2.1) and (2.2) covers many well-known stochastic volatility models, for instance setting $\alpha = 1, 1/2$ recovers the Hull and White (see Hull and White (1987)) and Heston models, respectively. However, we do not constrain ourselves to either of these worlds. Here, α can take arbitrary values. We will denote the correlation between the two Brownian motions W^S and W^V by ρ .

We shall first consider a call option, but our results can be extended to a put option in a straightforward manner. Our first step is to find the associated partial differential equation which describes this model. We do so by applying the Feynman-Kac formula, see e.g Hull (1997), which states that for the price of a call option, as defined by the model dynamics in Equations (2.1) and (2.2) is given by:

$$\frac{\partial C}{\partial t} + \sum_{i=1} \mu_i(t, x) \frac{\partial C}{\partial x_i} + \frac{1}{2} \sum_{i,j=1} \rho_{ij} \sigma_i(t, x) \sigma_j(t, x) \frac{\partial^2 C}{\partial x_i \partial x_j} - rC = 0, \quad (2.3)$$

where C is the price of a call. Using this formula, one obtains (see Srikant (1998))

$$\frac{\partial C}{\partial t} + rS \frac{\partial C}{\partial S} + \frac{1}{2} V S^2 \frac{\partial^2 C}{\partial S^2} + (\lambda + \mu V) \frac{\partial C}{\partial V} + \rho \xi V^{1/2+\alpha} S \frac{\partial^2 C}{\partial S \partial V} + \xi^2 V^{2\alpha} \frac{\partial^2 C}{\partial V^2} - rC = 0, \quad (2.4)$$

where $\lambda = \kappa\theta$, $\mu = -\kappa$. The call price $C = C(S, V, t)$ depends on the time t , the price of the underlying S and the time dependent volatility $V = V(t)$. In this model, there are four free parameters ρ , λ , μ and α . As we explained previously, existing solutions to this partial differential equation are numerical ones which have been obtained using Monte Carlo methods (except for the case of $\rho = 0$, where an analytical solution exists, see Srikant (1998)). Note that the put price $P = P(S, V, t)$ fulfills the same differential equation, but it is subject to a different boundary condition.

2.3 Reduction to the Symmetrical Model and Perturbative Solution to the Merton-Garman Model

By studying the partial differential equation given in Equation (2.4), it quickly becomes clear that the difficulty in finding an analytical solution to Equation (2.4) is due to the lack of symmetry between the different terms of the partial differential equation. It is useful to study the dimensions of the different terms and constants in this partial differential equation. The price of the call is obviously given in a specific currency which we shall take to be the USD or \$. The remaining dimensions follow from this. We have:

- $[C] = \$$,
- $[\frac{\partial C}{\partial t}] = \$/\text{time}$,
- $[rS\frac{\partial C}{\partial S}] = [r]\$$ thus $[r] = 1/\text{time}$,
- $[\frac{1}{2}VS^2\frac{\partial^2 C}{\partial S^2}] = [V]\$$ thus $[V] = 1/\text{time}$,
- $[(\lambda + \mu V)\frac{\partial C}{\partial V}] = ([\lambda] + [\mu]1/\text{time})$ thus $[\lambda] = 1/\text{time}^2$ and $[\mu] = \frac{1}{\text{time}}$,
- $[\rho\xi V^{1/2+\alpha}S\frac{\partial^2 C}{\partial S\partial V}] = [\rho][\xi](1/\text{time})^{1/2+\alpha}\$ \text{ time} = \$/\text{time}$ thus $[\rho][\xi] = \text{time}^{\alpha-3/2}$,
- $[\xi^2 V^{2\alpha}\frac{\partial^2 C}{\partial V^2}] = [\xi]^2(1/\text{time})^{2\alpha-2}\$ = \$/\text{time}$ thus $[\xi] = \text{time}^{\alpha-3/2}$ and ρ is dimensionless.

It is instructive to see that S and V have different dimensions. Nevertheless, our goal is to treat S and V as symmetrically as possible to make a global Galilean invariance in 2+1 manifest (see Section 2.5). This can be achieved by adequate variable transformations and by identifying the terms in the differential equation that violate this symmetry.

2.3.1 Symmetrical Model

Our aim is to derive a differential equation that is symmetrical in S and V . With this aim in mind, let us introduce an averaged volatility σ^2 which is constant. As in the case of the

Black-Scholes model, different definitions for the averaged volatility are possible, the specific choice will not impact our methodology and results.

Inspecting the differential equation (2.4), it is clear that we need to pick $\alpha = 1$ (corresponding to the Hull and White (1987) model) to emphasize the symmetry between S and V . We thus consider

$$\frac{\partial C}{\partial t} + rS \frac{\partial C}{\partial S} + \frac{1}{2}\sigma^2 S^2 \frac{\partial^2 C}{\partial S^2} + \mu V \frac{\partial C}{\partial V} + \rho \xi_0 V^{3/2} S \frac{\partial^2 C}{\partial S \partial V} + \xi_0^2 V^2 \frac{\partial^2 C}{\partial V^2} = rC. \quad (2.5)$$

We need to keep in mind that we will need to reintroduce $\frac{1}{2}V S^2 \frac{\partial^2 C}{\partial S^2}$, $\lambda \frac{\partial C}{\partial V}$ and the terms corresponding to deviations from 1 for α . Note that ξ_0 is different from ξ , in particular they do not have the same dimensions. Finally, we see that there is a mixed derivative term which needs to be eliminated. We thus set $\rho = 0$ and we will reintroduce this term as symmetry breaking term. We thus end up with:

$$\frac{\partial C}{\partial t} + rS \frac{\partial C}{\partial S} + \frac{1}{2}\sigma^2 S^2 \frac{\partial^2 C}{\partial S^2} + \mu V \frac{\partial C}{\partial V} + \xi_0^2 V^2 \frac{\partial^2 C}{\partial V^2} = rC. \quad (2.6)$$

This partial differential equation can be massaged with standard substitutions into a 2+1 dimensional heat equation (see Appendix 6.1) in which case the symmetry in S and V becomes manifest. In order to do so, we introduce

$$x = \log(S/K), \quad (2.7)$$

$$y = \log(V/V_0), \quad (2.8)$$

and

$$C(x, y, \tau) = K \phi(x, y, \tau) \psi_0(x, y, \tau), \quad (2.9)$$

where K is the strike price and V_0 is some constant with units of 1/sec. The function $\phi(x, y, \tau)$ and the rescaled time τ are defined in Appendix 6.1. Standard manipulations described in Appendix 6.1 lead to

$$\frac{\partial \psi_0}{\partial \tau} = \frac{\partial^2 \psi_0}{\partial x^2} + \frac{\partial^2 \psi_0}{\partial y^2}. \quad (2.10)$$

which is manifestly symmetrical in x and y . We will thus refer to the model described by the differential equation (2.6) as the symmetrical model. Another reason for massaging the symmetrical model into a heat equation is that this equation is easy to solve analytically.

We impose the standard boundary condition for the call price:

$$C(S, V, T) = \left(S(T) - K \right)^+. \quad (2.11)$$

For a put option we have

$$P(S, V, T) = \left(K - S(T) \right)^+. \quad (2.12)$$

2.3.2 Solution of the Symmetrical Model

Details of the derivation of the analytical solution of the symmetrical model, i.e., of the 2+1 dimensional heat equations, are given in Appendix 6.2. We find

$$C_0(S, V, t) = S\mathcal{N}(d_1) - Ke^{-r(T-t)}\mathcal{N}(d_2), \quad (2.13)$$

where

$$\mathcal{N}(d) = \frac{1}{\sqrt{2\pi}} \int_{-\infty}^d \exp\left(-\frac{z^2}{2}\right) dz, \quad (2.14)$$

and

$$d_1 = \frac{x}{\sqrt{2\tau}} + \frac{\sqrt{2\tau}}{2}(R_1 + 1) = \frac{\log(S/K) + (r + \sigma^2/2)(T - t)}{\sigma\sqrt{T - t}}, \quad (2.15)$$

$$d_2 = \frac{x}{\sqrt{2\tau}} + \frac{\sqrt{2\tau}}{2}(R_1 - 1) = d_1 - \sigma\sqrt{T - t}. \quad (2.16)$$

Remarkably, because of the boundary condition that only depends on S , it is identical to the Black-Scholes solution. For a put option, the very same procedure leads to

$$P_0(S, V, t) = C_0(S, V, t) - S + Ke^{-r(T-t)}, \quad (2.17)$$

equally we could have applied put-call parity. In the next subsection, we shall restore the symmetry breaking terms and discuss the full Merton-Garman model.

2.3.3 Symmetry Breaking Terms and Solution to the Merton-Garman Model

We are now in a position to solve the full Merton-Garman model using perturbation theory around the symmetrical solution $C_0(S, V, t)$. We organise perturbation theory as an expansion in terms the coefficients of the symmetry breaking terms. We first need to restore the full model by re-introducing the symmetry breaking terms

$$\begin{aligned} \frac{\partial C}{\partial t} + rS \frac{\partial C}{\partial S} + \frac{1}{2}\sigma^2 S^2 \frac{\partial^2 C}{\partial S^2} + \frac{c_1 S^2}{2} \left(V - \sigma^2 \right) \frac{\partial^2 C}{\partial S^2} + \mu V \frac{\partial C}{\partial V} + c_2 \lambda \frac{\partial C}{\partial V} + \xi_0^2 V^2 \frac{\partial^2 C}{\partial V^2} \\ + c_3 \left(\xi^2 V^{2\alpha} - \xi_0^2 V^2 \right) \frac{\partial^2 C}{\partial V^2} + c_4 \rho \xi V^{\alpha+1/2} S \frac{\partial^2 C}{\partial S \partial V} - rC = 0. \end{aligned} \quad (2.18)$$

Note that we have introduced dimensionless coefficients c_i which denote the strength of the symmetry breaking terms. In the limit $c_i = 1$ one recovers the original Merton-Garman model. These coefficients are simply introduced as a bookkeeping trick to keep track of which terms correspond to a deviation of the 2+1 Galilean invariant theory. In the end of the day, we set $c_i = 1$. We now do perturbation theory around the symmetrical solution

$C_0(S, V, t)$ (see Appendix 6.3 for details) and obtain

$$\begin{aligned}
 C_1(S, V, t) = & -K \frac{\left(\frac{S}{K}\right)^{\frac{1}{2} - \frac{r}{\sigma^2}} e^{\left(\frac{4 \log^2\left(\frac{S}{K}\right) + (2r + \sigma^2)^2 (t-T)^2\right)}{8\sigma^2(t-T)}}}{4\sqrt{2\pi} \left(\frac{\sqrt{2}\gamma}{\sigma\xi_0} + 1\right) \sqrt{\sigma^2(T-t)}} \\
 & \times \left(\frac{1}{2} \sigma^4 \left(\frac{\sqrt{2}\gamma}{\sigma\xi_0} + 1 \right) (t-T) + V \left(e^{\frac{1}{2}\sigma^2 \left(\frac{\sqrt{2}\gamma}{\sigma\xi_0} + 1 \right) (T-t)} - 1 \right) \right), \tag{2.19}
 \end{aligned}$$

where we have set $c_1 = 1$. We expect that our approximation should work well when λ and ρ are small, when α is close to one and when the variation of V around its average value σ^2 is not too large. In the limit when V is large, σ^2 is large as well and we expect that, as in the Black-Scholes case, the price of the call becomes the price of the underlying S .

It may appear surprising that the leading order correction does not depend on the symmetry breaking terms parametrised by c_2 , c_3 and c_4 . It can easily be shown (see Appendix 6.3) that the boundary condition (6.13) insures that only the contribution from the c_1 term survives. The boundary condition implies that the contributions of c_2 , c_3 and c_4 vanish to leading order in the perturbation theory. These symmetry breaking terms will, however, contribute to higher order corrections. Higher precision, if required, can be obtained by going to higher order in perturbation theory. However, as demonstrated in Section 2.4 the error from going to first order only is sufficiently small. In a simulated calibration exercise the error, as measured by implied volatility RMSE, is of the order of $< 2\%$, in comparison to the Merton-Garman model (see Tables 2.2 and 2.3). This magnitude of error is well within expected option implied volatility bid-ask spreads. Option prices can be calculated extremely rapidly using this formalism. Note that, in principle, if we resummed perturbation theory to all order in c_i , the dependence on σ and ξ_0 would vanish. It is also worth noticing that our results are independent on V_0 which is only introduced to match the dimension of V .

It is straightforward to show that we obtain the same result for a put option

$$\begin{aligned}
P_1(S, V, t) = & -K \frac{\left(\frac{S}{K}\right)^{\frac{1}{2} - \frac{r}{\sigma^2}} e^{\left(\frac{4 \log^2\left(\frac{S}{K}\right) + (2r + \sigma^2)^2 (t-T)^2}{8\sigma^2(t-T)}\right)}}{4\sqrt{2\pi} \left(\frac{\sqrt{2}\gamma}{\sigma\xi_0} + 1\right) \sqrt{\sigma^2(T-t)}} \\
& \times \left(\frac{1}{2} \sigma^4 \left(\frac{\sqrt{2}\gamma}{\sigma\xi_0} + 1 \right) (t-T) + V \left(e^{\frac{1}{2} \sigma^2 \left(\frac{\sqrt{2}\gamma}{\sigma\xi_0} + 1 \right) (T-t)} - 1 \right) \right).
\end{aligned} \tag{2.20}$$

The prices obtained to leading order in perturbation theory for a call and put option are thus given by

$$C(S, V, t) = C_0(S, V, t) + C_1(S, V, t), \tag{2.21}$$

and

$$P(S, V, t) = P_0(S, V, t) + P_1(S, V, t). \tag{2.22}$$

In the next section we shall compare our results to solutions obtained numerically.

2.4 Comparison with Numerical Simulations

In this section we investigate how the approximative solution compares to a Monte Carlo simulation of the full Merton-Garman model. It is well-known that the Merton-Garman model is not solvable analytically, however it can be solved using numerical methods. The first step is a comparison of static cross sections of options. Then we compare using a simulated time series calibration exercise.

2.4.1 Static Cross-Section Comparison

The first step in evaluating the performance of the leading order perturbative solution is to compare to a multitude of simulated data of the Merton-Garman model to ensure that the approximation is sufficient to fit a range of different options at one time. We start by

describing the data simulation of the Merton-Garman model, then move onto the calibration procedure for the approximative solution and discuss the results.

We choose a standard Monte Carlo framework, using stratified sampling and antithetic variables, simulating seven million paths with a time step of one-tenth of a day for option maturities. We choose a spot underlying price of $S = \$100$ and a strike range of $K \in [90, 110]$ to give a moneyness range of $K/S \in [0.9, 1.1]$ to simulate call options throughout the spectrum of moneyness ¹. We simulate the Merton-Garman model with the structural parameter vector: $\Theta^{MG} = \{1.5, 0.08, 1.5, -0.5, 1\}$ ² and initial volatility $V(t = 0) \in [10, 35]\%$.

The leading order perturbative solution is independent on the symmetry breaking terms, characterized by c_2 , c_3 and c_4 . The solution is thus independent of the parameters ρ and θ . However, as a by-product of the perturbation theory we introduced the following parameters: ξ_0, σ . From inspection we fix ξ_0 using $\xi_0 = \xi \sigma^{2(\alpha-1)}$ which guarantees that it has the right dimensions. While σ remains to be determined by calibration. Yielding the parameter vector: $\Theta^{pert.} = \{\kappa, \xi, \alpha, \sigma\}$.

The parameter σ is determined by calibration. When fitting σ to these simulations, there is a risk of overfitting expensive out of the money (OTM) options. For that reason, it is best to consider the implied volatility objective function (this is noted in Christoffersen et al. (2014)) which is given by:

$$IVRMSE = \sqrt{\frac{1}{N} \sum_{i=1}^N (IV_{MC_i} - IV_{pert.i})^2}, \quad (2.23)$$

where IV_{MC_i} stands for the implied volatility of the i^{th} -option simulated using the Monte Carlo and $IV_{pert.i}$ is the implied volatility of the i^{th} -option calculated using the leading order perturbative solution. The calibration exercise is extremely fast as our formula for option prices is an approximative analytical solution. Figure (2.1) and Table 2.1 demonstrates the results of the static calibration exercise for a 30 day maturity horizon ³. It should be noted

¹In this exercise we simulate call options only, as put prices can be calculated from the put-call parity.

²We also simulate for $\alpha \in [0.75, 1.5]$ and $\kappa \in [1.5, 5]$, $\rho \in [-0.5, -0.9]$, $\theta \in [0.08, 0.15]$. However, the results of the calibration exercise are represented by the choice of parameters made above.

³While we simulated time horizons between 5-100 days, this is representative of our results and we drop the other results for brevity.

that in Figure (2.1) the price Panels are the difference of the log, this is needed to observe any difference in prices as the two methods produce prices which are very similar. However, from these Panels it is clear that smaller moneyness, i.e $K/S < 1$ see extremely small errors where $\log(C_{MC}/C_{pert.}) \sim 0$. Also apparent is that at some moneyness level the leading order perturbative solution will over price options, the level of moneyness at which this occurs is inversely proportional to volatility. Lastly, the range of under to over pricing is also inversely proportional to volatility. However, as the prices from both methods are very similar it is more informative to look at the implied volatility cross-section in the even Panels of Figure (2.1) along with the IVRMSE column of Table 2.1. These show throughout the range of volatility regimes the leading order perturbative solution is able to approximate successfully a low IVRMSE, with a maximum occurring from the low volatility regime (Panel 8) of 1.57%.

Figure 2.1: Static Price and Implied Volatility Fits

Comparative fit of four different initial volatilities. Odd numbered panels (1,3,5,7) display the natural logarithm difference of the prices, while even numbered panels (2,4,6,8) display the implied volatility curves. For the implied volatility panels the solid black line represents the Monte Carlo implied volatility and the dashed line is that of the leading order perturbative solution. We pick an option maturity of 30 days.

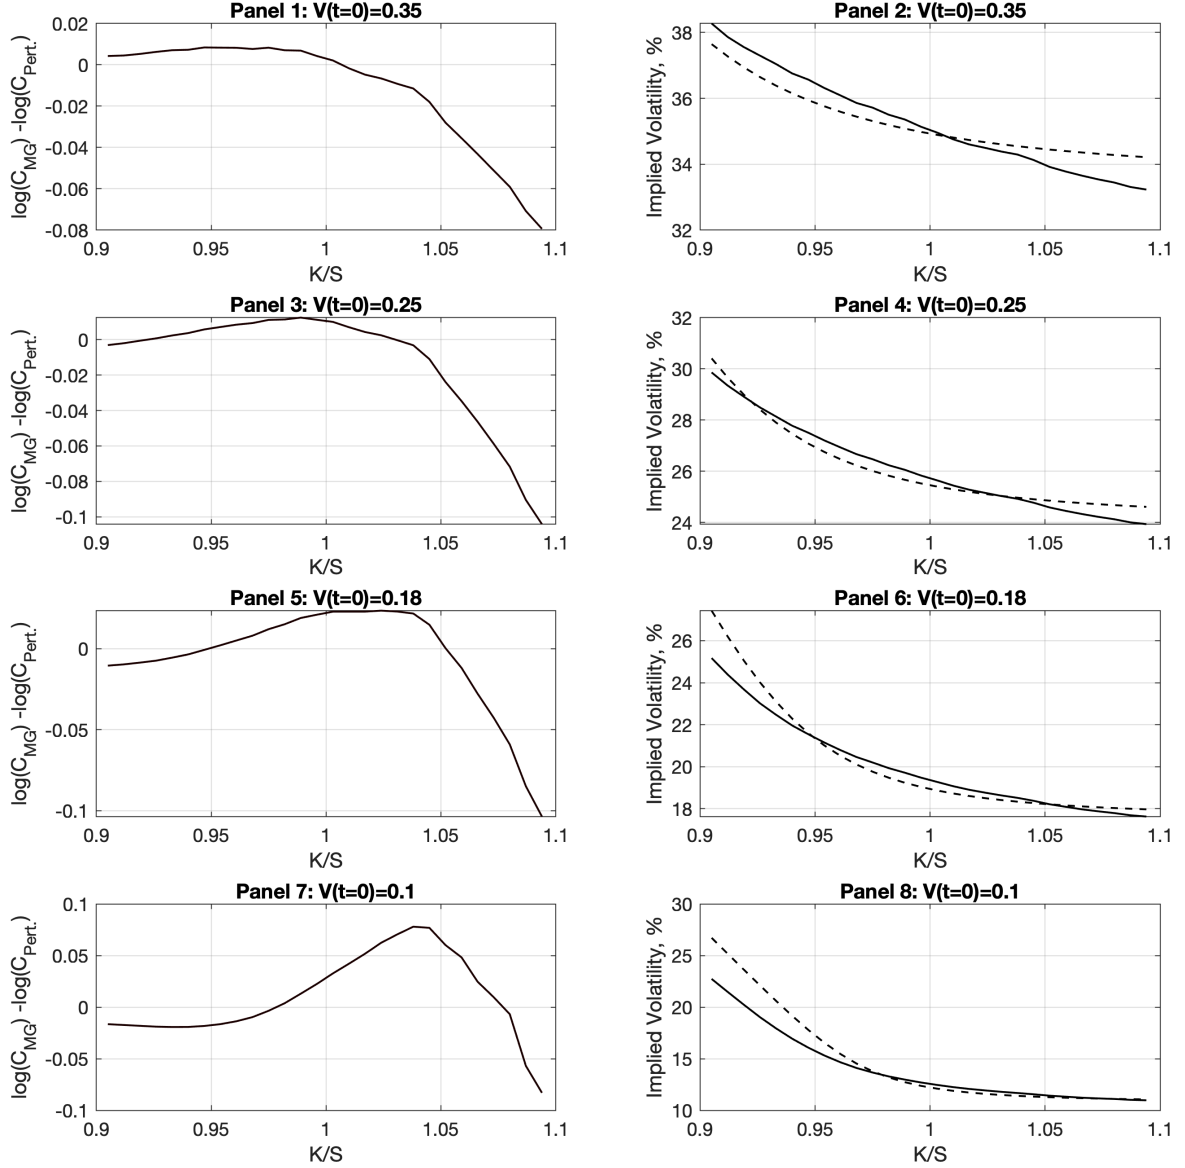


Table 2.1: Static cross-sectional calibration errors

This table reports the results of static cross-sectional calibration exercises between our leading order perturbative solution compared to a Monte Carlo simulation of the Merton-Garman model. We simulate four different initial volatilities (V), 10, 18, 25, 35%, for a time horizon of 30 days. The column denoted σ reports results from the calibration, using the IVRMSE objective function of Equation 2.23, as to the calibrated value of the average volatility parameter. The final column reports the IVRMSE of each simulation. All value are given in decimal percentages.

V	σ	IVRMSE
0.3500	0.3254	0.0055
0.2500	0.2361	0.0037
0.1800	0.1730	0.0070
0.1000	0.1069	0.0157

Another way to test the consistency of the perturbative expansion is to consider the ratio $C_1/(C_0 + C_1)$ as a function of moneyness, K/S , where C_0 is the contribution to the price of the symmetrical solution and C_1 is the leading order correction in perturbation theory. Figure (2.2) shows these ratios for several cases. Clearly $C_1 \ll C_0$ even when the volatility is large. This demonstrates nicely the validity of the perturbative expansion even in the case where the volatility is large. While this exercise confirms that for fixed scenarios the perturbative solution is a very good approximation to the actual Merton-Garman solution, it is essentially a multiple curve fitting exercise, a more thorough analysis is needed to be able to gauge the reliability of the perturbative approximation. This is what we shall focus on next.

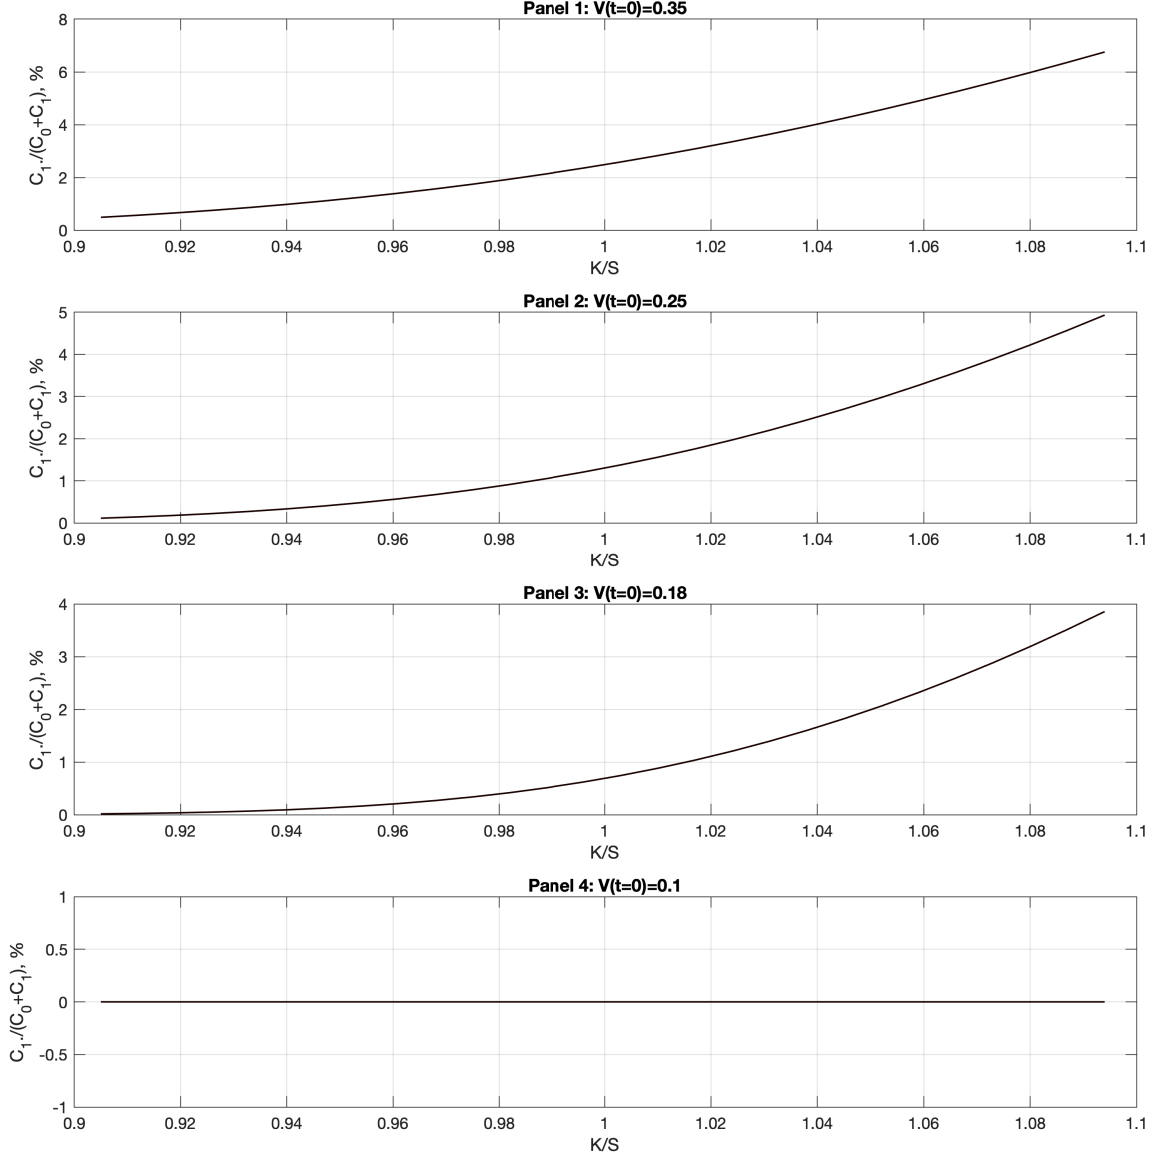
2.4.2 Simulated Time Series Calibration

The second step in evaluating the performance of the leading order perturbative solution is to estimate it against a time series simulation of the Merton-Garman model. The time series uses 100 different Monte Carlo paths to simulate the asset price and variance paths including 6 unique maturities within $[7, 180]$ days, for details see Appendix 6.4.

The benefits of the stress test is twofold: firstly it is particularly pertinent to run a number of different simulations for different parameter values, specifically investigating the

Figure 2.2: Pricing corrections of first order approximation

Test of the validity of the perturbation theory. These panels depict the ratios $C_1/(C_0 + C_1)$ in % where C_0 is the contribution to the price of the symmetrical solution and C_1 the leading order correction in perturbation theory. One has $C_1 \ll C_0$ even when the volatility is large. This demonstrates the validity of the perturbative expansion. The four cases correspond respectively to panels (1,3,5,7) of Figure 1.



effect different θ, ρ has on performance, as the other parameters in the $\Theta^{\text{pert.}}$ vector will have to attempt to absorb the information contained in the absent parameters. Secondly, it also

provides a first estimate into the applicability of the perturbative solution to different types of options markets. We simulate four different data sets with varying parameter vectors, described below.

- *data set 1*: $\Theta^{MG} = \{1.1768, 0.0823, 0.3000, -0.5459, 1.0000\}$, with a negative correlation which is a reasonable choice for modeling equity options, such as the S&P 500.
- *data set 2*: $\Theta^{MG} = \{1.1768, 0.0823, 0.3000, 0.0000, 1.0000\}$, with a correlation of zero.
- *data set 3*: $\Theta^{MG} = \{1.1768, 0.0823, 0.3000, +0.5459, 1.0000\}$, with a positive correlation coefficient; this is used to gain inference about modeling VIX options, see Park (2015).
- *data set 4*: $\Theta^{MG} = \{1.1768, 0.1250, 0.3000, -0.5459, 1.0000\}$, with a high(er) central tendency we investigate an equity style option market with a central tendency which is significantly higher than the initial variance value of: 0.08 and thus tests how the leading order perturbative solution handles significant change in the variance path.

For the following simulated calibration exercises, it is imperative to note the difference in IVRMSE and parameters between the data sets as this will highlight the following: firstly, parameter regions where the leading order perturbative solution might breakdown. Secondly, potential difficulties in estimating certain parameters of the model. Thirdly, where information contained in θ, ρ might be absorbed. These results can be found in Tables 2.2-2.3. Table 2.2 contains the results of the parameter vector estimates for data sets 1-4 along with summary statistics comparing to the Merton-Garman parameter vector. Table 2.3 contains results of the IVRMSE and standard deviations for each data set. The results of the data sets are described below:

- *Data set 1*: from Table 2.2 Panel 1 parameters κ, ξ appear to be challenging to estimate, being significantly larger and with quite high standard deviations, while α appears to be stable. While $\Theta^{\text{Pert.}}$ does not contain θ a significant amount of this missing information is absorbed by σ and ξ . Table 2.3 demonstrates the leading perturbative solution does

very well in approximating the Merton-Garman model with an IVRMSE of 1.2977% and standard deviation of 0.3474%.

- *Data set 2*: from Table 2.2 Panel 2 it starts to become clear some of the information contained in the correlation coefficient is absorbed by both ξ, α , particularly the latter. With the value of α reducing significantly while the standard deviation approximately doubles (relative to data set 1). Although it does appear that this regime is slightly easier to estimate κ, ξ . Table 2.3 reports a significant decrease (relative to data set 1) in IVRMSE with a moderate decrease in standard deviation.
- *Data set 3*: Table 2.2 Panel 3 demonstrates the absence of the information contained in the correlation coefficient has an effect on the ability to estimate κ, ξ, α in a similar manner to that of Panel 1, observing very similar biases and standard deviations. Perhaps suggesting that it is more the non-zero nature of the correlation coefficient which the leading perturbative solution struggles with. Furthermore, Table 2.3 demonstrates very similar errors to data set 1.
- *Data set 4*: from Table 2.2 Panel 4 the increase in central tendency also clearly has an impact in ξ, α the two variance related parameters of the leading perturbative solution. This is due to the increased difference in magnitude between the central tendency and initial variance. It suggests that both parameters absorb missing information contained in θ . This difference manifests itself in Table 2.3 with a resulting IVRMSE of 1.7199%, the largest across all data sets.

In summary, on the estimation side κ could certainly be a challenge to estimate, although this is not unique to our approach. For substantial difference between variance and central tendency it appears that α, ξ could also be a challenge, other than this estimating α is generally inconsiderable. Regarding the matter of information absorption we note that it is clear that σ absorbs significant amount of the information contained in θ , with contribution from α for large disparity between initial variance and central tendency. The parameters ξ, α also seem to share the majority of the information contained in ρ . Table 2.3 indicates that across data sets the leading perturbative solution does well in approximating the Merton-Garman

Table 2.2: Simulated time series parameter estimates

This table reports the results of our simulated time series calibration exercise against a Monte Carlo simulation of the Merton-Garman model. We use 100 different sample paths of weekly returns and latent variance over one year. At each observation we simulate six maturities within 7 to 180 days to maturity across a moneyness range of 0.9 to 1.1, using ten strikes. Each option price is computed using a Monte Carlo framework with 50,000 simulations and a time-step of $1/20^{\text{th}}$ of a trading day. We simulate four different data sets, each investigating a different initial parameter vector. The first rows (MG) presents the initial parameters used by the Monte Carlo simulation. The second rows (Pert.) presents the mean final parameter vector of our perturbative solution. The third rows (Bias) presents the difference between the Monte Carlo and perturbative solution parameter vectors. The fourth rows (Std. Error) presents the standard deviation of each parameter from the calibration exercises, across the 100 different calibrations for each data set.

		κ	θ	ξ	ρ	α	σ^2
data set 1	MG	1.1768	0.0823	0.3000	-0.5459	1.0000	-
	Pert.	4.1905	-	0.8772	-	0.9698	0.0855
	Bias	3.0137	-	0.5772	-	0.0302	-
	Std.						
	Error	3.6179	-	0.4163	-	0.1717	0.0120
data set 2	MG	1.1768	0.0823	0.3000	0.0000	1.0000	-
	Pert.	2.5454	-	0.6326	-	0.6178	0.0810
	Bias	1.3686	-	0.3326	-	0.3822	-
	Std.						
	Error	3.3240	-	0.4805	-	0.3584	0.0133
data set 3	MG	1.1768	0.0823	0.3000	0.5459	1.0000	-
	Pert.	4.4657	-	0.8806	-	0.9585	0.0851
	Bias	3.2889	-	0.5806	-	0.0415	-
	Std.						
	Error	3.3320	-	0.3094	-	0.1448	0.0120
data set 3	MG	1.1768	0.1250	0.3000	-0.5459	1.0000	-
	Pert.	5.4417	-	1.0104	-	1.6696	0.1022
	Bias	4.2649	-	0.7104	-	0.6696	-
	Std.						
	Error	4.0231	-	0.4529	-	1.4916	0.0145

model with fairly consistent errors, given the standard deviations, with an approximate error range of $1.2 - 1.7\%$.

Table 2.3: Simulated time series error estimates

This table displays the resulting IVRMSE and standard deviations to the simulated time series calibration exercise. We use 100 different sample paths of weekly returns and latent variance over one year. At each observation we simulate six maturities within 7 to 180 days to maturity across a moneyness range of 0.9 to 1.1, using ten strikes. Each option price is computed using a Monte Carlo framework with 50,000 simulations and a time-step of $1/20^{\text{th}}$ of a trading day. We simulate four different data sets, each investigating a different initial parameter vector. The first row (IVRMSE) presents the average IVRMSE between the Monte Carlo simulation and our perturbative solution. The second row (Std. Error) presents the standard deviation in IVRMSE across the 100 different calibrations for each data set.

	data set 1	data set 2	data set 3	data set 4
IVRMSE	0.0130	0.0117	0.0132	0.0172
Std. Error	0.0035	0.0031	0.0036	0.0061

2.5 Model building in finance, symmetries and group theory

In this section, we shall first discuss Galilean invariance, see e.g., (Bose (1995) and Lévy-Leblond (1967)), in the context of mathematical finance before explaining how new option pricing models can be constructed using the concept of symmetries. We shall assume that the dimensionless option price $\psi(x, y, t)$ (where x and y represent the dimensionless price and volatility variables as defined in Equations 2.7 and 2.8 and t is the time), from which the usual price of the option is derived, is the fundamental quantity. If we posit that $\psi(x, y, t)$ is a measurable quantity, it should not depend on the coordinate system \mathcal{P} that is being used to measure it. We could use \mathcal{P} parametrised by (x, y, t) or \mathcal{P}' parametrised by (x', y', t') and obtain the same dimensionless price, assuming that the two coordinate systems are related by a transformation which we shall take to be a Galilean transformation, knowing that $\psi(x, y, t)$ is a solution to the 2+1 heat equation:

$$\frac{\partial \psi(x, y, t)}{\partial t} = \frac{\partial^2 \psi(x, y, t)}{\partial x^2} + \frac{\partial^2 \psi(x, y, t)}{\partial y^2}. \quad (2.24)$$

A Galilean transformation can be decomposed as the composition of a rotation, a translation and a uniform motion in the space (x, y, t) where $\vec{x} = (x, y)$ represents a point in two-dimensional space, and t a point in one-dimensional time. A general point in this space can be ordered by a pair (\vec{x}, t) . Let us first consider the case where the two coordinate systems are moving away from each other in a uniform motion fixed by a two-dimensional constant vector \vec{w} . A uniform motion with vector \vec{w} , is given by

$$(\vec{x}, t) \mapsto (\vec{x} + t\vec{w}, t). \quad (2.25)$$

A translation is given by

$$(\vec{x}, t) \mapsto (\vec{x} + \vec{a}, t + s), \quad (2.26)$$

where \vec{a} is two-dimensional vector and s a real number. Finally a rotation is given by

$$(\vec{x}, t) \mapsto (G\vec{x}, t), \quad (2.27)$$

where G is 2×2 orthogonal transformation. It is easy to see that the partial differential equation obeyed by $\psi(x, y, \tau)$

$$\left(\frac{\partial}{\partial \tau} - \frac{\partial^2}{\partial x^2} - \frac{\partial^2}{\partial y^2} \right) \psi(x, y, \tau) = 0, \quad (2.28)$$

is covariant, in the sense that it is form invariant, under two-dimensional Galilean transformations $\text{Gal}(2)$.

Our Monte Carlo investigation of the Merton-Garman model demonstrates that the symmetry $\text{Gal}(2)$ is a good symmetry of the model as the leading perturbation theory around the symmetrical solution works very well. This demonstrates the importance of the symmetry between S and V which is manifest when expressed in the appropriate variables. To a very good approximation and for a range of parameters relevant to financial applications, the Merton-Garman model possesses a hidden $\text{Gal}(2)$ symmetry that is only softly broken. This is very clearly illustrated by the results obtained in Figure (2.2) which demonstrates

that the larger the volatility fluctuations are, the more the symmetry breaking terms become important. In panel 4 where the volatility is very close to being constant, the leading order correction is essentially vanishing and the symmetrical solution is very close to that obtained with the original Merton-Garman model. This is suggestive of an alternative way for building option pricing models. Instead of starting from stochastic processes, we could simply have derived the symmetrical model by positing that the option price should depend on the price of the underlying, a time dependent volatility and time. By requiring that the dimensionless option price follows a differential equation that is Gal(2) covariant, we would immediately have obtained the 2+1 dimensional heat equation. The choice of Gal(2) symmetry is imposed by the model we are considering, in different models there could be Lorentz invariance, see below. The very small deviations from the Gal(2) covariance, can be accounted for by symmetry breaking terms as explained above. This led us to an approximative perturbative analytical solution of the original Merton-Garman differential equation. We explicitly break the Gal(2) symmetry to reproduce the original Merton-Garman model. However, as a point for further research it would be instructive to find a controlled way to break these symmetries as in spontaneous symmetry breaking.

There is another interesting consequence of having a symmetry group as a fundamental building block. Namely, one can classify how the different objects in the model will transform under this symmetry. This is a very well developed field of mathematics called representation theory. In the case above, $\psi(x, y, t)$ is a scalar under Galilean transformations. A scalar representation of a symmetry group means that it is invariant under the group transformations. The differential equation is covariant under such transformations. Besides the scalar transformations, there are vector representations such as e.g., $\partial\psi(x, y, t)/\partial x$, $\partial\psi(x, y, t)/\partial y$ or $\partial\psi(x, y, t)/\partial t$ which are nothing but the Greeks. They appear from a very different perspective than is usually the case in finance.

These ideas open up new directions in model building for option pricing. One could for example consider a 3+1 dimensional model (x, y, z, t) where the z -direction could describe a time dependent interest rate or additionally this could be used to model a two-factor volatility process with stochastic central tendency, see e.g., Bardgett et al. (2019). The

differential equation for the price would be of the type

$$\frac{\partial \psi(x, y, z, t)}{\partial t} = \frac{\partial^2 \psi(x, y, z, t)}{\partial x^2} + \frac{\partial^2 \psi(x, y, z, t)}{\partial y^2} + \frac{\partial^2 \psi(x, y, z, t)}{\partial z^2}, \quad (2.29)$$

for which it is easy to find solutions:

$$G(x, y, \tau; x', y', z', t) = \Theta(\tau - t) \frac{1}{(4\pi(\tau - t))^{3/2}} \exp\left(-\frac{(x - x')^2 + (y - y')^2 + (z - z')^2}{4(\tau - t)}\right). \quad (2.30)$$

One could also consider “relativistic” extensions of the Merton-Garman model treating time on the same footing as the underlying price and the volatility:

$$\frac{\partial^2 \psi(x, y, t)}{\partial t^2} = \frac{\partial^2 \psi(x, y, t)}{\partial x^2} + \frac{\partial^2 \psi(x, y, t)}{\partial y^2}. \quad (2.31)$$

It is well known that this model is covariant under Lorentz transformations. Clearly identifying the right symmetry group for a given financial system is of paramount importance. Making use of symmetries to model physical system has been extremely successful in all fields of physics. Applying these ideas to option pricing models opens up new perspectives for model building in finance using the concept of symmetry groups and representation theory.

2.6 Conclusion

In this chapter we have introduced a perturbative method to obtain analytical approximative solutions to models such as the Merton-Garman model. The key idea consists in treating the price of the underlying and the volatility in a symmetrical way. This leads to a model which has an exact Galilean invariance in two-dimensions as it is described by the two-dimensional heat equation which has an analytical solution. By folding this solution with the boundary condition leading to the correct price at maturity for a call option, we obtained an analytical

symmetrical solution for this model which corresponds to the Black-Scholes solution, despite being derived from a very different perspective and in a framework with a time dependent volatility.

The Merton-Garman model is recovered by introducing symmetry breaking terms and we have calculated the leading order correction to the symmetrical solution. We have shown that our perturbative solution works very well by comparing it to a Monte Carlo simulation of the Merton-Garman model for a range of parameters which are relevant from a financial point of view. The moneyness curves of the two prices are so overlapping that we had to plot implied volatility curves to be able to discuss in a quantitative manner the differences between the two solutions.

We argued that the fact that the symmetrical model works so well is a sign that the Merton-Garman model has a hidden two-dimensional Galilean symmetry which is softly broken for the relevant parameter ranges. We have explained that the concept of symmetry, groups and representation theory could be extremely useful in building pricing models in financial mathematics. This clearly needs to be explored further. From this point of view, our work is opening up a new perspective on model building in mathematical finance.

Chapter 3

Macroeconomic Announcements and Volatility Seasonalities, a High-Frequency Approach to Option Pricing

The central research question in this chapter is the following. Given the move in recent times towards high-frequency, intradaily, data what are the important factors and processes needed to accurately model high-frequency options data? This voluminous research question is broken up into three sub-questions, which are as follows. Firstly, are standard factors and techniques (as used in Christoffersen et al. (2009), Eraker et al. (2003) and Eraker (2004) to name a few) which have been demonstrated to be relevant to lower frequency data, such as daily or weekly data still important? Secondly, how much does investors bias of volatility accrual, as observed in the weekend and overnight effects, impact option prices? Thirdly, what is the effect of macroeconomic announcements on option prices?¹

The use of our high-frequency data set affords us the opportunity to study effects which would have been previously impossible or at the very best only able to gain very noisy infer-

¹Please note that the notation is not necessarily the same as the previous chapter and is self-consistent and contained within this chapter.

ence: for instance, incorporation of the overnight and weekend effects into traditional option pricing models. This involves investigating volatility accrual intra and interday, for which we require at the very least quotes at the open and close. Muravyev and Ni (2019) find average returns for S&P 500 index options are negative (approximately -0.7%) per day, comprised of positive intraday and negative interday returns (referred to as the overnight effect). The authors hypothesise this is a manifestation of option market makers (OMMs) bias in ignoring the daily volatility seasonality that the trading day is more volatile than the overnight period (a well known seasonality). Muravyev and Ni (2019) find OMMs actually price options as if each period was as volatile as the other. We address the issue of incorporating this bias into traditional option pricing models and provide a further empirical test of Muravyev and Ni's hypothesis as to the cause of the overnight effect. In addition, using a similar methodology, we attempt to capture one of the earliest and most researched seasonalities, the weekend effect. Option returns have been found to be significantly lower over weekends, Friday close to Monday close (referred to as the weekend effect), than any other day of the week (documented in Jones and Shemesh (2018)). Jones and Shemesh (2018) document a potential explanation of the weekend effect being down to the convention of measuring volatility accrual using conventional calendar time, thus over a weekend two calendar days have elapsed with no trading (Friday close to Monday open) causing option implied volatilities to be too high prior to the weekend. Secondly, on the subject of our investigation into macroeconomic announcements we choose to investigate FOMC meetings. We choose FOMC meetings as Stroud and Johannes (2014) find this is one of the most significant macroeconomic announcements. Eight times a year the FOMC meet to determine near-term monetary policy. These announcements occur during the trading day, thus without high-frequency data researchers could only gain very noisy inference (such as Chen and Clements (2007)), while we are able to study their effect on options using minute data. Specifically we wish to determine if these announcements, whose timing is known ex-ante, generate fundamentally different risks compared to the traditional Brownian or jump driven risks in option pricing models, due to their predictable timing. While there is a substantial literature base on these topics to the best of our knowledge we are the first to incorporate these effects into option pricing models.

The work is ordered as follows: Section 3.1 provides a literature review and discussion. Sections 3.2.1 - 3.2.4 provides an introduction and discussion into the estimators of uncertainty around FOMC announcements, calibration methodology, VIX construction and preliminary data analysis, while Sections 3.2.5 - 3.2.8 provide summary statistics on the uncertainty estimates around FOMC announcements and option pricing implications of FOMC announcements. Sections 3.3.1 - 3.3.2 undergo a more in-depth discussion of the main research which influences our approach into volatility seasonalities, the model framework and notation for the work into the overnight and weekend effects. Sections 3.3.3 to 3.3.4 describe the empirical results from these investigations into the overnight and weekend effects, respectively.

3.1 Literature Review and Discussion

Empirical option pricing received a substantial increase in interest with the Black-Scholes (1973) model, which led mathematical legitimacy to option pricing. The model gives a theoretical estimate of the price of European-style options, which is unique regardless of the expected return of the underlying asset. However, one main drawback of the Black-Scholes model is the assumption of constant volatility regardless of moneyness or time to maturity of the option. For instance Macbeth and Merville (1979) find implied volatilities are not flat and often exhibit a "volatility smile" or "smirk" pattern. The empirical shortcomings of the Black-Scholes model prompted research along dimensions relaxing the restrictive assumptions. For instance univariate diffusion models, such as Cox and Ross (1976) who investigate a constant elasticity of variance. Geske (1979) incorporates leverage effects into option pricing, thus the variance of the rate of return on the underlying is non-constant.

A very substantial body of literature exists utilising stochastic volatility models, which allow the instantaneous volatility of asset returns to evolve stochastically, with respect to time. This usually takes the form of a diffusion or jump-diffusion process (such as in Merton (1976) or Duffie et al. (2000) which is what we shall focus on) but could incorporate a regime-switching framework as in Duan et al. (2002), who consider that the conditional

variance of the logarithmic return has a number of distinct volatility regimes, in each period there is a probability the volatility will move into a new regime.

The inclusion of standard independent and identically distributed jumps in diffusion models has received a lot of attention in the literature. Bates (1991) (initially) uses Poisson jumps in the price process to model extreme events such as the crash in 1987. Models with Poisson jumps match volatility smiles well for a single maturity, however, less well for multiple maturities. Jump models with this standard assumption converge towards Black-Scholes model option prices at longer tenor maturities, in stark contradiction to observed smiles at these maturities. Tompkins (2001) provides details of this pattern for stock index, bond, currency and interest rate options for a number of different countries, therefore demonstrating the feature is persistent.² Eraker et al. (2003) study the effect of jumps in returns and volatility on stochastic volatility models and find strong support for jumps in both processes. However, another important consideration is how to model the jump processes. Kaeck (2013) investigates the performance of alternative jump size distributions in the price process, using models which have separate distributions for positive and negative price jumps. Kaeck (2013) concludes a double-gamma jump specification performs best. This type of specification arises when the standard deviation of the negative jump is not restricted to be equal to the mean, but this restriction applies for positive jumps.

While single factor jump-diffusion models are very popular in the time series literature there is evidence which suggests a multi-factor volatility framework is needed, with one highly persistent and one quickly mean-reverting factor. Andersen et al. (2001) and Andersen et al. (2003) undergo a daily and intradaily study using stock return volatilities, both find evidence in need of a long-memory process, which could be filled by a second latent stochastic volatility factor. Additional evidence comes from the study of exchange rate data, from Alizadeh et al. (2002), who also find the evidence strongly suggests the need for one highly persistent factor and one quickly mean-reverting factor. In addition to Chernov et al. (2003) who study multiple continuous time specifications using a 37 year period of daily Dow Jones returns. More recently, Christoffersen et al. (2009) use a two-factor volatility

²Some observed phenomenon are purely market specific. For instance Kaeck and Alexander (2013) investigate the dynamics of European equity indices and demonstrate they have significant differences from the S&P 500 index. Finding the European indices to be much simpler to model.

model to fit the slope and level of the volatility smile, which move largely independently of each other. Their empirical findings demonstrate significant improvement over a benchmark Heston model.

In response to the above evidence there has been research into models which combine these myriad of different factors. For instance Bardgett et al. (2019) consider a two-factor stochastic volatility model with a persistent latent central tendency and a quickly mean reverting instantaneous latent volatility. In addition to which they model for contemporaneous random Poisson jumps in the price and instantaneous volatility in addition to random Poisson jumps in the central tendency. The intensities of the jumps also vary stochastically.

3.1.1 Volatility Seasonalities

Many pieces of research document systematic patterns in stock price volatility and volume of trade. French and Roll (1986) find stock volatility is higher during trading hours mainly due to private information accrued during periods of non-trading. Kaplanski and Levy (2015) show weekend, holiday and overnight trading breaks generate excessive perceived risk in options markets, presumably from asymmetric information, which encourages uninformed option traders to postpone trading until later in the week. Wood et al. (1985) and Harris (1986) document the well known volatility intradaily U-shape pattern, with volatility being higher at the open (Foster and Viswanathan (1993)) and close. This same pattern is also observed in stock bid-ask spreads, however options do not follow the same pattern, see Chan et al. (1995). This pattern is also linked to trading volume, Stoll and Whaley (1990) find during trading hours the relationship between higher volatility and volume is proportional. However, according to evidence from Foster and Viswanathan (1993), the intraday correlation becomes negative interday. There is also evidence from Jain and Joh (1988) to suggest volume has a weekly seasonality, being lower on Mondays and Fridays.

We investigate and contribute to the literature on the weekend effect, perhaps one of the earliest known seasonalities (reported by Fields (1931)). The effect has been shown to be prevalent in U.S. equities since 1885 (see Bessembinder and Hertz (1993)), and is one of

the most researched seasonalities (although some argue the effect is simply a result of data mining, see Sullivan et al. (2001)). The weekend effect has been observed in a myriad of different markets, for instance Jaffe and Westerfield (1985) observe evidence of the effect in international equity indices. Gibbons and Hess (1981) find evidence of the weekend effect in government bonds. However, there is also a strand of conflicting literature calling the robustness of the weekend effect into question. For instance Chang et al. (1993) study the robustness of day of the week effect in international markets. The authors test robustness using individual sample size and/or error term adjustments. The study finds sample size and/or error term adjustments yield U.S. day of the week effects statistically insignificant. However, five countries in Europe survive the double combination of sample size and error term adjustment. In most countries, where the effects are robust, they are statistically significant in not more than two weeks of the month. This evidence is inconsistent with explanations of the day of the week effect based on institutional differences or arrival of new information. Chang et al. (1993) suggest the effect disappeared somewhere between 1970-1980. Kamara (1997) shows the introduction of equity derivatives and institutionalisation of equity markets affects the weekend seasonality in the S&P 500. The weekend effect declines significantly over the sample period from 1962 to 1993, which is found to be positively related to the ratio of institutional to individual trading. However, the seasonality for small stocks does not decline during this period. Kamara (1997) hypothesises this is because for small stocks, institutions receive comparable trading costs to individuals, hence keeping the seasonality alive. Chen and Singal (2003) argue short sellers affect prices in a significant and systematic way, hypothesising the weekend effect has origins in speculative short selling before the weekend, suggesting these investors close out their positions on Friday and re-open them Monday morning. This causes the price to rise on Fridays and fall on Mondays. Further, reporting the weekend effect seemed to vanish for firms which had options traded on them, suggesting options might be a potential explanation of Kamara's results. Jones and Shemesh (2018) investigate delta hedged and un-hedged returns on options on individual equities. In response to Chen and Singal (2003) they reason a weekend effect should be seen in call and put options. The reasoning for which is if investors are averse to holding positions with unbounded risk over the weekend, then those who have written call options

should attempt to close out their positions. Whereas put option writers, while not facing unbounded downside risk, still face the possibilities of losses. As hypothesised, the authors find evidence in support of a weekend effect for individual equity options, but inconclusive results for options written on the S&P 500.

Our research also contributes to the modelling of the overnight effect. Muravyev and Ni (2019) suggest the cause of the overnight effect is essentially one of behavioural bias from market makers. There has not been a lot of research into behavioural factors in derivatives markets (Shefrin (2010) provides a good summary of behavioural finance). Stein (1989) investigates the term structure of option implied volatilities on S&P 100 options. Finding long maturity options tend to over react to implied volatility changes of shorter term options. Poteshman (2001) finds option implied volatilities under-react to individual daily changes in variance of the underlying asset returns and overreacts to mostly increasing or decreasing daily changes in variance. Han (2008) investigates investor sentiments and if this affects prices of S&P 500 options. Finding evidence to suggest changes in investor sentiment about the market has influence on the slope of the option implied volatility smile and also the risk-neutral skewness. There exists a growing literature base investigating the overnight effect in equity returns. For example, Cliff et al. (2008) demonstrates all the equity risk premium is a result of overnight returns. Their sample includes individual stocks, equity indexes and futures contracts. Bogousslavsky (2018) investigates intraday and overnight return anomalies in stock returns. We are aware of two papers which investigate option returns in an intradaily manner, in chronological order they are: Sheikh and Ronn (1994) and Muravyev and Ni (2019). Sheikh and Ronn (1994) use short term at-the-money (ATM) options on individual equity options (30 different stocks) for a 21 month period ending prior to the 1987 crash. They report observations of the overnight effect but find the results are not statistically significant. Muravyev and Ni (2019) suggest this inconclusive result is perhaps due to their sample size being too small. Sheikh and Ronn (1994) actually do not discuss overnight versus intraday returns but rather focus on returns towards the end of the trading day. Muravyev and Ni (2019) investigate delta-hedged average returns for S&P 500 index options. Finding they are negative and large per day. This daily return is broken up into positive intraday returns and negative interday. They find this is due to OMMs pricing

an equal amount of volatility accrual in both the overnight and trading day periods, thus ignoring a well known seasonality, more volatility accrues during the day than overnight.

3.1.2 Macroeconomic and Political Announcements

The notion of uncertainty as a fundamental motivating force of asset prices is not novel and has been well documented. Influential theoretical contributions include the work of Bloom (2009) who conclude uncertainty undergoes a discrete jump upward after major shocks. Further evidence from Pastor and Veronesi (2012) analyse the effect of government policy on stock prices, declining at the time of the announcement. The magnitude of decline is directly proportional to the uncertainty around the government policy. Furthermore, policy changes should increase volatilities and correlations amongst stocks, with jump risk premiums being positive on average.

There is also a substantial strand of literature which examines the effect of information in trading patterns. For example, Berry and Howe (1994) find a positive proportional relationship between news and stock volume (for further evidence see McQueen and Roley (1993), Mitchell and Mulherin (1994)). This also ties in with work done by Andersen et al. (2007) who investigate high-frequency stock, bond and exchange rate dynamics across the U.S., Germany and UK, finding a close link with fundamental surprises in news. Continuing this theme Scholtus et al. (2014) investigate the effects of macroeconomic news on algorithmic trading.

Theoretical contributions also come from authors utilising options, for instance Bekaert and Engstrom (2017) introduce a consumption-based asset pricing model which allows for realistic properties of equity index option prices and their co-movements with macroeconomic outlooks. When option implied volatility is high consumption growth is more negatively skewed. Furthermore, Kelly et al. (2016) study international equity option quotes to demonstrate that political uncertainty from national elections and global summits is priced by investors with the result that options spanning important events are on average more expensive. Lastly, Dew-Becker et al. (2018) use options contracts on 19 different markets

(covering a range of different components of the economy, including financial conditions, inflation and the price of real assets) to study the pricing of shocks to uncertainty and realised volatility. The entire research methodology revolves around the idea that instead of undergoing complex and sophisticated regressions based on uncertainty, the researchers let the investors communicate their views on uncertainty via option prices. The authors construct two types of straddle portfolios: one directly hedging innovations in uncertainty, and one hedging realised volatility. For both types of portfolio it is possible to learn about risk premia from their average returns. The authors find hedges against uncertainty earn positive returns and are thus not viewed as bad states of the economy. On the other hand portfolios hedging realised volatility earn statistically and economically significantly negative returns, indicating investors view shocks to fundamentals as undesirable.

Specifically related to our research there has been a thread of the literature interested in the impact of FOMC announcements. For instance, Lucca and Moench (2015) investigate returns accrued around FOMC announcements. Their results indicate since its beginning in 1994, U.S. stock returns have been on average thirty times higher on announcement days, compared to other days. Furthermore, the abnormal returns are largely accrued twenty-four hours before the announcement, with more than 80% of equity premium in this period. This feature is documented over a 17 year period from 1994 to 2011. The authors explore a few risk based theories, none of which can explain the abnormal returns. In conjunction with evidence from Savor and Wilson (2013) who investigate the difference between stock market average returns on days when macroeconomic announcements have taken place, versus non-announcements days. They find the average announcement day excess return from 1958 to 2009 is 11.4 basis points versus 1.1 for all other days. Suggesting over 60% of the cumulative annual equity risk premium is due to announcement days. There is also a large body of literature reporting the effects of macroeconomic announcements utilising options as a source of data. For instance Chen and Clements (2007) investigate the pattern of the VIX leading up to and after FOMC announcements between 3 January 1996 and 1 September 2006. They find no evidence to suggest a predictable movement in the VIX in the 5 days prior to and after an FOMC announcement. However, they find the VIX falls approximately 2% on the day of the FOMC meeting, which cannot be attributed to mean reversion. Nikkinen

and Sahlström (2004) investigate the impact of FOMC announcements (among others) on stock market uncertainty. They measure the implied volatility of the S&P 100 (the VIX at this time was based off S&P 100 options), concluding implied volatility increases prior to the scheduled news and drops after the announcement. Vähämaa and Äijö (2011) study the effect of FOMC announcements on volatility of the S&P 500. They find stock market volatility generally decreases after FOMC meetings, with a positive relationship between target rate surprise and market uncertainty. Krieger et al. (2012) also find the VIX declines significantly on scheduled meeting dates. They attribute the decline to resolution of uncertainty regarding future interest rates. Nikkinen and Sahlström (2004) and Savor and Wilson (2013) theorise the drop in VIX is mechanical. Krieger et al. (2012) mention the VIX can be thought of as a portfolio of one-day conditional volatilities. Therefore, when a high volatility day (such as an FOMC day) is removed from the portfolio of volatilities, a drop is to be expected. More recently Fernandez-Perez et al. (2017) study the intraday effect of FOMC announcements on the VIX and VIX futures. They find at the time of the announcement the VIX and VIX futures decline significantly. Fernandez-Perez et al. (2017) also report the decline after the announcement is not instantaneous but gradually declines for approximately 45 minutes after the announcement. This is consistent with work done by Patell and Wolfson (1984) who study the effect of earnings announcements occurring during trading hours, finding the majority of the effect occurs within the first few minutes. Amengual and Xiu (2018) model the term structure of variance in a general non-affine framework and conclude most of the downward volatility jumps, measured using daily VIX changes, are associated with FOMC announcements. Thus, confirming the findings from other studies concluding a decrease in volatility is a result of policy uncertainty resolution. In summary although there is some ambiguity as to when the effect of the announcement manifests, the undeniable conclusion is macroeconomic news is an important driver of uncertainty. Further this has been shown to be evident in Bond markets by Nakamura and Steinsson (2018). The authors demonstrate that an unexpected change in interest rates in a 30-minute window surrounding a scheduled FOMC announcement arise from new about monetary policy. Further suggesting that macroeconomic news has a large part to play in forecasts about output growth.

There has also been substantial research in recent times on the potential information

asymmetry of macroeconomic announcements. Bernanke and Kuttner (2005) analyse the impact of unanticipated cuts in the federal funds rate target on equity prices. This is done with two aims: to estimate the size of a typical reaction, to unanticipated rate target cuts, and further understand the reasons for the markets response. The authors find the CRSP value-weighted index registers a one-day gain of roughly 1 percent in response to a 25 basis point easing. In relation to what explains equity prices' response: the market reacts minimally, if at all, to the component of rate changes which are anticipated by markets participants. The authors find only a minimal portion of the effect is attributed to the changes on dividend forecasts and less to the effect on forecasts of real interest rates. The majority of the effect of funds rate surprises comes through their effect on expectations of future excess returns. Onan, Salih and Yasar (2014) investigate the impact of macroeconomic announcements on the high-frequency behaviour of observed implied volatility skew of the S&P 500 index options and VIX. The authors use tick-by-tick data of SPX options contracts for 250 trading days in 2006. As suggested by Aït-Sahalia et al. (2005) sampling too frequently might not be optimal in the presence of market microstructure noise. Onan, Salih and Yasar (2014) therefore choose 30 minute frequency data. Onan, Salih and Yasar (2014) find macroeconomic announcements affect the VIX level significantly and the slope to a lesser extent. Positive and negative news significantly and asymmetrically alters the implied volatility slope and the VIX. Gospodinov and Jamali (2012) examine the effect from expected and surprise target rate changes, from FOMC announcements, on realised and implied volatility. Their findings suggest surprise changes in the target rate significantly increase volatility. The results also indicate the expected component of a target rate change, along with the target rate change itself, do not significantly affect volatility. Asymmetric informational effects, such as described above, are also observed in currency markets as shown by Lobo et al. (2006). However, in contradiction to the above, Füss et al. (2018) introduce a new information density indicator (IDI) with the aim of understanding the effects of news on prices, specifically to better understand the effect upon price jumps. The IDI, which measures the abnormal flow of information before scheduled macroeconomic announcements, is significantly related to the likelihood of price jumps and is independent of the magnitude of news related surprises and pre-announcement trading activity. The authors interpret this

variable as an additional source of uncertainty in the market, inducing diffusive beliefs among investors. These beliefs are then resolved through the relating macroeconomic event(s) as facts.

However, it could be that there is no asymmetry as such, rather reactions are a press coverage issue. Gu et al. (2018) find evidence to suggest average stock returns after FOMC announcements are positive if accompanied by the release of Summary Economic Projections (SEP) and a press conference. The authors show several measures of uncertainty are significantly higher on days of FOMC announcements accompanied by SEP and press conferences than on announcement days without. Du et al. (2018) examines the informational content of equity options around FOMC announcements. The authors findings demonstrate information contained in option trades prior to FOMC rate change announcements, measured by implied volatility spread,³ predicts firm stock returns to a greater extent than volatility spread on non-meeting days. The authors go on to document that information is reflected in options prior to stock prices. The evidence cited above provides significant motivation to investigate macroeconomic announcements on option dynamics. Furthermore, of high importance is understanding how agents form expectations of the short-term interest rate in real-time is vital. Neglecting or miss-estimating risk premiums could lead to miss-perceptions of the market's outlook and could very possibly lead to sudden unexpected moves which were not anticipated. Cieslak (2018) studies how agents form expectations of the short-term interest rate. She finds persistent and significant differences between the expected real rate perceived by agents and the full-sample counterpart estimated by econometricians. These results are particularly pertinent when entering recessions, agents systematically overestimate the real rate and underestimate future unemployment and the amount of monetary easing.

While there is a substantial strand of literature, discussed above, which documents the effect of macroeconomic news releases there has also been research conducted along very different dimensions to predict Federal funds target rates, for example: Cieslak and Vissing-Jorgensen (2020), who study the impact of the stock market on the Federal Reserve's mone-

³Implied volatility spread is measured as the weighted difference in implied volatility between put and call options with the same strike and maturity.

tary policy. The authors find that low stock returns tend to predict accommodating monetary policy, to such an extent that statistically stock returns are a more powerful predictor of the Federal funds target rate, than standard macroeconomic news releases. The authors use a method of textual analysis of FOMC minutes and transcripts to argue that stock returns lead Federal Reserve policy. However, this relationship is not one sided, the Federal reserve causes meaningful impact on the stock market similarly. Cieslak et al. (2019) studies how much of realised stock returns can be attributed to the Federal Reserve. The authors find evidence to suggest that monetary policy news may not arrive only via FOMC announcement dates, but also from discount rate meetings, held bi-weekly. The authors document that the equity premium is earned entirely even weeks starting from the last FOMC meeting. Evidence is found to suggest systematic informal communication from Fed officials to the media and general financial sector as a channel through which news of the Feds intentions to provide needed stimulus, thus reducing the equity premium. This two way relationship is also corroborated by Stein and Sunderam (2018), who demonstrate that the central bank, in clear possession of private information regarding its long-run target rate, is averse to bond market volatility. As a consequence the central bank gradually impounds changes to its central rate, while actively taking steps to learn what market participants think the Fed is thinking. Prior to each FOMC meeting the Fed performs a detailed survey of primary market participants, asking detailed questions about what they think the Fed will do.

In modelling the response of macroeconomic news many researchers have found the news to induce a jump or discontinuity. Johannes (2004) and Barndoff-Nielsen and Shephard (2006) find macroeconomic news can induce jumps in interest rate models. Beber and Brandt (2006) examine the effect of regularly scheduled macroeconomic announcements on treasury bond futures option prices closely surrounding times of announcements. Their findings show the announcements reduce the uncertainty regardless of the content of the news. Piazzesi (2005) investigate bond yields in response to FOMC announcements, modelling the policy news as a pure jump process. This influences the approach to our empirical work, prior research clearly suggests it is reasonable to model the uncertainty due to macroeconomic announcements as a discontinuous jump in the price process, which we also adopt.

3.2 Incorporating FOMC Announcements in Index Price Models

This section augments FOMC announcement risks into standard continuous-time stochastic volatility models. The first point of discussion is how to model FOMC announcements and their impact on index prices. We assume FOMC announcements induce a jump in the price path. The announcement releases information regarding monetary policy, the jump sizes transforms this information into discontinuities in the price process. Thus the jump distribution acts as a model of how information affects prices. We are not alone in choosing to model announcement effects in such a way (see Piazzesi (2005) and Dubinsky et al. (2018)) who also use the jump model. Furthermore, it is consistent with statistical evidence which suggests announcements are the cause of jumps in jump-diffusion models (Johannes (2004), Barndoff-Nielsen and Shephard (2006)). As the FOMC jump times are known in advance we use a counter, N_t^d to count the number of FOMC announcements prior to maturity at time T : $N_t^d = \sum_j \mathbb{1}_{[t_j \leq T]}$, where $\mathbb{1}$ is the indicator function and t_j is an increasing sequence of predictable announcement times. FOMC announcements occur during trading hours, usually around 2:00 p.m., we anticipate the effects of the announcement to be absorbed within an hour. An assumption which is supported by prior work, for instance Fernandez-Perez et al. (2017) study intradaily effects of FOMC announcements on the VIX and VIX futures, finding the decline after the announcement is not instantaneous but lasts for approximately 45 minutes. Further evidence from Patell and Wolfson (1984) find, when studying earnings announcements which occur during the trading day, that the information is absorbed within minutes.

We consider an affine, mean reverting square root two-factor stochastic volatility model,

where prices and variance processes solve, under the risk-neutral measure \mathbb{Q} :

$$\begin{aligned} dY_t = & -(\lambda\bar{\mu} + \frac{1}{2}V_{t-})dt + \sqrt{V_{t-}}dW_t^Y(\mathbb{Q}) \\ & + (e^{\xi_t^y} - 1)dZ_t(\mathbb{Q}) - \lambda\mu_y^{\mathbb{Q}}dt + (e^{\xi_t^{FOMC}} - 1)d\bar{Z}_t(\mathbb{Q}), \end{aligned} \quad (3.1)$$

$$dV_t = \kappa_v^{\mathbb{Q}}(m_{t-} - V_{t-})dt + \sigma_v\sqrt{V_{t-}}dW_t^v(\mathbb{Q}) + \xi_t^v dZ_t(\mathbb{Q}), \quad (3.2)$$

$$dm_t = \kappa_m^{\mathbb{Q}}(\theta_m^{\mathbb{Q}} - m_{t-})dt + \sigma_m\sqrt{m_{t-}}dW_t^m(\mathbb{Q}). \quad (3.3)$$

Where for a stochastic jump process X_t , we denote X_{t-} as the value of X prior to any jump at time t . Denote $Y_t = \ln(F_t(\tau))$, with $F_t(\tau)$ being the forward price, at time t with maturity at time T and thus time to maturity $\tau = T - t$. V_t is the instantaneous latent variance process of Y_t with parameter σ_v , and $\kappa_v^{\mathbb{Q}}$ is the rate of mean reversion to m_t , the central tendency (with volatility parameter σ_m and $\kappa_m^{\mathbb{Q}}$ being its rate of mean reversion to the constant parameter $\theta_m^{\mathbb{Q}}$). We assume the three Brownian motions: W_t^Y, W_t^V, W_t^m have the following correlations:

$$d\langle W_t^Y, W_t^V \rangle = \rho, \quad (3.4)$$

$$d\langle W_t^Y, W_t^m \rangle = 0, \quad (3.5)$$

$$d\langle W_t^V, W_t^m \rangle = 0. \quad (3.6)$$

No arbitrage conditions require the volatility of volatilities, σ_v and σ_m and the correlation to be equal under both the historical and risk-neutral measures. We define Z_t as a Poisson jump process, with constant intensity λ (and with finite activity).⁴ In modelling the random jumps, we assume they occur only in the index and its instantaneous variance, Bardgett et al. (2019) use similar risk-neutral dynamics and also model for random jumps in the central tendency but conclude these are unnecessary, based on their work we do not include them. Furthermore, jumps are contemporaneous, occur at random times depending on the Poisson process Z_t . Following the standard literature (see Eraker et al. (2003), Eraker (2004) or Broadie et al. (2007)) we assume the variance jump sizes are exponentially distributed

⁴Better model fit could be achieved by allowing for a non-constant jump intensity which varies with the level of the variance processes, such as is used in Bardgett et al. (2019). However, this is unlikely to detract substantially from our work and the conclusions regarding announcement effects or volatility seasonalities.

with $\xi_t^v \sim E(\mu_v^Q)$. The return jumps are normally distributed with $\xi_t^y \sim \mathcal{N}(\mu_y^Q, \sigma_y^Q)$, such that $\bar{\mu} = \exp(\mu_y^Q + \frac{1}{2}(\sigma_y^Q)^2) - 1 - \mu_y^Q$.⁵ The jump \bar{Z}_j captures the price movement in response to information released in FOMC announcements. We assume jump magnitudes are state independent and conditionally normally distributed under the risk-neutral measure, \mathbb{Q} : $\xi_t^{FOMC} \sim \mathcal{N}(-\frac{1}{2}(\sigma_j^Q)^2, (\sigma_j^Q)^2)$, where σ_j^Q represents the volatility of the FOMC announcement occurring at time t_j . This approach to modelling predictably timed price jumps is in line with previous work in Dubinsky et al. (2018). The addition of deterministic jumps only adds one additional parameter, per FOMC announcement, in order to keep the expected growth of the underlying at the risk-free rate. Appendix 6.5 presents a derivation of the characteristic function. Estimating σ_j^Q is a basal focus of this work. The model dynamics in 3.1-3.3 nest several models which we will use for various estimation exercises. Firstly the single-factor stochastic volatility model of Heston (1993), defined with a constant central tendency and no jumps, denoted SV. Secondly, the single-factor stochastic volatility model augmented with random price jumps of Bates (2000), denoted SVJ. Thirdly, the two-factor stochastic volatility model with no jumps, denoted SV2. The two-factor stochastic volatility model augmented with random contemporaneous jumps in the price and variance process, denoted SVCJ2. Finally the two-factor jump model SVCJ2 augmented with deterministic FOMC jumps, denoted SVCJ2D.

3.2.1 Anticipated Uncertainty Estimators

The estimation and interpretation of σ_j^Q is a fundamental piece for this work, as such we first introduce a simple model to gain intuition and insight as to the information contained in and how to estimate σ_j^Q . Following the framework of Dubinsky et al. (2018) consider a special case of the general stochastic volatility model with constant volatility and price

⁵Note as far as risk premia are concerned we follow Pan (2002) and Eraker (2004) and impose the intensities of jumps be the same under the risk-neutral and historical measures. As a manifestation of such an assumption it is imposed this jump-timing risk premium is absorbed into the mean jump size risk premium. Details of this assumption do not affect our work as we do not estimate any historical measure parameters.

jumps on FOMC announcements:

$$F_t = F_0 \exp \left[-\frac{1}{2} \sigma^2 t + \sigma W_t(\mathbb{Q}) + \sum_{j=1}^{N_t^d} \bar{Z}_j(\mathbb{Q}) \right], \quad (3.7)$$

where $\bar{Z}_j(\mathbb{Q}) = -\frac{1}{2}(\sigma_j^{\mathbb{Q}})^2 + \sigma_j^{\mathbb{Q}}\epsilon_j^{\mathbb{Q}}$ is responsible for the price discontinuities caused by FOMC announcements with $\epsilon_j^{\mathbb{Q}} \sim \mathcal{N}(0, 1)$. Additionally, N_t^d counts FOMC announcements prior to time t . As the Brownian motion, $W_t^{\mathbb{Q}}$, and FOMC jumps are normally distributed the prices follow a log-normal distribution. As such a European call option with maturity at time T will be given by the Black-Scholes formula:

$$C(F_t, t) = e^{-r(T-t)} \left[F_t \mathcal{N} \left(\frac{\ln(F_t/K) + (\sigma^2/2)(T-t)}{\sigma \sqrt{T-t}} \right) - K \mathcal{N} \left(\frac{\ln(F_t/K) - (\sigma^2/2)(T-t)}{\sigma \sqrt{T-t}} \right) \right], \quad (3.8)$$

where $\mathcal{N}(\cdot)$ is the cumulative normal distribution function and K is the strike price.⁶ However, with a modified volatility parameter to account for the FOMC jump:

$$\sigma_{t, T-t}^2 = \sigma^2 + (T-t)^{-1} \sum_{j: t < t_j \leq t+T} (\sigma_j^{\mathbb{Q}})^2, \quad (3.9)$$

this model provides the main test hypotheses we will investigate in our work for predetermined announcements for option prices. Firstly, for an option with one announcement prior to maturity, the change in implied variance (IV) just before and after the announcement is: $(T-t)^{-1}(\sigma_j^{\mathbb{Q}})^2$, implying IV drops discontinuously immediately after the announcement. Second, implied variance increases at a rate $(T-t)^{-1}$. Thirdly, fixing the number of FOMC announcements (and therefore jumps) constant the term structure of implied variance slopes downwards.

As noted in Dubinsky et al. (2018) there are two estimation methodologies: from the term structure and the times series. First we explain the term structure estimate. Given two options with maturities T_1 and T_2 ($T_1 < T_2$) and one FOMC announcement prior to matu-

⁶Note to calculate the value of a put option one would invoke the put-call parity.

ity, then from Equation (3.9) the following must hold: $\sigma_{t,T_1-t}^2 > \sigma_{t,T_2-t}^2$ and $\sigma_j^{\mathbb{Q}}$ is estimated by:

$$(\sigma_j^{\mathbb{Q}})^2 = \frac{\sigma_{t,T_1-t}^2 - \sigma_{t,T_2-t}^2}{(T_1 - t)^{-1} - (T_2 - t)^{-1}}. \quad (3.10)$$

The time series estimator uses information from changes in implied variance over the announcement. Let there be an FOMC announcement on date t , further let there be only one announcement prior to maturity, then the post-announcement IV is σ . The time series estimator is thus given by:

$$(\sigma_j^{\mathbb{Q}})^2 = (T - t)(\sigma_{t-,T-t}^2 - \sigma_{t+\delta t,T-t-\delta t}^2), \quad (3.11)$$

where due to the high-frequency data $\delta t = 60$ minutes. For details on the discussion of robustness to stochastic volatility (the model outlined in Equation (3.7) assumes constant volatility) and problematic areas of each methodology see Appendix 6.6.

3.2.2 Calibration Methodology

This section provides details on the calibration methodology used in our calibration exercises, paying particular attention to the objective function and the method of latent variance estimation.

We consider a number of nested versions of the two-factor stochastic volatility model developed in Section 3.2. All model option prices are achieved by using the Fourier Cosine Expansion of Fang and Oosterlee (2009) (for details of the procedure and how this relates to the two-factor models see Appendix 6.7). For all models we estimate the latent variances using an adapted method of Duan and Yeh (2010). The work of Duan and Yeh (2010) establishes a theoretical link between the VIX and the latent variance for a one-factor stochastic volatility model with jumps in the price process. We extend this theoretical methodology to encompass our model framework (the details of which are presented in Appendix 6.8). Let $O_{t,i}^{mod}(F_t, V_t, m_t, T, K, \Theta^{\mathbb{Q}})$ denote the model based price of option i , put or call, struck at K

at time t , with maturity at time T and let

$$\Theta^{\mathbb{Q}} = \{\kappa_v^{\mathbb{Q}}, \theta_m^{\mathbb{Q}}, \sigma_v, \rho, \kappa_m^{\mathbb{Q}}, \sigma_m, \lambda, \mu_y^{\mathbb{Q}}, \sigma_y^{\mathbb{Q}}, \mu_v^{\mathbb{Q}}, \{\sigma^{\mathbb{Q}}\}_{j=1}^{N_{\text{FOMC}}}\}, \quad (3.12)$$

denote the structural parameter vector of the SVCJ2D model. Where the last element, $\{\sigma^{\mathbb{Q}}\}_{j=1}^{N_{\text{FOMC}}}$, represents a vector of FOMC announcement volatility parameters, with $N_{\text{FOMC}} = 8$, one for each FOMC announcement, each year. Ideally we would use the implied volatility RMSE (IVRMSE) for the objective function and comparing model performance. However, estimating the model parameters using IVRMSE is numerically very intensive as the Black-Scholes inversion must be done for every set of model option prices queried by the optimisation function. Instead we use a first order approximation to the IVRMSE. Denote $\sigma_{t,i}^{mkt}, \sigma_{t,i}^{mod}$ as the market and model implied volatility of option i at time t . To compute the implied volatilities we invert each option price using the Black-Scholes (1973) model. A first order expansion in the price of the option with respect to the implied volatility is:

$$O_{t,i}^{mkt} - O_{t,i}^{mod} \approx \frac{\partial O_{t,i}^{mkt}}{\partial \sigma_{t,i}^{mkt}} (\sigma_{t,i}^{mkt} - \sigma_{t,i}^{mod}), \quad (3.13)$$

where $O_{t,i}^{mkt}$ denotes the market price of option i (put or call), at time t . Denote $\partial O_{t,i}^{mkt} / \partial \sigma_{t,i}^{mkt}$ as the Black-Scholes option vega computed at the market implied volatility, denoted as $BSV_{t,i}^{mkt}$. Thus, an approximation to the difference in implied volatilities between the market and model can be given as:

$$\sigma_{t,i}^{mkt} - \sigma_{t,i}^{mod} \approx \frac{O_{t,i}^{mkt} - O_{t,i}^{mod}}{BSV_{t,i}^{mkt}}. \quad (3.14)$$

Which gives our objective function as the Vega-Weighted RMSE (VWRMSE henceforth) of Trolle and Schwartz (2009):

$$\text{VWRMSE} = \frac{1}{N} \sum_{t=1}^N \sqrt{\frac{1}{n_t} \sum_{j=1}^{n_t} \left(\frac{O_{t,j}^{mkt} - O_{t,j}^{mod}}{BSV_{t,j}^{mkt}} \right)^2}, \quad (3.15)$$

where N denotes the total number of cross-sectional observations in the data (i.e., the total number of 30 minute observations) and n_t denotes the total number of unique strikes and

maturities across all option maturities available at time t .

Throughout this work we are concerned with incorporating the effects of FOMC announcements and volatility seasonalities into traditional option pricing models. For each of these two effects we choose to use a two-step calibration process. Whereby we calibrate the base model (SV, SVJ, SV2 or SVCJ2) to data from January to August, inclusive, leaving September to December as an out-of-sample period. In the second step we calibrate the parameters relating to each effect we are incorporating, in that specific test, e.g., FOMC announcements or volatility seasonalities. During the second step, the parameters from the base model are held fixed at their optimal values found in step one. In choosing this calibration methodology we are able to explicitly measure the advantage of including each individual effect in an option pricing framework. The reason for which is that, as the parameters from step one are fixed, there can be no interplay between parameters of the original model and the new parameters responsible for the new effect we are measuring. When studying stochastic volatility models there is the challenge of having to jointly estimate the model's structural parameters as well as the latent volatilities. As a result there are a myriad of different approaches used in the literature. What follows is a non-exhaustive list of some of the studies which employ different methodologies. One approach is to simply treat the volatilities as simply another parameter which is re-estimated daily, such as in Bakshi et al. (1997). Another approach consists of filtering the latent volatility using a time series of underlying returns. This can be done in a number of different ways, for example Eraker (2004) does this in a Bayesian setting. Alternatively Pan (2002) uses the Generalized Method of Moments while Carr and Wu (2007) use a Kalman filter. Alternatively Christoffersen et al. (2009) use an iterative two-step procedure (similar to Bates (2000)) by estimating the structural parameters and spot volatilities using option data. Where on the first step for a given set of structural parameters the latent volatilities are optimised, the second step uses the set of optimised spot volatilities from step one to solve one aggregate sum of squared pricing error optimisation problems for the structural parameters. A further study done by Bardgett et al. (2019) uses a particle filter to estimate the conditional densities of un-observable latent processes such as the volatility and jump process at every point in time. It is combined with a maximum likelihood procedure for parameter and standard error estimations. Due to the

many different approaches available we feel our two-step procedure is reasonable.

3.2.3 VIX Construction

The VIX is needed in order to estimate the latent variance and central tendency, using our adapted approach of the Duan and Yeh (2010) method, for use in calibration exercises. The necessity of having to construct the VIX is purely down to dealing with intradaily data we need estimates for the VIX, across multiple time horizons, to match our observation frequency.⁷ Whereas previous research which deals with daily or weekly data at the close could simply obtain the recorded values from the CBOE. The VIX index is defined in the White paper of the CBOE (2009) and is calculated practically using a mixture of S&P 500 out-of-the-money (OTM) options with maturities adjacent to 30 days. To suit our purpose we calculate the quantity:

$$VIX_t(\tau) = \sqrt{\frac{2}{\tau} e^{r\tau} \left[\int_0^{F_t(\tau)} \frac{P_t(K, \tau)}{K^2} dK + \int_{F_t(\tau)}^{\infty} \frac{C_t(K, \tau)}{K^2} dK \right]}, \quad (3.16)$$

where $P_t(K, \tau)$ and $C_t(K, \tau)$ denote the prices of put and call options, respectively, at time t with time to maturity τ (i.e., maturity at time T , yielding $\tau = T - t$) and strike K . The risk-free rate is denoted as r . Lastly, $F_t(\tau)$ denotes the forward price at time t , with maturity at time T . Therefore, the quantity $VIX_t(\tau)$ denotes the calculated VIX value at time t which spans the time-horizon τ . In practice however, as the price distribution against moneyness is sharply peaked and decays to zero we may truncate the upper bound. We restrict the upper bound to $nF_t(\tau)$, where n is some integer sufficiently large such that a call option with strike $nF_t(\tau)$ is worth zero. In practice we choose $n = 6$, we thus approximate the VIX by:

$$VIX_t(\tau) \sim \sqrt{\frac{2}{\tau} e^{r\tau} \left[\int_0^{F_t(\tau)} \frac{P_t(K, \tau)}{K^2} dK + \int_{F_t(\tau)}^{nF_t(\tau)} \frac{C_t(K, \tau)}{K^2} dK \right]}. \quad (3.17)$$

⁷The use of estimates of the VIX across multiple time horizons is required in order to estimate both the latent variance and central tendency.

The next issue to tackle is in calculating the option prices $C_t(K, \tau)$ and $P_t(K, \tau)$. Due to having to integrate over these strikes we are required to be able to provide a continuum of $P_t(K, \tau)$ and $C_t(K, \tau)$ for the strike range $[0, nF_t(\tau)]$. In order to do so we use an interpolation method to provide prices at all strikes. The first step is to generate a least squares regressed implied volatility smile from the raw data with strike vector k , using a polynomial of degree four:

$$IV_R(k) = \hat{\beta}_0 + \hat{\beta}_1 \frac{k}{F} + \hat{\beta}_2 \left(\frac{k}{F} \right)^2 + \hat{\beta}_3 \left(\frac{k}{F} \right)^3 + \hat{\beta}_4 \left(\frac{k}{F} \right)^4 + \epsilon. \quad (3.18)$$

The β coefficients are then used to generate a query implied volatility smile, IV_Q :

$$IV_Q(K) = \hat{\beta}_0 + \hat{\beta}_1 \frac{K}{F} + \hat{\beta}_2 \left(\frac{K}{F} \right)^2 + \hat{\beta}_3 \left(\frac{K}{F} \right)^3 + \hat{\beta}_4 \left(\frac{K}{F} \right)^4. \quad (3.19)$$

The query strikes, K cover a range much greater than that of the data. As a result we set any data outside the observed range $[k_{\max}, k_{\min}]$ to the regressed maximum/minimum values obtained from the regression on k :

$$IV_Q(K > k_{\max}) = \hat{\beta}_0 + \hat{\beta}_1 \frac{k_{\max}}{F} + \hat{\beta}_2 \left(\frac{k_{\max}}{F} \right)^2 + \hat{\beta}_3 \left(\frac{k_{\max}}{F} \right)^3 + \hat{\beta}_4 \left(\frac{k_{\max}}{F} \right)^4, \quad (3.20)$$

$$IV_Q(K < k_{\min}) = \hat{\beta}_0 + \hat{\beta}_1 \frac{k_{\min}}{F} + \hat{\beta}_2 \left(\frac{k_{\min}}{F} \right)^2 + \hat{\beta}_3 \left(\frac{k_{\min}}{F} \right)^3 + \hat{\beta}_4 \left(\frac{k_{\min}}{F} \right)^4. \quad (3.21)$$

This mitigates potential reliability damage, due to noisy data, from options with very low moneyness, the same process is applied in: Carr and Wu (2008), Rehman and Vilkov (2012) and Neumann and Skiadopoulos (2013). In the low moneyness region the implied volatility smile can start to dip down again. By setting these extreme regions of moneyness to constant values it creates a stable environment for interpolation and it is possible to back out reliable prices using the Black-Scholes model. We are aware of other methods of interpolating the implied volatility surface. For instance we could have used linear extrapolation in the tails (as used in Jiang and Tian (2007) and Aït-Sahalia and Lo (1998)). The reason we chose horizontal extrapolation is because of the evidence found in Amman and Feser (2019). The authors examine the approaches mentioned above with the view of determining the most

reliable. They find as the domain of available strikes declines linear out performs horizontal extrapolation in the tails. However, they find horizontal extrapolation still performs well (and out performs linear extrapolation) in the presence of high micro-structure noise. As we have a relatively dense and wide spanning strike range and we analyse intradaily data we feel it is more important to guard against micro-structure noise. The resulting shape of prices against moneyness is a sharply peaked almost Christmas tree like distribution around at-the-money strikes and zero in the tails.⁸ To calculate the integral in Equation (3.17) we use a numerical adaptive quadrature based on a 7-point Gauss rule with a 15-point Kronrod rule. This quadrature formula is preferred over others such as a low order adaptive recursive Simpson's rule or a high order adaptive Gaussian Lobatto quadrature rule. The reason for this is the Gauss-Kronrod we use has a better resolution and is thus more robust (see Shampine (2009) for details). The above procedure generates VIX values over every time horizon in each cross-section. In order to standardise the resultant values we linearly interpolate the VIX_t^2 between 30 and 360 days in 30 day increments. We choose to interpolate the VIX_t^2 as this is more reliable, due to variance being linear in time. We report summary statistics and results of the VIX time series in Section 3.2.4.

3.2.4 Preliminary Data Analysis

We obtain intradaily European options data from the CBOE. We have access to data from Jan 2004 to July 2017. Following previous work (Onan, Salih and Yasar (2014) and Aït-Sahalia et al. (2005)) we sample data every 30-minutes during the trading day to avoid market microstructure noise. Due to the computational burdens of investigating intradaily data, it is not feasible to use the entire data set. Further the motivation to include the recent financial crisis and more recent data lead to using alternate years starting in 2008 and ending in 2016. This time series spans periods of relative calm and very turbulent periods

⁸In the above procedure, it would be a reasonable question to ask why we appear to create more work for ourselves in calculating the implied volatilities, interpolating them and then converting back to prices when we had the raw price data all along. The answer to which lies in the interpolation phase. Due to the distribution of prices against moneyness it is significantly easier to interpolate around the volatility smile instead.

with extreme events, which is especially relevant in the exercise of estimating frequency and magnitude of jumps. Thus, for investigation into the FOMC announcements we consider only scheduled announcements starting in 2008 taking every other year (e.g., we consider scheduled announcements in 2008, 2010, 2012, 2014 and 2016, only making for a total of 40 announcements). We choose to neglect the first 15 minutes of trading and take our first observation at 9:45 a.m., Chan et al. (1995) show (confirmed by Muravyev and Ni (2019)) option quotes are intermittent with large bid-ask spreads during the initial open period. Furthermore, by delaying fifteen minutes it is possible to take the last observation of each day at exactly the close.

The maturity profile we have available to us changes significantly over the time span of our study. In 2010 the CBOE introduced Friday weekly expiry options (symbol SPXW), in conjunction to their existing monthly and quarterly expiration profile. The Friday weekly options increase in traded volume dramatically during our sample period, where in 2010 they comprised approximately 10% of trading volume by mid-2015 this grows to approximately 30% (as noted in Andersen et al. (2017)). Later in our sample the CBOE introduced an additional two weekly expiration's, on Wednesday and Monday, which were introduced on February 23 and August 15 of 2016. As these weekly expiries were only introduced in one year and they contain duplicated information with the Friday expiration we remove them from our sample. The justification of including the Friday weekly expiries into our data set is twofold: firstly it will increase the stability of the VIX estimation and in turn that of the latent variance and central tendency. Without the weekly expiries we would have to interpolate the 30-day VIX with only one data point less than 30 days. By including the weekly expiries we increase the number of option expiries before 30 days, thus increasing the stability of the VIX term structure. Secondly, it should be easiest to observe the effects of FOMC announcements and volatility seasonalities on short tenor options. In terms of maturity filters we only consider options between 4 and 550 days to maturity. The longer maturity options (those with greater than a year to expiry) are rich in information content for determining the mean-reversion rate and variance of the central tendency process, as such are important for the two-factor models in the calibration exercise. We also delete all quotes which have a bid less than 50 cents. Furthermore, we delete all in-the-money (ITM) options

since they are illiquid compared to their out-of-the-money counterparts. Furthermore, we keep only those OTM options which are within the moneyness range: 0.8 and 1.1. We do not delete quotes with zero volume, making the interpolation on the VIX more stable. If options with zero volume were removed, the strike range would increase throughout the day as more options are traded. This increasing strike range would lead to more stable inference on the VIX towards the end of the day, but crucially would lead to more extrapolation needed at the beginning of the day and thus a less stable VIX towards the beginning of the day.

We obtain interest rate data from OptionMetrics. Another issue we need to contend with is our interest rate data is only quoted at a daily level at the close. The data is comprised of three columns: a quote time, days to maturity and rate. As the interest rates are quoted at the close for all calculations using interest rates we use the previous days data. Also note, each day does not necessarily match the maturity profile we require. Thus, we linearly interpolate to generate intradaily interest rates at the required maturities per cross-section. It is important to note in the interest rate data the front maturity may not be less than the front maturity of the cross-section. For instance, years 2008 to 2012 the front maturity ranges between 6 to 7 days. While years 2014 to 2016 the front maturity is between 1-7 days. Henceforth, some of the front maturities are calculated using a linear extrapolation. However, the need for extrapolation is not an issue for the last maturity as the interest rate data always quotes a rate for a minimum of a 3561 day horizon and we consider a maximum maturity of 550 days.⁹

Next, we infer the forward price using the at-the-money put-call pair (price of the options are based on the best bid and ask price midpoints).¹⁰ This avoids two problems: firstly, making estimations of future dividends. Secondly, the issue of using forward prices which are not synchronised with option prices also traditionally forward contracts are only quoted quarterly so would not be available for all maturities. Therefore, we consider the underlying of the options is the Futures index and not the actual index itself (this method is similar to Bardgett et. al, (2019)).¹¹

⁹We choose not to add stochastic interest rates into our model framework. Following the results of Bakshi et al. (1997) who find for out-of-sample pricing performance there is no gain for stochastic volatility models.

¹⁰However, similar results are obtained using an averaging method of all the options in each cross-section.

¹¹This is feasible for the following reasons. At maturity the spot and forward price are equal. Furthermore,

Table 3.1: Number of options (binned by year)

This table provides information on the numbers of out-of-the-money options used in our data set. The table displays the number of option used in each year (Year). The table also displays the average number of options included in the data set each week, binned by year, (Weekly). The average number of options included in the data set per day, binned by year, (Daily). The average number of options quoted by observation, as we sample every 30-minutes starting at 9:45 a.m. we have 14 samples per day, binned by year, (Observation).

	2008	2010	2012	2014	2016
Year	731202	852620	1257787	2002583	2209493
weekly	14062	16397	24188	38511	42490
Daily	2902	3383	4991	7947	8768
Observation	207	242	357	568	626


The results of the above filtering procedure on the number of options comprising the data set can be seen in Table 3.1. The first row of Table 3.1 reports the total number of options contained in each year comprising the data set. The total number of options grows substantially over time, with 2008 having 731,202 quotes and ending with 2016 having 2,209,493, making up nearly one third of the total data set in one year! This increase is most likely caused by growth on two fronts, one the number of quoted maturities (due to the introduction and increased popularity of weekly expiry options) and two, the density of strikes. In terms of strike density, on average in 2008 there are 28 strikes quoted per maturity, this grows to 59 strikes per maturity for 2016. The subsequent four rows of Table 3.1 report average number of options for each year based on incrementally smaller time aggregation methods, ranging from the number of options available, on average, per week to per 30-minute cross-section. As a benchmark Bardgett et. al, (2019) use a data set with a mixture of SPX and VIX options from March, 2006 to October, 2010 having a total of 427,061 options in the entire data set. In comparison our total data set is comprised of 7,053,685 options (i.e., approximately seventeen times that of Bardgett et. al, (2019)). This gives a flavour for the computational burden we undertake by using intradaily data.

Figure 3.1 displays the returns of the S&P 500 index, sampled every 30 minutes, binned by year. Table 3.2 reports the first four moments on the returns. The first column of Table 3.2 demonstrates the means of each year. Due to the observation frequency (every 30 minutes)

assuming a constant risk-free rate and dividend the forward must follow the same process as the spot, thus the only difference is in initial condition.

Figure 3.1: Returns of S&P 500 index

This figure displays the returns of the S&P 500 index, using a sampling frequency of 30 minutes. The figure reports returns on years 2008 to 2016 sampling every other year.



the magnitude is expected to be approximately zero (note using daily returns measured at the close this yields mean returns between -0.24% and 0.04% which is approximately of an expected order of magnitude). However, as expected 2008 is the only year with negative mean returns due to the financial crisis towards the end of the year, with the S&P 500 index falling from a high of \$1,471 on January 02 to to \$903 on December 31 at the close with a low of \$747 on November 21. The second column depicts the standard deviations of returns. With 2008 standing out, having approximately double the magnitude of standard deviation (0.5983% or 42.82% annualised using daily returns at the close), compared to the next highest value seen in 2010 (0.2835% or 18.29% annualised), with 2012 to 2016 values ranging between 0.1733% and 0.2146% or 11.07% to 13.20% in annualised terms. The third moment (skewness) is displayed in column three. All years are negatively skewed, suggesting tail event movements of a large magnitude occur negatively. The fourth moment (kurtosis) is shown in column

four. All years demonstrate consistently high kurtosis and are leptokurtic, with a minimum value of 13.8 in 2012 and a maximum value of 19.4 in 2016. The fatness of the return tail of 2008 is only the second largest to 2016 with values of 14.4 and 19.4 respectively. All of which suggesting the presence of abnormal significant movements (inferred as kurtosis is a measure of fat tails). From knowledge of events which occurred in 2008 during the financial crisis we might expect to see significantly larger jump related parameters when calibrating stochastic volatility (SV) jump models. Figure 3.1 demonstrates although there is some market turbulence towards the beginning of the year around January and March, the first two-thirds of the year (until around October with the likes of Lehman Brothers and the Royal Bank of Scotland going bankrupt around this time) are reasonably tranquil. Thus it could be that we might not expect jump parameters to be too different from other years, especially as we calibrate to the first two-thirds of the year and leave the last third as an out-of-sample region.

Table 3.2: Summary statistics of returns on S&P 500 index

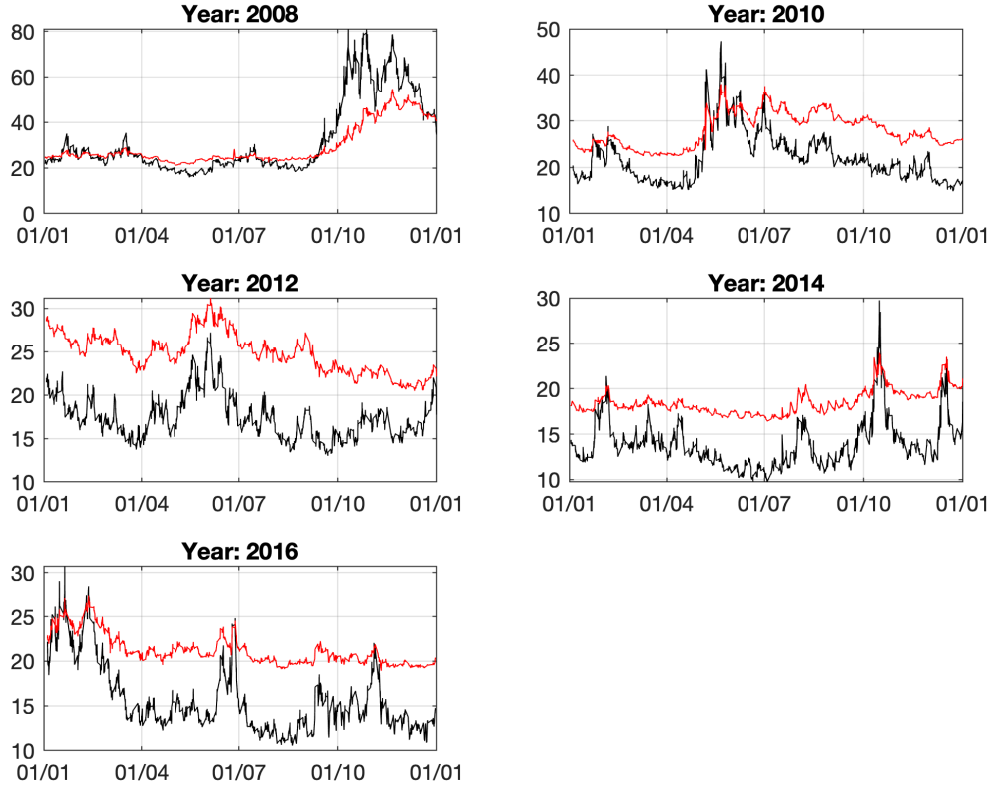
This table reports summary statistics for the returns of S&P 500, with a sampling frequency of 30 minutes. The mean and standard deviation (Std) are given as a percentage.

	Mean	Std	Skewness	Kurtosis
2008	-0.0180	0.5983	-0.3573	14.4177
2010	0.0022	0.2835	-0.5807	13.7791
2012	0.0026	0.2046	-0.4352	13.7774
2014	0.0029	0.1733	-0.1378	12.0321
2016	0.0028	0.2146	-0.8893	19.3517

The results of the above VIX reconstruction time series are displayed in Figure 3.2 with Panels 1 to 5 representing years 2008 to 2016. Table 3.3 reports the first four moments of the relative difference between 30 minute cross-sections of the VIX over time horizons 30 to 360 days, i.e., $\delta VIX_{t,t+\delta t}(\tau) = (VIX_{t+\delta t}(\tau) - VIX_t(\tau))/VIX_t(\tau)$. We choose to only present statistics over the time horizons of 30 and 360 days as these are the VIX time horizons we use in calibration exercises. The Mean column of Table 3.3 displays the mean difference between cross-sections in $\delta VIX_{t,t+\delta t}(\tau)$ over the two time horizons. The interpretation of the mean difference is as follows. The VIX value at the beginning of the year multiplied by the mean growth over the entire year should be approximate to the final VIX value. For example, in 2010 the initial VIX over a 30-day horizon is 20.4% and the final VIX is 16.9%. The mean change in VIX over 30-minutes is -0.00054% , thus: $20.4 \times (1 - 0.000054)^{(252 \times 14)} = 16.9\%$, where the 14 represents the number of 30-minute observations in a trading day. The relative difference is of the order $0.001 - 0.02\%$, with both time horizons for 2008 being one order of magnitude larger than the other years except the 30 day horizon for 2016. Secondly, on average, the volatility (irrespective of time horizon) increased in 2008 substantially as expected. The second column displays the standard deviations in the relative difference of the VIX. Across all years the standard deviation of the relative difference in the VIX is greater across the 30-day horizon than over the 360-day horizon, confirming our expectation of VIX term structure that over longer time horizons the VIX is smoother. The skewness is displayed in the third column, this is where we start to get a feel for the magnitude of turmoil in the market during 2008. For all years, 2008 included, the skewness over a 30 day period is positive. Notably, there is very strong positive skewness for years 2008 and 2012, with values significantly above unity, whereas the other years in the sample only possess slight positive

Figure 3.2: VIX time series (30 and 360 day horizon)

This figure displays the estimated VIX time series, as calculated by Equation (3.17), over a 30 and 360 day horizon for all years in the sample. The solid black line represents the 30 day horizon (denoted $VIX(30)$), the solid red line represents the 360 day horizon (denoted $VIX(360)$). Values are reported as a percentage.



skewness in the range: 0.19 to 0.47, suggesting the distribution is relatively symmetrical. However, over the longer time horizon of 360 days, with the exception of 2008, the skewness of $\delta VIX_{t,t+\delta t}(360)$ is negatively skewed. The fourth column displays the kurtosis, across all sample years the kurtosis is large suggesting significant movements in the volatility overall time horizons. Indicating the need for some form of discontinuous process in the variance to capture such extreme movements.

Table 3.3: Summary statistics for VIX estimates

This table reports the first four moments for difference in predicted VIX, $\delta VIX_{t,t+\delta t}(\tau)$, from time t to $t + \delta t$ where δt is 30 minutes and τ denotes the time horizon in days, either 30 or 360 days. Statistics are reported for the two time horizons (30 and 360 days) that we use in calibration exercises. The mean is presented as a percentage. We abbreviate the standard deviation to Std, which is also presented as a percentage.

		Mean	Std	Skewness	Kurtosis
2008	VIX(30)	0.0154	2.1129	1.5986	24.7550
	VIX(360)	0.0139	0.8310	0.4280	22.0175
2010	VIX(30)	-0.0054	2.2036	0.1911	30.0276
	VIX(360)	0.0017	0.6340	0.0921	18.1290
2012	VIX(30)	-0.0056	1.7966	1.6469	25.4846
	VIX(360)	-0.0071	0.6469	-0.0908	19.9859
2014	VIX(30)	0.0076	2.1512	0.4678	12.9105
	VIX(360)	0.0051	0.5094	-0.2598	11.1105
2016	VIX(30)	-0.0123	2.3870	0.1963	20.3028
	VIX(360)	-0.0029	0.6573	-0.2994	15.3451

3.2.5 Summary Statistics

Table 3.4 provides summary statistics on the implied volatility of the front at-the-money maturity option on the first observation (at 9:45 a.m.) of each day. These statistics are categorised depending on if the day contains an FOMC announcement. The first two columns present averaged estimates of FOMC against non-FOMC opening implied variances for the front ATM option for each year. With the exception of 2010 all years have higher implied variances on FOMC announcement days. To illustrate the effect of FOMC announcements on option implied variance our results imply on average approximately 3.7% of the total annualised variance occurs on FOMC announcement days.¹² FOMC implied variances are significantly higher during the 2008 recession. However, during 2008 a similar calculation shows total FOMC variance is approximate to the average level. Demonstrating during the crisis the macroeconomic effect of FOMC announcements are still relevant. Assuming volatility is constant across business days, FOMC announcements would account for $8/252 \simeq 3.2\%$

¹²For example, taking the mean results for FOMC and Non FOMC volatility, dividing the variance on FOMCs, 8×19.85^2 , by the total non FOMC variance, $8 \times 19.85^2 + 244 \times 18.25^2$ is 3.73%

Table 3.4: Summary statistics for FOMC and non-FOMC implied volatilities

This table provides summary statistics for the opening (first observation taken at 9:45 a.m.) implied volatility of the shortest maturity at-the-money option. Statistics are categorised depending on FOMC or non-FOMC days. We report volatility (Vol), skewness (Skew), kurtosis (Kurt) as well as the ratio (Var Ratio) between FOMC and non-FOMC variance. The last column (Num FOMC) provides the number of FOMC announcements in each year.

	FOMC Vol	Non- FOMC Vol	Var Ratio	FOMC Skew	Non- FOMC Skew	FOMC Kurt	Non- FOMC Kurt	Num FOMC
2008	35.43	31.90	1.11	0.56	1.46	1.72	4.33	8.00
2010	18.97	19.22	0.99	0.24	1.29	2.24	5.28	8.00
2012	15.50	15.12	1.02	0.01	0.82	1.93	3.68	8.00
2014	14.06	11.52	1.22	1.48	1.91	3.91	8.62	8.00
2016	15.27	13.51	1.13	0.49	1.32	1.92	4.23	8.00
Avg	19.85	18.25	1.09	0.55	1.36	2.34	5.23	8.00

of the total annualised variance. Thus suggesting FOMC announcements are responsible for only a second order effect of overall volatility generation. The third column ("Var Ratio") reports the ratio of average volatility on FOMC and non-FOMC days, from which the largest disparity occurs in 2014 where FOMC variance makes up approximately 4.7% of the total annualised variance. The skewness values being small for most years indicates for models with random price jumps near zero jump means. In line with Dubinsky et al. (2018) who model the impact of earnings announcements on individual equity options we assume FOMC announcements generate a discontinuity in the sample path of equity prices.¹³

¹³See Dubinsky et al. (2018) for discussion of alternative assumptions.

3.2.6 Non-Parametric Tests

In this section we test the main effects of FOMC announcements on option prices, namely: does IV increase prior to an FOMC announcement? Is the term structure of IV's downward sloping prior to the announcement? Finally, does IV decrease after the announcement? In order to test these questions it is important to choose a test which can accommodate the small data sample. The Wilcoxon signed rank test for-fills this criteria. Generally the Wilcoxon test compares two related samples to asses whether the medians differ. In this application it is used to test if the median is statistically different (using a 5% significance level) to a median of zero. Regarding the two questions of if the IVs increase prior to the announcement and decrease post announcement, as we have hypothesised the direction of the move we use a one sided version of the Wilcoxon test. Dubinsky et al. (2018) use the same test and implementation, so do Patell and Wolfson (1979,1981).¹⁴ For both the increase prior to and decrease after FOMC announcement tests we use the front maturity option and require the maturity be greater than three days post announcement. To test the IV increase prior to announcement we measure the change in ATM IV over a week preceding the announcement. However, our results are not robust to a two week measurement period. For the decrease in IV post announcement we use the change in IV from the last observation prior to the announcement and the last observation at the close of the day of the FOMC announcement. For the term structure estimate we the use the ATM options for the shortest two maturities, using the last observation prior to the announcement.

Table 3.5 reports the results of the questions raised above. We see for the increase in IVs prior to an announcement the null hypothesis is only rejected in years 2010 and 2014. Suggesting for the years: 2008, 2012 and 2016 there was no statistically significant increase in IV from a week prior to the announcement. However, when aggregated as a whole the null hypothesis is rejected, indicating there is a statistically significant increase in IV prior to the announcement. The result that three out of the five years we study do not have a statistically significant increase, examined individually, coupled with the lack

¹⁴Although the Patell and Wolfson (1979,1981) implementation is different because they use volatilities.

Table 3.5: Wilcoxon test (by calendar year)

This table presents the p -values for the Wilcoxon signed rank non-parametric test, grouped by calendar year. We use a one-sided version to test the increase in implied variance in the one week prior to an FOMC announcement (Increase Prior to FOMC), and the decrease in implied variance after the FOMC announcement (Decrease after FOMC). For the decreasing term structure of implied variance before the FOMC announcements (Term Structure on FOMC) we use the two-sided version.

	Increase Prior to FOMC	Term Structure on FOMC	Decrease after to FOMC
2008	0.11	0.95	0.89
2010	0.01	0.01	0.99
2012	0.58	0.32	0.77
2014	0.02	1.00	1.00
2016	0.66	1.00	0.95
Avg	0.00	0.95	1.00

of robustness across different measurement periods suggests perhaps the effect is only of second order importance. For the decreasing term structure the only year with a statistically significant decreasing term structure is 2010. Nor is the aggregate statistically significant either, adding evidence to our suspicion that the effect of FOMC announcements might only be of second order importance. This evidence continues when we investigate the decrease post FOMC announcement as none of the years reject the null hypothesis that the median is zero, indicating no decrease post announcement. The evidence documented in Table 3.5 does not suggest the effect of FOMC announcements is statistically overwhelmingly strong. In order to deal with this potentially weak effect we choose to use time varying jump volatilities, rather than constant as used by Dubinsky et al. (2018).

3.2.7 Characterising Anticipated FOMC Volatility

Table 3.6 provides details on individual calculations of the term structure and time series estimates of each FOMC announcement in our sample. While Tables 3.7 and 3.8 provide summary statistics of the two estimates.

Table 3.6 displays the first two maturities τ_1 and τ_2 and the associated IVs of the ATM options on the last observation before an FOMC announcement (depending on the time of

the announcement this is either 15 or 30 minutes prior) which are used to calculate the term structure estimate. Table 3.6 also reports the IVs of the ATM front maturity directly before and after an FOMC announcement, $(IV_{\tau_1-}, IV_{\tau_1})$, which is used to calculate the time series estimate. With the change in relevant IVs being denoted with δIV_{term} , δIV_{time} for the term structure and time series estimates, respectively. In the event where either of these changes in IVs is negative the respective estimate cannot be calculated and is set to zero. Summary statistics on the frequency of how often the estimates cannot be calculated and by what magnitude are reported in Tables 3.7 and 3.8. While the theory of considering a special case of the general stochastic volatility model with constant volatility and price jumps on FOMC announcements (see Section 3.2.1), would predict the two different estimation methodologies should yield approximate results Table 3.6 demonstrates there might not be an acceptable level of agreement. For example on 30th January 2008 (1st FOMC announcement) the term structure estimate is 4.99% while the times series estimate is close to double at 9.07%. Based on the other estimates of σ_j^Q for the other years these are by far the largest (although on 28th October 2008 the term structure estimate predicts $\sigma_{term} = 9.20\%$ but we do not include this observation in our current conversation as the time series estimate is not calculable). To offer some economic insight as to why these estimates are so large on 22nd January 2008, following an unscheduled meeting citing concerns about a weakening economy, the Federal Reserve cut both the federal funds rate and the discount rate by three-quarters of a percent. This was the first unscheduled meeting cut, only eight days before a scheduled meeting, since a cut of half a percent the day after the terrorist attacks in 2001. In conjunction with the rate cut being the largest by the Federal Reserve since October 1984. At the time many market participants felt the Federal Reserve would cut rates further in the upcoming meeting the next week causing large uncertainty in the market. Another concern about the estimation methods is the potential evidence of conflicting information. For instance on 28th October 2008 (7th FOMC announcement) the change in term structure IV estimate yields $\delta IV_{term} = 7.16\%$ while the time series is $\delta IV_{time} = -0.45\%$ implying the term structure estimate is calculable and the time series is not. The reason why this example is perhaps noteworthy is not only that one estimate is not-calculable but also the magnitude of the term structure estimate is significantly different from zero (unlike in the case on the 25th June

2008 (4th FOMC announcement) where the time series is calculable and the term structure is not). Tables 3.7 and 3.8 report summary statistics on this matter, which we explain below.

Table 3.7 demonstrates an average term structure estimate across years of 1.5%. Table 3.8 reports the time series counterpart being 1.17%, in close agreement with Chen and Clements (2007). Across years the term structure estimate mean is always greater than the median, indicating positive skewness. The same is true for the time series estimate except for 2012 where the mean is slightly less than the median (0.73, 0.91%, respectively). We also provide characterising statistics on the number and magnitude of hypothesis violation under columns Err_1 and Err_2 . Under the Err_1 columns we report the number of FOMC announcements on which the hypothesis of a decreasing term structure/implied variance after the announcement is violated, grouped by year. While Err_2 counts the number of FOMC announcements on which the violations were more than 1.5% (using implied variances). Both estimation methodologies seem to struggle with reliability. For instance in 2010 and 2012 the term structure estimate predicts violation in seven and six cases out of eight, respectively. While for these years the time series estimator is more reliable, when aggregated across years the results for violations are 3 and 3.2 for the time series and term structure estimators. Thus, implying there is approximately a 38% ($3/8 = 37.5\%$) probability of the hypothesis of a decreasing term structure/implied variance after the announcement being violated. The number of error dates are most likely caused by low FOMC volatility. In order to compare to a normal, non-announcement period, Table 3.9 displays summary statistics on the 1-hour variability in implied variance (IV). We choose a 1-hour interval as when calculating the time series estimator we use information in the observations before and after the announcement which is a 45 or 60 minute interval, depending on the announcement time. Unsurprisingly, the highest average variability is in 2008 due to the financial crisis. However, even with this, the 1-hour variability in IV is between 0.45% and 0.78%, with a mean of 0.55%. The range of IV difference in the 1-hour window surrounding FOMC announcements varies between 0.06% and 2.14%, with a mean of 0.93%. Thus, the 1-hour IV difference is approximately twice as variable over the FOMC announcement as a normal window.

Next we discuss the error occurrences in Tables 3.7 and 3.8 in more detail. The main goal is to provide some theoretical explanation behind the idea that errors can be linked to

stochastic volatility and other market-microstructure effects. To understand the effects of stochastic volatility on the estimators we will first begin with the term structure estimator. From Equation (3.10) a low level of stochastic volatility would cause σ_{t,T_1}^2 to be small, but due to the low level there would be a large amount of mean reversion of variance causing σ_{t,T_2}^2 to be larger, producing a downward sloping estimate. For the time series estimator the main cause would likely come from the level of volatility-of-volatility as higher volatility-of-volatility adds further noise to the change of IV on an FOMC announcement making detection of an FOMC jump harder. Another contributing factor to the occurrence of errors is the variance ratio between FOMC and non-FOMC days reported in Table 3.4. Theory suggests a proportional relationship between years with variance ratios close to unity and number of error occurrences, because an FOMC jump is much easier to identify if the jump size standard deviation is large relative to the average day-to-day variation in returns.

A further effect to mention which plays a role on the signal-to-noise ratio is the maturity of the option. The signal-to-noise ratio is inversely proportional to maturity as for short tenor options the annualised return variance will be significantly impacted by the FOMC jump volatility. In summary, more errors are expected for years with low variance ratios and longer front maturity options. This realisation gives evidence to why the term structure estimator might struggle so much in years 2010 and 2012. From Table 3.7 for these years the errors are 7 and 6 respectively. Table 3.4 reports these years have notably low variance ratios of 0.99, 1.02 when the mean across years is 1.09. Further Table 3.6 reports 2010 has the longest average front maturity options of 16 and 44 days.

The two estimates do seem to approximately capture the same information, using a correlation of the mean estimates per year the results is 60%.¹⁵ In summary, the term structure and time series estimates vary between 0% and 9.20%, 0% and 9.07%, respectively, with averages of 1.5% and 1.17%. Chen and Clements (2007) find evidence to suggest that the VIX falls by approximately 2% on days of FOMC announcements. Thus the term structure and time series estimates predict a slightly smaller decline, on average, but are still in agreement. In general we find the FOMC announcements seem to have a lesser impact on S&P 500 index options than on the actual index itself (see Lucca and Moench (2015)).

¹⁵Using the Spearman's rho for robustness against outliers. The result for Pearson's correlation is very

Table 3.6: Individual and average uncertainty estimates

This table provides individual and average estimates of σ_j^Q for both the term structure and time series estimates (reported as $\sigma_{term}, \sigma_{time}$, respectively). The table also displays the constituent parts necessary for these calculations. The front two maturities (τ_1 and τ_2) used for the term structure estimate, also the IV of the first two maturity at-the-money options prior to the announcement (IV_{τ_1-} and IV_{τ_2-}), these are used for the term structure estimates. For the time series estimates we report the IV of the front ATM option immediately prior to and after the announcement (IV_{τ_1-} and IV_{τ_1}). Maturities are given in days and IV's and FOMC jump estimates are given as percentages.

	τ_1	IV_{τ_1-}	IV_{τ_1}	τ_2	IV_{τ_2}	δIV_{term}	δIV_{time}	σ_{term}	σ_{time}
2008	16.43	29.05	10.79	51.43	25.28	3.77	18.26	4.99	9.07
	3.47	27.65	28.24	31.47	26.17	1.48	-0.59	1.26	0.00
	16.47	18.43	18.71	51.47	18.95	-0.53	-0.28	0.00	0.00
	23.47	19.53	19.52	51.47	20.27	-0.74	0.01	0.00	0.24
	10.47	21.02	20.97	45.47	20.91	0.11	0.05	0.63	0.38
	3.47	41.75	42.77	31.47	29.70	12.05	-1.02	3.59	0.00
	23.47	61.63	62.07	51.47	54.46	7.16	-0.45	9.20	0.00
	3.45	55.46	54.30	31.45	50.85	4.60	1.15	2.21	1.04
Avg	12.59	34.31	32.17	43.21	30.83	3.49	2.14	2.73	1.34
2010	23.47	21.32	21.49	51.47	21.73	-0.40	-0.16	0.00	0.00
	3.47	14.53	15.00	31.47	15.36	-0.84	-0.48	0.00	0.00
	23.47	18.29	16.06	51.47	18.54	-0.25	2.23	0.00	3.79
	23.47	24.05	23.75	58.47	24.42	-0.37	0.30	0.00	1.39
	10.47	20.30	20.11	38.47	20.33	-0.03	0.19	0.00	0.74
	24.47	18.44	18.95	59.47	20.19	-1.75	-0.51	0.00	0.00
	16.47	19.84	19.08	44.47	19.68	0.16	0.76	1.08	1.85
	3.43	14.46	14.22	17.43	15.02	-0.56	0.23	0.00	0.47
Avg	16.09	18.90	18.58	44.09	19.41	-0.50	0.32	0.14	1.03
2012	23.39	14.77	14.31	51.39	16.64	-1.87	0.45	0.00	1.71
	3.49	11.49	12.14	17.49	13.00	-1.50	-0.65	0.00	0.00
	23.43	14.33	14.59	51.43	15.39	-1.07	-0.26	0.00	0.00
	9.53	18.37	19.21	16.53	16.54	1.83	-0.84	3.36	0.00
	9.53	18.38	18.06	16.53	16.50	1.88	0.32	3.41	0.91
	8.47	12.47	11.68	15.47	13.04	-0.57	0.79	0.00	1.35
	9.53	16.96	16.59	16.53	17.35	-0.39	0.37	0.00	0.98
	9.47	12.93	12.61	16.47	14.08	-1.15	0.33	0.00	0.92
Avg	12.10	14.96	14.90	25.23	15.32	-0.36	0.06	0.85	0.73
2014	9.43	17.17	14.69	16.43	16.06	1.11	2.48	2.59	2.53
	2.43	11.87	11.59	9.43	11.66	0.22	0.28	0.44	0.43
	9.47	11.70	11.71	16.47	11.51	0.19	-0.01	1.08	0.00
	2.47	11.87	12.12	9.47	10.01	1.86	-0.25	1.30	0.00
	9.47	11.81	11.58	16.47	11.43	0.39	0.23	1.53	0.77
	2.47	12.97	11.53	9.47	10.85	2.12	1.44	1.39	0.99
	9.47	14.25	15.47	16.47	13.44	0.81	-1.22	2.22	0.00
	2.43	21.11	17.18	9.43	17.38	3.73	3.93	1.83	1.62
Avg	5.95	14.09	13.23	12.95	12.79	1.30	0.86	1.55	0.79
2016	9.43	22.52	22.93	16.43	20.85	1.67	-0.41	3.18	0.00
	16.43	14.34	11.57	23.43	14.13	0.21	2.77	1.77	3.53
	16.43	12.45	11.24	23.43	12.36	0.08	1.21	1.12	2.33
	9.43	18.36	17.04	16.43	17.88	0.48	1.31	1.70	1.84
	16.43	10.81	9.09	23.43	10.87	-0.06	1.72	0.00	2.78
	16.43	14.79	11.53	23.43	13.96	0.83	3.25	3.54	3.83
	16.43	19.34	19.79	23.43	17.44	1.91	-0.45	5.36	0.00
	9.43	11.04	10.44	16.43	10.72	0.32	0.61	1.40	1.25
Avg	13.80	15.45	14.20	20.80	14.78	0.68	1.25	2.26	1.95

Table 3.7: Anticipated uncertainty (term structure estimator, by year)

Table provides average estimates of anticipated uncertainty σ_j^Q using the term structure estimator. We report summary statistics overall years 2008-2016 for all FOMC announcements pooled by calendar year. We report the mean, (Mean), median, (Median), the standard error (Std Err), and the lower and upper quantile (25% and 75%) of all observations without errors. Err₁ counts the number of FOMC announcements on which the hypothesis of a decreasing term structure is violated. Err₂ counts the number of FOMC announcements on which the violations were more than 1.5% (using implied volatilities).

	Mean	Median	Std Err	25%	75%	Err ₁	Err ₂
2008	2.73	1.73	3.15	0.31	4.29	2.00	0.00
2010	0.14	0.00	0.38	0.00	0.00	7.00	1.00
2012	0.85	0.00	1.57	0.00	1.68	6.00	2.00
2014	1.55	1.46	0.67	1.19	2.02	0.00	0.00
2016	2.26	1.74	1.68	1.26	3.36	1.00	0.00
Avg	1.50	0.99	1.49	0.55	2.27	3.20	0.60

Table 3.8: Anticipated uncertainty (time series estimator, by year)

This table provides average estimates of anticipated uncertainty σ_j^Q using the time series estimator. We report summary statistics overall years 2008-2016 for all FOMC announcements pooled by calendar year. We report the mean, (Mean), median, (Median), the standard error (Std Err), and the lower and upper quantile (25% and 75%) of all observations without errors. Err₁ counts the number of FOMC announcements on which the hypothesis of decreasing implied volatility after the announcement is violated. Err₂ counts the number of FOMC announcements on which the violations were more than 1.5% (using implied volatilities).

	Mean	Median	Std Err	25%	75%	Err ₁	Err ₂
2008	1.34	0.12	3.14	0.00	0.71	4.00	0.00
2010	1.03	0.60	1.31	0.00	1.62	3.00	0.00
2012	0.73	0.91	0.66	0.00	1.17	3.00	0.00
2014	0.79	0.60	0.91	0.00	1.30	3.00	0.00
2016	1.95	2.09	1.46	0.63	3.16	2.00	0.00
Avg	1.17	0.87	1.50	0.13	1.59	3.00	0.00

3.2.8 Option Pricing Implications

We consider a number of stochastic volatility models augmented with deterministic jumps on FOMC announcements: a diffusive stochastic volatility model SV, which augmented with FOMC jumps denoted as SVD. SVJ is the model with randomly timed jumps, with a constant intensity in the price process and SVJD is the SVJ model augmented with FOMC

Table 3.9: Average variability of front maturity ATM implied variance

This table provides summary statistics on average 1-hour variability on the implied variance of the at-the-money shortest maturity option, excluding 1-hour periods which contain an FOMC announcement. The Mean column displays the mean average percentage difference. The following columns give the i -th percentile.

Year	Mean	5%	25%	50%	75%	95%
2008	0.78	0.04	0.19	0.45	0.93	3.03
2010	0.63	0.03	0.15	0.37	0.77	2.30
2012	0.45	0.02	0.12	0.28	0.58	1.42
2014	0.45	0.02	0.10	0.28	0.58	1.54
2016	0.47	0.02	0.12	0.31	0.57	1.54
Avg	0.55	0.03	0.14	0.34	0.68	1.97

jumps. SVCJ2 is the two-factor volatility model with randomly timed contemporaneous Poisson jumps in the price and variance process, SVCJ2D denotes the FOMC deterministic counterpart to the SVCJ2 model.¹⁶ Essentially we consider three models and their counterparts with the addition of deterministic FOMC announcement jumps. The calibration methodology is as described in Section 3.2.2, thus we initially calibrate the base models (SV, SVJ and SVCJ2). For the second step of the calibration we fix the base model parameters from step one and add the FOMC jump volatilities, σ_j^Q , to the parameter vector. Consider FOMC jump volatilities σ_j^Q are not constant over time, there is a vector of σ_j^Q 's, one for each FOMC announcement. Hence, each deterministic jump model requires the addition of eight parameters (one for each FOMC announcement) to the original base model. Note, as there are FOMC announcements throughout the year we no longer have an out-of-sample period as we consider all announcements in the second calibration step.

Table 3.10 presents the parameter estimates for each of the models: SV, SVJ, SV2 and SVCJ2. As far as parameter error estimates are concerned we use an asymptotic standard errors for maximum likelihood estimates approach, as in Hamilton (1994).¹⁷ Overall we find significant evidence for randomly timed jumps in the price and variance processes and multiple stochastic volatility factors. Intuitively jumps in returns can generate significant moves

¹⁶We also calibrated the SV2D model. However, we exclude these results, for brevity, as they are qualitatively no different.

¹⁷As a side issue, in order to make sure our code is correct for the more complex two-factor models we undertake a Monte Carlo exercise as is done in Bates (2006) and Fulop et al. (2014). Concluding the numerical methods employed produce very low errors ($< 0.17\%$) for all maturities used, see Appendix 6.9 for further details.

in the price, such as the crash of 1987 (evidence as to the need of price jumps provided by Bates (1991)). The volatility of returns, without jumps, is driven by a Brownian motion, thus it can only increase gradually via a sequence of normally distributed increments. Evidence from Bates (2000) and Duffie et al. (2000) has shown empirically that the conditional volatility of returns can increase rapidly. Eraker et al. (2003) demonstrate that jumps in volatility allow the volatility of returns to increase suddenly. The need for multiple volatility factors is due to the relationship between the level and slope of the volatility smile being independent. Thus, there are days with low volatility and a steep slope and also days with a flat slope but high volatility. A single-factor volatility model can generate either a steep or flat slope at a fixed volatility level, but cannot handle varying both for a fixed set of structural parameters. The introduction of a second volatility factor alleviates this restriction. This finding is consistent with previous research (Bardgett et al. (2019), Christoffersen et al. (2009) and Kaeck and Alexander (2012) to name a few). Table 3.10 also presents model errors, measured by the VWRMSE (Err.). For instance, in 2008 the in-sample VWRMSE of the standard SV model is 1.99%, with the addition of randomly timed price jumps the VWRMSE drops to 1.67%, a reduction of approximately 16% relative to the SV model. The addition of a second variance process modelling the central tendency yields a VWRMSE of 1.58% which is a fall of approximately 20%. Indicative of a significant role for randomly timed jumps and a larger role of multiple stochastic volatility factors. This trend continues to hold for all years in our sample. In terms of structural parameters, the estimates of κ_v^Q are fairly similar between models. For instance the SV estimates are between 2.43 and 5.73, SVJ estimates are approximately equal. There is some variation in 2008 where the estimate of κ_v^Q is lower for SVJ than SV and in 2014 where the SVJ estimate is higher. However, the size of variations in the rate of mean reversion will not significantly affect the pricing of options. Estimates from the multi-factor volatility SV2 and SVCJ2 models are mostly similar to the SV and SVJ counterparts. Note for 2010, 2012 and 2016 the SV2 and SVCJ2 estimated rate of mean reversion to the central tendency is higher than the SV. This is not entirely out of line with theoretical expectation, for all three years there is multiple periods of significantly high volatility which in turn effects the level of the central tendency. Due to the SV2 and SVCJ2 models varying the level of mean reversion a higher rate can be

acceptable as the model struggles less to fit both high and low volatility periods (something which can only be done in one-factor volatility models with a low rate of mean reversion which explains the lower $\kappa_v^{\mathbb{Q}}$'s of years 2010, 2012 and 2016). The estimates for $\theta_m^{\mathbb{Q}}$ provide plausible long-run volatility averages ranging from 16.73% in 2014 to 36.06% in 2010. However, theoretically in the one and two-factor models the $\theta_m^{\mathbb{Q}}$ does not quite represent the same quantity, but the values are reasonably similar. The values for σ_v are mostly determined from the time-series of the volatility. Estimates of σ_v are similar across models for each year. Although there is some contradiction to expectations regarding the inter-model comparison of σ_v . Our expectation is that the SVJ and SVCJ2 σ_v is less than predicted by the SV model, i.e., $\sigma_v^{\text{SVJ,SVCJ2}} < \sigma_v^{\text{SV}}$, our reasoning is as follows. For the SV model the volatility is purely determined by its diffusive contribution while the SVJ and SVCJ2 models have, in addition, the contribution from the randomly timed jumps.¹⁸ This is in line with Eraker et al. (2003). This relationship is broken for the SVJ model in years 2012 and 2016. However, the relationship holds for the SVCJ2 model for all years, showing greater reliance. However, the estimates are generally in line with prior research (see discussion in Broadie et al. (2007) who mention option based estimates tend to have large $\sigma_v^{\mathbb{Q}}$ and small ρ).

The parameters relating to the two-factor volatility models (SV2 and SVCJ2) are $\kappa_m^{\mathbb{Q}}$ and σ_m . These parameters are mostly determined by long maturity options. For the SV2 model, the rate of mean reversion in the central tendency, $\kappa_m^{\mathbb{Q}}$ are between 0.003 and 0.320. However, estimates for 2008 and 2012 are quite low at 0.003 and 0.005, respectively. Estimates for the SVCJ2 model are lower (with the exception of 2008 and 2012) with estimates ranging from 0.001 to 0.275, although years 2008, 2010 and 2014 represent levels of mean reversion only effective for option tenors much longer than included in our sample. These parameters are estimated under the risk-neutral measure \mathbb{Q} , from the discussion in Bardgett et al. (2019) who find using S&P 500 returns, VIX index levels and S&P 500 options (similar to our data set, only at a lower frequency) the two-factor model has a lower rate of mean reversion in the historical measure \mathbb{P} , i.e., $\kappa_m^{\mathbb{Q}} > \kappa_m^{\mathbb{P}}$. Thus, it appears under neither measure is there a significant amount, if any, mean reversion in the central tendency. Although for many years the two-factor models seem to mean revert, if at all, over very long time spans a two-factor

¹⁸In the jump models the conditional volatility of returns is: $V_t + \lambda((\mu_y^{\mathbb{Q}})^2 + (\sigma_y^{\mathbb{Q}})^2)$.

model with only one variance process mean reverting is still perfectly valid. For instance, Christoffersen et al. (2009) find a two-factor model with only one mean reverting process performs better out of sample. In summary, we feel small κ_m^Q is not a concern provided σ_m^Q is not also approximately zero. In such a case if both κ_m^Q and σ_m^Q are approximately zero then the two-factor model essentially collapses to a one-factor model (this happens only once in our sample when estimating the SVCJ2 model in 2014). The estimates for σ_m range between 0.025 and 0.753 for the SV2 model and estimates range between 0.002 and 0.803 for the SVCJ2 model. With the exception of 2014 the estimates for both models are similar. Although some potentially low estimates for these parameters are not unsurprising as suggested from our Monte Carlo calibration exercise in Appendix 6.9 displayed in Table 6.3, which implies κ_m^Q and σ_m are difficult to estimate, with standard deviations of 0.060 and 0.087, respectively. Even given some of the theoretically low values for κ_m^Q and σ_m we find they are in line with previous research, Bardgett et al. (2019) find κ_m^Q to be between 0.383 and 0.905, and σ_m between 0.113 and 0.206. Thus, our estimates are below for κ_m^Q and slightly above for σ_m . Besides, the estimates for σ_m being lower than the σ_v estimate is consistent with expectation as the central tendency process should be smoother than that of the variance. The parameters relating to the randomly timed price jumps in the SVJ model imply between 3 to 17 jumps a year (approximately), with average jump sizes close to zero between -0.006% and -0.13% and jump size volatilities of 1.43% to 3.55% . For the SVCJ2 model there are significantly less jumps per year between 0.07 to 3 jumps per year, however, the average price jump size ranges from -19.3% to 3.13% with jump size volatility between 0.5% and 12.2% . For all years except 2014 we see a smaller jump size volatility for the SVCJ2 model than the SVJ.¹⁹ The literature, from calibration to option prices, is a little inconclusive when it comes to estimating the number of jumps per year. For instance, Bakshi et al. (1997) expect less than one jump a year, however the jump in returns is of the order of -5% . On the other hand Duan and Yeh (2010) estimate more than 100 jumps per year with jump size of approximately 0.3% . Thus, previous research supports either large rare jump

¹⁹Although direct comparison between the number and size of jumps between the SVJ and SVCJ2 models is not entirely valid as the SVCJ2 model contains an extra variance factor. It could be due to the misspecification of the SVJ model the jumps in this model are picking up information which is somehow contained in the additional variance factor.

events or small and frequent, intuitively either makes sense, when considering the overall contribution from jumps. Given this wide range our estimates are in line with the literature. The mean volatility jump size estimates range from 0.010 to 0.093, in line with previous research, for instance Bardgett et al. (2019) estimates the mean volatility jump size to be between 0.010 and 0.320. As far as out-of-sample error estimates are concerned we observe an expected rise, compared to the in-sample counterparts. However, as anticipated the more complex SVCJ2 model still out performs the other models out-of-sample. In summary our estimates provide economically plausible parameters and the change to high-frequency observations does not seem to significantly affect the parameter estimates of well-established models.

Table 3.10: Parameter estimates (full maturity data set)

This table reports the parameter estimates of the SV, SVJ, SV2 and SVCJ2 models. The models are estimated with the VWRMSE objective function defined in Equation (3.15). Latent volatility estimates are computed using an extended approach of Duan and Yeh (2010) given in Appendix 6.8. The parameters are estimated on the full maturity data set $4 \leq \text{DTM} \leq 550$. Models are calibrated to an in-sample period January to August, each year. The parenthesis report asymptotic standard errors. The Err. column denotes the model in-sample Vega-Weighted RMSE (VWRMSE). The OS-Err. column denotes the model out-of-sample VRMSE.

Year	Model	κ_v^Q	θ^Q	σ_v	ρ	κ_m^Q	σ_m	λ	$\mu_y^Q(\%)$	$\sigma_y^Q(\%)$	μ_v^Q	Err.	OS-Err.
2008	SV	5.059 (0.006)	0.065 (0.000)	0.892 (0.001)	-0.776 (0.001)							1.989	2.581
	SVJ	2.000 (0.000)	0.075 (0.000)	0.852 (0.000)	-1.000 (0.000)			14.191 (0.001)	-0.130 (0.000)	2.200 (0.000)		1.675	1.923
	SV2	2.001 (0.006)	0.116 (0.000)	0.862 (0.001)	-0.847 (0.001)	0.003 (0.004)	0.753 (0.002)					1.583	1.872
	SVCJ2	2.208 (0.011)	0.093 (0.000)	0.857 (0.002)	-0.869 (0.002)	0.004 (0.007)	0.803 (0.005)	3.375 (0.387)	0.562 (0.264)	0.500 (0.695)	0.010 (0.001)	1.575	1.847
2010	SV	2.937 (0.006)	0.104 (0.000)	1.489 (0.003)	-0.674 (0.001)							1.773	1.919
	SVJ	2.030 (0.002)	0.113 (0.000)	1.026 (0.000)	-0.886 (0.000)			9.741 (0.014)	-0.438 (0.000)	3.554 (0.000)		1.603	1.825
	SV2	6.062 (0.007)	0.082 (0.000)	1.712 (0.002)	-0.702 (0.001)	0.214 (0.003)	0.243 (0.002)					1.442	1.738
	SVCJ2	4.233 (0.005)	0.083 (0.000)	1.303 (0.003)	-0.703 (0.001)	0.001 (0.001)	0.218 (0.005)	2.875 (0.021)	-6.078 (0.018)	0.500 (0.153)	0.093 (0.001)	1.380	1.681
2012	SV	2.425 (0.002)	0.087 (0.000)	0.981 (0.001)	-0.792 (0.000)							1.044	1.579
	SVJ	2.518 (0.000)	0.088 (0.000)	1.155 (0.000)	-0.941 (0.000)			16.639 (0.000)	-0.048 (0.000)	2.022 (0.000)		0.915	1.527
	SV2	2.460 (0.001)	0.100 (0.000)	1.138 (0.001)	-0.732 (0.000)	0.005 (0.003)	0.275 (0.001)					0.900	1.416
	SVCJ2	4.201 (0.002)	0.052 (0.000)	0.901 (0.001)	-0.790 (0.000)	0.275 (0.005)	0.110 (0.002)	0.266 (0.001)	3.125 (0.201)	0.500 (0.038)	0.0106 (0.004)	0.893	1.333
2014	SV	5.728 (0.008)	0.032 (0.000)	1.082 (0.001)	-0.683 (0.000)							1.103	1.782
	SVJ	8.879 (0.015)	0.028 (0.000)	1.021 (0.001)	-0.722 (0.001)			2.727 (0.040)	-0.567 (0.000)	2.499 (0.000)		1.008	1.491
	SV2	5.591 (0.000)	0.033 (0.000)	1.288 (0.001)	-0.670 (0.000)	0.156 (0.000)	0.025 (0.006)					1.003	1.468
	SVCJ2	3.498 (0.001)	0.031 (0.000)	0.788 (0.001)	-0.670 (0.000)	0.004 (0.001)	0.002 (0.097)	0.070 (0.000)	-19.312 (0.021)	12.219 (0.018)	0.0102 (0.001)	1.000	1.424
2016	SV	3.615 (0.002)	0.057 (0.000)	1.198 (0.000)	-0.767 (0.000)							1.126	1.372
	SVJ	3.184 (0.001)	0.058 (0.000)	1.200 (0.000)	-0.843 (0.000)			9.837 (0.002)	-0.006 (0.000)	1.430 (0.000)		1.099	1.315
	SV2	4.391 (0.000)	0.052 (0.000)	1.253 (0.001)	-0.765 (0.000)	0.320 (0.001)	0.091 (0.002)					1.080	1.337
	SVCJ2	4.411 (0.000)	0.053 (0.000)	1.055 (0.001)	-0.760 (0.000)	0.251 (0.001)	0.131 (0.003)	2.615 (0.005)	-0.130 (0.039)	0.520 (0.037)	0.011 (0.002)	1.050	1.269

Table 3.11 presents the calibration results of the FOMC jump augmented models, using the two-step calibration procedure outlined in Section 3.2.2 (as such we use the original structural parameters from Table 3.10). The table presents calibrations to four different maturity categories: short ($4 \leq \text{DTM} \leq 30$), medium ($31 \leq \text{DTM} \leq 90$) and long ($91 \leq \text{DTM} \leq 550$) and all ($4 \leq \text{DTM} \leq 550$). As a general comment there is not very strong agreement between models on the exact value of σ_j^Q and holding the model constant σ_j^Q still varies between the maturity categories, as expected. The most notable FOMC volatility jump is seen on the 7th announcement of 2010, which correspond to the November, 3rd FOMC announcement. It was after this announcement that the FOMC made public its intention to purchase up to \$600 billion of Treasury securities by the end of the second quarter 2011. This announcement stands out as the market appeared to price in such a purchasing program in the weeks prior to the announcement this is thought to be due to comments made by, then chairman Bernanke, at the Kansas City Fed's Jackson Hole conference (see Fuhrer and Olivei (2011)). However, the exact magnitude and timing surrounding the purchasing program could only be guessed prior to the announcement, still maintaining the importance of the announcement information. Our model picks up this event information with varying volatility jumps ranging from 0.69% to 5% depending on the model and maturity category. The estimates from Table 3.6 on the FOMC jump volatility, using the term structure and time series estimation methods, give an FOMC jump volatility of between 1.08% and 1.85%, which is closely in line with that predicted by our model which gives estimates of 1.42, 1.34, 1% for the SVD, SVJD and SVCJ2D models, respectively (it is only really applicable to compare the estimates of Table 3.6 to the short maturity category as the term structure and time series estimates use information only from short maturity options). In the case of the SVCJ2D model, for 2010, we see error reductions of between 8% and 33% for including FOMC jumps in the base SVCJ2 model, which suggests at least for some years and particular announcements the effect is very strong. In this example the most significant error reduction is seen for the long maturity category (and least significant in the short category), this is likely explained as these are the options which contain the effects of each σ_j^Q for longest and will also be influenced by multiple FOMC announcements, thus making the aggregated effect greater. The phenomenon of the longer maturity category displaying greater error reduction holds

in years 2008, 2010 and 2014. However, in 2012 all values are approximately the same, with both the short and medium categories representing a decrease of 1%, while the long maturity category has no reduction (this is likely down to most jump volatilities being very small this year). Chen and Clements (2007) estimate a 2% drop in the VIX on the day of an FOMC announcement (there is no overlap with our sample, Chen and Clements use data from 1996 to 2006). For the SVCJ2D model²⁰ the average (denoted Avg.) rows from Table 3.11, for all option maturities, demonstrate the FOMC jump volatilities range from 0.08% to 1.51% with a mean across years of 0.47%. Although, our results estimate a smaller drop than Chen and Clements (2007) for 2010 where there were some significant improvements made by the inclusion of FOMC deterministic jumps we see an average jump of 1.51% which is of an approximate order to that of Chen and Clements (2007). By way of model improvement the short category has error reductions of: 0% to 8%, the medium of: 0% to 23%, the long of: 0% to 33% and the all maturity category of: 0% to 20%. To put the results from Table 3.11 into perspective we find that the addition of FOMC jumps in the all maturity category can lead to significant improvements, as meaningful as the addition of a second volatility factor (in comparison with the results from Table 3.10, where the addition of a second volatility factor is between 4% and 20%, relative to the SV model). However, the improvements from FOMC jump augmentation are very sporadic and do not produce a consistent error reduction.

Next, consider the behaviour of average option pricing errors as a measure of overall intradaily fit on announcement days. Table 3.12 provides the mean VWRMSE pricing errors (aggregated across years) for the minutes leading up to and after FOMC announcements (going from 120 minutes prior to 30 minutes post). Table 3.12 also groups errors into the four maturity categories used above. Taking consideration of FOMC announcements provides a moderate pricing improvement in the hours before an FOMC announcement, especially for the short maturity category (options with less than 30 days to maturity). Considering the short maturity category only we see the VWRMSE falls from 1.79, 1.84, 1.79, 1.78 in the SV model to 1.72, 1.76, 1.72, 1.70 in the SVD model, respectively. A similar reduction pattern is observed in the other models. In the short maturity category the ratio between the FOMC

²⁰We discuss only the SVCJ2D model as it is the most sophisticated and therefore least likely to have reduced errors due to implementation of additional parameters.

jump and non-FOMC jump models is unity from the time of the announcement on-wards. The reason for which is due to the maturity profile not spanning two announcements, thus as soon as the announcement takes place the contribution from the FOMC jump falls away and we are left with the original model. In the medium and long maturity categories the addition of the FOMC announcement jumps makes little to no improvement in the hours surrounding an announcement. In fact in some cases the FOMC jump augmentation performs worse, when evaluated over the FOMC announcement, than the original model. As these maturity categories can span multiple announcements there is multiple event information. Intuitively it makes sense that longer term options will benefit little from the FOMC jump effect in the hours leading up to the announcement as the annualised variance will not be significantly impacted by the FOMC jump volatility, unlike for short tenor options. The all maturity category demonstrates the effect of accounting for FOMC announcements considering all maturity options (between 4 and 550 days). The models still demonstrate some improvement before the announcement, approximately 1 – 2.5% for the SVD, SVJD and SVCJ2D models.

Table 3.11: Individual FOMC volatility estimates

This table reports estimates of the FOMC announcement jump volatility σ_j^Q (given in percentages) for the SVD, SVJD and SVCJ2D models. The models are estimated with the VWRMSE objective function defined in Equation (3.15). Latent volatility estimates are computed using an extended approach of Duan and Yeh (2010) given in Appendix 6.8. With the structural parameters for each model fixed at the values in Table 3.10. The jump volatilities are estimated over four different maturity categories: All (maturity in $4 \leq \text{DTM} \leq 550$), short (maturity in $4 \leq \text{DTM} \leq 30$), medium (maturity in $31 \leq \text{DTM} \leq 90$) and long (maturity in $91 \leq \text{DTM} \leq 550$). The Avg. column denotes the average σ_j^Q given for that model, year and maturity category. The Err. row denotes the error ratio between the FOMC jump augmented model and the original model.

		All			Short			Medium			Long		
Year		SVD	SVJD	SVCJ2D	SVD	SVJD	SVCJ2D	SVD	SVJD	SVCJ2D	SVD	SVJD	SVCJ2D
2008	$\sigma_{j,1}$	0.01	0.01	1.35	0.01	0.01	1.14	0.01	0.01	1.49	0.01	0.01	1.25
	$\sigma_{j,2}$	0.01	0.06	0.00	0.01	0.01	1.28	0.01	0.01	0.00	0.01	0.01	0.00
	$\sigma_{j,3}$	0.01	0.06	0.00	0.11	0.06	1.11	0.01	0.06	0.78	0.01	0.06	0.00
	$\sigma_{j,4}$	0.01	0.01	0.01	0.01	0.08	0.79	0.01	0.08	0.80	0.01	0.08	0.01
	$\sigma_{j,5}$	0.01	0.02	0.82	0.01	0.37	0.75	0.01	0.37	1.15	0.01	0.37	0.00
	$\sigma_{j,6}$	0.01	1.46	0.00	0.66	0.49	1.55	0.29	0.49	0.59	0.01	0.49	0.00
	$\sigma_{j,7}$	0.01	0.06	0.00	0.01	0.08	0.09	0.01	1.16	2.68	0.01	1.16	0.00
	$\sigma_{j,8}$	0.01	0.01	0.01	0.01	0.01	0.66	0.01	0.01	0.66	0.01	0.01	0.01
	Avg.	0.01	0.21	0.27	0.10	0.14	0.92	0.05	0.27	1.02	0.01	0.27	0.16
	Err.	1.00	1.00	0.99	1.00	1.00	0.96	1.00	1.00	0.77	1.00	0.98	1.00
2010	$\sigma_{j,1}$	0.01	0.01	0.43	0.01	0.01	0.13	0.57	0.01	0.43	0.01	0.01	0.56
	$\sigma_{j,2}$	0.01	0.01	1.07	0.76	0.55	1.21	1.27	0.55	1.05	0.01	0.01	1.20
	$\sigma_{j,3}$	0.01	0.01	0.71	0.55	0.01	0.13	0.65	0.01	0.53	0.01	0.01	1.45
	$\sigma_{j,4}$	0.01	0.01	0.32	0.01	0.01	0.01	0.01	0.01	0.11	0.01	0.01	0.43
	$\sigma_{j,5}$	0.01	0.01	0.60	0.01	0.78	0.16	1.04	1.45	0.30	0.01	0.01	1.51
	$\sigma_{j,6}$	3.23	0.01	1.03	0.00	0.01	0.00	2.73	1.79	1.50	3.87	0.30	1.07
	$\sigma_{j,7}$	1.38	1.20	5.00	1.42	1.34	1.00	2.53	1.34	5.00	0.69	4.68	5.00
	$\sigma_{j,8}$	0.01	0.46	2.97	0.46	0.40	1.24	0.85	0.40	1.24	0.01	0.01	2.00
	Avg.	0.59	0.22	1.51	0.40	0.39	0.49	1.21	0.70	1.27	0.58	0.63	1.65
	Err.	0.96	0.88	0.80	1.00	0.88	0.92	0.93	0.90	0.91	0.95	0.85	0.67
2012	$\sigma_{j,1}$	0.78	0.80	0.01	0.01	0.01	1.08	1.31	1.54	0.01	1.83	0.01	0.05
	$\sigma_{j,2}$	0.63	0.46	0.01	0.01	0.01	0.11	0.79	0.46	0.01	1.88	0.01	0.01
	$\sigma_{j,3}$	0.91	0.01	0.01	0.66	0.01	0.52	1.02	0.04	0.29	0.01	0.69	0.92
	$\sigma_{j,4}$	0.01	0.84	0.00	0.99	1.18	0.67	0.01	0.26	0.00	0.01	1.18	0.00
	$\sigma_{j,5}$	0.96	0.77	0.56	0.99	0.96	0.50	0.01	0.84	0.82	0.01	0.96	0.01
	$\sigma_{j,6}$	0.01	0.01	0.01	0.01	0.01	0.82	1.16	0.01	0.01	0.78	0.01	0.01
	$\sigma_{j,7}$	0.01	0.01	0.01	0.01	0.01	0.73	0.77	0.01	0.60	0.84	0.01	0.01
	$\sigma_{j,8}$	0.01	0.01	0.00	0.01	0.01	0.79	0.01	0.01	0.79	0.01	0.01	0.00
	Avg.	0.42	0.37	0.08	0.34	0.28	0.65	0.64	0.40	0.32	0.67	0.36	0.13
	Err.	0.98	1.00	1.00	0.97	1.00	0.99	0.94	1.00	0.99	0.95	1.00	1.00
2014	$\sigma_{j,1}$	0.26	0.46	0.29	0.29	0.55	0.01	0.01	0.01	0.62	0.48	0.23	0.72
	$\sigma_{j,2}$	0.01	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.21	0.29	0.44	0.44
	$\sigma_{j,3}$	0.54	0.71	0.01	0.68	0.83	0.01	0.01	0.01	0.44	0.79	1.20	0.01
	$\sigma_{j,4}$	0.44	0.42	0.38	0.51	0.66	0.01	0.48	0.01	0.68	0.24	0.01	0.01
	$\sigma_{j,5}$	0.01	0.02	0.37	0.01	0.02	0.01	0.01	0.01	0.77	0.08	0.01	0.95
	$\sigma_{j,6}$	0.01	0.04	0.01	0.01	0.01	0.10	0.01	0.01	0.15	0.01	0.01	0.01
	$\sigma_{j,7}$	0.01	0.08	0.01	0.22	0.22	0.02	0.01	0.01	0.33	0.01	0.01	0.01
	$\sigma_{j,8}$	1.02	0.26	0.01	1.00	0.01	0.07	0.02	0.06	0.07	1.01	1.55	0.01
	Avg.	0.29	0.25	0.14	0.34	0.29	0.03	0.07	0.02	0.41	0.37	0.43	0.27
	Err.	1.00	1.00	1.00	1.00	1.00	1.00	1.00	0.99	0.96	0.98	1.00	0.94
2016	$\sigma_{j,1}$	0.74	1.44	0.29	0.74	1.44	0.29	0.74	0.04	0.29	0.74	0.74	0.29
	$\sigma_{j,2}$	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.39	0.00	0.00	0.00	0.00
	$\sigma_{j,3}$	0.83	0.83	0.68	0.83	0.83	0.68	0.83	0.01	0.68	0.83	0.83	0.68
	$\sigma_{j,4}$	0.00	0.00	0.51	0.00	0.00	0.51	0.00	0.01	0.51	0.00	0.00	0.51
	$\sigma_{j,5}$	0.00	0.00	0.01	0.00	0.00	0.01	0.00	0.01	0.01	0.00	0.00	0.01
	$\sigma_{j,6}$	0.00	0.00	0.01	0.00	0.00	0.01	0.00	4.99	0.01	0.00	0.00	0.01
	$\sigma_{j,7}$	0.00	0.00	0.22	0.00	0.00	0.22	0.00	2.12	0.22	0.00	0.00	0.22
	$\sigma_{j,8}$	0.00	0.00	1.00	0.00	0.00	1.00	0.00	0.46	1.00	0.00	0.00	1.00
	Avg.	0.20	0.28	0.34	0.20	0.28	0.34	0.20	1.00	0.34	0.20	0.20	0.34
	Err.	1.00	1.00	1.00	0.99	1.00	1.00	0.99	1.00	1.00	1.00	1.00	1.00

Table 3.12: Average pricing errors around FOMC announcements

This table reports the mean Vega-Weighted RMSE (VWRMSE) pricing errors for the SV, SVD, SVJ, SVJD, SVCJ2 and SVCJ2D models, aggregated over all years in the sample. Pricing errors are grouped into four different maturity categories all (maturity in $4 \leq \text{DTM} \leq 550$), short (maturity in $4 \leq \text{DTM} \leq 30$), med (maturity in $31 \leq \text{DTM} \leq 90$) and long (maturity in $91 \leq \text{DTM} \leq 550$). The columns labelled -120 to +30 provide the mean VWRMSE for options from -120 to +30 minutes from an FOMC announcement, all VWRMSE values are displayed as a percentage.

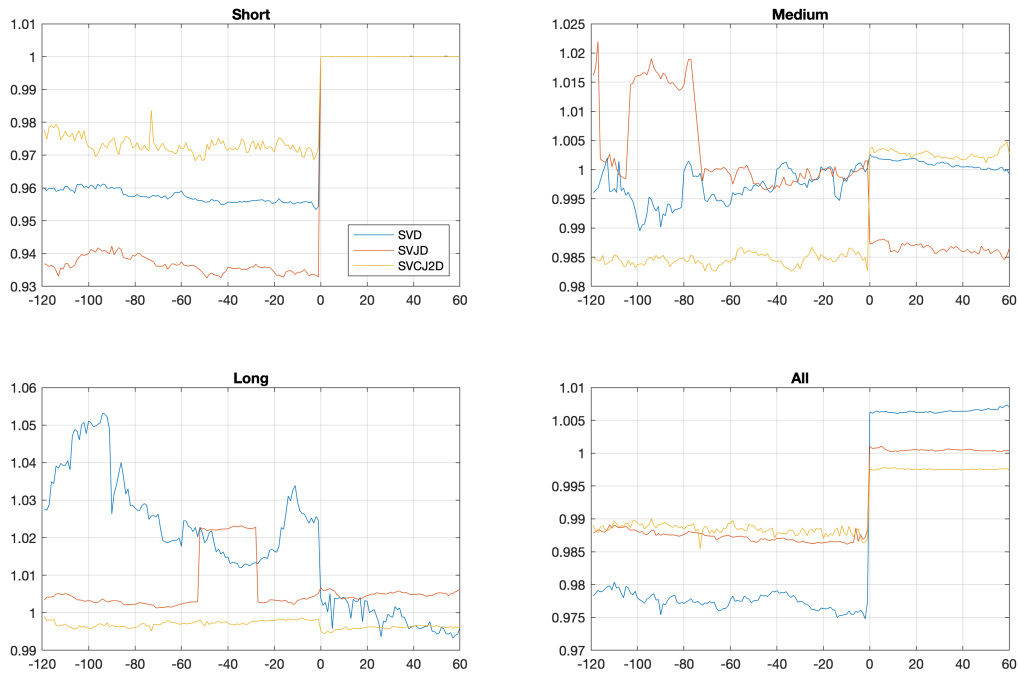
Cat.	Model	-120	-60	-30	-5	0	30
short	SV	1.79	1.84	1.79	1.78	1.71	1.75
	SVD	1.72	1.76	1.72	1.70	1.71	1.75
	SVJ	1.60	1.64	1.60	1.58	1.52	1.56
	SVJD	1.50	1.54	1.50	1.47	1.52	1.56
	SVCJ2	1.54	1.59	1.54	1.54	1.56	1.56
	SVCJ2D	1.50	1.55	1.50	1.50	1.56	1.56
med	SV	0.78	0.80	0.77	0.77	0.81	0.77
	SVD	0.78	0.80	0.77	0.77	0.81	0.77
	SVJ	0.83	0.84	0.82	0.83	0.84	0.83
	SVJD	0.84	0.84	0.82	0.83	0.83	0.82
	SVCJ2	1.22	1.23	1.21	1.21	1.22	1.22
	SVCJ2D	1.20	1.21	1.19	1.19	1.22	1.22
long	SV	0.70	0.73	0.71	0.70	0.65	0.70
	SVD	0.72	0.74	0.72	0.71	0.65	0.70
	SVJ	1.15	1.16	1.17	1.17	1.16	1.16
	SVJD	1.16	1.16	1.20	1.17	1.17	1.17
	SVCJ2	1.29	1.30	1.29	1.29	1.30	1.29
	SVCJ2D	1.29	1.30	1.29	1.29	1.29	1.28
all	SV	1.17	1.20	1.18	1.17	1.19	1.16
	SVD	1.15	1.17	1.15	1.14	1.20	1.17
	SVJ	1.16	1.19	1.17	1.16	1.15	1.16
	SVJD	1.15	1.18	1.16	1.15	1.15	1.16
	SVCJ2	1.44	1.47	1.45	1.44	1.44	1.45
	SVCJ2D	1.42	1.46	1.44	1.42	1.44	1.45

However, it could be the price improvements detected in the one-factor models are simply down to misspecification of the original model (SV or SVJ in this case). Therefore, any error reduction observed could be attributed to missing information being detected by the FOMC jumps. Although, The results for the SVCJ2D model receive a similar, albeit slightly smaller, error reduction around the FOMC announcement, reducing the impact of potential suspected misspecification, as our results are reasonably robust regardless of model.

Figure 3.3 displays the mean VWRMSE ratio between each FOMC jump augmented model and the original model from two hours prior to the announcement to one hour post,

Figure 3.3: Time series of pricing errors around FOMC announcements

This figure displays Vega-Weighted RMSE (VWRMSE) pricing error ratios for the SVD, SVJD and SVCJ2D models aggregated across FOMC announcements and years. The figure also displays the error time series grouped by different maturity categories. All (maturity in $4 \leq \text{DTM} \leq 550$), short (maturity in $4 \leq \text{DTM} \leq 30$), medium (maturity in $31 \leq \text{DTM} \leq 90$) and long (maturity in $91 \leq \text{DTM} \leq 550$). Errors are observed for each FOMC day between two hours prior and one hour after the announcement, using minute frequency data. Error ratios are calculated as the VWRMSE of the FOMC augmented model against the original model. The blue line presents the ratios from the SVD related models, orange the SVJD and yellow the SVCJ2D.



using one minute frequency data. As in Table 3.12 we split the maturity categories. While there is some information overlap between Table 3.12 and Figure 3.3 two additional points are noteworthy: in the short maturity Panel for all models a downward drift can be observed in the minutes leading up to the announcement, indicating the effect of the FOMC jumps becomes more pronounced closer to the announcement. Second, in all panels after the announcement there is an instantaneous move towards an error ratio of unity. This move is due to the shortest tenor options in each panel no longer possessing a σ_j^Q , thus any effect left is due to longer maturity options for which the effect is much less significant. In the case of the short maturity panel this move causes the error ratio to collapse to unity as the FOMC jump augmented models essentially collapse back to their original models.

3.2.9 Summary

In this Section we investigate models which incorporate FOMC announcements for pricing options. We take into consideration the timing of FOMC announcements and develop a pricing approach including deterministic jumps on FOMC announcements. We explore two different types of estimators of the uncertainty surrounding FOMC announcements and document their reliability. To quantify the impact on option prices, we calibrate a two-factor volatility model, with contemporaneous jumps in the price and variance process along with deterministic FOMC jumps (SVCJ2D). Empirically we find somewhat sporadic results for the importance of the inclusion of deterministic FOMC jumps. For all option maturities in our sample the inclusion of FOMC jumps can have a significant impact on error reduction, in 2010 we see a fall of 20%, due to a few important announcements. However, the jumps do not always demonstrate a reduction in errors. The importance of FOMC jumps in the minutes surrounding an announcement are most relevant for short tenor options, where a reduction of a few percent is observed, when aggregated across years.

3.3 Volatility Seasonality Adjustments

This section attempts to gauge the importance of various puzzling factors reported in the option returns literature, focusing on the overnight and weekend seasonalities, in option pricing models. We attempt to reconcile these seasonalities into standard option pricing frameworks, via a time re-ordering distribution approach. The motivation behind considering a re-ordering of time distribution is that the time to maturity is a vital piece of information in determining the price of an option. As the time to maturity increases, more volatility will be accrued during the lifetime of the option, thus making it more expensive. Therefore, in the investigation into which parts of the week accrue more volatility the idea of altering the time weights of each part of the week, to measure such seasonalities, is intuitively appealing. The overnight seasonality is prevalent in both equity and index option contracts and holds

irrespective of moneyness or maturity category. The nature of the effect is the decomposition of option returns into day (intraday) and overnight (interday) periods. The overall returns on S&P 500 options are negative and large, however, this is comprised of positive intraday and negative interday returns. Prior research hypothesises this return asymmetry is a manifestation of option market maker's (OMMs) incorrectly pricing the volatility seasonality between day and night (it is well known the day is more volatile than the night). The weekend seasonality is defined as the lower returns of options over the weekend period, relative to any other day of the week. The effect is not attributed to the underlying security, being observed in both delta-hedged and un-hedged positions. The effect is wanted across moneyness, maturity and option type (either put or call). We wish to be able to capture these apparently puzzling seasonalities by incorporating additional factors into traditional option pricing models.

Muravyev and Ni (2019) investigate the overnight effect and show option returns are negative and large, approximately -0.7% per day. This is comprised of positive intradaily (approximately 0.3%) and negative interdaily returns (approximately -1%). Muravyev and Ni (2019) hypothesise the change of sign of option returns from day to night is a direct result of OMMs failing to price the variation of volatility from day to night, OMMs fail to take into account the volatility seasonality that the trading day is more volatile than the overnight period. In order to gauge how much more volatile the trading day is compared to the overnight period Muravyev and Ni (2019) measure the standard deviation of intra/inter daily returns, on the underlying asset, over the preceding 60 days to calculate σ_i , where the subscript i is used to denote either the trading day or overnight period. The ratio of the standard deviations $\gamma = \sigma_{day}/\sigma_{night}$ is found to range from 1.5 to 2.0. That is the data suggests the trading day is up to twice as volatile as the overnight period. However, when Muravyev and Ni (2019) introduce the notion of different volatilities in the Black-Scholes model and simulate different levels of under-reaction to the volatility seasonality they find the ratio of accrued volatility which fits the option return patterns in the data (that is the change of sign of option returns between day and night) best is the case corresponding to equal volatility accrual in both periods. This comes as somewhat of a surprise as the standard deviation in returns of the underlying suggests this is clearly not the case. From

this investigation Muravyev and Ni (2019) conclude OMMs completely ignore this volatility seasonality. Our intuition behind the empirical test to include the overnight effect into option pricing models is as follows. The quantity σ_i can be thought of as the amount of volatility accrual during period i , thus it is a product of the amount of time elapsed and the amount of volatility accrued per unit time. If the hypothesis concluding that OMMs ignore the day-night volatility seasonality is correct, we would expect a decrease in the trading period,²¹ thus increasing the overnight period. Deduced due to the following reasoning: as the trading day carries higher volatility, per unit time, than the overnight period there should be an increase in the product of overnight volatility, per unit time, and the weight of this period in order to balance the total volatility accrual during each period.

Jones and Shemesh (2018) investigate the weekly return pattern of options, with the objective of determining why options exhibit significantly lower returns over the weekend (defined by them as Friday-Monday close-close) relative to any other close-close period (coined the weekend effect). They categorise their results according to hedged/unhedged returns, on individual equities and S&P 500 options. Their results demonstrate a clear existence of the weekend effect in both hedged/unhedged returns and implied volatilities on individual equity options. However, they find unconvincing results for the effect in S&P 500 options. Jones and Shemesh (2018) also confirm French and Roll's (1986) conclusion that Friday-Monday close-close is not substantially different from any other close-close, for S&P 500 volatility. The intuition behind the overnight test is as follows. If the observations by French and Roll (1986) and Jones and Shemesh (2018) are indeed true, that the weekend period is not as volatile as suggested by the elapse of two calendar days, then we would expect a shift from the weekend time weight into the week.

3.3.1 Time Notation

To avoid confusion for what is to come we shall discuss the nomenclature. Using the regular calendar time convention a week is comprised of three unique periods: market trading (de-

²¹See Section 3.3.1 below for details regarding the time re-ordering distribution.

noted by: t_{tr}), the weekday overnight period (denoted by: t_{ov}) and the weekend (denoted by: t_{wknd}). The trading period occurs five times a week (for an average week exclusive of weekday holidays) from 9.30 a.m. to 16.15 p.m. Monday - Friday. We define the overnight period as the time between trading periods, e.g., Monday close to Tuesday open. This period occurs four times a week Monday close - Thursday close. Therefore the weekend period is defined as being Friday close to Monday open. Using this time convention the time until maturity, from time t to expiry at time T , is given by $\tau = T - t$:

$$\tau = (N_{tr} + \epsilon_t)t_{tr} + N_{ov}t_{ov} + N_{wknd}t_{wknd}, \quad (3.22)$$

where $N_i \in \mathbb{N}$ and $i \in \{tr, ov, wknd\}$ denotes the number of each unique time periods which elapse between time t and maturity at T . The ratio of time remaining during the current trading day is denoted by ϵ_t . We keep t_i to denote the time weight of period i as measured by the calendar time convention and use $t_{adj,i}$ to denote the time weight of period i which is a result of the calibrations and thus not measured by the calendar time convention. Thus, in the time adjusted notation the time to maturity is denoted by:

$$\tau_{adj} = (N_{tr} + \epsilon_{adj,t})t_{adj,tr} + N_{ov}t_{adj,ov} + N_{wknd}t_{adj,wknd}. \quad (3.23)$$

It is more intuitive to think about ratios of the time weights rather than the actual values. As such we define $R_i = t_{adj,i}/t_i$ where $i \in \{tr, ov, wknd\}$ to represent the ratio of period i . Thus, the quantity R_i reflects the ratio of change of time period i under the time-adjusted regime. The logic behind finding the ratio of adjustment easier to work with is because it gives a direct handle on the expected ratio change of volatility accrual during period i .

3.3.2 Model Framework

We investigate two standard stochastic volatility models, a single factor stochastic volatility model (SV) and a single factor model augmented with random Poisson jumps in the price process (SVJ). The idea of this work is not to demonstrate which model, augmented with

many factors produces the lowest errors, rather to demonstrate which components are now relevant to intradaily option pricing. In unreported results we investigate these effects using a two-factor volatility model and achieve qualitatively similar results. As there is no further additional insights into what biases are made by OMMs or the magnitude of correction we only display the one-factor models.

The addition of the time weights increases the parameter vector by one parameter because all the tests we implement introduce two constraint equations leaving only one parameter to be determined by calibration. One such constraint which applies irrespective of the seasonality is the total calendar time of a week (therefore also a year) being fixed. This fixed time condition, of a week, may be expressed as follows:

$$5t_{adj,tr} + 4t_{adj,ov} + t_{adj,wknd} = \frac{1}{52}. \quad (3.24)$$

The additional constraint equation is seasonality dependent and will be described in each relevant section (and are given in Equations (3.37) and (3.44)). We denote these so called time adjusted models as M_i where $M \in \{SV, SVJ\}$ and $i \in \{tr, ov, wknd\}$, denotes the unique time period which is augmented into the original parameter vector. Given the constraint Equation (3.24), this makes the new time to maturity from Equation (3.23):

$$\tau_{adj} = \frac{1}{52}N_{wks} + t_{adj,tr}(N_{days} + \epsilon_t) + t_{adj,ov}N_{days}, \quad (3.25)$$

where N_{days} denotes the number of complete calendar days from time t until the nearest Friday. This change of time to maturity affects option prices from traditional models on two fronts. Upon initial thought it may seem moving from τ to τ_{adj} requires no model alterations and we can simply use the new times to maturity to generate prices. While this is true, there is also a slightly more subtle effect which needs to be considered, the effect on the latent volatility detection. As previously discussed in Section 3.2.2 we use the VIX to determine the latent volatility in our calibrations. The VIX values we obtain as outlined in Section 3.2.3 are interpolated over integer multiples of 30 days. When calibrating one-factor models we use the 30-day VIX value to infer the latent volatility. Therefore, in the time adjusted

models a calendar time of 30 days is now slightly different (τ_{adj}). Consequently, the VIX is mapped from a calendar time $VIX_t^2(\tau)$ to a transformed time $VIX_t^2(\tau_{adj})$. As the VIX is the square root of the expected integrated variance over the next 30 days, a general statement would be the VIX over the period $\tau = T - t$, is given by:

$$VIX_t^2(\tau) = \frac{1}{T-t} \mathbb{E}_t^{\mathbb{Q}} \left[\int_t^T V_s ds \right], \quad (3.26)$$

it is then possible to find a model dependent expression for the integrated variance. For simplicity consider a one-factor diffusion model, whereby application of Itô's lemma yields:

$$d\ln(S_t) = (r - \frac{1}{2}V_t)dt + \sqrt{V_t}dW_t, \quad (3.27)$$

with r denoting the risk-free rate. Integration and taking expectation values yields:

$$\mathbb{E}_t^{\mathbb{Q}}[\ln(S_T/S_t)] = r(T-t) - \frac{1}{2} \mathbb{E}_t^{\mathbb{Q}} \left[\int_t^T V_s ds \right], \quad (3.28)$$

by use of the forward price this yields an expression for the expected variance as:

$$\frac{1}{T-t} \mathbb{E}_t^{\mathbb{Q}} \left[\int_t^T V_s ds \right] = -\frac{2}{T-t} \left(\mathbb{E}_t^{\mathbb{Q}}[\ln(S_T/S_t)] - \ln(F_t(\tau)/S_t) \right). \quad (3.29)$$

By the payoff decomposition theorem we have:

$$\mathbb{E}_t^{\mathbb{Q}}[\ln(S_T/S_t)] = \ln(x/S_t) - \frac{x - F_t(\tau)}{x} - e^{r\tau} \int_0^x \frac{P_t(K)}{K^2} dK - e^{r\tau} \int_x^\infty \frac{C_t(K)}{K^2} dK, \quad (3.30)$$

where $P_t(K)$ and $C_t(K)$ define the prices of European put and call options, respectively, with strike K at time t with maturity at time T . For the regular $VIX_t(\tau)$ calculation we choose $x = F_t(\tau)$, thus cancelling the second term. Upon substitution the first term from Equation (3.30) cancels with the second term of Equation (3.29), leaving:

$$VIX_t^2(\tau) = \frac{2}{\tau} \left(e^{r\tau} \int_0^{F_t(\tau)} \frac{P_t(K)}{K^2} dK + e^{r\tau} \int_{F_t(\tau)}^\infty \frac{C_t(K)}{K^2} dK \right). \quad (3.31)$$

For the adjusted $VIX_t(\tau_{adj})$ we have the same, only with a different forward price $F_t(\tau_{adj})$, i.e.,

$$VIX_t^2(\tau_{adj}) = \frac{-2}{\tau_{adj}} \left(\mathbb{E}_t^{\mathbb{Q}}[\ln(S_T/S_t)] - \ln(F_t(\tau_{adj})/S_t) \right). \quad (3.32)$$

Similar use of the payoff decomposition theorem yields:

$$\mathbb{E}_t^{\mathbb{Q}}[\ln(S_T/S_t)] = \ln(x/S_t) - \frac{x - F_t(\tau_{adj})}{x} - e^{r\tau_{adj}} \int_0^x \frac{P_t(K)}{K^2} dK - e^{r\tau_{adj}} \int_x^\infty \frac{C_t(K)}{K^2} dK, \quad (3.33)$$

here clearly a sensible choice would be to choose: $x = F_t(\tau_{adj})$, however, we are forced to choose: $x = F_t(\tau)$ if we wish to use the calendar time $VIX_t(\tau)$ to calculate $VIX_t(\tau_{adj})$, leaving:

$$VIX_t^2(\tau_{adj}) = \frac{-2}{\tau_{adj}} \left(\ln(F_t(\tau)/F_t(\tau_{adj})) - \frac{F_t(\tau) - F_t(\tau_{adj})}{F_t(\tau)} - e^{r\tau_{adj}} \int_0^{F_t(\tau)} \frac{P_t(K)}{K^2} dK - e^{r\tau_{adj}} \int_{F_t(\tau)}^\infty \frac{C_t(K)}{K^2} dK \right). \quad (3.34)$$

However, as τ is very similar to τ_{adj} then $F(\tau) \sim F(\tau_{adj})$, thus for all maturities the first two terms in Equation (3.34) will always be small yielding:

$$VIX_t^2(\tau_{adj}) \approx \frac{e^{r(\tau_{adj}-\tau)}\tau}{\tau_{adj}} \left[\frac{2}{\tau} \left(e^{r\tau} \int_0^{F_t(\tau)} \frac{P_t(K)}{K^2} dK + e^{r\tau} \int_{F_t(\tau)}^\infty \frac{C_t(K)}{K^2} dK \right) \right]. \quad (3.35)$$

However, the expression in the square brackets of Equation (3.35) is simply the $VIX_t^2(\tau)$ from Equation (3.31), giving the final result:

$$VIX_t^2(\tau_{adj}) \approx \frac{\tau}{\tau_{adj}} e^{r(\tau_{adj}-\tau)} VIX_t^2(\tau). \quad (3.36)$$

Over a 30 day horizon the VIX values will change very little as we impose the constraint Equation (3.24) and 30 days is only slightly over four weeks we must have $\tau_{adj} \sim \tau$.

3.3.3 Overnight Effect

This subsection investigates if the model feature of a time re-ordering distribution can be used to capture the overnight effect. During this test only the trading/overnight period becomes a parameter in the model. Imposing the condition that both the overnight and trading day periods still sum to one day, the maturity is not altered:

$$\tau_{\text{day}} + \tau_{\text{night}} = \frac{1}{7 \times 52}. \quad (3.37)$$

Making the two constraint Equations (3.24) and (3.37) which signify a week must still be a week and the overnight and day periods must still sum to one day. Inferring any increase in one time period is reflected by a decrease in the other. Therefore, for this investigation the parameter vector only grows by one additional parameter ($t_{adj,tr}$ or $t_{adj,ov}$), requiring the weekend be fixed to its calendar time: $t_{adj,wknd} = t_{wknd}$. The choice of calibrating to either $t_{adj,tr}$ or $t_{adj,ov}$ is arbitrary and we only report the results for $t_{adj,tr}$, although similar results are obtained for $t_{adj,ov}$ demonstrating the effect is robust.

In order to interpret the calibration results some bookkeeping is in order first. Define the ratio of volatility accrual in each period, the level of bias parameter as:

$$\gamma = r_{\text{vol}} \frac{t_{tr}}{t_{ov}}, \quad (3.38)$$

where the quantity r_{vol} represents the ratio of volatility, per unit time, between the day and overnight periods. Following the finding by Muravyev and Ni (2019) in a fully rational world γ is between 1.6 and 1.8. This ratio is computed using the calendar time periods, t_{day} and t_{night} , i.e., how many hours the market is/is not open. This yields:

$$r_{\text{vol}} = \gamma_{\text{rational}} \frac{t_{ov}}{t_{tr}}, \quad (3.39)$$

as the calibration does not affect this ratio we may use the same value in determining the new adjusted volatility accrual ratio:

$$\begin{aligned}\gamma &= r_{\text{vol}} \frac{t_{\text{adj},tr}}{t_{\text{adj},ov}}, \\ &= \gamma_{\text{rational}} \frac{t_{ov}}{t_{tr}} \frac{t_{\text{adj},tr}}{t_{\text{adj},ov}}.\end{aligned}\tag{3.40}$$

We look at only very short maturity categories for all of our time adjusted tests, as these will contain the strongest signals of investor biases. Our logic is as follows: for long maturity options the exact distribution of time weight between two periods e.g., the trading and overnight periods will have a very weak effect on the value of the objective function as there are many weeks until maturity. Making the objective function surface increasingly flat with longer maturity options. Thus, we examine only two maturity categories an ultra-short category (between 4 and 7, inclusive, days to maturity), short (between 8 and 20, inclusive, days to maturity). While these categories are somewhat arbitrary we find for much longer maturities and even for some years for the short category the surface becomes too flat.²² Table 3.13 reports the parameter estimates of the SV_{tr} and SVJ_{tr} models (see Appendix 6.10 for initial results for the SV and SVJ models calibrated to the short maturity data sets. Once we have the original structural parameters we employ the two-step calibration methodology of Section 3.2.2 and calibrate to the new time weight of the trading day, $t_{\text{adj},tr}$). Table 3.13 presents a clear in-sample error reduction when compared to the original models (displayed in the Err. Ratio column). For instance, for the SV_{tr} model in 2008 using the ultra-short maturity category we see a reduction of 16%. While this case is the most significant the average error reduction, across years, for the SV_{tr} model in the ultra-short category is 6.6% (with a range of 1% to 16%, per year). The SVJ_{tr} counterpart average error reduction is 4.6% (with a range of 1% to 11%, per year). The short maturity counterparts are 6.4% (with a range of 3% to 8%, per year) and 4.4% (with a range of 1% to 9%, per year) for the SV_{tr} and SVJ_{tr} models respectively. Table 3.13 also demonstrates the out-of-sample error ratios with the original models (displayed in the OS Err. Ratio column). The out-of-sample error reductions depict very meaningful model improvement for taking into account market

²²See Appendix 6.11 for details on the effect of a flat objective function.

maker bias to volatility accrual between the day and night periods. For example, the average out-of-sample reduction for the SV_{tr} model in the ultra-short category is 8.2% (ranging from 4% to 16%). The SVJ_{tr} counterpart average error reduction is 8.8% (with a range of 5% to 14%, per year). The short maturity counterparts are 5.8% (with a range of 3% to 9%, per year) and 4.2% (with a range of 1% to 7%, per year) for the SV_{tr} and SVJ_{tr} models respectively. The γ column reports the implied γ resulting from R_{tr} given by each model using Equation (3.40). There is generally stable agreement between models on the value of γ implying the effect of OMM bias is model independent, consistent with our expectations. Secondly, values range from 0 to 1.23 when using the short maturity category (although this is likely the effect of longer maturity options as previously mentioned, see Appendix 6.11 for details) and from 0.42 to 1.23 using only the ultra-short maturities. The number of cases for which $\gamma < 1$ is only 4 out of 10 for the ultra-short and 8 out of 10 for the short category. As a value of $\gamma < 1$ implies OMMs price more volatility accrual during the overnight period than the trading day, these results are placed down to maturity related effects and only the results of the ultra-short category are discussed.

With the exception of 2014 and 2016 the estimates of γ are between 1.04 and 1.23 indicating a ratio well below the rational ranges suggested by Muravyev and Ni (2019) who suggest in a rational world a range between 1.6 and 1.8. Thus, the calibration results provide model-based evidence in support of Muravyev and Ni's (2016) hypothesis.

In summary accounting for the overnight effect leads to a clear and significant out-of-sample error reduction of approximately 7%. This effect seems to be heightened in times of market turmoil and short tenor options, both seeing further error reductions. The effect is not model dependent. We also support the hypothesis of the overnight effect being caused by OMMs bias in ignoring the volatility seasonality between day and night, as noted in Muravyev and Ni (2019). With the increased popularity of weekly expiry option contracts market participants have improved ability to acquire or reduce exposure to diffusive and jump priced risks. While Andersen et al. (2017) demonstrate the importance of using weekly OTM options to gain inference on jump risks our research demonstrates the importance for market participants to understand the risks of misspecifying the volatility accrual process, especially at very short time horizons.

3.3.4 Weekend Effect

This subsection is designed to test if the model described in Section 3.3.2 can also explain the weekend option return seasonality. The empirical test is constructed by splitting the week into two periods: week and weekend. Defined as such:

$$\text{week} = 5t_{tr} + 4t_{ov}, \quad (3.41)$$

$$\text{weekend} = t_{wknd}, \quad (3.42)$$

for this test we require that the week and weekend still sum to one week, hence we have the constraint equation:

$$\text{week} + \text{weekend} = \frac{1}{52}. \quad (3.43)$$

For this test we are not interested in where the time weight accrues during the trading week, only in the total time weight of the trading week. Thus, we shall keep the same ratios between trading time and non-trading time during the week, yielding an additional constraint equation:

$$\frac{t_{adj,tr}}{t_{adj,ov}} = \frac{t_{tr}}{t_{ov}}, \quad (3.44)$$

resulting in direct calibration between week and weekend. Once again the choice of calibration to either $t_{adj,wk}$ or $t_{adj,wknd}$ is arbitrary and we only report results for $t_{adj,wk}$, although similar results are obtained for $t_{adj,wknd}$. In this work the weekend is defined slightly differently from Jones and Shemesh (2018) (who include one trading day on Monday), our current definition includes no trading time. As a consequence our expectation is less volatility will accrue over the weekend than predicted using a calendar time convention. Volatility accrual over the weekend can be expressed as:

$$\text{vol} = \int_t^{t+\delta t} V_s ds, \quad (3.45)$$

where δt represents the passage of time over a weekend. As this quantity is small (~ 0.0074 years.²³) and V_t will be relatively constant the volatility accrual can be approximated by:

$$\text{vol} = V_t \delta t. \quad (3.46)$$

As the weekend includes no trading time we would expect the total volatility accrued to be small, approximately that of the overnight period ($V_{t, wknd} t_{adj, wknd} \approx V_{ov} t_{ov}$). In the language of Equation (3.46)

$$t_{adj, wknd} \approx \frac{V_{ov} t_{ov}}{V_{t, wknd}}. \quad (3.47)$$

Although Equation (3.47) is not very informative as currently stated, it seems theoretically reasonable to assume a relationship along the lines of: $V_{ov} \approx V_{wknd}$, as in both periods no trading occurs, leaving the relation: $t_{adj, wknd} \approx t_{adj, ov}$.²⁴

Table 3.14 displays the results of the SV_{wk} and SVJ_{wk} models. Firstly, note an error reduction of 0 – 3% in the ultra-short maturity data set with the exception of the SV_{wk} model in 2010 which has a reduction of 8%. For the short maturity category we see a reduction of 0 – 3% with no outliers. Indicating, controlling for the weekend seasonality is of less importance than the overnight seasonality as shown in Section 3.3.3 and Table 3.13. As observed when testing for the overnight effect we also observe no model dependence with regard to error reduction.

For the ultra-short maturity category ratio values, R_{wk} , ranges from 0 to 0.94 (~ 4 days). For the short maturity case values range from 0 to 1. Upon initial inspection the results look very erratic. However, most results can be categorized as either: $R_{wk} \sim 1$ or hitting the lower bound. In this case where R_{wk} hits the lower bound we suspect this is likely down to maturity related effects as previously mentioned. However, even only considering the

²³A weekend, measured in the calendar time convention is 2.7188 days, this is comprised of the two days from Saturday and Sunday plus the overnight time left from market close on Friday and the overnight time from Monday morning until market open (which accumulates to 17.25 hours or 71.88% of a day). Thus $2.7188/365=0.0074$.

²⁴When considering the above argument relating the volatility accrual of the weekend to that of the overnight period, it is not clear if we should be using the calendar time convention t_{ov} or the adjusted time $t_{adj, ov}$. However, as demonstrated by Table 3.13 the estimates for $\delta t_{adj, ov}$ are reasonably close together, varying between 0.8 and 0.98 days (a range of 4 hours). Where the adjusted time of the overnight period is calculated as: $t_{adj, ov} = \frac{1}{4}(7 - t_{wknd} - 5(R_{tr} t_{tr}))$, and taking the two most extreme R_{tr} of 0.26 and 0.77 from the ultra-short maturity category in years 2016 and 2012 respectively.

results which do not hit the lower bound none of the results align with our hypothesis that weight leaves the weekend and goes into the week, as indicated by no results for $R_{wk} > 1$, leading to inconclusive results for incorporating the weekend effect into traditional option pricing models. This is not unforeseen as commented upon by Jones and Shemesh (2018) who only find evidence of a weekend effect for individual equity and not S&P 500 options. Intuitively, it might seem reasonable to expect index and individual equity options to be driven by the same systematic risk factors. However, index options are options formed on a portfolio whereas equity options are a portfolio of options. As such variance risk on S&P 500 options will be decomposed of two parts, individual variance risk and correlation risk, while clearly individual equity options are only exposed to individual variance risk. The work of Driessen et al. (2009) is the first to study the effects of expected option return under market-wide correlation shocks. Driessen et al. (2009) propose a model which encapsulates the effects of correlations on the instantaneous variance of index options. Suppose the price of stock i , at time t , is $S_i(t)$, which follows an Itô process with instantaneous variance $V_i^2(t)$, which also follows an Itô process. Thus, the instantaneous correlation between the Wiener processes that drive stocks i and j is $\rho_{ij}(t)dt$, for $i \neq j$. For the purposes of this conversation it is not necessary to place any particular restrictions on the process of $\rho_{ij}(t)$, only that it is non-constant in time. As the S&P 500 is the weighted average of the individual stocks, given a set of weights $\{\omega_i\}$ the instantaneous variance of the index $V_{S\&P500}(t)^2$ at time t is expressed as:

$$V_{S\&P500}(t)^2 = \sum_{i=1}^{500} \omega_i^2 V_i(t)^2 + \sum_{i=1}^{500} \sum_{j \neq i}^{500} \omega_i \omega_j V_i(t) V_j(t) \rho_{ij}(t). \quad (3.48)$$

Therefore, it is clear that variance changes are also driven by shocks to correlations $\rho_{ij}(t)$, in addition to individual variances $V_i(t)^2$. Thus, it could be that risk factors that drive correlations are responsible for not observing a weekend effect in S&P 500 options. As hypothesised by Driessen et al. (2009) it could be that investors' desire to hedge against correlation risk by using options written on the index. This is in line with evidence found by Jones and Shemesh (2018) and Gârleanu et al. (2009) who find that end-users are net long index options.

3.3.5 Conclusion

We investigate a number of factors, effects and seasonalities which are relevant for high-frequency intradaily option pricing. We find empirical evidence to suggest multiple volatility factors and jumps in both the price and variance process, as suggested by more traditional option pricing models which use daily and or weekly data, are still important.

We take into consideration the timing of FOMC announcements and develop a pricing approach including deterministic jumps on FOMC announcements. We explore two different types of estimators of the uncertainty surrounding FOMC announcements and document their reliability. To quantify the impact on option prices, we calibrate a two-factor volatility model augmented with random jumps in the price and variance process. Empirically we find evidence to suggest FOMC announcements are an important part of option prices. We find that accounting for FOMC announcements is meaningful for short tenor options only on the day of the announcement, this helps reduce intradaily pricing errors by approximately 1 – 2%. While accounting for FOMC jumps has a larger impact throughout the sample due to the accumulated effect of accounting for these deterministic price jumps, this can yield error reductions of up to 20%.

We also investigate the weekend and overnight volatility seasonalities. We develop a novel framework to incorporate such effects into traditional option pricing models. We are able to effectively capture the overnight effect leading to reduction in short term option pricing errors of approximately 7%, out-of-sample. This effect seems to become more pressing in times of market turmoil. We also investigate the weekend effect in our option pricing model framework, but obtain inconclusive results for the presence of this effect, in line with prior research such as Jones and Shemesh (2018).

Chapter 4

Earnings Announcements and Expected Option Returns

4.1 Introduction

The main test hypothesis of this work is: does a positive relationship between delta-hedged option returns and earnings announcement variance of the underlying stock exist? We do indeed find this to be the case. To measure earnings announcement variance we use the square of the four-day excess return around earnings announcements.¹ There is extensive literature which demonstrates earnings announcements impact stock returns (for instance, Bernard and Thomas (1989), Chan et al. (1996), Novy-Marx (2015), Sadka (2006) amongst many others). There is also a growing strand of literature which documents the effects of volatility on expected option returns (for instance, Cao and Han (2013), Duarte and Jones (2007), Hu and Jacobs (2016)). However, to the best of our knowledge there is no work documenting that earnings announcements are an important source of volatility for

¹We are aware that there are other ways to measure earnings announcement surprise, such as the standardised unexpected earnings (SUE) as used by Xing et al. (2010). Multiple reasons motivate our choice behind choosing a return-based measure such as CAR. By definition an earnings-based measure compares announced earnings to analyst forecasted earnings, which is a noisy proxy of the measure. Secondly, the degree of persistence affects the information content in any given earnings surprise, this is not captured by SUE but is by a return-based measure does.

expected option returns.

Our main contribution to the literature is to demonstrate that options on firms with high earnings announcement variance earn, on average, higher returns than options on low earnings announcement variance firms. This is the key novel finding of our work. This finding is robust to: option type (call or put), portfolio formation, and also the length of time for which the portfolio is held. We also investigate a number of different measurements of earnings announcement variance and find our results are robust.

To test our hypothesis we examine a cross-section of at-the-money options on individual stocks each month. For each optionable stock, in each month, we evaluate the return over the following month of a portfolio that buys one call (or put), delta-hedged with the underlying stock.² We choose the strategy of delta-hedging option returns, instead of studying raw returns, such that our results are not driven by sensitivities in stock price changes. Therefore, we study the part of option returns which are related to volatility change.

Using Fama-MacBeth cross-sectional regressions from 1996 to 2017, we find high statistical significance on our measure for earnings announcement variance in predicting expected delta-hedged option returns. Our results are robust to several measures of stock, jump and arbitrage-related characteristics. We are particularly careful to include an array of volatility-related characteristics in our Fama-MacBeth regressions, as it could be that our results are driven by one or more combinations of already known powerful explanatory variables. For example, it is readily conceivable that firms with high earnings announcement variance simply have very high total variance anyway and by extension might have high volatility around earnings announcements. To control for this, we use a total variance specification (similar to Cao and Han (2013)), finding total variance does not explain our results. We further complement the findings of Hu and Jacobs (2016) on the negative relation between expected delta-hedged call option return and underlying stock volatility. Furthermore, as prior research documents there is a strong negative relation between delta-hedged option returns and idiosyncratic volatility (see Cao and Han (2013)), we complement this finding. Logic suggests, and we confirm, earnings announcements are a high source of idiosyncratic

²Please note that the notation is not necessarily the same as the previous chapters and is self-consistent and contained within this chapter.

volatility. As option market makers do not know the outcome of the announcement, it is hard to hedge against such events, thus they may demand a higher premium. However, including idiosyncratic volatility in our regression framework does not take away statistical significance from our earnings announcement variance measure. Also conceivable is that earnings announcement variance is merely a product of jump risk. For example, option market makers might observe a jump in returns of an underlying stock in response to an earnings announcement and, as a consequence, adjust their jump risk premium going forward, potentially completely subsuming any earnings announcement variance statistical significance. To account for this we follow the jump measure as given in Bali et al. (2011), using several different specifications, however the strong statistical significance on earnings announcement variance remains. We find that our results are robust to controlling for all volatility, arbitrage and jump-related measures.

We further create an option trading strategy which buys delta-hedged call options on firms ranked in the top earnings announcement variance quintile, and sells delta-hedged options which are ranked in the bottom earnings announcement variance quintile.³ This strategy earns approximately 0.31% per month, using an option value-weighting scheme. We also investigate time series regressions of the returns to our option trading strategy against the Fama-French (2015) five-factors, momentum and the idiosyncratic volatility strategy from Cao and Han (2013). Our option trading strategy has a statistically significant positive alpha of approximately 0.44% per month, thus our results cannot be explained by idiosyncratic volatility and other common factors.

Ultimately, the profitability of our earnings announcement variance-based option trading strategy depends on option trading costs. If we assume the effective option spread is equal to 25% of the quoted spread, then the average return of our option trading strategy is rendered insignificant; however, a significant time series regression alpha remains of 0.26% per month. Following the work of Muravyev and Pearson (2019), who shows option effective spread is under 30%, this suggests our option trading strategy is reasonably robust to trading costs.

³It is for this reason that we consider the CAR_4^2 and not CAR, it is not because CAR has less predictive power, rather to do with the shape of the function. Using CAR one would obtain a U-shaped profitability as both firms with very high and very low CAR have high earnings announcement variance, thus making it very challenging to construct a trading strategy.

Our results demonstrate that idiosyncratic volatility contains a positive component in relation to earnings announcements, while the overall relationship with delta-hedged option returns is negative. Although this seems like an apparent contradiction, one potential explanation is a demand-based model. We find that investors have diminishing demand for options with increasing earnings announcement variance, therefore this elevates pressure on option market makers in relation to such firms, potentially lowering the premium as market makers are net long individual equity options (see Christoffersen et al. (2017)). This economic reasoning is in line with previous research on demand-based option pricing, e.g., Gârleanu et al. (2009). However, we find that controlling for open interest does not drive our results.

The work is structured as follows. We describe the data in Section 4.2 and detail the measurement of the delta-hedged option returns and earnings announcement variance measure. Section 4.3 presents summary statistics on the data and the main regression-based results. Section 4.4 presents the portfolio-sorting results and studies the option trading strategy, taking into account trading costs. Section 4.5 provides robustness checks on our results. Section 4.6 discusses a possible economic interpretation of the sign relation between delta-hedged option returns and earnings announcement variance. Section 4.7 concludes the work.

4.1.1 Literature Review

Although there has been tremendous growth in the equity options market in recent decades, there is still room for growth in knowledge about the determinants of expected returns in this market. Partly, the issue is that index and individual equity options behave differently, thus there is less known about equity option expected returns as there has been emphasis by some to investigate index option returns (for instance, Coval and Shumway (2001), Bakshi and Kapadia (2003)). In contrast, there is an extremely substantial literature on the cross-section of expected stock returns. This work attempts to bridge some of the gap.

Bakshi and Kapadia (2003) investigate delta-hedged option returns in S&P 500 index

options. The motivations for their research is to investigate whether a volatility risk premium is priced into options, and if so, what this signifies? Their results indicate that the volatility risk premium exists and is negative. Second the delta-hedged returns of S&P 500 index options are negative, on average. The authors control for jump risks by assuming (in line with Bates (2000) and Bakshi et al. (2003)) that the skewness and kurtosis of the risk-neutral distribution can be used as a proxy for the mean jump size and the jump intensity, respectively. Thus, they include the skew and kurtosis as variables in their regression framework, using the model-free approach of Bakshi et al. (2003). However, even accounting for jump risks, the volatility risk premium still significantly affects delta-hedged returns. The work on volatility risk premia is continued by Duarte and Jones (2007), who analyse volatility risk premia using a large cross-section of option returns on individual equities. Their results strongly suggest the presence of a volatility risk premium which is increasing in the level of overall market volatility. This risk premium provides compensation for risk from the underlying asset and from characteristics of the option contract. As the work of Duarte and Jones (2007) demonstrates, option returns can be affected by option implied volatility and volatility from the underlying stock. However, does option implied volatility have any predictive power for stock returns? An et al. (2014) demonstrate stocks with large increases in call option implied volatilities over the previous month tend to have high future returns. The authors also find the reverse is true for put options. Sorting stocks ranked into decile portfolios by past call implied volatilities, produces spreads in average returns of approximately 1% per month. The authors also document that this finding is highly persistent, with return differences lasting for up to six months. Continuing the theme of predictive power from stock volatility characteristics on option returns, Hu and Jacobs (2016) investigate the relation between expected option returns and the volatility of the underlying security. In the cross-section of stock option returns, they find returns on call option portfolios are negatively related with underlying stock volatility. The authors also document a positive relation between put options and underlying stock volatility. Continuing the topic of volatility and option returns, Vasquez (2017) investigates the relationship between implied volatility term structure and option returns, finding a positive relationship.

There is also a strand of literature dealing with the effects of skewness; for instance, Boyer

and Vorkink (2014) investigate the relationship between option ex-ante total skewness and returns on equity options. The authors' results suggest that total skewness demonstrates a strong negative relationship with average option returns. Their results indicate that large option premiums are required to compensate intermediaries for bearing unhedgeable risk when accommodating investors' desire for lottery-like options. Although the work done by Boyer and Vorkink (2014) is very interesting, it does not capture the option's lottery-like characteristics caused by the underlying stock's lottery-like behaviour. Byun and Kim (2016) investigate such a relationship, demonstrating that call options also present a strong negative relationship between option returns and underlying stock lottery-like behaviour. Blau et al. (2016) investigate whether the gambling behaviour of investors affects volume and volatility in options markets. The authors find that the ratio of call option volume to total option volume is greater for stocks with return distributions that mimic lottery-like behaviour.

Amongst the voluminous literature on stock cross-sectional returns, there is a subsection which uses informational content from options to predict stock returns. Such papers include Pan and Poteshman (2006) who present strong evidence suggesting option trading volume is a predictor for future stock returns. Cao et al. (2005) investigate mergers and find that information is received into the options market before the stock market. Bali and Hovakimian (2009) investigate whether realised and implied volatilities of individual stocks can predict cross-sectional variation in expected returns. They find there is a significant relation between volatility spreads and expected stock returns.

There has also been a large amount of work done in the area of earnings announcements and their effect on the cross-section of stock returns. Notable on this last point is the work of Chan et al. (1996) who find past returns and past earnings surprises each predict large drifts in future returns, even after controlling for the other. The point that earnings announcements are of extremely high importance, on an individual firm level, is one which is stressed in the finding of Novy-Marx (2015). Novy-Marx (2015) documents that earnings momentum (i.e., the tendency of stocks that have recently announced positive earnings to outperform those which have announced weak earnings) can explain the performance of strategies based on price momentum (i.e., the tendency of stocks that have performed well

over the prior year to outperform, going forward, stocks that have performed badly over the prior year). That is, earnings surprise measure engulfs past performance in the cross-sectional regressions of returns on firm characteristics (in contradiction to the results of Chan et al. (1996)). Novy-Marx goes on to document that controlling for earnings surprises, when constructing price momentum strategy, significantly reduces their profitability. There has also been work done on option volatility characteristics which can predict earnings shocks. For instance, Xing et al. (2010) investigate the shape of the volatility smile and its effect on cross-sectional predictive power for future stock returns. Xing et al. (2010) find that stocks which exhibit the steepest smile in their options outperform stocks with the least pronounced smile. The predictability is persistent up to six months, and firms with the steepest volatility smiles experience the most severe earnings shocks in the next quarter. This finding is then harnessed in Atilgan (2014) who supports Xing et al. (2010) in finding that the predictability of equity returns by volatility spread is stronger during earnings announcements. It is also found that option volatility spreads can be used to predict stock returns, as in Cremers and Weinbaum (2010) who study how deviations from put-call parity contain information about future stock returns. They use the difference in implied volatility between pairs of call and put options, finding that stocks with relatively expensive calls perform better than stocks with expensive puts by 50 basis points per week. This observation cannot be explained by short sale constraints or by stocks which are hard to borrow.

We find an economic interpretation of our results lies in that of demand-based models. For instance, Gârleanu et al. (2009) model demand pressure effects on option prices. Their model shows that demand pressure in an option increases its price by an amount directly proportional to the unhedgeable part of the variance of the option. Along a similar path the authors find demand pressure increases the price of any other option by an amount which is directly proportional to the covariance of the unhedgeable parts of the two options. Along a similar theme Bollen and Whaley (2004) investigate the relation between net buying pressure and the shape of the implied volatility smile for index and individual stock options. The authors find changes in implied volatility are directly related to net buying pressures, with S&P 500 options most strongly affected by buying pressure from index puts, while changes in implied volatility from stock options are dominated by call option demand. In

response to the findings of Gârleanu et al. (2009) who document that end-users are net sellers (supported by Lakonishok et al. (2007)) Christoffersen et al. (2017) document a positive option illiquidity premium in expected equity option returns. The authors' results indicate a positive premium, i.e., expected hedged returns on less liquid options are higher than those on more liquid options. The authors find returns of 3.4% per day for at-the-money calls and 2.5% for at-the-money puts. These findings are consistent with evidence that suggests market-makers hold large, risky net long positions, and a positive illiquidity premium compensates them for this. Furthermore, the idea that demand-pressure and order flow can impact prices is cemented by the work of Muravyev (2016). Muravyev (2016) investigates inventory risk faced by market-makers, which he shows has a first-order effect on option prices. Muravyev also documents that order imbalances have a greater predictive power, than any other commonly used predictor, on option returns.

Even given the literature setting detailed above, to the best of our knowledge, we are the first to study the effect of earnings announcement variance on equity option expected returns.

4.2 Data and Delta-Hedged Option Returns

This section introduces the data used in the empirical work and describes the measurement of the key variables of interest: delta-hedged option returns, the earnings announcement surprise and our metric of earnings announcement variance.

4.2.1 Data

We use equity option and stock market data for the period January 1996 to December 2017. We obtain data on the U.S. individual stock options from the Ivy DB database provided by OptionMetrics. The data fields we use include: daily closing bid and ask quotes, trading volume and open interest of each option, implied volatility, as well as the option's delta

(which is computed by OptionMetrics). These are all obtained from the price OptionMetrics file (opprcd). We also obtain interest rate data from OptionMetrics using the zerocd file, where the interest rate is calculated from a collection of continuously compounded zero-coupon interest rates at various maturities, known as the zero curve. We also retrieve daily and monthly split-adjusted stock returns, stock prices and trading volume from the Center for Research in Security Prices (CRSP). We focus our analysis on the options of common stocks only (CRSP share codes of 10 and 11). We further get the daily and monthly Fama-French factor returns from Kenneth French's data library.⁴ By way of filtering methodology, we follow the standard practice in Cao and Han (2013) and Cao et al. (2015). At the end of each month, for each optionable stock, we collect one call and one put option which are closest to being at-the-money (where moneyness is defined as stock price divided by strike price, i.e., we choose a put and call pair with a ratio of stock price to strike price closest to unity) and have the second shortest maturity (i.e., we choose options with the shortest maturity which is greater than a month, in line with Cao and Han (2013) and Cao et al. (2015)). The option filters we apply are as follows. For our analysis, we choose options whose stocks do not have an ex-dividend date prior to option maturity (i.e., we will exclude an option if the firm pays a dividend during the remaining life of the option). We obtain dividend date and amount information from the OptionMetrics Distribution file (distrd).⁵ In order to mitigate the effects of market microstructure noise, we apply the following additional filters. We retain only options whose trading volume and open interest (in line with Choy and Wei (2020)) implied volatility and bid quotes are positive, with the bid price strictly smaller than the ask price; additionally, the midpoint of the bid and ask prices is at a minimum of \$1/8. We also require that the option has standard settlement (i.e., 100 shares of the underlying security are delivered at exercise, denoted in the opprcd file by a zero in the SS flag field), in line with Hu and Jacobs (2016). We also require that the last trade date of the option matches the underlying security price date (this step is in line with accepted

⁴The data library is available at <http://mba.tuck.dartmouth.edu/pages/faculty/ken.french/>.

⁵Using short maturity options, the early exercise premium is small and we confirm that including options whose underlying stock pays a dividend before maturity does not qualitatively affect our results. However, by excluding the dividend, we effectively consider only European-type calls, as the call is not optimal to exercise early.

data filtering techniques, such as used in Cao and Han (2013)).⁶ Furthermore, we exclude options with moneyness (the ratio of underlying stock price over option strike price, given as a percentage) less than 80 or higher than 120. While this moneyness range does include other types and not strictly at-the-money options (which technically are those with a value of exactly 100) we consider deep in and out-of-the money options (as defined by Chung et al. (2016)) so as to get good coverage of firms. However, as demonstrated by Table 4.1 the mean option moneyness used for our analysis is at-the-money, although we do use some in and out-of-the money options, thus increasing the robustness of our findings. Our reason for using at-the-money options only is that they will be the most sensitive to changes in volatility.

This leaves us with a cross-section of options which are approximately at-the-money with a common short-term maturity. Our final sample contains 1,423 firms, on average each month. The total data has 354,379 observations for delta-hedged call returns and 286,328 observations for delta-hedged put returns.

4.2.2 Delta-Hedged Option Returns

We define the delta-hedged gain (our definition follows that of Bakshi and Kapadia (2003)) as the change in the value of a self-financing portfolio consisting of a long call position, hedged by a short position in the underlying stock (to ensure the portfolio is not sensitive to the movement of the underlying stock price), whereby the overall investment earns the risk-free rate. Let us consider a portfolio of a call option that is hedged discretely N times over a period $[t, t + \tau]$, with the hedge being re-balanced at each of the dates t_n , with $n \in \{0, 1, 2, \dots, N - 1\}$, where we define $t_0 = t$ and $t_N = t + \tau$, with $t_0 < t_1, \dots, < t_N$. Thus,

⁶Our results are also qualitatively similar if we exclude firms with an initial stock price of under \$5.

over the period $[t, t + \tau]$ the call option gain is defined as:

$$\begin{aligned} \Pi(t, t + \tau) = & C_{t+\tau} - C_t - \sum_{n=0}^{N-1} \Delta_{C,t_n} [S(t_{n+1}) - S(t_n)] \\ & - \sum_{n=0}^{N-1} (e^{r_{t_n}(t_n - t_{n-1})/365} - 1) [C(t_n) - \Delta_{C,t_n} S(t_n)], \end{aligned} \quad (4.1)$$

where the price of the call option on date t_n is denoted as $C(t_n)$ and Δ_{C,t_n} is the delta of the call option on date t_n . The price of the underlying stock on date t_n is denoted as $S(t_n)$. The continuously compounded risk-free rate is r_{t_n} . The definition for the put delta-hedged gain is the same as Equation (4.1), except with put-related price and delta replacing their call counterparts. Equation (4.1) represents the excess dollar return of a delta-hedged call option. As the option price is a function of the underlying stock price, $\Pi(t, t + \tau)$ is proportional to the initial price of the underlying. In order to be able to compare across stocks, we need to use a scaling function. Cao and Han (2013) solve this issue by dividing the call delta-hedged gain by $(\Delta_t S_t - C_t)$.⁷ However, they also note their results are qualitatively similar when scaling by the initial stock or option price. We shall scale by $(\Delta_t S_t - C_t)$ and we refer to the delta-hedged call return as $\Pi(t, t + \tau)/(\Delta_t S_t - C_t)$.

4.2.3 Earnings Surprises

For earnings surprises, we use a measure commonly employed (see e.g., Chan et al. (1996), Daniel et al. (2020)): the cumulative four-day excess returns (CAR_4). Following Chan et al. (1996), we measure the earnings surprise over a four-day window, starting two days prior to the earnings announcement date (EAD), as the cumulative excess returns around the most recent quarterly earnings announcement date.⁸ Define the earnings surprise for firm i as $CAR_{4,i}$, given by:

$$CAR_{4,i} = \sum_{d=-2}^{d=1} (R_{i,d} - R_{m,d}), \quad (4.2)$$

⁷The put delta-hedged gain is divided by $(P_t - \Delta_t S_t)$.

⁸This refers to the RDQ field in COMPUSTAT.

with $R_{i,d}$ denoting the return on day d of stock i and $R_{m,d}$ denoting the return of the market on day d , relative to the earnings announcement (thus $d = 0$ is the announcement day). As in Daniel et al. (2020), we require that there be at least two days with valid returns over the four-day horizon. Furthermore, it is necessary that the earnings announcement date be at least two trading days prior to the end of the month. When forming the portfolios, if there is no announcement date in the previous six months, we exclude firm i .

We use a function of the earnings surprise measure, CAR_4 , as a metric to gauge earnings announcement variance. With the present construction of observing earnings surprise, firms which have very large positive returns will have a large CAR_4 , while firms which have very large negative returns will have a very negative CAR_4 . However, both of the above categories of firms will have large earnings announcement variance; thus, for our empirical analysis we wish to construct a measure wherein both types of firms will be categorised by the same metric. To satisfy this criteria we choose to use the CAR_4^2 , thus both extremely high and low excess returns will both have large CAR_4^2 . We will use this measure throughout to denote earnings announcement variance. However, our results are qualitatively similar when we use different measures of earnings announcement variance. Instead of squaring the final excess return value (CAR_4) we achieve similar results if we use the sum of squared results. Secondly, our results are qualitatively similar when we use $|CAR_4|$ as our measure of earnings announcement variance instead. Thirdly, we achieve similar results by constructing a specific earnings announcement volatility-based measure constructed as the standard deviation of excess returns around the EAD.

4.3 Empirical Results

Section 4.3.1 presents an investigation into the summary statistics with different lengths of holding the portfolio. Section 4.3.2 reports our results of Fama-MacBeth regressions. We investigate the effects of including stock, volatility and jump characteristics as well as investigating limits to arbitrage. Furthermore, we include a total Fama-MacBeth regression test, in which we include an intersection of the previously investigated independent variables.

Table 4.1: Summary statistics of individual equity options data

This table reports the descriptive statistics of delta-hedged option returns. The option sample period is from January 1996 to December 2017. Delta-hedged gain is the change over the month (or until maturity) in the value of a portfolio consisting of one contract of long option and a proper amount of the underlying stock, re-hedged daily. The call option delta-hedged gain is scaled by $(\Delta S - C)$, where Δ is the Black-Scholes delta, S is the price of the underlying stock and C is the price of the call option. The put option delta-hedged gain is scaled by $(P - \Delta S)$, where P denotes the price of the put option. All return measures are given as a percentage. Maturity is given as the number of calendar days until maturity. Moneyness is the ratio of stock price over option strike price, given as a percentage.

Variable	Mean	Median	Std	10%	25%	75%	90%
Panel A. Calls							
Delta-hedged gain until maturity	-1.16	-1.56	17.48	-10.16	-5.06	1.59	6.79
Delta-hedged gain until month-end	-0.65	-1.10	8.34	-7.21	-3.57	1.34	5.77
Days to maturity	48.17	48.00	2.28	45.00	46.00	50.00	51.00
Moneyness	98.87	99.20	5.65	91.94	95.83	101.95	105.20
Panel B. Puts							
Delta-hedged gain until maturity	-0.89	-1.59	12.20	-8.88	-4.60	1.21	5.80
Delta-hedged gain until month-end	-0.84	-1.15	8.93	-6.43	-3.29	0.92	4.55
Days to maturity	48.15	48.00	2.27	45.00	46.00	50.00	51.00
Moneyness	100.22	100.17	5.82	93.54	97.32	103.21	107.02

4.3.1 Summary Statistics

Table 4.1 displays summary statistics on the data. Table 4.1 shows the average moneyness, defined as the price of the underlying divided by the strike price given as a percentage i.e., $100 \times S/K$ for call options is 98.25 and 100.26 for puts, with standard deviations of 5.66 and 5.82 respectively. The time to maturity ranges from 45 to 51 calendar days across our sample, with the average being 48. Table 4.1 Panels A and B also provide mean and median delta-hedged option returns for individual firms. Both Panels demonstrate that the average option returns are negative for puts and calls. This is not affected by the period over which the option return is measured, nor is it affected if maturity is included. To illustrate, the mean delta-hedged at-the-money call option return is -0.65% over the next month and -1.16% if held until maturity (which is on average 48 calendar days). The median delta-hedged call option return is -1.10% (-1.56%) over the next month (until maturity). The mean delta-hedged put option return is -0.84% (-0.89%) over the next month (until maturity). The median is -1.15% (-1.59%) over the next month (until maturity).

4.3.2 Fama-MacBeth Regressions

Table 4.2 reports Fama-MacBeth regressions controlling for stock-related characteristics. The dependent variable is call option delta-hedged gain until month-end, scaled by $(\Delta S - C)$ at the beginning of the period (delta-hedged returns). We re-balance daily the delta-hedges to minimise the effect of change in underlying stock price. The common time to maturity is approximately one and a half months. Model 1 runs the regression of the CAR_4 , measured as the firms' excess return around the most recent earnings announcement, and CAR_4^2 our earnings announcement variance measure. The beta for CAR_4 and CAR_4^2 both demonstrate strong statistical significance with respective betas of 0.0070, 0.0659 and t -statistics of 2.2735 and 2.6745. Even though delta-hedges are re-balanced daily, it could be that the strong link between delta-hedged option return and earnings announcement variance is attributable to some imperfections in the delta-hedges. Thus, the results of Model 1 could potentially be attributed to some known pattern in the cross-section of expected stock returns. To address this, the regressions reported in Models 2 - 6 include several stock characteristics which are known predictors of the cross-section of stock returns; we include: size of the underlying firm, past stock returns and various jump risk measures. The CAR_4^2 coefficient remains highly statistically significant in all regressions of Table 4.2. In Models 2 and 3, we control for past stock returns from the previous month (Model 2) and stock returns over the previous year (skipping the prior month to avoid diluting price momentum by short-term reversals, as suggested in Novy-Marx (2015)). We find no statistically significant coefficients for the past returns, thus the coefficients on CAR_4^2 change very little, ranging from 0.0635 to 0.0639, with t -statistics 2.7078 and 2.5889 respectively. In Model 4, we control for any size-related effects. The size variable (ME) is the product of monthly closing stock prices and the number of outstanding shares, in the previous month. The coefficient for size is strongly positive (in line with prior research, see for example: Christoffersen et al. (2017) and Xing et al. (2010)), with a beta of 0.0014 and t -statistic of 3.8792. However, the earnings announcement variance beta is still significant, with t -statistic of 3.3920. Although thus far, stock characteristics do not explain the strong positive relation on earnings announcement variance, it could be that firms

with high earnings announcement variance are just those which have high jump risk and there is nothing unique about the EAD. To account for such an effect, model 5 includes a measure for jumps in the underlying stock price. To proxy for jumps, we follow the work of Bali et al. (2011) and include a measure of extreme returns ($\text{MAX}(\text{Ret}_{0,12})$) which is calculated as the maximum absolute daily return of the underlying stock from the previous 12 months.⁹ Model 5 depicts a statistically significant negative relation between delta-hedged option return and maximum daily return (this is in line with the relation between maximum daily return and expected stock returns, as found in Bali et al. (2011), and is also in line with prior research on expected option returns, see Byun and Kim (2016)). However, more importantly, inclusion of a jump measure does not take away statistical significance from our earnings announcement variance measure as model 5 has a t -statistic of 3.9994 on CAR_4^2 , concluding that jumps do not explain our results. However, there is other work which investigates the preference for lottery-like payoffs from investors. For instance, the work of Boyer and Vorkink (2014), who use a measure of total skewness for option returns, is executed by assuming a lognormal distribution of stock prices. Their results also demonstrate a negative relation between options with high jump risk and expected option returns. However, their measure does not capture the options' jump characteristics, caused by jump risk in the underlying stock, due to the simplifying assumption of lognormal stock prices. Therefore, their option skewness measure is a function of only the underlying stock's expected return and volatility for given option parameters. It is for this prior reason (the lack of consideration of the underlying stock's jump risk), that we do not use this measure. Although, as Byun and Kim (2016) find complementary results for option returns, using the maximum return method of Bali et al. (2011) and the skewness measure of Boyer and Vorkink (2014), we feel confident that our results would be qualitatively robust to a skewness-based measure of jump risk. The positive coefficient on the earnings announcement variance measure is not affected by controlling for all independent variables in Models 2-5 in one regression. Model 6 presents a total regression model which still loads positively on CAR_4^2 , with a coefficient of 0.1045 and t -statistic of 4.6236. In summary, the significant positive relation between delta-hedged option returns and the earnings announcement variance measure is robust to accounting for

Table 4.2: Fama-Macbeth regressions stock-related characteristics

This table reports the average coefficients from monthly Fama-MacBeth cross-sectional regressions; variables are formed at the end of each month. The dependent variable is call option delta-hedged gain until month-end scaled by $(\Delta S - C)$ at the beginning of the period. Earnings surprise is measured using the cumulative four-day excess returns around the most recent earnings announcement (CAR_4), taken from the previous month. We also include our measure for earnings announcement variance (CAR_4^2), taken from the previous month. The stock return in the prior month is denoted by $r_{0,1}$. Past performance is represented using the individual monthly stock returns, measured over the previous year, skipping the most recent month ($r_{1,12}$). Size is accounted for in the ME variable, which is calculated as the product of monthly closing price and the number of outstanding shares in the previous month. We also include a jump risk proxy ($MAX(Ret_{0,12})$), which is calculated as the absolute value of the largest daily return from the previous 12 months. To adjust for serial correlation, robust Newey-West (1987) t -statistics are reported in brackets. All independent variables are winsorized each month at the 0.5% level. The sample period is from January 1996 to December 2017.

Variables	Model 1	Model 2	Model 3	Model 4	Model 5	Model 6
Intercept	-0.0036 (-2.3325)	-0.0038 (-2.5444)	-0.0039 (-2.5577)	-0.0243 (-3.8599)	-0.0010 (-0.7213)	-0.0233 (-3.9352)
CAR_4	0.0070 (2.2735)	0.0069 (2.4177)	0.0067 (2.2725)	0.0066 (2.2283)	0.0090 (2.8770)	0.0072 (2.5826)
CAR_4^2	0.0659 (2.6745)	0.0635 (2.7078)	0.0639 (2.5889)	0.0917 (3.3920)	0.0959 (3.9994)	0.1045 (4.6236)
$r_{0,1}$		-0.0011 (-0.3958)				-0.0020 (-0.7198)
$r_{2,12}$			-0.0006 (-0.7741)			-0.0012 (-1.5978)
$\text{Ln}(\text{ME})$				0.0014 (3.8792)		0.0014 (3.9442)
$\text{MAX}(\text{Ret}_{0,12})$					-0.0205 (-5.2285)	-0.0182 (-2.1159)

various stock-related characteristics.

Table 4.3 investigates Fama-MacBeth regressions controlling for volatility-related characteristics. The dependent variable is still call option delta-hedged returns until month-end. There is a strand of literature which investigates the cross-sectional characteristics of volatility measures on delta-hedged option returns, such as: Duan and Wei (2008) and Cao and Han (2013). Firstly, we wish to confirm our data sample and filtering methodology is reasonable by complementing key findings of these works. Secondly, we wish to confirm that our results are not simply driven by firms which have high volatility. Thirdly, we wish to investigate the effect of earnings announcements on idiosyncratic and systematic volatility. Cao and

⁹We also investigate two other specifications of jump risk: we take the maximum absolute value of daily returns over the past month and 6 months, achieving similar results.

Han (2013) posit a strong negative relationship between delta-hedged option returns and idiosyncratic volatility. The economical justification for their results is that firms with high idiosyncratic volatility attract attention from speculators, however these firms are difficult to hedge against, consequently market makers charge a higher premium, yielding lower returns. Tying into this argument, our expectation is that earnings announcement dates are a source of high idiosyncratic volatility, relative to each firm. Therefore, if the dates surrounding the EAD are removed, then a substantial part of that month's idiosyncratic volatility should also vanish. Fourth and finally, we wish to understand if idiosyncratic or systematic volatility measures drive our results. Model 1 includes independent variables for CAR_4^2 and total variance (VAR), as measured by the square of the standard deviation of daily stock returns over the previous month. The coefficient on CAR_4^2 remains positive and statistically significant, with beta of 0.0900 and t -statistic 4.6208. Model 1 also complements the results of Cao and Han (2013) (as well as Hu and Jacobs (2016), although they study raw option returns) in finding a statistically significant negative coefficient for variance of the underlying, regressed on delta-hedged option returns. Therefore, we have demonstrated that our results are not driven by highly volatile firms. Models 2 and 3 investigate the effect on idiosyncratic and systematic variance of removing the dates surrounding the EAD (two days prior to, and one day post, each EAD). The idiosyncratic volatility (IVOL) is measured as the standard deviation of the residuals of the Fama and French three-factor model, estimated using the daily stock returns over the previous month and is the same definition as used in Ang et al. (2006). The systematic volatility (SysVOL) is measured as: $\sqrt{VOL^2 - IVOL^2}$ and is the same definition as that used in Duan and Wei (2008). Model 2 investigates the measures using all trading days each month, for which the beta on idiosyncratic variance is -0.0050 with a significant t -statistic of -2.0344 . Model 2 also reports that the beta on systematic variance is 0.0038 with an insignificant t -statistic of 0.7060 , thus further complementing the findings of Cao and Han (2013), who conclude there is a negative idiosyncratic volatility relation with delta-hedged option returns. Model 3 investigates the effect on idiosyncratic variance with the exclusion of the dates surrounding the EAD for firm i (if firm i under-went an earnings announcement that month). This demonstrates two key points. Firstly, removal of the dates around the EADs causes the idiosyncratic variance to become non-significant, with a t -statistic of

-1.2489 and a non-significant beta on systematic variance ($\text{SysVOL}_{\text{exc. EAD}}^2$). Secondly, the beta on idiosyncratic variance becomes more negative, in line with the notion that earnings announcement variance is directly proportional to delta-hedged option returns.¹⁰ With the understanding that every firm has to have four earnings announcements every year, thus approximately only one-third of firms will have dates removed each month and this has a marked effect on the statistical significance, we conclude that earnings announcements are a high source of idiosyncratic volatility. Models 4 and 5 investigate the effect of including our earnings announcement variance measure, CAR_4^2 , into the regressions of models 2 and 3. Both models demonstrate a statistically significant beta on our earnings announcement variance measure, indicating our results are robust to idiosyncratic and systematic variance. Our intuition as to the positive sign of earnings announcement variance is also supported by the results in Models 4 and 5. For instance, in both pairs (Models 2 and 4, Models 3 and 5) the addition of a separate beta to model the earnings announcement variance component renders the betas for idiosyncratic variance (IVOL^2 for Model 4 and $\text{IVOL}_{\text{exc. EAD}}^2$ for Model 5) more negative in order to balance the significant positive beta from the EAD. In summary, our results are not driven by high variance firms or firms with high idiosyncratic and/or systematic variance. We also show that earnings announcements are a high source of idiosyncratic variance with a positive component in relation to delta-hedged option returns.

Table 4.4 reports regression results while controlling for limits to arbitrage. We do so because arbitrage between stocks and options is more challenging when transaction costs in options are high and/or when the underlying stock is illiquid. Furthermore, as these cases are associated with high stock volatility, there is a possibility that some of these measures might explain the large movements around earnings announcements and therefore leave the earnings announcement variance measure insignificant. In order to investigate this, we include as independent variables: relative option bid-ask spread to proxy for transaction costs in trading options, and option demand, measured by option's open interest scaled by monthly stock trading volume (the same definition as used in Cao and Han (2013)). To

¹⁰For this table we report the idiosyncratic variance not volatility as variance is a linear measure and it is thus easier to see the effect of removal of the EAD dates.

Table 4.3: Fama-Macbeth regressions volatility-related characteristics

This table reports the average coefficients from monthly Fama-MacBeth cross-sectional regressions; variables are formed at the end of each month. The dependent variable is call option delta-hedged gain until month-end scaled by $(\Delta S - C)$ at the beginning of the period. We include our measure for earnings announcement variance (CAR_4^2), taken from the previous month. VAR is the square of the standard deviation of daily stock returns over the previous month. Idiosyncratic variance ($IVOL^2$) is the square of the standard deviation of the residuals of the Fama-French three-factors model, estimated using the daily stock returns over the previous month. Systematic variance ($SysVOL^2$) is the square root of $(VOL^2 - IVOL^2)$. The variables with the subscript: "exc. EAD" refer to the original variables ($IVOL^2$ and $SysVOL^2$) calculated with the exclusion of the CAR_4 dates. To adjust for serial correlation, robust Newey-West (1987) t -statistics are reported in brackets. All independent variables are winsorized each month at the 0.5% level. The sample period is from January 1996 to December 2017.

Variables	Model 1	Model 2	Model 3	Model 4	Model 5
Intercept	-0.0031 (-2.3180)	-0.0029 (-2.1912)	-0.0031 (-2.4449)	-0.0032 (-2.4389)	-0.0035 (-2.7160)
CAR_4^2	0.0900 (4.6208)			0.0910 (4.6312)	0.0719 (3.2311)
VAR	-0.0045 (-2.5799)				
$IVOL^2$		-0.0050 (-2.0344)		-0.0071 (-3.1585)	
$SysVOL^2$		0.0038 (0.7060)		0.0016 (0.2780)	
$IVOL_{exc. EAD}^2$			-0.0054 (-1.2489)		-0.0058 (-2.3117)
$SysVOL_{exc. EAD}^2$			0.0024 (0.4209)		0.0005 (0.0888)

proxy for stock illiquidity, we use the Amihud (2002) illiquidity measure.¹¹ Model 1 provides the univariate regression with only CAR_4^2 which remains positive and statistically significant in the absence of CAR_4 . Model 2 includes our measure for demand as an independent variable, which is strongly negatively loaded against delta-hedged returns, with a beta of -0.0750 and t -statistic -7.0651 . Thus demonstrating that delta-hedged option returns are more negative with high in-demand options (consistent with findings in Cao and Han (2013) and demand-based models of option pricing). However, controlling for demand does not appreciably diminish the significance of our earnings announcement variance measure.

¹¹The Amihud (2002) measure of illiquidity for stock i at month t is defined as:

$$IL_{i,t} = \frac{1}{D_t} \sum_{d=1}^{D_t} |R_{i,d}| / VOLUME_{i,d},$$

where D_t is the number of trading days in month t and $R_{i,d}$ is stock i 's daily return, while $VOLUME_{i,d}$ is the stocks' trading volume in day d of month t .

Model 3 includes the proxy for option transaction costs, relative option bid-ask spread. The coefficient on relative option bid-ask spread is 0.0050, with a significant t -statistic of 3.3566. Inclusion of transaction costs does not dampen any of the significance of the earnings announcement variance measure. Model 4 includes the natural logarithm of the Amihud (2002) measure of illiquidity. The regression reports a highly statistically significant and negative relation between stock illiquidity and delta-hedged option returns, with beta of -0.0013 and t -statistic -3.5447 , demonstrating that delta-hedged option returns are more negative with less liquid stocks. There is a slight increase in beta and significance of the earnings announcement variance measure when a regressor for illiquidity is included, but not a momentous amount. However, this does demonstrate that options are not redundant, as the negative relation, between delta-hedged option returns and stock illiquidity, is in distinct contrast to expected stock returns and stock illiquidity, with many studies in the equity literature finding a positive relationship (first observed in Amihud and Mendelson (1986)). This point confirms that our results are not just some distorted reflection of the already known and highly researched results on the cross-section of expected stock returns. Model 5 reports the total regressions of all the limits to arbitrage variables, even for which there is no significance taken out of the earnings announcement variance measure. Table 4.4 also complements the results found in Cao and Han (2013) regarding the negative relations between delta-hedged option returns and the limits to arbitrage variables.

Table 4.4: Fama-MacBeth regressions limits to arbitrage

This table reports the average coefficients from monthly Fama-MacBeth cross-sectional regressions; variables are formed at the end of each month. The dependent variable is call option delta-hedged gain until month-end scaled by $(\Delta S - C)$ at the beginning of the period. CAR_4^2 is our measure for earnings announcement variance. Option open interest is the number of option contracts open at the beginning of the period. Stock volume is the stock trading volume over the previous month. Relative option bid-ask spread is the ratio of bid-ask spread to midpoint, at the beginning of the period. Illiquidity is the average of the daily Amihud (2002) stock illiquidity measure over the previous month. To adjust for serial correlation, robust Newey-West (1987) t -statistics are reported in brackets. All independent variables are winsorized each month at the 0.5% level. The sample period is from January 1996 to December 2017.

Variables	Model 1	Model 2	Model 3	Model 4	Model 5
Intercept	-0.0035 (-2.3251)	-0.0013 (-0.8615)	-0.0045 (-3.3041)	-0.0312 (-3.6961)	-0.0361 (-4.8726)
CAR_4^2	0.0650 (2.8286)	0.0600 (2.5916)	0.0646 (2.8243)	0.0867 (3.5699)	0.0824 (3.2979)
(Option open interest/stock volume)* 10^3		-0.0750 (-7.0651)			-0.0685 (-9.8331)
Relative option bid-ask spread			0.0050 (3.3566)		0.0084 (6.7450)
Ln(illiquidity)				-0.0013 (-3.5447)	-0.0016 (-4.9011)

Table 4.5 reports several robustness checks on our results. In previous regression tables, the dependent variable has been delta-hedged option returns measured until month-end. In Table 4.5, we also include results for if the holding period is changed to be until maturity. Models 1 and 2 of Table 4.5 report regression results using monthly delta-hedged returns, while Models 3 and 4 report the delta-hedged returns until maturity. Furthermore, the Table also reports regression results using put options instead of calls, Models 2 and 4. In all regression results the coefficient for CAR_4^2 remains positive (varying between 0.0780 and 0.0371) and statistically significant. Thus, not even controlling for stock, volatility or limits to arbitrage measures can explain the significant positive relation between delta-hedged option returns and our earnings announcement variance measure (CAR_4^2).

Table 4.5: Fama-Macbeth regressions, alternate measures of delta-hedged returns

This table reports the average coefficients from monthly Fama-MacBeth cross-sectional regressions; variables are formed at the end of each month. The dependent variable is call option delta-hedged gain until month-end scaled by $(\Delta S - C)$ at the beginning of the period for Model 1. For Model 2 the dependent variable is the put option delta-hedged gain until month-end scaled by $(P - \Delta S)$ at the beginning of the period. For Model 3 the dependent variable is call option delta-hedged gain until maturity scaled by $(\Delta S - C)$ at the beginning of the period. For Model 4 the dependent variable is the put option delta-hedged gain until maturity scaled by $(P - \Delta S)$ at the beginning of the period. All independent variables are the same as in Tables 4.2 to 4.4. To adjust for serial correlation, robust Newey-West (1987) t -statistics are reported in brackets. All independent variables are winsorized each month at the 0.5% level. The sample period is from January 1996 to December 2017.

Variables	Model 1	Model 2	Model 3	Model 4
Intercept	-0.0390 (-5.6556)	-0.0490 (-9.1520)	-0.1055 (-10.0755)	-0.1122 (-11.3499)
CAR_4^2	0.0760 (3.0995)	0.0780 (2.8024)	0.0371 (2.4939)	0.0476 (2.3585)
$IVOL_{exc. EAD}^2$	-0.0031 (-1.1997)	-0.0070 (-2.9483)	-0.0090 (-1.7084)	-0.0076 (-1.9374)
$SysVol_{exc. EAD}^2$	0.0028 (0.9303)	-0.0032 (-0.3379)	0.0030 (0.7417)	0.0076 (1.1762)
$r_{0,1}$	0.0013 (0.5082)	-0.0035 (-1.2319)	0.0016 (0.4876)	-0.0006 (-0.1312)
$r_{2,12}$	-0.0003 (-0.4486)	-0.0015 (-2.3344)	-0.0007 (-0.6406)	0.0038 (2.6932)
$\ln(ME)$	-0.0001 (-0.1114)	-0.0033 (-5.3060)	-0.0066 (-6.2536)	-0.0035 (-3.8157)
(Option open interest/stock volume)* 10^3	-0.0666 (-8.7257)	-0.0623 (-7.1261)	-0.0857 (-15.7616)	-0.1079 (-14.2921)
Relative option bid-ask spread	0.0088 (7.1549)	0.0278 (6.6863)	0.0043 (1.0834)	-0.0035 (-0.9279)
$\ln(\text{illiquidity})$	-0.0017 (-2.3005)	-0.0043 (-7.7667)	-0.0094 (-8.9176)	-0.0077 (-8.8591)

4.4 Earnings Announcement Option Trading Strategy

This section studies the relation between delta-hedged option returns and our earnings announcement variance measure (CAR_4^2) using a time series regression and the portfolio sorting approach. We confirm the results obtained by the Fama-MacBeth regressions, introduce an earnings announcement variance-based trading strategy and examine the impact of transaction costs on our option trading strategy. Our option trading strategy is formed as follows, at the end of each month, we rank firms with traded options into five quintiles based on their CAR_4^2 . The strategy then buys the delta-hedged options with firms ranked in the top quintile (quintile number five) and sells the delta-hedged options with firms ranked in the bottom quintile (quintile number one). As before, the delta-hedges are re-balanced daily. As in Section 4.3, for each optionable stock we choose only one call and put option that is closest to being at-the-money and has a common time to maturity of, on average, 48 calendar days. Equation (4.1) defines the long position of delta-hedged gain. The short position involves selling one contract of call option against a long position of Δ shares (where Δ is the Black-Scholes call option delta) of the underlying stock. We measure the performance over the next month.¹²

4.4.1 Time Series Regressions

The purpose of the time series regressions is essentially to see whether, after controlling for common risk factors and other known option trading strategies, the alpha on our option earnings announcement variance strategy is significant. We do so by regressing the returns of the option trading strategy onto the returns of explanatory strategies, including the Fama and French five-factors, momentum and idiosyncratic volatility trading strategy of Cao and Han (2013). Table 4.6 reports the results from the time series regressions. For the time series regression, we use the option value-weighted trading strategy (weighting by the market

¹²Section 4.5 investigates robustness to using returns held until maturity for which we achieve qualitatively similar results.

value of total option open interest multiplied by the midpoint). Model 1 includes only the market excess return (in excess of the risk-free rate); we measure the excess return (MKT-RF) as the value-weighted portfolio of all U.S.-based common stocks in the CRSP database, minus the risk-free rate.¹³ Model 1 demonstrates the alpha for our option earnings announcement variance trading strategy is statistically significant, with a positive alpha of 0.3850% and t -statistic 2.4300. Model 2 is a time series regression, using the Fama and French five-factors. Including the Fama and French five-factors has no powerful effect on the statistical significance of our trading strategy with an alpha of 0.4082% and t -statistic 2.5157. Model 3 includes the momentum strategy, loading negatively on momentum with statistical significance; however, there is little change to the magnitude or significance of the alpha. Model 4 reports results for controlling for the idiosyncratic volatility strategy of Cao and Han (2013) along with the Fama and French five-factors. The option trading strategy of Cao and Han (2013) is constructed as follows: at the end of each month, firms are ranked into five quintiles based on their idiosyncratic volatility; the strategy then buys the delta-hedged call options on firms ranked in the bottom quintile, and sells the delta-hedged call options on firms ranked in the top quintile. We find the idiosyncratic volatility trading strategy has little statistical significance, in contrast to Cao and Han (2013). However this is in line with research by Karakaya (2014) who, using an extended sample to Cao and Han (2013), find that the idiosyncratic volatility trading strategy is completely explained by the Fama and French three-factor model. Therefore including the idiosyncratic volatility trading strategy does not subsume our earnings announcement variance-based trading strategy. Model 5 reports a regression controlling for all strategies. Even with this large array of strategies implemented, it does not take away any significance from our trading strategy, with the alpha being statistically significant at 0.4409% with t -statistic of 2.7319.

¹³In practice, we obtain the daily excess returns for this portfolio from Ken French's website.

Table 4.6: Time series regressions

This table reports the results of monthly time series regressions on the return of our option trading strategy, using returns held until month-end, in the option value-weighting scheme. Our option trading strategy consists of selling delta-hedged calls on low earnings announcement variance firms (measured by CAR_4^2) and buying delta-hedged options on high earnings announcement firms. For this time series regression we use the option value-weighted scheme. The regressions include Fama-French (2015) risk factors (MKT-RF, SMB, HML, RMW, CMA), the momentum strategy (UMD), and the idiosyncratic volatility trading strategy of Cao and Han (2013) (IVOL). To adjust for serial correlation, robust Newey-West (1987) t -statistics are reported in brackets. All independent variables are winsorized each month at the 0.5% level. The sample period is from January 1996 to December 2017.

Variables	Model 1	Model 2	Model 3	Model 4	Model 5
Intercept	0.3850 (2.4300)	0.4082 (2.5157)	0.4435 (2.7631)	0.4037 (2.4760)	0.4409 (2.7319)
MKT-RF	-0.0938 (-2.6707)	-0.0797 (-1.9045)	-0.1054 (-2.4897)	-0.0800 (-1.9073)	-0.1055 (-2.4849)
SMB		-0.0864 (-1.5318)	-0.0665 (-1.1860)	-0.0862 (-1.5260)	-0.0665 (-1.1835)
HML		-0.0030 (-0.0437)	-0.0576 (-0.8079)	-0.0023 (-0.0336)	-0.0569 (-0.7964)
RMW		-0.1221 (-1.6144)	-0.0975 (-1.2976)	-0.1237 (-1.6296)	-0.0985 (-1.3049)
CMA		0.2762 (2.7369)	0.2955 (2.9626)	0.2744 (2.7109)	0.2945 (2.9415)
UMD			-0.0836 (-2.6738)		-0.0832 (-2.6496)
IVOL				0.0510 (0.3622)	0.0275 (0.1972)
Adj. R^2	0.0898	0.148	0.163	0.117	0.181

4.4.2 Average Portfolio Returns

Table 4.7 reports the average returns of five option value-weighted portfolios sorted on earnings announcement variance; each portfolio consists of buying delta-hedged calls. Table 4.7 also reports the "5-1" portfolio as the difference in the average returns of the top and bottom option value-weighted earnings announcement variance quintile portfolios, which is by definition the return of our earnings announcement variance-based trading strategy. Table 4.7 demonstrates that the average return of buying delta-hedged calls is negative (as demonstrated by Table 4.1) and increases with earnings announcement variance, corresponding to the significant positive relation between earnings announcement variance and delta-hedged option returns. Over the entire sample period, we find our option-based trading strategy has statistically and economically significant returns of 0.31% per month (with a t -statistic of

Table 4.7: Average monthly returns of option trading strategy

This table reports the average returns of the call option earnings announcement variance trading strategy, using returns until month-end. At the end of each month we rank all firms with options traded into five groups by their CAR_4^2 . Our option trading strategy consists of selling delta-hedged calls on low earnings announcement variance firms (measured by CAR_4^2), and buying delta-hedged options on high earnings announcement firms. We weight by the market value of the option open interest at the beginning of the period. All values are expressed as a percentage. To adjust for serial correlation, robust Newey-West (1987) t -statistics are reported in brackets. All independent variables are winsorized each month at the 0.5% level. The sample period is from January 1996 to December 2017.

Quintile	1-low	2	3	4	5-high	5-1
Average returns	-0.44 (-2.82)	-0.32 (-2.49)	-0.23 (-1.35)	-0.13 (-0.76)	-0.13 (-0.71)	0.31 (2.05)

2.05).¹⁴ In terms of placing the economic size of this effect with previous research consider the following. Table 4.6 demonstrates a monthly alpha of 0.44%, a monthly return of 0.44% corresponds to an annual return of 5.28% without compounding, which is approximately a third of the effect from the bid-ask spread (e.g., 13.44% as reported by Cao and Wei (2010)). Figure 4.1 displays a time series plot of the returns to our option trading strategy.

¹⁴Section 4.5 investigates the returns of our trading strategy using different weighting schemes and concludes our trading strategy is robust.

Figure 4.1: Time series of returns

This figure displays the monthly time series of the difference in the option value-weighted average returns of selling delta-hedged calls on firms ranked in the bottom earnings announcement variance quintile (as measured by CAR_4^2), and buying delta-hedged calls on firms ranked in the top quintile. The delta-hedged option positions are re-balanced daily to be delta-neutral. The earnings announcement variance CAR_4^2 is the square of the four-day excess return of the stock following an earnings announcement and is taken from the previous month. Our sample consists of short-term at-the-money call options on individual stocks from January 1996 to December 2017.




Table 4.8 investigates the effect of various double sorts on the profitability and significance of our trading strategy, using option value-weighted returns. The table also reports the intercept (α) from time series regressions using the Fama and French five-factor model, momentum and the idiosyncratic volatility trading strategy of Cao and Han (2013). The size row investigates further the effect size has on the profitability of our option trading strategy. Each month we first sort firms into five quintiles by their market capitalisation and

then, within each size quintile, we sort firms into one of five earnings announcement variance quintiles (as measured by CAR_4^2). We then take the average across all the size quintiles for each of the five earnings announcement variance quintiles as well as the 5-1 portfolio. The average 5-1 portfolio sorted on size is still statistically and economically significant at 0.34% per month, with a statistically significant time series regression alpha of 0.37% and t -statistic 2.46. The average profitability is similar to that of Table 4.7 (which reports an average return of 0.31% per month), suggesting that size does not drive our results. The second row sorts by total volatility and earnings announcement variance. The returns in all categories of earnings announcement variance, except the last, are statistically different from zero. The 5-1 portfolio is economically and statistically significant with 0.38% return per month, with a significant alpha of 0.48% and t -statistic of 3.27. The third row sorts by idiosyncratic volatility and earnings announcement variance. The returns in all but the last two categories of earnings announcement variance are statistically different from zero. The average 5-1 portfolio is still statistically and economically significant at 0.36% per month, with a statistically significant time series regression alpha of 0.46% and t -statistic 3.37. The fourth row sorts by relative option bid-ask spread and earnings announcement variance. The returns at all categories of earnings announcement variance are statistically different from zero. Furthermore, the average 5-1 portfolio is still statistically and economically significant at 0.19% per month, with a statistically significant time series regression alpha of 0.27% and t -statistic 2.09. Comparing the profitability from the average of the 5-1 portfolio to that of Table 4.7, we see a marked decrease with returns going from 0.31 down to 0.19 (decreasing by approximately 40%), demonstrating trading costs are highly relevant to the profitability of our option trading strategy. The fifth row sorts by stock illiquidity (using the Amihud (2002) measure of illiquidity) and earnings announcement variance. The average 5-1 portfolio is still statistically and economically significant at 0.43% per month, with a statistically significant time series regression alpha of 0.54% and t -statistic 3.07. Comparison with Table 4.7 suggests the profitability of our trading strategy is also highly sensitive to stock illiquidity. Thus, in line with Cao and Han (2013) we find from the investigation into stock illiquidity and option bid-ask spreads that limits to arbitrage from both the option contract and the underlying are highly relevant to profitability. The sixth row sorts by option open interest (formed at

the beginning of the period) and earnings announcement variance. The average 5-1 portfolio is still statistically and economically significant at 0.20% per month, with a statistically significant time series regression alpha of 0.24% and t -statistic 2.74. Demonstrating option open interest is an important part of the profitability to our trading strategy, however our results are not driven by option open interest.

In summary, the double sorts indicate little difference in profitability when accounting for the size of the firm, in line with the Fama-MacBeth regressions in Tables 4.2 and 4.5. We also confirm controlling for both total and idiosyncratic volatility has little effect on profitability, further supporting the conclusion from Tables 4.3 and 4.5 that our results are not driven by highly volatile firms. Table 4.8 also confirms that limits to arbitrage proxies play an important role on profitability, with both stock and option limits to arbitrage seeming highly relevant. Table 4.8 further shows that option open interest is of high relevance to the profitability of our option trading strategy

Table 4.8: Average monthly returns, double portfolio sorts

This table reports the average return of buying short-maturity at-the-money call options on firms with traded options, using the option value-weighting scheme. For each stock we buy one call option contract against a short position of Δ shares of the underlying stock (where Δ is the Black-Scholes call option delta). The delta-hedges are re-balanced daily. We compound the daily returns of the re-balanced delta-hedged call option positions until month-end to arrive at a return over the next month. At the end of each month we rank all firms into five groups by their CAR_4^2 . We then sort based on either the market capitalisation, total volatility, idiosyncratic volatility, relative option bid-ask spread, stock illiquidity (using the Amihud (2002) measure of illiquidity) or option open interest. All values are expressed as a percentage. To adjust for serial correlation, robust Newey-West (1987) t -statistics are reported in brackets. All independent variables are winsorized each month at the 0.5% level. The sample period is from January 1996 to December 2017.

Quintile	1-low	2	3	4	5-high	5-1	α
AVG. Size	-0.74 (-4.12)	-0.68 (-4.97)	-0.69 (-4.45)	-0.20 (-0.90)	-0.40 (-2.66)	0.34 (2.17)	0.37 (2.46)
AVG. VOL	-0.49 (-2.90)	-0.35 (-2.75)	-0.33 (-2.38)	-0.17 (-1.98)	-0.11 (-0.84)	0.38 (2.38)	0.48 (3.27)
AVG. IVOL	-0.47 (-2.79)	-0.27 (-1.99)	-0.32 (-1.99)	-0.21 (-1.52)	-0.11 (-0.74)	0.36 (2.49)	0.46 (3.37)
AVG. Bid-ask spread	-0.56 (-3.78)	-0.54 (-4.27)	-0.41 (-2.85)	-0.44 (-2.29)	-0.37 (-2.30)	0.19 (2.04)	0.27 (2.09)
AVG. Ln(illiquidity)	-0.82 (-5.34)	-0.69 (-4.40)	-0.64 (-4.47)	-0.34 (-1.72)	-0.39 (-2.32)	0.43 (2.58)	0.54 (3.07)
AVG. Open interest	-0.48 (-3.31)	-0.41 (-3.04)	-0.39 (-2.91)	-0.28 (-1.86)	-0.29 (-1.95)	0.20 (2.50)	0.24 (2.74)

4.4.3 Controlling for Trading Costs

For all of the prior results, we have made the standard assumption that options can be traded at the midpoint of the bid and ask price quotes. However, realising this is an idealised assumption, Table 4.9 investigates the impact of transaction costs on the profitability of our earnings announcement variance option trading strategy. To simulate the costs of trading options, we assume the effective option spread is equal to 10% or 25% of the quote bid-ask spread. We define effective spread as twice the difference between the execution price and the market quote at the time of order entry, i.e., the midpoint corresponds to an effective spread of 0%. The row "MidP" of Table 4.9 represents the zero effective spread returns, as

in all previous tables. Table 4.9 displays the average return of our trading strategy using the option value-weighting scheme.¹⁵ The final column of Table 4.9 also displays the alpha from time series regressions using the Fama and French five-factors, momentum and the idiosyncratic volatility strategy of Cao and Han (2013). Due to the first row using zero effective spread, it is the same as that from Table 4.7 regarding the returns. At an effective spread level of 10%, the profitability is slightly under statistical significance; however, the alpha remains economically and statistically significant at 0.37% per month, with a t -statistic of 2.48. Even when the effective spread is increased to 25% the alpha remains statistically significant at 0.26% per month, with a t -statistic of 1.98 (significant at the 5% level). While the statistical significance on the alpha of our trading strategy will clearly not be strong at much of a higher effective spread level, Muravyev and Pearson (2019) demonstrate that, at high frequencies, option traders are able to time their executions to buy when the option fair value is close to the ask, and sell close to the bid. Muravyev and Pearson (2019) find that option traders who exploit this predictability can reduce their effective spread to just under 30%, thus the fact that our trading strategy is significant at 25% indicates reasonable robustness to trading costs. Figure 4.2 displays a time series plot of the returns to our option trading strategy, using the option value-weighted scheme with effective spread levels of 10% and 25%.

4.5 Robustness Checks

This section details several robustness checks on our results. We investigate the effects of different weighting schemes on our option trading strategy. Secondly, we have only considered the profitability of our option trading strategy using options held until month-end (in line with Cao and Han (2013)) we investigate the further holding pattern of returns held until maturity. This section demonstrates that our results are robust to different weighting schemes and portfolio holding periods.

¹⁵Section 4.5 investigates the effect of trading costs using two further weighting schemes and achieves qualitatively similar results.

Table 4.9: Controlling for trading costs


This table reports the impact of trading costs of stock options on the profitability of our option trading strategy based on earnings announcement variance, using the option value-weighted scheme. In the first row (MidP), we assume the options are transacted at the midpoint of the bid and ask quotes (effective spread zero). The other two rows report varying levels of effective spread for the 10% and 25% levels. At the end of each month we rank all firms with options traded into five groups by their CAR_4^2 . Also reported is the intercept (α reported as a percentage) of a time series regression using the Fama-French (2015) five-risk factors, momentum and the idiosyncratic volatility trading strategy of Cao and Han (2013). To adjust for serial correlation, robust Newey-West (1987) t -statistics are reported in brackets. All independent variables are winsorized each month at the 0.5% level. The sample period is from January 1996 to December 2017.

Quintile	1-low	2	3	4	5-high	5-1	α
MidP	-0.44 (-2.82)	-0.32 (-2.49)	-0.23 (-1.35)	-0.13 (-0.76)	-0.13 (-0.71)	0.31 (2.05)	0.44 (2.73)
10%	-0.55 (-3.55)	-0.43 (-3.40)	-0.35 (-2.07)	-0.26 (-1.58)	-0.31 (-1.74)	0.24 (1.62)	0.37 (2.48)
25%	-0.71 (-4.68)	-0.60 (-4.77)	-0.52 (-3.17)	-0.47 (-2.81)	-0.59 (-3.29)	0.12 (1.09)	0.26 (1.98)

Table 4.10 reports results for two different weighting schemes of our earnings announcement variance-based trading strategy: equal-weighted and stock value-weighted (by the market capitalisation value of the underlying), using returns until month-end. The table also investigates the impact of transaction costs on the profitability of our earnings announcement variance option trading strategy using the new weighting schemes. Panel A reports the equal-weighting scheme. There is a linear decrease in the profitability of our trading strategy for increasing effective spread. At zero effective spread the strategy has profitability of 0.19%, with a t -statistic of 2.52, a time series regression alpha of 0.21% and t -statistic of 2.78. This demonstrates that our trading strategy is robust to using an equal weighting scheme instead of option value-weighted. However, this decreases to a profitability of 0.08% and alpha of 0.09%, with neither being statistically significant, for an effective spread of 10%. This demonstrates that the effect is not as robust to trading costs using an equal-weighting scheme rather than the option value-weighting scheme. Panel B reports the stock-value weighting scheme. There is again a linear decrease in the profitability of our trading strategy for increasing effective spread. At zero effective spread, the strategy has profitability of

Figure 4.2: Time series of returns with effective spread

This figure displays the monthly time series of the difference in the option value-weighted average returns of selling delta-hedged calls on firms ranked in the bottom earnings announcement variance quintile (as measured by CAR_4^2), and buying delta-hedged calls on firms ranked in the top quintile. The delta-hedged option positions are re-balanced daily to be delta-neutral. The earnings surprise measure CAR_4^2 is the square of the four-day excess return of the stock following an earnings announcement and is taken from the previous month. The first panel uses an effective spread of 10%, the second panel uses an effective spread of 25%. Our sample consists of short-term at-the-money call options on individual stocks from January 1996 to December 2017.



0.14%, with a significant t -statistic of 2.15, a statistically significant time series regression alpha of 0.23% and t -statistic of 2.66. For an effective spread level of 10%, the alpha remains statistically significant, at the 5% level, at 0.17%, with a t -statistic of 1.98. However, for an effective spread level of 25%, neither the profitability nor the alpha are statistically significant. In summary, we have demonstrated that our earnings announcement variance-based option trading strategy is robust to different weighting schemes.

Table 4.10: Controlling for trading costs, alternate weighting schemes

This table reports the effect of using different weighting schemes while investigating the impact of trading costs of stock options on the profitability of our option trading strategy, based on earnings announcement variance, using returns until month-end. In the first row (MidP), we assume the options are transacted at the midpoint of the bid and ask quotes (effective spread zero). The other two rows report varying levels of effective spread for the 10% and 25% levels. At the end of each month, we rank all firms with options traded into five groups by their CAR_4^2 . Also reported is the intercept (α reported as a percentage) of a time series regression using the Fama-French (2015) five-risk factors, momentum and the idiosyncratic volatility trading strategy of Cao and Han (2013). To adjust for serial correlation, robust Newey-West (1987) t -statistics are reported in brackets. All independent variables are winsorized each month at the 0.5% level. The sample period is from January 1996 to December 2017.

Quintile	1-low	2	3	4	5-high	5-1	α
Panel A. Equal weight							
MidP	-0.55 (-3.02)	-0.54 (-3.28)	-0.57 (-3.15)	-0.51 (-2.71)	-0.36 (-1.57)	0.19 (2.52)	0.21 (2.78)
10%	-0.71 (-4.77)	-0.71 (-5.25)	-0.75 (-5.03)	-0.72 (-4.74)	-0.63 (-3.93)	0.08 (1.05)	0.09 (1.23)
25%	-1.11 (-7.39)	-1.11 (-8.21)	-1.18 (-7.84)	-1.19 (-7.75)	-1.20 (-7.47)	-0.10 (-1.30)	-0.09 (-1.16)
Panel B. Stock value-weight							
MidP	-0.15 (-1.25)	-0.14 (-1.27)	-0.09 (-0.68)	-0.03 (-0.20)	-0.01 (-0.08)	0.14 (2.15)	0.23 (2.66)
10%	-0.24 (-1.99)	-0.23 (-2.10)	-0.18 (-1.44)	-0.14 (-1.06)	-0.17 (-1.12)	0.07 (1.76)	0.17 (1.98)
25%	-0.37 (-3.12)	-0.37 (-3.37)	-0.32 (-2.61)	-0.31 (-2.35)	-0.40 (-2.71)	-0.03 (-0.31)	0.06 (0.74)

Next, we move onto investigating the profitability of our option trading strategy when forming portfolios at the beginning of the month and holding until maturity, using the option value-weighting scheme. Table 4.11 reports results from time series regressions. Model 1 demonstrates that the alpha for our option earnings announcement variance trading strategy is statistically significant, with a positive alpha of 0.3717% and t -statistic 2.3501. Model 2 is a time series regression using the Fama and French five-factors. Including the Fama and French five-factors reduces the alpha slightly to 0.3173% and is still statistically significant, with t -

statistic 2.4204. Model 3 includes the momentum strategy, loading negatively on momentum with a statistically significant beta of -0.0795 and t -statistic -2.1342 (in agreement with Table 4.6). Even with inclusion of momentum, there is still statistical significance in the alpha at 0.3461% , with t -statistic 2.3181 . Model 4 reports results which control for the idiosyncratic volatility strategy of Cao and Han (2013) along with the Fama and French five-factors. Once again, we do not find that the earnings announcement variance-based trading strategy is subsumed by the idiosyncratic volatility trading strategy of Cao and Han (2013). Model 5 reports a regression controlling for all strategies. Even with this vast array of strategies implemented, it does not take away any significance from our trading strategy, with the alpha being statistically significant at 0.3473% , with t -statistic of 2.9517 .

Table 4.11: Time series regressions, held until maturity

This table reports the results of monthly time series regressions on the return of our option trading strategy, using returns held until maturity, in the option value-weighting scheme. Our option trading strategy consists of selling delta-hedged calls on low earnings announcement variance firms (measured by CAR_4^2) and buying delta-hedged options on high earnings announcement firms. For this time series regression we use the option value-weighted scheme. The regressions include Fama-French (2015) risk factors (MKT-RF, SMB, HML, RMW, CMA), the momentum strategy (UMD), and the idiosyncratic volatility trading strategy of Cao and Han (2013) (IVOL). To adjust for serial correlation, robust Newey-West (1987) t -statistics are reported in brackets. All independent variables are winsorized each month at the 0.5% level. The sample period is from January 1996 to December 2017.

Variables	Model 1	Model 2	Model 3	Model 4	Model 5
Intercept	0.3717 (2.3501)	0.3173 (2.4204)	0.3461 (2.3181)	0.3189 (2.1751)	0.3473 (2.9517)
MKT-RF	-0.0318 (-0.7894)	-0.0069 (-0.1375)	-0.0279 (-0.5475)	-0.0076 (-0.1500)	-0.0283 (-0.5542)
SMB		0.0150 (0.2284)	0.0269 (0.4103)	0.0135 (0.2038)	0.0257 (0.3899)
HML		-0.0172 (-0.2110)	-0.0706 (-0.8317)	-0.0171 (-0.2084)	-0.0703 (-0.8264)
RMW		-0.0087 (-0.0967)	0.0085 (0.0940)	-0.0087 (-0.0956)	0.0085 (0.0939)
CMA		0.1948 (1.6747)	0.2222 (1.9131)	0.1930 (1.6535)	0.2208 (1.8937)
UMD			-0.0795 (-2.1342)		-0.0793 (-2.1219)
IVOL				-0.0388 (-0.2750)	-0.0289 (-0.2060)
Adj. R^2	0.0015	0.0010	0.0014	0.0048	0.0097

Table 4.12 investigates the impact of trading costs on the profitability of our earnings announcement variance option trading strategy, using returns held until maturity and the option value-weighting scheme. The final column of Table 4.12 also displays the alpha from time series regressions using the Fama and French five-factors, momentum and the idiosyncratic volatility strategy of Cao and Han (2013). There is a linear decrease in the profitability of our trading strategy for increasing effective spread. At zero effective spread, the strategy has profitability of 0.35% with a t -statistic of 2.40, a time series regression alpha of 0.35% and t -statistic of 2.95; both the profitability and alpha are economically and

Table 4.12: Controlling for trading costs, held until maturity

This table reports the impact of trading costs of stock options on the profitability of our option trading strategy based on earnings announcement variance, using returns until maturity. The first row (MidP) we assume the options are transacted at the midpoint of the bid and ask quotes (effective spread zero). The other two rows report varying levels of effective spread for the 10% and 25% levels. At the end of each month we rank all firms with options traded into five groups by their CAR_4^2 . Also reported is the intercept (α reported as a percentage) of a time series regression using the Fama-French (2015) five-risk factors, momentum and the idiosyncratic volatility trading strategy of Cao and Han (2013). To adjust for serial correlation, robust Newey-West (1987) t -statistics are reported in brackets. All independent variables are winsorized each month at the 0.5% level. The sample period is from January 1996 to December 2017.

Quintile	1-low	2	3	4	5-high	5-1	α
MidP	-0.73 (-4.44)	-0.54 (-3.50)	-0.52 (-2.68)	-0.54 (-2.79)	-0.38 (-1.92)	0.35 (2.40)	0.35 (2.95)
10%	-0.79 (-4.79)	-0.60 (-3.88)	-0.58 (-3.01)	-0.61 (-3.15)	-0.48 (-2.40)	0.31 (2.14)	0.31 (2.46)
25%	-0.88 (-5.33)	-0.69 (-4.45)	-0.67 (-3.50)	-0.72 (-3.70)	-0.62 (-3.11)	0.26 (1.74)	0.26 (2.14)

statistically significant. This decreases to a profitability of 0.31% (t -statistic 2.14) and alpha of 0.31% (t -statistic 2.46) for an effective spread of 10%; thus at this level, the profitability and alpha are still statistically significant. At an effective spread level of 25%, only the alpha remains statistically significant at 0.26%, with t -statistic of 2.14. In summary, our trading strategy is robust to using different portfolio holding periods and transaction costs.

4.6 Discussion

The aim of this section is to offer a possible economical explanation as to why the relationship between earnings announcement variance and delta-hedged option returns is positive. The sign of this result is potentially at odds with economic instinct as previous studies have found there is a negative relationship between stock volatility and delta-hedged option returns (Hu and Jacobs (2016) for example), furthermore studies have shown (and we confirm) that total idiosyncratic volatility also has a negative relationship with delta-hedged option returns,

e.g., Cao and Han (2013).

It could be that this positive relationship is caused by demand-pressure effects. Table 4.4 demonstrates (consistent with Cao and Han (2013)) that delta-hedged option returns decrease with option open interest, supporting the idea that, due to limits to arbitrage, option market makers charge higher premiums for options with high end-user demand. However, as Table 4.13 demonstrates, cross-sectional averages of option open interest sorted by earnings announcement variance decreases with earnings announcement variance by approximately 35%.¹⁶ Therefore, investors demonstrate less demand for options on high earnings announcement variance firms. Thus, consistent with the negative relationship between delta-hedged option returns and option open interest it is self-consistent that earnings announcement variance is proportional to delta-hedged option returns. Consistent with demand-pressure effects (for instance, Gârleanu et al. (2009)). Furthermore, Christoffersen et al. (2017) demonstrate that market makers are net long positions in equity options, suggesting that a plausible economical explanation of our results is that market makers charge a lower price for options on high earnings announcement variance firms due to the lack of demand and given their net position is long they would prefer to buy for a price below the "fair value". However, as Table 4.8 demonstrates, our results are not purely driven by open interest, therefore while demand-pressure effects seem reasonable they do not explain our results.

As further research it would be instructive to also study the relation between our earnings announcement variance trading strategy with an alternative strategy based on the price of volatility risk. Duarte and Jones (2007) demonstrate that the conditional volatility risk premium varies with market volatility. Furthermore, the implied volatility distribution is majorly effected by earnings announcements causing non-normality (see Dubinsky and Johannes (2006)). It could be that this would cause a difference between implied and historical volatility which is being captured by the earnings announcement variance measure. However, the issue with this type of approach, considering the volatility risk premium, is one of liquidity. As in order to construct the variance risk premium one has to compute the expected return variation, this integral is calculated numerically, to ensure reliability it is

¹⁶We achieve qualitatively similar results if we use the option open interest stock volume ratio, as used in Table 4.4.

Table 4.13: Investor demand sorted by earnings announcement variance

This table reports the average option open interest, sorted by earnings announcement variance. Option open interest is the number of option contracts open at the beginning of the period. To adjust for serial correlation, robust Newey-West (1987) t -statistics are reported in brackets. All independent variables are winsorized each month at the 0.5% level. The sample period is from January 1996 to December 2017.

Quintile	1-low	2	3	4	5-high
Option open interest/1000	1.14	1.12	1.02	0.90	0.75
	(12.86)	(12.99)	(13.78)	(13.81)	(14.73)

standard to use two in-money options, one at-the-money and two out-of-the-money options. This data requirement means an approximate halving of firms which can be considered (see Cao and Han (2013)).

4.7 Conclusion

This work provides a comprehensive study of delta-hedged individual stock option returns. The key novel finding is that the average delta-hedged option return is increasing with an increase in the earnings announcement variance of the underlying stock. Our finding is robust to call and put options. It is statistically and economically significant. Furthermore, our trading strategy is robust to different portfolio formation dates and lengths of holding. For example, when using delta-hedged call option returns over the next month, options ranked in the top earnings announcement variance quintile, on average, out-perform the portfolio of delta-hedged call options on firms ranked in the bottom earnings announcement variance quintile by 0.31% per month, using an option value-weighted scheme. A time series regression exercise demonstrates our trading strategy yields a robust and statistically significant alpha of approximately 0.44%, per month, after controlling for the Fama-French five-factors, momentum and the idiosyncratic volatility trading strategy of Cao and Han (2013). We suggest that the positive relationship between delta-hedged option returns and earnings announcement variance is consistent with demand-pressure effects. Given market makers are

net long positions in equity options they thus charge a lower premium, coupled with the decreased demand for options with high earnings announcement variance potentially explains the positive sign between the idiosyncratic volatility component to earnings announcements and delta-hedged option returns.

Chapter 5

Conclusion

In the second chapter of this thesis we concern ourselves with aspects of theoretical option pricing, specifically non-affine model specifications. We consider a general one-factor stochastic volatility model with a non-square root dependence on volatility in the variance process. Using perturbation theory we construct a first order approximation solution to this non-affine model. In a simulated Monte Carlo time series estimation exercise our solution performs very well and is able to generate accurate prices with very low computational burden.

Tying into model specification concerns from the second chapter, the third chapter undertakes an empirical investigation into novel aspects of high-frequency option pricing model specifications. The task of pricing options in a high-frequency intradaily setting opens the door to a plethora of new phenomena, one of which is market maker biases to volatility accrual. We develop a novel approach to traditional option pricing by incorporating a mechanism to account for the change in sign of option returns from day to night. We find our model reduces errors in short term options by approximately 7% percent out-of-sample and seems to carry more emphasis in times of market turmoil. An additional potentially important aspect of intradaily option pricing model specification is the effect intradaily macroeconomic news announcements have on options. As the timing of such news events is known ex-ante we account for such an effect by inclusion of a deterministic jump process in our model specification to empirically test the importance of FOMC announcements. Our empirical

results show FOMC announcements are an important part of option pricing. We also find additional evidence that indicates FOMC announcements are particularly important during times of economic distress. We conclude that macroeconomic announcements are potentially highly relevant to option pricing models.

The fourth chapter of this thesis links to the above work concerning the importance of news announcements on option pricing, we examine the impact of earnings announcements on individual equity options. We study delta-hedged option returns within a six month aftermath of earnings announcements and find that, on average, firms with high earnings announcement volatility out-perform firms with low earnings announcement volatility, over the period 1996 to 2017. We perform a time series regression exercise which demonstrates our trading strategy yields a robust alpha of approximately 0.44%, per month, after controlling for the Fama-French five-factors, momentum and the idiosyncratic volatility trading strategy of Cao and Han (2013). We suggest the positive relationship between delta-hedged option returns and earnings announcement volatility is consistent with demand-based models and market makers being net long positions in individual equity options, thus charging a lower premium.

I will now provide some details for further research. Evidence in Chapter 2 clearly demonstrates that the concept of perturbation theory is a fruitful methodology for option pricing models. Further research might be conducted along the dimension of considering different base-case models. In Chapter 2 we consider a base-case of the Black-Scholes (1973) model, however as noted previously this model has the significant short coming of constant volatility. Thus further improvements could be made by considering a stochastic volatility base case model, for example the Heston (1993) model, this would be achieved by setting $\alpha = 1/2$ in the Merton-Garman model and using Taylor series expansions for the two α dependent terms.¹

Regarding Chapter 3 it would be instructive to consider different modelling approaches

¹Of course by setting $\alpha = 1$ one could do a similar approach using the Hull and White (1987) model as a base-case. However, the Heston (1993) model has the advantage of an analytical solution to the price of the option, with only numerical methods required to compute the integral, whereas the Hull and White model is more challenging.

towards the effect of macroeconomic news announcements, specifically FOMC announcements. There is evidence from the literature which suggests the Federal Reserve is averse to surprising investors (e.g., Stein and Sunderam (2018)). To combat this potential surprise on FOMC announcement date the Federal Reserve also uses discount rate meetings to guide investors' expectations and to also affirm that the Fed will choose the path investors prefer (e.g., Cieslak et al. (2019)). As such one idea to deal with the small-sample issues would be to use the same modelling methodology but to consider the biweekly discount rate meetings as well. Separately, one could consider an alternative modelling choice entirely. Asset prices will move to the surprise news component, implying the existence of news-related jumps and their magnitude could be state-dependent (i.e., how surprised markets are by the news). An alternative modelling methodology could investigate jump volatility dependent on a function of the news surprise (see Kuttner (2001)). Although this new methodology brings with it a major downside that one would lose the nature of the deterministic approach used in this research.

In relation to Chapter 4 it would be interesting to investigate the effect short sale constraints might have on the earnings announcement variance strategy. There is a substantial debate around the effects of short selling enhancing price efficiency around earnings announcements. This debate initially started with Ball and Brown (1968) and sees sporadic evidence on both sides. For example, Cao et al. (2007) finds weak evidence to suggest short sellers reduce drift after earnings announcements. However, McKenzie (2012) finds short selling increases after negative earnings shocks, however concludes it does not remove long-term PEAD, measured over the following quarter. However, using daily data on shorting flows Boehmer and Wu (2013) demonstrate that short sellers promote efficient pricing, thus removing the drift over the next quarter for negative surprises. Based on evidence provided by Christoffersen et al. (2017), that market makers hold net long positions in equity options, it could be they are buying calls from end users and hedging by short selling the underlying. There exists the possibility that shorting the underlying stock is more difficult for portfolios with low earnings announcement variance, meaning that a temporary misalignment between stocks and option embedded information might be a contributing factor to the returns of the strategy.

Bibliography

.
. .
. .
. .

- Aguilar, J.-P. (2017). A series representation for the black-scholes formula. *arXiv preprint arXiv:1710.01141*.
- Ai, Y., Kimmel, R. Et al. (2007). Maximum likelihood estimation of stochastic volatility models. *Journal of financial economics*, 83(2), 413–452.
- Ait-Sahalia, Y., & Lo, A. W. (1998). Nonparametric estimation of state-price densities implicit in financial asset prices. *The Journal of Finance*, 53(2), 499–547.
- Ait-Sahalia, Y., Mykland, P. A., & Zhang, L. (2005). How often to sample a continuous-time process in the presence of market microstructure noise. *The review of financial studies*, 18(2), 351–416.
- Alizadeh, S., Brandt, M. W., & Diebold, F. X. (2002). Range-based estimation of stochastic volatility models. *The Journal of Finance*, 57(3), 1047–1091.
- Amengual, D., & Xiu, D. (2018). Resolution of policy uncertainty and sudden declines in volatility. *Journal of Econometrics*, 203(2), 297–315.
- Amihud, Y. (2002). Illiquidity and stock returns: Cross-section and time-series effects. *Journal of financial markets*, 5(1), 31–56.
- Amihud, Y., & Mendelson, H. (1986). Asset pricing and the bid-ask spread. *Journal of financial Economics*, 17(2), 223–249.

- Ammann, M., & Feser, A. (2019). Robust estimation of risk-neutral moments. *Journal of Futures Markets*, 39(9), 1137–1166.
- An, B.-J., Ang, A., Bali, T. G., & Cakici, N. (2014). The joint cross section of stocks and options. *The Journal of Finance*, 69(5), 2279–2337.
- Andersen, T. G., Benzoni, L., & Lund, J. (2002). An empirical investigation of continuous-time equity return models. *The Journal of Finance*, 57(3), 1239–1284.
- Andersen, T. G., Bollerslev, T., Diebold, F. X., & Ebens, H. (2001). The distribution of realized stock return volatility. *Journal of financial economics*, 61(1), 43–76.
- Andersen, T. G., Bollerslev, T., Diebold, F. X., & Labys, P. (2003). Modeling and forecasting realized volatility. *Econometrica*, 71(2), 579–625.
- Andersen, T. G., Bollerslev, T., Diebold, F. X., & Vega, C. (2007). Real-time price discovery in global stock, bond and foreign exchange markets. *Journal of international Economics*, 73(2), 251–277.
- Andersen, T. G., Fusari, N., & Todorov, V. (2017). Short-term market risks implied by weekly options. *The Journal of Finance*, 72(3), 1335–1386.
- Ang, A., Hodrick, R. J., Xing, Y., & Zhang, X. (2006). The cross-section of volatility and expected returns. *The Journal of Finance*, 61(1), 259–299.
- Atilgan, Y. (2014). Volatility spreads and earnings announcement returns. *Journal of Banking & Finance*, 38, 205–215.
- Baaquie, B. E. (1997). A path integral approach to option pricing with stochastic volatility: Some exact results. *Journal de Physique I*, 7(12), 1733–1753.
- Baaquie, B. E., Coriano, C., & Srikant, M. (2003). Quantum mechanics, path integrals and option pricing: Reducing the complexity of finance, In *Nonlinear physics: Theory and experiment ii*. World Scientific.
- Bachelier, L. (1900). *Théorie de la spéculation*. Gauthier-Villars.
- Bakshi, G., Cao, C., & Chen, Z. (1997). Empirical performance of alternative option pricing models. *The Journal of finance*, 52(5), 2003–2049.
- Bakshi, G., & Kapadia, N. (2003). Delta-hedged gains and the negative market volatility risk premium. *The Review of Financial Studies*, 16(2), 527–566.

- Bali, T. G., Cakici, N., & Whitelaw, R. F. (2011). Maxing out: Stocks as lotteries and the cross-section of expected returns. *Journal of Financial Economics*, 99(2), 427–446.
- Bali, T. G., & Hovakimian, A. (2009). Volatility spreads and expected stock returns. *Management Science*, 55(11), 1797–1812.
- Ball, R., & Brown, P. (1968). An empirical evaluation of accounting income numbers. *Journal of accounting research*, 159–178.
- Bardgett, C., Gourier, E., & Leippold, M. (2019). Inferring volatility dynamics and risk premia from the s&p 500 and vix markets. *Journal of Financial Economics*, 131(3), 593–618.
- Barndorff-Nielsen, O. E., & Shephard, N. (2006). Econometrics of testing for jumps in financial economics using bipower variation. *Journal of financial Econometrics*, 4(1), 1–30.
- Bates, D. S. (1991). The crash of 87: Was it expected? the evidence from options markets. *The journal of finance*, 46(3), 1009–1044.
- Bates, D. S. (1996). 20 testing option pricing models. *Handbook of statistics*, 14, 567–611.
- Bates, D. S. (2000). Post-'87 crash fears in the s&p 500 futures option market. *Journal of Econometrics*, 94(1-2), 181–238.
- Bates, D. S. (2006). Maximum likelihood estimation of latent affine processes. *The Review of Financial Studies*, 19(3), 909–965.
- Beber, A., & Brandt, M. W. (2006). The effect of macroeconomic news on beliefs and preferences: Evidence from the options market. *Journal of Monetary Economics*, 53(8), 1997–2039.
- Bekaert, G., & Engstrom, E. (2017). Asset return dynamics under habits and bad environment—good environment fundamentals. *Journal of Political Economy*, 125(3), 713–760.
- Benzoni, L. (2002). *Pricing options under stochastic volatility: An empirical investigation* (tech. rep.). working paper, University of Minnesota.
- Berkman, H., & McKenzie, M. D. (2012). Earnings announcements: Good news for institutional investors and short sellers. *Financial Review*, 47(1), 91–113.
- Bernanke, B. S., & Kuttner, K. N. (2005). What explains the stock market's reaction to federal reserve policy? *The Journal of finance*, 60(3), 1221–1257.

- Bernard, V. L., & Thomas, J. K. (1989). Post-earnings-announcement drift: Delayed price response or risk premium? *Journal of Accounting research*, 27, 1–36.
- Berry, T. D., & Howe, K. M. (1994). Public information arrival. *The Journal of Finance*, 49(4), 1331–1346.
- Bessembinder, H., & Hertzel, M. G. (1993). Return autocorrelations around nontrading days. *The Review of Financial Studies*, 6(1), 155–189.
- Black, F. (1975). Fact and fantasy in the use of options. *Financial Analysts Journal*, 31(4), 36–41.
- Black, F., & Scholes, M. (1973). The pricing of options and corporate liabilities. *Journal of political economy*, 81(3), 637–654.
- Blau, B. M., Bowles, T. B., & Whitby, R. J. (2016). Gambling preferences, options markets, and volatility. *Journal of Financial and Quantitative Analysis*, 51(2), 515–540.
- Blazhyevskiy, L., & Yanishevsky, V. (2011). The path integral representation kernel of evolution operator in merton-garman model. *arXiv preprint arXiv:1106.5143*.
- Bloom, N. (2009). The impact of uncertainty shocks. *econometrica*, 77(3), 623–685.
- Boehmer, E., & Wu, J. (2013). Short selling and the price discovery process. *The Review of Financial Studies*, 26(2), 287–322.
- Bogousslavsky, V. (2018). The cross-section of intraday and overnight returns. *Available at SSRN 2869624*.
- Bollen, N. P., & Whaley, R. E. (2004). Does net buying pressure affect the shape of implied volatility functions? *The Journal of Finance*, 59(2), 711–753.
- Bose, S. (1995). Representations of the $(2+1)$ -dimensional galilean group. *Journal of Mathematical Physics*, 36(2), 875–890.
- Boyer, B. H., & Vorkink, K. (2014). Stock options as lotteries. *The Journal of Finance*, 69(4), 1485–1527.
- Broadie, M., Chernov, M., & Johannes, M. (2007). Model specification and risk premia: Evidence from futures options. *The Journal of Finance*, 62(3), 1453–1490.
- Byun, S.-J., & Kim, D.-H. (2016). Gambling preference and individual equity option returns. *Journal of Financial Economics*, 122(1), 155–174.

- Cao, C., Chen, Z., & Griffin, J. M. (2005). Informational content of option volume prior to takeovers. *The Journal of Business*, 78(3), 1073–1109.
- Cao, J., & Han, B. (2013). Cross section of option returns and idiosyncratic stock volatility. *Journal of Financial Economics*, 108(1), 231–249.
- Cao, M., & Wei, J. (2010). Option market liquidity: Commonality and other characteristics. *Journal of Financial Markets*, 13(1), 20–48.
- Carr, P., & Wu, L. (2007). Stochastic skew in currency options. *Journal of Financial Economics*, 86(1), 213–247.
- Carr, P., & Wu, L. (2008). Variance risk premiums. *The Review of Financial Studies*, 22(3), 1311–1341.
- CBOE, C. B. O. E. (2009). The cboe volatility index–vix. *White paper*.
- Chan, K., Chung, Y. P., & Johnson, H. (1995). The intraday behavior of bid-ask spreads for nyse stocks and cboe options. *Journal of Financial and Quantitative Analysis*, 30(3), 329–346.
- Chan, L. K., Jegadeesh, N., & Lakonishok, J. (1996). Momentum strategies. *The Journal of Finance*, 51(5), 1681–1713.
- Chang, E. C., Pinegar, J. M., & Ravichandran, R. (1993). International evidence on the robustness of the day-of-the-week effect. *Journal of Financial and quantitative Analysis*, 28(4), 497–513.
- Chen, H., & Singal, V. (2003). Role of speculative short sales in price formation: The case of the weekend effect. *The Journal of Finance*, 58(2), 685–705.
- Chen, J., & Clements, A. (2007). S&p 500 implied volatility and monetary policy announcements. *Finance Research Letters*, 4(4), 227–232.
- Chernov, M. (2007). On the role of risk premia in volatility forecasting. *Journal of Business & Economic Statistics*, 25(4), 411–426.
- Chernov, M., Gallant, A. R., Ghysels, E., & Tauchen, G. (2003). Alternative models for stock price dynamics. *Journal of Econometrics*, 116(1-2), 225–257.
- Chourdakis, K., & Dotsis, G. (2011). Maximum likelihood estimation of non-affine volatility processes. *Journal of Empirical Finance*, 18(3), 533–545.

- Choy, S. K., & Wei, J. (2020). Liquidity risk and expected option returns. *Journal of Banking & Finance*, 111, 105700.
- Christoffersen, P., Feunou, B., Jacobs, K., & Meddahi, N. (2014). The economic value of realized volatility: Using high-frequency returns for option valuation. *Journal of Financial and Quantitative Analysis*, 49(3), 663–697.
- Christoffersen, P., Goyenko, R., Jacobs, K., & Karoui, M. (2017). Illiquidity premia in the equity options market. *The Review of Financial Studies*, 31(3), 811–851.
- Christoffersen, P., Heston, S., & Jacobs, K. (2009). The shape and term structure of the index option smirk: Why multifactor stochastic volatility models work so well. *Management Science*, 55(12), 1914–1932.
- Christoffersen, P., Jacobs, K., & Mimouni, K. (2010). Volatility dynamics for the s&p500: Evidence from realized volatility, daily returns, and option prices. *The Review of Financial Studies*, 23(8), 3141–3189.
- Cieslak, A. (2018). Short-rate expectations and unexpected returns in treasury bonds. *The Review of Financial Studies*, 31(9), 3265–3306.
- Cieslak, A., Morse, A., & Vissing-Jorgensen, A. (2019). Stock returns over the fomc cycle. *The Journal of Finance*, 74(5), 2201–2248.
- Cieslak, A., & Vissing-Jorgensen, A. (2020). *The economics of the fed put* (tech. rep.). National Bureau of Economic Research.
- Cliff, M., Cooper, M. J., & Gulen, H. (2008). Return differences between trading and non-trading hours: Like night and day. *SSRN eLibrary*, 32.
- Coval, J. D., & Shumway, T. (2001). Expected option returns. *The journal of Finance*, 56(3), 983–1009.
- Cox, J. C., & Ross, S. A. (1976). The valuation of options for alternative stochastic processes. *Journal of financial economics*, 3(1-2), 145–166.
- Cremers, M., & Weinbaum, D. (2010). Deviations from put-call parity and stock return predictability. *Journal of Financial and Quantitative Analysis*, 45(2), 335–367.
- Daniel, K., Hirshleifer, D., & Sun, L. (2020). *Short-and long-horizon behavioral factors* (tech. rep.). National Bureau of Economic Research.

- Dew-Becker, I., Giglio, S., & Kelly, B. T. (2018). Hedging macroeconomic and financial uncertainty and volatility. *Available at SSRN*.
- Diamond, D. W., & Verrecchia, R. E. (1987). Constraints on short-selling and asset price adjustment to private information. *Journal of Financial Economics*, 18(2), 277–311.
- Driessen, J., Maenhout, P. J., & Vilkov, G. (2009). The price of correlation risk: Evidence from equity options. *The Journal of Finance*, 64(3), 1377–1406.
- Du, B., Fung, S., & Loveland, R. (2018). The informational role of options markets: Evidence from fomc announcements. *Journal of Banking & Finance*, 92, 237–256.
- Duan, J.-C., Popova, I., Ritchken, P., Et al. (2002). Option pricing under regime switching. *Quantitative Finance*, 2(116-132), 209.
- Duan, J.-C., & Wei, J. (2008). Systematic risk and the price structure of individual equity options. *The Review of Financial studies*, 22(5), 1981–2006.
- Duan, J.-C., & Yeh, C.-Y. (2010). Jump and volatility risk premiums implied by vix. *Journal of Economic Dynamics and Control*, 34(11), 2232–2244.
- Duarte, J., & Jones, C. S. (2007). The price of market volatility risk, In *Afa 2009 san francisco meetings paper*.
- Dubinsky, A., Johannes, M., Kaeck, A., & Seeger, N. J. (2018). Option pricing of earnings announcement risks. *The Review of Financial Studies*, 32(2), 646–687.
- Dubinsky, A., Johannes, M. Et al. (2006). Fundamental uncertainty, earning announcements and equity options, In *Working paper*.
- Duffie, D., Pan, J., & Singleton, K. (2000). Transform analysis and asset pricing for affine jump-diffusions. *Econometrica*, 68(6), 1343–1376.
- Eraker, B. (2004). Do stock prices and volatility jump? reconciling evidence from spot and option prices. *The Journal of Finance*, 59(3), 1367–1403.
- Eraker, B., Johannes, M., & Polson, N. (2003). The impact of jumps in volatility and returns. *The Journal of Finance*, 58(3), 1269–1300.
- Fama, E. F., & French, K. R. (1993). Common risk factors in the returns on stocks and bonds. *Journal of financial economics*, 33(1), 3–56.
- Fama, E. F., & French, K. R. (2015). A five-factor asset pricing model. *Journal of financial economics*, 116(1), 1–22.

- Fang, F., & Oosterlee, C. W. (2009). Pricing early-exercise and discrete barrier options by fourier-cosine series expansions. *Numerische Mathematik*, 114(1), 27.
- Fernandez-Perez, A., Frijns, B., & Tourani-Rad, A. (2017). When no news is good news—the decrease in investor fear after the fomc announcement. *Journal of Empirical Finance*, 41, 187–199.
- Fields, M. J. (1931). Stock prices: A problem in verification. *The Journal of Business of the University of Chicago*, 4(4), 415–418.
- Foster, F. D., & Viswanathan, S. (1993). Variations in trading volume, return volatility, and trading costs: Evidence on recent price formation models. *The Journal of Finance*, 48(1), 187–211.
- French, K. R., & Roll, R. (1986). Stock return variances: The arrival of information and the reaction of traders. *Journal of financial economics*, 17(1), 5–26.
- Fuhrer, J. C., Olivei, G. P. Et al. (2011). The estimated macroeconomic effects of the federal reserve’s large-scale treasury purchase program. *Public Policy Briefs*, 2.
- Fulop, A., Li, J., & Yu, J. (2014). Self-exciting jumps, learning, and asset pricing implications. *The Review of Financial Studies*, 28(3), 876–912.
- Füss, R., Grabellus, M., Mager, F., & Stein, M. (2018). Something in the air: Information density, news surprises, and price jumps. *Journal of International Financial Markets, Institutions and Money*, 53, 50–75.
- Garleanu, N., Pedersen, L. H., & Poteshman, A. M. (2008). Demand-based option pricing. *The Review of Financial Studies*, 22(10), 4259–4299.
- Garman, M. B. Et al. (1976). *A general theory of asset valuation under diffusion state processes* (tech. rep.). University of California at Berkeley.
- Geske, R. (1979). The valuation of compound options. *Journal of financial economics*, 7(1), 63–81.
- Gibbons, M. R., & Hess, P. (1981). Day of the week effects and asset returns. *Journal of business*, 579–596.
- Gospodinov, N., & Jamali, I. (2012). The effects of federal funds rate surprises on s&p 500 volatility and volatility risk premium. *Journal of Empirical Finance*, 19(4), 497–510.

- Gu, C., Kurov, A., & Wolfe, M. H. (2018). Relief rallies after fomic announcements as a resolution of uncertainty. *Journal of Empirical Finance*, 49, 1–18.
- Hamilton, J. D. (1994). *Time series analysis* (Vol. 2). Princeton university press Princeton, NJ.
- Han, B. (2008). Investor sentiment and option prices. *The Review of Financial Studies*, 21(1), 387–414.
- Harris, L. (1986). A transaction data study of weekly and intradaily patterns in stock returns. *Journal of financial economics*, 16(1), 99–117.
- Hawkes, A. G. (1971). Point spectra of some mutually exciting point processes. *Journal of the Royal Statistical Society: Series B (Methodological)*, 33(3), 438–443.
- Heston, S. L. (1993). A closed-form solution for options with stochastic volatility with applications to bond and currency options. *The review of financial studies*, 6(2), 327–343.
- Hu, G., & Jacobs, K. (2016). Volatility and expected option returns. *Journal of Financial and Quantitative Analysis*, 1–77.
- Huang, J.-z., & Wu, L. (2004). Specification analysis of option pricing models based on time-changed lévy processes. *The Journal of Finance*, 59(3), 1405–1439.
- Hull, J., & White, A. (1987). The pricing of options on assets with stochastic volatilities. *The journal of finance*, 42(2), 281–300.
- Hull, J. C. (1997). Options, futures and other derivatives. -prentice hall, upper saddle river, new jersey.
- Jaffe, J., & Westerfield, R. (1985). The week-end effect in common stock returns: The international evidence. *The journal of finance*, 40(2), 433–454.
- Jain, P. C., & Joh, G.-H. (1988). The dependence between hourly prices and trading volume. *Journal of Financial and Quantitative Analysis*, 23(3), 269–283.
- Jiang, G. J., & Tian, Y. S. (2007). Extracting model-free volatility from option prices: An examination of the vix index. *The Journal of Derivatives*, 14(3), 35–60.
- Johannes, M. (2004). The statistical and economic role of jumps in continuous-time interest rate models. *The Journal of Finance*, 59(1), 227–260.

- Johannes, M. S., Polson, N. G., & Stroud, J. R. (2009). Optimal filtering of jump diffusions: Extracting latent states from asset prices. *The Review of Financial Studies*, 22(7), 2759–2799.
- Jones, C. S. (2003). The dynamics of stochastic volatility: Evidence from underlying and options markets. *Journal of econometrics*, 116(1-2), 181–224.
- Jones, C. S., & Shemesh, J. (2018). Option mispricing around nontrading periods. *The Journal of Finance*, 73(2), 861–900.
- Kaeck, A. (2013). Asymmetry in the jump-size distribution of the s&p 500: Evidence from equity and option markets. *Journal of Economic Dynamics and Control*, 37(9), 1872–1888.
- Kaeck, A., & Alexander, C. (2012). Volatility dynamics for the s&p 500: Further evidence from non-affine, multi-factor jump diffusions. *Journal of Banking & Finance*, 36(11), 3110–3121.
- Kaeck, A., & Alexander, C. (2013). Stochastic volatility jump-diffusions for european equity index dynamics. *European Financial Management*, 19(3), 470–496.
- Kamara, A. (1997). New evidence on the monday seasonal in stock returns. *Journal of Business*, 63–84.
- Kaplanski, G., & Levy, H. (2015). Trading breaks and asymmetric information: The option markets. *Journal of Banking & Finance*, 58, 390–404.
- Karakaya, M. M. (2014). *Characteristics and expected returns in individual equity options*. University of Chicago, Division of the Social Sciences, Department of Economics.
- Kelly, B., Pástor, L., & Veronesi, P. (2016). The price of political uncertainty: Theory and evidence from the option market. *The Journal of Finance*, 71(5), 2417–2480.
- Kleinert, H., & Korbel, J. (2016). Option pricing beyond black–scholes based on double-fractional diffusion. *Physica A: Statistical Mechanics and its Applications*, 449, 200–214.
- Krieger, K., Mauck, N., & Chen, D. (2012). Vix changes and derivative returns on fomc meeting days. *Financial Markets and Portfolio Management*, 26(3), 315–331.
- Kuttner, K. N. (2001). Monetary policy surprises and interest rates: Evidence from the fed funds futures market. *Journal of monetary economics*, 47(3), 523–544.

- Lakonishok, J., Lee, I., Pearson, N. D., & Poteshman, A. M. (2007). Option market activity. *The Review of Financial Studies*, 20(3), 813–857.
- Lévy-Leblond, J.-M. (1967). Nonrelativistic particles and wave equations. *Communications in Mathematical Physics*, 6(4), 286–311.
- Lobo, B. J., Darrat, A. F., & Ramchander, S. (2006). The asymmetric impact of monetary policy on currency markets. *Financial Review*, 41(2), 289–303.
- Lord, R., Koekkoek, R., & Dijk, D. V. (2010). A comparison of biased simulation schemes for stochastic volatility models. *Quantitative Finance*, 10(2), 177–194.
- Lucca, D. O., & Moench, E. (2015). The pre-fomc announcement drift. *The Journal of Finance*, 70(1), 329–371.
- MacBeth, J. D., & Merville, L. J. (1979). An empirical examination of the black-scholes call option pricing model. *The journal of finance*, 34(5), 1173–1186.
- McQueen, G., & Roley, V. V. (1993). Stock prices, news, and business conditions. *The Review of Financial Studies*, 6(3), 683–707.
- Merton, R. C. (1973). Theory of rational option pricing. *The Bell Journal of economics and management science*, 141–183.
- Merton, R. C. (1976). Option pricing when underlying stock returns are discontinuous. *Journal of financial economics*, 3(1-2), 125–144.
- Mitchell, M. L., & Mulherin, J. H. (1994). The impact of public information on the stock market. *The Journal of Finance*, 49(3), 923–950.
- Muravyev, D. (2016). Order flow and expected option returns. *The Journal of Finance*, 71(2), 673–708.
- Muravyev, D., & Ni, X. C. (2019). Why do option returns change sign from day to night? *Journal of Financial Economics*.
- Muravyev, D., & Pearson, N. D. (2019). Option trading costs are lower than you think. *Available at SSRN 2580548*.
- Nakamura, E., & Steinsson, J. (2018). High-frequency identification of monetary non-neutrality: The information effect. *The Quarterly Journal of Economics*, 133(3), 1283–1330.

- Neumann, M., & Skiadopoulos, G. (2013). Predictable dynamics in higher-order risk-neutral moments: Evidence from the s&p 500 options. *Journal of Financial and Quantitative Analysis*, 48(3), 947–977.
- Newey, W. K., & West, K. D. (1987). Hypothesis testing with efficient method of moments estimation. *International Economic Review*, 777–787.
- Nikkinen, J., & Sahlström, P. (2004). Impact of the federal open market committee’s meetings and scheduled macroeconomic news on stock market uncertainty. *International Review of Financial Analysis*, 13(1), 1–12.
- Novy-Marx, R. (2015). *Fundamentally, momentum is fundamental momentum* (tech. rep.). National Bureau of Economic Research.
- Onan, M., Salih, A., & Yasar, B. (2014). Impact of macroeconomic announcements on implied volatility slope of spx options and vix. *Finance Research Letters*, 11(4), 454–462.
- Pan, J. (2002). The jump-risk premia implicit in options: Evidence from an integrated time-series study. *Journal of financial economics*, 63(1), 3–50.
- Pan, J., & Poteshman, A. M. (2006). The information in option volume for future stock prices. *The Review of Financial Studies*, 19(3), 871–908.
- Park, Y.-H. (2015). *Price dislocation and price discovery in the s&p 500 options and vix derivatives markets* (tech. rep.). Working Paper). Federal Reserve Board.
- Pastor, L., & Veronesi, P. (2012). Uncertainty about government policy and stock prices. *The journal of Finance*, 67(4), 1219–1264.
- Patell, J. M., & Wolfson, M. A. (1979). Anticipated information releases reflected in call option prices. *Journal of accounting and economics*, 1(2), 117–140.
- Patell, J. M., & Wolfson, M. A. (1981). The ex ante and ex post price effects of quarterly earnings announcements reflected in option and stock prices. *Journal of Accounting Research*, 434–458.
- Patell, J. M., & Wolfson, M. A. (1984). The intraday speed of adjustment of stock prices to earnings and dividend announcements. *Journal of Financial Economics*, 13(2), 223–252.
- Piazzesi, M. (2005). Bond yields and the federal reserve. *Journal of Political Economy*, 113(2), 311–344.

- Poteshman, A. M. (2001). Underreaction, overreaction, and increasing misreaction to information in the options market. *The Journal of Finance*, 56(3), 851–876.
- Rehman, Z., & Vilkov, G. (2012). Risk-neutral skewness: Return predictability and its sources. *Available at SSRN 1301648*.
- Sadka, R. (2006). Momentum and post-earnings-announcement drift anomalies: The role of liquidity risk. *Journal of Financial Economics*, 80(2), 309–349.
- Savor, P., & Wilson, M. (2013). How much do investors care about macroeconomic risk? evidence from scheduled economic announcements. *Journal of Financial and Quantitative Analysis*, 48(2), 343–375.
- Scholtus, M., van Dijk, D., & Frijns, B. (2014). Speed, algorithmic trading, and market quality around macroeconomic news announcements. *Journal of Banking & Finance*, 38, 89–105.
- Shampine, L. (2009). Quadrature error estimates for integrands with modest differentiability. *Can Appl Math Q*, 17(3), 571–583.
- Shefrin, H. Et al. (2010). Behavioralizing finance. *Foundations and Trends® in Finance*, 4(1–2), 1–184.
- Sheikh, A. M., & Ronn, E. I. (1994). A characterization of the daily and intraday behavior of returns on options. *The Journal of Finance*, 49(2), 557–579.
- Srikant, M. (1998). *Option pricing with stochastic volatility* (tech. rep.). <http://srikant.org/professional/personal/thesis/thesis.html>
- Stein, J. (1989). Overreactions in the options market. *The Journal of Finance*, 44(4), 1011–1023.
- Stein, J. C., & Sunderam, A. (2018). The fed, the bond market, and gradualism in monetary policy. *The Journal of Finance*, 73(3), 1015–1060.
- Stoll, H. R., & Whaley, R. E. (1990). Stock market structure and volatility. *The Review of Financial Studies*, 3(1), 37–71.
- Stroud, J. R., & Johannes, M. S. (2014). Bayesian modeling and forecasting of 24-hour high-frequency volatility. *Journal of the American Statistical Association*, 109(508), 1368–1384.

- Sullivan, R., Timmermann, A., & White, H. (2001). Dangers of data mining: The case of calendar effects in stock returns. *Journal of Econometrics*, 105(1), 249–286.
- Tompkins, R. G. (2001). Implied volatility surfaces: Uncovering regularities for options on financial futures. *The European Journal of Finance*, 7(3), 198–230.
- Trolle, A. B., & Schwartz, E. S. (2009). Unspanned stochastic volatility and the pricing of commodity derivatives. *The Review of Financial Studies*, 22(11), 4423–4461.
- Utama, B., & Purqon, A. (2016). Feynman path integral application on deriving black-scholes diffusion equation for european option pricing, In *Journal of physics: Conference series*. IOP Publishing.
- Vähämaa, S., & Äijö, J. (2011). The fed’s policy decisions and implied volatility. *Journal of Futures Markets*, 31(10), 995–1010.
- Vasquez, A. (2017). Equity volatility term structures and the cross section of option returns. *Journal of Financial and Quantitative Analysis*, 52(6), 2727–2754.
- Wood, R. A., McInish, T. H., & Ord, J. K. (1985). An investigation of transactions data for nyse stocks. *The Journal of Finance*, 40(3), 723–739.
- Xing, Y., Zhang, X., & Zhao, R. (2010). What does the individual option volatility smirk tell us about future equity returns? *Journal of Financial and Quantitative Analysis*, 45(3), 641–662.
- Zhang, L., Mykland, P. A., & Ait-Sahalia, Y. (2005). A tale of two time scales: Determining integrated volatility with noisy high-frequency data. *Journal of the American Statistical Association*, 100(472), 1394–1411.

Chapter 6

Appendix

6.1 Heat Equation

In this Appendix, we show how to massage the partial differential equation corresponding to the symmetrical model:

$$\frac{\partial C}{\partial t} + rS \frac{\partial C}{\partial S} + \frac{1}{2} \sigma^2 S^2 \frac{\partial^2 C}{\partial S^2} + \mu V \frac{\partial C}{\partial V} + \xi_0^2 V^2 \frac{\partial^2 C}{\partial V^2} = rC, \quad (6.1)$$

into the heat equation in 2+1 dimensions. To do so, we now make the standard substitutions for the underlying and variance, transforming them to dimensionless variables:

$$x = \log(S/K), \quad (6.2)$$

$$y = \log(V/V_0), \quad (6.3)$$

where K is the strike price and V_0 is some constant with units of 1/sec. Re-casting Equation (6.1) in terms of x and y , we obtain

$$\frac{\partial C}{\partial t} + \omega \frac{\partial C}{\partial x} + \gamma \frac{\partial C}{\partial y} + \frac{1}{2} \sigma^2 \frac{\partial^2 C}{\partial x^2} + \xi_0^2 \frac{\partial^2 C}{\partial y^2} = rC, \quad (6.4)$$

where, $\omega = r - \frac{1}{2} \sigma^2$ and $\gamma = \mu - \xi_0^2$.

In order to remove the constant term, $\frac{1}{2}\sigma^2$ in front of the second derivative with respect to x , we make the time transformation: $\tau = \frac{\sigma^2}{2}(T - t)$, yielding:

$$-\frac{\partial C}{\partial \tau} + (R_1 - 1)\frac{\partial C}{\partial x} + \frac{2\gamma}{\sigma^2}\frac{\partial C}{\partial y} + \frac{\partial^2 C}{\partial x^2} + \eta\frac{\partial^2 C}{\partial y^2} - R_1 C = 0, \quad (6.5)$$

where $R_1 = \frac{2r}{\sigma^2}$ and $\eta = \frac{2\xi_0^2}{\sigma^2}$. We then proceed with the further substitution: $y = \frac{1}{\sqrt{\eta}}y$, transforming the coefficient of the second derivative with respect to y to unity:

$$-\frac{\partial C}{\partial \tau} + (R_1 - 1)\frac{\partial C}{\partial x} + \sqrt{2}\frac{\gamma}{\sigma\xi_0}\frac{\partial C}{\partial y} + \frac{\partial^2 C}{\partial x^2} + \frac{\partial^2 C}{\partial y^2} - R_1 C = 0. \quad (6.6)$$

Next, we make the price transformation:

$$C(x, y, \tau) = K\phi(x, y, \tau)\psi_0(x, y, \tau), \quad (6.7)$$

with:

$$\phi(x, y, \tau) = e^{ax+by+c\tau}. \quad (6.8)$$

The constants a, b and c are chosen by inspection, after substitution into Equation (6.6) we see that the choice:

$$a = -\frac{1}{2}(R_1 - 1), \quad (6.9)$$

$$b = -\frac{1}{2}(R_2 - 1), \quad (6.10)$$

$$c = -\frac{1}{4}\left((R_2 - 1)^2 + (R_1 + 1)^2\right), \quad (6.11)$$

where: $R_2 = 1 + \sqrt{2}\gamma/\sigma\xi_0$, leads to the heat equation in 2+1 dimensions:

$$\frac{\partial \psi_0}{\partial \tau} = \frac{\partial^2 \psi_0}{\partial x^2} + \frac{\partial^2 \psi_0}{\partial y^2}. \quad (6.12)$$

There are well known techniques to solve this partial differential equation analytically. It is also well known that the heat equation has a symmetry based on the Galilean group in 2+1 dimensions. This symmetry is now manifest. In particular, the variable x and y are

interchangeable as advertised previously. Before we can solve the heat equation, we need to specify an appropriate boundary equation.

Guided by our intuition formed by the Black-Scholes model, we propose a boundary condition of the form for a call option:

$$\psi_0(x, y, 0) = e^{\frac{1}{2}(R_2-1)y} \left(e^{(R_1+1)x/2} - e^{(R_1-1)x/2} \right)^+. \quad (6.13)$$

To verify that this boundary condition makes sense from an option pricing point of view, we work out the implied boundary conditions for the call price in the original variables:

$$C(x, y, 0) = K e^{ax+by} \psi_0(x, y, 0), \quad (6.14)$$

and thus

$$C(x, y, 0) = K \left(e^x - 1 \right)^+. \quad (6.15)$$

Substitution back to original variable, using: $S = K e^x$, we find

$$C(S, V, T) = \left(S(T) - K \right)^+. \quad (6.16)$$

This is the standard payoff of a call option. We have thus found an appropriate boundary condition for the heat equation and can thus proceed to solving this partial differential equation. For a put option we have

$$\psi_0(x, y, 0) = -e^{\frac{1}{2}(R_2-1)y} \left(e^{(R_1+1)x/2} - e^{(R_1-1)x/2} \right)^+, \quad (6.17)$$

which leads to

$$P(S, V, T) = \left(K - S(T) \right)^+. \quad (6.18)$$

6.2 Solution of the Symmetrical Model

In this Appendix, we show how to solve the heat equation in 2+1 dimensions:

$$\frac{\partial \psi_0}{\partial \tau} = \frac{\partial^2 \psi_0}{\partial x^2} + \frac{\partial^2 \psi_0}{\partial y^2}. \quad (6.19)$$

So far the function $\psi_0(x, y, \tau)$ is only defined for $\tau > 0$, however by introducing the Heaviside function $\Theta(\tau)$ we may extend the definition domain to the range $\tau < 0$

$$\bar{\psi}_0(x, y, \tau) = \Theta(\tau) \psi_0(x, y, \tau), \quad (6.20)$$

Thus we get an inhomogeneous differential equation:

$$\left(\frac{\partial}{\partial \tau} - \frac{\partial^2}{\partial x^2} - \frac{\partial^2}{\partial y^2} \right) \bar{\psi}_0(x, y, \tau) = \bar{\psi}_0(x, y, 0) \delta(\tau). \quad (6.21)$$

This equation is solved by

$$\bar{\psi}_0(x, y, \tau) = \int \bar{\psi}_0(x', y', 0) G(x, y, \tau | x', y', 0) dx' dy'. \quad (6.22)$$

This is the partial differential equation for the Gaussian propagator of the heat equation in 2+1 dimensions:

$$G(x, y, \tau | X, Y, 0) = \frac{1}{4\pi\tau} e^{-\frac{(x-X)^2}{4\tau} - \frac{(y-Y)^2}{4\tau}} \Theta(\tau). \quad (6.23)$$

Combining the above two results, the solution can be written in the form

$$\bar{\psi}_0(x, y, \tau) = \frac{1}{4\pi\tau} \int \bar{\psi}_0(X, Y, 0) e^{-\frac{(x-X)^2}{4\tau} - \frac{(y-Y)^2}{4\tau}} dX dY, \quad (6.24)$$

$$\bar{\psi}_0(x, y, \tau) = \frac{1}{4\pi\tau} \int e^{(R_2-1)Y/2} \left(e^{(R_1+1)X/2} - e^{(R_1-1)X/2} \right)^+ e^{-\frac{(x-X)^2}{4\tau} - \frac{(y-Y)^2}{4\tau}} dX dY, \quad (6.25)$$

which leads to

$$\psi_0(x, y, \tau) = \frac{1}{4\pi\tau} \int \psi_0(X, Y, 0) e^{-\frac{(x-X)^2}{4\tau} - \frac{(y-Y)^2}{4\tau}} dX dY, \quad (6.26)$$

and

$$\psi_0(x, y, \tau) = \frac{1}{4\pi\tau} \int e^{(R_2-1)Y/2} \left(e^{(R_1+1)X/2} - e^{(R_1-1)X/2} \right)^+ e^{-\frac{(x-X)^2}{4\tau} - \frac{(y-Y)^2}{4\tau}} dX dY. \quad (6.27)$$

Note that the two integrals can be separated:

$$\begin{aligned} \psi_0(x, y, \tau) = \frac{1}{\sqrt{4\pi\tau}} \int_0^\infty dX \left(e^{(R_1+1)X/2} - e^{(R_1-1)X/2} \right)^+ e^{-\frac{(x-X)^2}{4\tau}} \\ \times \frac{1}{\sqrt{4\pi\tau}} \int_{-\infty}^\infty dY e^{(R_2-1)Y/2} e^{-\frac{(y-Y)^2}{4\tau}}, \end{aligned} \quad (6.28)$$

where the first of these integrals precisely corresponds to the 1+1 dimensional Black-Scholes model. We thus finally obtain

$$\begin{aligned} \psi_0(x, y, \tau) = e^{\frac{1}{2}(R_2-1)(\frac{\tau}{2}(R_2-1)+y)} \left[e^{(R_1+1)x/2} + (R_1+1)^2\tau/4\mathcal{N}(d_1) \right. \\ \left. - e^{(R_1-1)x/2+(R_1-1)^2\tau/4}\mathcal{N}(d_2) \right], \end{aligned} \quad (6.29)$$

where

$$\mathcal{N}(d) = \frac{1}{\sqrt{2\pi}} \int_{-\infty}^d \exp\left(-\frac{z^2}{2}\right) dz, \quad (6.30)$$

and

$$d_1 = \frac{x}{\sqrt{2\tau}} + \frac{\sqrt{2\tau}}{2}(R_1+1) = \frac{\log(S/K) + (r + \sigma^2/2)(T-t)}{\sigma\sqrt{T-t}}, \quad (6.31)$$

$$d_2 = \frac{x}{\sqrt{2\tau}} + \frac{\sqrt{2\tau}}{2}(R_1-1) = d_1 - \sigma\sqrt{T-t}. \quad (6.32)$$

We have obtained an analytical solution to the symmetrical model. Remarkably, because of

the boundary condition that only depends on S , it is identical to the Black-Scholes solution. Going back to the original variables we find:

$$C_0(S, V, t) = S\mathcal{N}(d_1) - Ke^{-r(T-t)}\mathcal{N}(d_2). \quad (6.33)$$

6.3 Symmetry Breaking Terms and Solution to the Merton-Garman Model

In this Appendix, we give details of the derivation of the perturbative solution to the full Merton-Garman. We first need to restore the full model by re-introducing the symmetry breaking terms

$$\begin{aligned} \frac{\partial C}{\partial t} + rS\frac{\partial C}{\partial S} + \frac{1}{2}\sigma^2 S^2 \frac{\partial^2 C}{\partial S^2} + \frac{c_1 S^2}{2} \left(V - \sigma^2 \right) \frac{\partial^2 C}{\partial S^2} + \mu V \frac{\partial C}{\partial V} + c_2 \lambda \frac{\partial C}{\partial V} + \xi_0^2 V^2 \frac{\partial^2 C}{\partial V^2} \\ + c_3 \left(\xi^2 V^{2\alpha} - \xi_0^2 V^2 \right) \frac{\partial^2 C}{\partial V^2} + c_4 \rho \xi V^{\alpha+1/2} S \frac{\partial^2 C}{\partial S \partial V} - rC = 0. \end{aligned} \quad (6.34)$$

Note that we have introduced dimensionless coefficients c_i which denote the strength of the symmetry breaking terms. In the limit $c_i = 1$ one recovers the original Merton-Garman model. These coefficients are simply introduced as a bookkeeping trick to keep track of which terms correspond to a deviation of the 2+1 Galilean invariant theory. In the end of the day, we set $c_i = 1$. We can now apply the same variables transformations to Equation (6.34) that we had applied in the symmetric case and obtain

$$\left(\frac{\partial}{\partial \tau} - \frac{\partial^2}{\partial x^2} - \frac{\partial^2}{\partial y^2} + \mathcal{D}(x, y) \right) \psi(x, y, \tau) = 0, \quad (6.35)$$

where the operator $\mathcal{D}(x, y)$ is defined as:

$$\begin{aligned} \mathcal{D}(x, y) = & \left(\frac{c_1}{2} (V_0 e^y - \sigma^2) \left((a^2 - a) + (2a - 1) \frac{\partial}{\partial x} + \frac{\partial^2}{\partial x^2} \right) \right. \\ & + c_2 \frac{\lambda}{V_0} e^{-y} \left(\frac{\partial}{\partial y} + b \right) \\ & + c_3 \left(\xi^2 V_0^{2\alpha-2} e^{(2\alpha-2)y} - \xi_0^2 \right) \left((b^2 - b) + (2b - 1) \frac{\partial}{\partial y} + \frac{\partial^2}{\partial y^2} \right) \\ & \left. + c_4 \xi \rho V_0^{\alpha-\frac{1}{2}} e^{(\alpha-\frac{1}{2})y} \left(ab + b \frac{\partial}{\partial x} + a \frac{\partial}{\partial y} + \frac{\partial^2}{\partial x \partial y} \right) \right), \end{aligned} \quad (6.36)$$

where a and b are respectively given in Equation (6.9) and Equation (6.10). Note that $\mathcal{D}(x, y)$ is τ independent.

We now do perturbation theory around the symmetrical solution ψ_0 which was given in Equation (6.29). To leading order in c_i , we write $\psi = \psi_0 + \psi_1$ where ψ_1 is of order c_i . Keeping in mind that \mathcal{D} is order c_i , we find

$$\left(\frac{\partial}{\partial \tau} - \frac{\partial^2}{\partial x^2} - \frac{\partial^2}{\partial y^2} \right) \psi_1(x, y, \tau) = -\mathcal{D}(x, y) \psi_0(x, y, \tau). \quad (6.37)$$

This equation can be solved by the Green's function method, we obtain

$$\psi_1(x, y, \tau) = - \int_0^\tau dt \int_{-\infty}^\infty dx' \int_{-\infty}^\infty dy' G(x, y, \tau; x', y', t) \mathcal{D}(x', y') \psi_0(x', y', t), \quad (6.38)$$

where

$$G(x, y, \tau; x', y', t) = \frac{1}{4\pi(\tau - t)} \exp \left(-\frac{(x - x')^2 + (y - y')^2}{4(\tau - t)} \right). \quad (6.39)$$

These integrals can be performed analytically. Our result provides an approximative and analytical solution to the Merton-Garman model. We find:

$$\psi(x, y, \tau) = \psi_0(x, y, \tau) + \psi_1(x, y, \tau), \quad (6.40)$$

with

$$\psi_1(x, y, \tau) = \frac{c_1 \left(\sigma^2 \tau (\sqrt{2}\gamma + \sigma\xi_0) - \sigma\xi_0 V_0 e^y \left(e^{\frac{\sqrt{2}\gamma\tau}{\sigma\xi_0} + \tau} - 1 \right) \right) e^{\left(\frac{\gamma^2\tau}{2\sigma\xi_0^2} - \frac{x^2}{4\tau} + \frac{\gamma y}{\sqrt{2}\sigma\xi_0} \right)}}{4\sqrt{\pi\tau} (\sqrt{2}\gamma + \sigma\xi_0)}, \quad (6.41)$$

to leading order. In the original variables, we find:

$$\begin{aligned} C_1(S, V, t) = & -K \frac{\left(\frac{S}{K} \right)^{\frac{1}{2} - \frac{r}{\sigma^2}} e^{\left(\frac{4 \log^2 \left(\frac{S}{K} \right) + (2r + \sigma^2)^2 (t-T)^2}{8\sigma^2(t-T)} \right)}}{4\sqrt{2\pi} \left(\frac{\sqrt{2}\gamma}{\sigma\xi_0} + 1 \right) \sqrt{\sigma^2(T-t)}} \\ & \times \left(\frac{1}{2} \sigma^4 \left(\frac{\sqrt{2}\gamma}{\sigma\xi_0} + 1 \right) (t-T) + V \left(e^{\frac{1}{2}\sigma^2 \left(\frac{\sqrt{2}\gamma}{\sigma\xi_0} + 1 \right) (T-t)} - 1 \right) \right), \end{aligned} \quad (6.42)$$

where we have set $c_1 = 1$.

It may appear surprising that the leading order correction does not depend on the symmetry breaking terms parametrised by c_2 , c_3 and c_4 . To understand what is happening, one can organise the perturbation theory slightly different, but mathematically equivalent, fashion by looking at corrections to the Green's function $G(x, y, \tau|x', y', 0)$. The differential equation to solve is given by:

$$\left(\frac{\partial}{\partial \tau} - \frac{\partial^2}{\partial x^2} - \frac{\partial^2}{\partial y^2} + \mathcal{D} \right) G(x, y, \tau|x', y', 0) = \delta(\tau) \delta(x - x') \delta(y - y'). \quad (6.43)$$

Perturbation theory is organised by expanding around the Green's function of the symmetrical symmetry: $G(x, y, \tau|x', y', 0) = G_0(x, y, \tau|x', y', 0) + G_1(x, y, \tau|x', y', 0)$ to leading order. One obtains:

$$\left(\frac{\partial}{\partial \tau} - \frac{\partial^2}{\partial x^2} - \frac{\partial^2}{\partial y^2} \right) G_0(x, y, \tau|x', y', 0) = \delta(\tau) \delta(x - x') \delta(y - y'). \quad (6.44)$$

The solution to this partial differential equation was given above. The correction $G_1(x, y, \tau|x', y', 0)$ is obtained by solving:

$$\left(\frac{\partial}{\partial \tau} - \frac{\partial^2}{\partial x^2} - \frac{\partial^2}{\partial y^2} \right) G_1(x, y, \tau|x', y', 0) = -\mathcal{D}(x, y) G_0(x, y, \tau|x', y', 0), \quad (6.45)$$

which can be solved easily. One finds:

$$G_1(x, y, \tau | x', y', 0) = - \int_0^\tau dt' \int_{-\infty}^\infty dx'' \int_{-\infty}^\infty dy'' G_0(x, y, \tau | x'', y'', t') \mathcal{D}(x'', y'') \quad (6.46)$$

$$\times G_0(x'', y'', t' | x', y', 0).$$

It is easy to show that this correction depends on all four symmetry breaking terms. We find

$$G_{1,c_1}(x, y, \tau | x', y', 0) = \frac{c_1}{64\pi\tau^{5/2}} e^{-\frac{(x-x')^2 + (y-y')^2}{4\tau}} \left((-1 + R_1)^2 \tau^2 \right. \quad (6.47)$$

$$+ 2\tau(-1 + R_1(x - x')) + (x - x')^2$$

$$\times \left(2\sigma^2 \sqrt{\tau} - \sqrt{\pi} e^{\frac{\tau^2 + (y-y')^2 + 2\tau(y-y')}{4\tau}} V_0 \left(\text{Erf} \left(\frac{\tau + y - y'}{2\sqrt{\tau}} \right) \right. \right.$$

$$\left. \left. + \text{Erf} \left(\frac{\tau - y + y'}{2\sqrt{\tau}} \right) \right),$$

$$G_{1,c_2}(x, y, \tau | x', y', 0) = c_2 \frac{\lambda}{16\pi V_0 \tau^2} e^{-\frac{(x-x')^2 + y'^2 + y'^2 + 4\tau(y-y')}{4\tau}} \quad (6.48)$$

$$\times \left(2e^{yy'} 2\tau(e^y - e^{y'})\tau + e^{(\tau+y)^2 + 2\tau y' + y'^2} 4\tau \sqrt{\pi} (R_2 - 2)\tau^{3/2} \right.$$

$$\left. \times \left(\text{Erf} \left(\frac{\tau - y + y'}{2\sqrt{\tau}} \right) - \text{Erf} \left(\frac{-\tau - y + y'}{2\sqrt{\tau}} \right) \right) \right),$$

$$\begin{aligned}
G_{1,c_3}(x, y, \tau | x', y', 0) &= c_3 \frac{1}{64\tau^2\pi(1-\alpha)} e^{-\frac{(x-x')^2}{4\tau}} \\
&\times \left(2e^{-\frac{(y-y')^2}{4\tau}-2(x-x')} (-2(\alpha-1)\xi_0^2 e^{2(y+y')}) \right. \\
&\times ((R_2^2-1)\tau^2 + 2\tau(R_2y - R_2y' - 1) + (y-y')^2) \\
&+ \xi^2 V^{2\alpha-2} \left(-e^{2(\alpha y+y')} \right) (-2\tau(\alpha + R_2 - 1) - y + y') \\
&- \xi^2 V^{2\alpha-2} e^{2(y+\alpha y')} (2\tau(3\alpha + R_2 - 3) + y - y') \\
&+ \sqrt{\pi}\xi^2 \tau^{3/2} V^{2\alpha-2} (2\alpha + R_2 - 3)(2\alpha + R_2 - 1) e^{(\alpha-1)((\alpha-1)\tau+y+y')} \\
&\times \left(\text{Erf} \left(\frac{2(\alpha-1)\tau + y - y'}{2\sqrt{\tau}} \right) - \text{Erf} \left(\frac{-2(\alpha-1)\tau + y - y'}{2\sqrt{\tau}} \right) \right) \Bigg),
\end{aligned} \tag{6.49}$$

and

$$\begin{aligned}
G_{1,c_4}(x, y, \tau | x', y', 0) &= c_4 \frac{\xi \rho V_0^{\alpha-\frac{1}{2}} ((R_1-1)\tau + x - x')}{32\pi\tau^2(1-2\alpha)} e^{-\frac{\alpha\tau^2+y(\tau+y)+(x-x')^2+y'^2}{4\tau}} \\
&\times \left(4 \left(e^{\alpha y + \frac{y'}{2}} - e^{\frac{y}{2} + \alpha y'} \right) e^{\frac{1}{4}(\alpha\tau+y(\frac{2y'}{\tau}-1)-2y')} \right. \\
&- \sqrt{\pi\tau}(2\alpha + 2R_2 - 3) \\
&\times \left(\text{Erf} \left(\frac{-2\alpha\tau + \tau + 2y - 2y'}{4\sqrt{\tau}} \right) - \text{Erf} \left(\frac{2\alpha\tau - \tau + 2y - 2y'}{4\sqrt{\tau}} \right) \right) \\
&\times \exp \left(\frac{4\alpha^2\tau^2 + 8\alpha\tau(y+y') + (\tau-2y')^2 + 4y^2}{16\tau} \right) \Bigg).
\end{aligned} \tag{6.50}$$

It is easy to see that when folding these corrections to the Green's function with the boundary condition (6.13) that only the contribution from the c_1 term survives and one recovers the result of (6.41). The boundary condition thus implies that the contributions of c_2 , c_3 and c_4 vanish to leading order in the perturbation theory. These symmetry breaking terms will, however, contribute to higher order corrections. Higher precision, if required, can be obtained by going to higher order in perturbation theory. Option prices can be calculated extremely rapidly using this formalism. Note that, in principle, if we resummed perturbation theory to all order in c_i , the dependence on σ and ξ_0 would vanish. It is also worth noticing that our results are independent on V_0 which is only introduced to match the dimension of V .

6.4 Time Series Simulation Merton-Garman Model

The Merton-Garman model is defined by the coupled two-dimensional SDE:

$$dS_t = rS_t dt + \sqrt{V_t} S_t dW_t^S, \quad (6.51)$$

$$dV_t = \kappa(\theta - V_t)dt + \xi V_t^\alpha dW_t^V, \quad (6.52)$$

with the two Brownian motion components: W_t^S, W_t^V having correlation ρ . We use an Euler discretization scheme for the asset price and variance process, with a full truncation to tackle the issue of negative variances. Conditional on a time s for $t > s$ the discretization scheme for the asset price and variance processes are:

$$S_t = S_s + rS_t \Delta t + S_s \sqrt{V_t^+ \Delta t} z_S, \quad (6.53)$$

$$V_t = V_s - \kappa(\theta - V_s^+) \Delta t + \xi(V_s^+ \Delta t)^\alpha z_V, \quad (6.54)$$

where $\Delta t = t - s$, $z_V \sim \mathcal{N}(0, 1)$ and $z_S = \rho z_V + \sqrt{1 - \rho^2} z$ with $z \sim \mathcal{N}(0, 1)$. This scheme is used to generate 100 different sample paths of weekly returns and latent variance over one year, i.e 52 observations with $\Delta t = 7$ days. At each observation time we simulate six unique maturities within $[7, 180]$ days maturity, with a moneyness range of $K/S \in [0.9, 1.1]$ across ten strikes. Each option price is computed using the Monte Carlo framework with 50,000 simulations and a time-step of $1/20^{\text{th}}$ of a trading day.

The calibration process is done using the objective function defined in Equation (2.23). It should be noted that as initial conditions we start with the true parameter vector, i.e $\Theta_{\text{initial}}^{\text{pert.}} = \Theta_{\text{true}}^{MG}$ and for σ we start with the initial variance. We also pass the variance path at each time step.

6.5 Characteristic Function

This appendix provides details on the derivation of the characteristic function of the option pricing model from Equations (3.1)-(3.3). Initially we omit the FOMC deterministic jumps and focus only on a two-factor stochastic volatility model with random Poisson contemporaneous jumps (with finite activity) and a constant intensity.

Denote the characteristic function as: $\psi(t, Y, V, m)$, application of Itô's lemma yields:

$$\begin{aligned}
 d\psi(t, Y, V, m) = & \psi_t(t_-, Y_-, V_-, m_-)dt + \psi_Y(t_-, Y_-, V_-, m_-)dY_{t-}^c \\
 & + \psi_V(t_-, Y_-, V_-, m_-)dV_{t-}^c + \psi_m(t_-, Y_-, V_-, m_-)dm_{t-}^c \\
 & + \frac{1}{2}\psi_{YY}(t_-, Y_-, V_-, m_-)dY_{t-}^c dY_{t-}^c + \frac{1}{2}\psi_{VV}(t_-, Y_-, V_-, m_-)dV_{t-}^c dV_{t-}^c \\
 & + \frac{1}{2}\psi_{mm}(t_-, Y_-, V_-, m_-)dm_{t-}^c dm_{t-}^c + \psi_{YV}(t_-, Y_-, V_-, m_-)dY_{t-}^c dV_{t-}^c \\
 & + (\psi(t, Y, V, m) - \psi(t_-, Y_-, V_-, m_-))dZ_t,
 \end{aligned} \tag{6.55}$$

where X_{t-} denotes the value of X prior to any jump at time t . The superscript c denotes the continuous part of the process and the dZ_t term denotes the discontinuous jump related component. As the correlations between: W_t^Y, W_t^m and W_t^v, W_t^m are zero we neglect these derivatives. Substitution from the dynamics in Equations (3.1)-(3.3) yields:

$$\begin{aligned}
 d\psi(t, Y, V, m) = & [\psi_t(t_-, Y_-, V_-, m_-) - \psi_Y(t_-, Y_-, V_-, m_-)(\lambda\tilde{\mu} + \frac{1}{2}V_{t-}) \\
 & + \psi_V(t_-, Y_-, V_-, m_-)(\kappa_v^Q(m_t - V_{t-})) + \kappa_m^Q(\theta_m^Q - m_{t-})\psi_m(t_-, Y_-, V_-, m_-) \\
 & + \frac{1}{2}\psi_{YY}(t_-, Y_-, V_-, m_-)V_{t-} + \frac{1}{2}\sigma_V^2\psi_{VV}(t_-, Y_-, V_-, m_-)V_{t-} \\
 & + \frac{1}{2}\sigma_m^2\psi_{mm}(t_-, Y_-, V_-, m_-)m_{t-} + \rho\sigma_v\psi_{YV}(t_-, Y_-, V_-, m_-)V_{t-}]dt \\
 & + (\psi(t, Y, V, m) - \psi(t_-, Y_-, V_-, m_-))dZ_t.
 \end{aligned} \tag{6.56}$$

Where $\tilde{\mu} = \exp(\mu_y^Q + \frac{1}{2}(\sigma_y^Q)^2) - 1$. The Brownian motion terms have been excluded as the characteristic function is a martingale, thus its derivative with respect to time must be zero, as such we are only interested in the non-Brownian motion terms. By the affine structure of

the problem it follows the discounted transform for $u \in \mathbb{R}$ is affine (see Duffie et al. (2000)) the characteristic function is given by:

$$\psi(t, Y, V, m) = \exp(\alpha(\tau) + iuY_t + \beta(\tau)V_t + \gamma(\tau)m_t), \quad (6.57)$$

the discontinuous jump terms from above can be expressed as:

$$(\psi(t, Y, V, m) - \psi(t_-, Y_-, V_-, m_-))dZ_t = \psi(t_-, Y_-, V_-, m_-)(\exp(iu\xi^Y + \beta\xi^V) - 1)dZ_t. \quad (6.58)$$

Which yields, when taking expectation value:

$$\begin{aligned} \mathbb{E}_t^{\mathbb{Q}} \left[\psi(t_-, Y_-, V_-, m_-)(\exp(iu\xi^Y + \beta\xi^V) - 1)dZ_t \right] &= \lambda \mathbb{E}_t^{\mathbb{Q}} [\exp(iu\xi^Y + \beta\xi^V) - 1] \\ &\quad \times \psi(t_-, Y_-, V_-, m_-)dt. \end{aligned} \quad (6.59)$$

Substitution gives:

$$\begin{aligned} \mathbb{E}_t^{\mathbb{Q}} \left[\frac{d\psi}{dt} \right] &= [(\alpha'(\tau) + \beta'(\tau) + \gamma'(\tau)) - iu(\lambda\tilde{\mu} + \frac{1}{2}V_t) + \beta(\tau)\kappa_v^{\mathbb{Q}}(m_t - V_t) \\ &\quad + \kappa_m^{\mathbb{Q}}(\theta_m^{\mathbb{Q}} - m_t)\gamma(\tau) - \frac{1}{2}u^2V_t + \frac{1}{2}\sigma_v^2V_t\beta(\tau)^2 + \frac{1}{2}\sigma_m^2m_t\gamma(\tau)^2 + \rho\sigma_vV_t(iu\beta(\tau)) \\ &\quad + \lambda\mathbb{E}_t^{\mathbb{Q}}[\exp(iu\xi^Y + \beta(\tau)\xi^V) - 1] - \lambda\mathbb{E}_t^{\mathbb{Q}}[(\exp(iu\xi^Y) - 1)]]\psi(t_-, Y_-, V_-, m_-). \end{aligned} \quad (6.60)$$

As the characteristic function is a martingale: $\frac{d\psi}{dt} = 0$ and using Duffie et al. (2000) for the expectation values, yields the grouped ODE's:

$$\begin{aligned} -\alpha'(\tau) + \kappa_m^{\mathbb{Q}}\theta_m^{\mathbb{Q}}\gamma(\tau) - \lambda(\exp(\mu_y^{\mathbb{Q}} + \frac{1}{2}(\sigma_y^{\mathbb{Q}})^2) - 1)iu + \\ \lambda \left(\frac{\exp(iu\mu_y^{\mathbb{Q}} - \frac{1}{2}(\sigma_y^{\mathbb{Q}})^2u^2)}{1 - \mu_y^{\mathbb{Q}}\beta(\tau)} - 1 \right) = 0, \end{aligned} \quad (6.61)$$

$$-\beta'(\tau) - \frac{1}{2}iu - \kappa_v^{\mathbb{Q}}\beta(\tau) - \frac{1}{2}u^2 + \frac{1}{2}\sigma_v^2\beta(\tau)^2 + i\sigma_v\rho u\beta(\tau) = 0, \quad (6.62)$$

$$-\gamma'(\tau) + \kappa_v^{\mathbb{Q}}\beta(\tau) - \kappa_m^{\mathbb{Q}}\gamma(\tau) + \frac{1}{2}\sigma_m^2\gamma^2(\tau) = 0. \quad (6.63)$$

The extension to deterministic FOMC jumps is as follows. Due to the multiplication laws of logs:

$$\ln(F_t) = \ln(\tilde{F}_t) + \sum_{j=1}^{N_t^d} \bar{Z}_j(\mathbb{Q}), \quad (6.64)$$

where $\ln(\tilde{F}_t)$ is the standard affine component from above and the second term represents the deterministic jumps from FOMC announcements. Under the assumption the deterministic jumps are conditionally independent of the affine state variables the transform of $\ln(F_t)$ is just the product of the traditional affine transform and the transform of the deterministic jumps:

$$\psi^{FOMC}(t, Y, V, m) = \mathbb{E}_t^{\mathbb{Q}}[\exp(iu \ln(\tilde{F}_t))] \mathbb{E}_t^{\mathbb{Q}} \left[\exp \left(iu \sum_{j=1}^{N_t^d} \bar{Z}_j(\mathbb{Q}) \right) \right]. \quad (6.65)$$

Simplification results in:

$$\psi^{FOMC}(t, Y, V, m) = \exp(\alpha^*(\tau) + \beta(\tau)V_t + \gamma(\tau)m_t + iuY_t), \quad (6.66)$$

where:

$$\alpha^*(\tau) = \alpha(\tau) + \left[\exp \left(iu \sum_{j=1}^{N_t^d} \bar{Z}_j(\mathbb{Q}) \right) \right]. \quad (6.67)$$

Thus only the constant term in the exponential is changed.

6.6 Impact of Stochastic Volatility on FOMC Estimators

This appendix addresses the concerns of impacts from stochastic volatility and jumps on the FOMC announcement estimators.

The change from constant to stochastic volatility in addition to the augmentation of asymmetric shocks could very well result in systematic biases of the term structure and time series estimates, below we extend the conversation in Dubinsky et al. (2018) to two-factor volatility models with non-constant contemporaneous jumps.

Addressing the concerns around stochastic volatility. From Bates (1996) if discontinuities to volatility and returns are independent then the stochastic volatility option price can be expressed as the expectation of the Black-Scholes price, with the Black-Scholes implied variance being the expected integrated risk-neutral variance:

$$EIV_{t,(T-t)} = \mathbb{E}_t^{\mathbb{Q}} \left[\int_t^{t+T} V_s ds \right]. \quad (6.68)$$

Assuming the Black-Scholes implied variance is an accurate approximation for the expected risk-neutral variance, i.e.,

$$(\sigma_{t,T-t}^{\text{BS}})^2 \approx \mathbb{E}_t^{\mathbb{Q}} \left[\int_t^{t+T} V_s ds \right]. \quad (6.69)$$

The errors in assuming the above should generally be small for this analysis. Hull and White (1987) find, for ATM options, the errors are 1.6% when $\rho = -0.6$ (and less than 1% with no leverage effect). While our parameter estimates (see Table 3.10) indicate a smaller correlation value this effect will be reduced by our use of very short tenor options which we use in our empirical analysis.¹

Addressing concerns raised by jumps. Chernov (2007) investigates the effect of the approximation in models for index options with both jumps in prices and volatility, with a non-zero correlation. Chernov concludes for ATM options the bias is negligible. Therefore, we conclude assuming Equation (6.69) is acceptable and does not introduce any substantial bias.

In the following discussion assume two ATM options maturing at T_1, T_2 ($T_2 > T_1$) with a single FOMC announcement between time t and T_1 . For generality consider a two-factor, square root, model as considered in Equations (3.1)-(3.3). However, in this Appendix for completeness we consider the intensity varies with the level of the factors:

$$\lambda_t = \lambda_0 + \lambda_1 V_t + \lambda_2 m_t, \quad (6.70)$$

as in Bardgett et al. (2019), although we do not estimate a two-factor variance model with

¹Although the errors can be large for OTM options, we use only ATM options for our estimation of $\sigma_j^{\mathbb{Q}}$.

non-constant intensity we include it here for understanding.

Each estimation methodology is effected by the implied variance over option maturities. Thus, to understand how stochastic volatility affects the estimators we compute the expected integrated variance, over an option which matures at time T_i :

$$\begin{aligned}
EIV_{t,T_i-t} = & \frac{1 - e^{-(\kappa_v^Q - \lambda_1 \mu_v^Q)(T_i-t)}}{\kappa_v^Q - \lambda_1 \mu_v^Q} V_t \\
& + (\kappa_v^Q + \lambda_2 \mu_v^Q) \frac{\kappa_m^Q (1 - e^{-(\kappa_v^Q - \lambda_1 \mu_v^Q)(T_i-t)}) - (\kappa_v^Q - \lambda_1 \mu_v^Q) (1 - e^{-\kappa_m^Q(T_i-t)})}{\kappa_m^Q (\kappa_v^Q - \lambda_1 \mu_v^Q) (\kappa_m^Q - (\kappa_v^Q - \lambda_1 \mu_v^Q))} (m_t - \theta_m^Q) \\
& + \left(\frac{(\lambda_0 \mu_v^Q + \theta_m^Q (\kappa_v^Q + \lambda_2 \mu_v^Q))}{\kappa_v^Q - \lambda_1 \mu_v^Q} \right) \left((T_i - t) - \frac{1 - e^{-(\kappa_v^Q - \lambda_1 \mu_v^Q)(T_i-t)}}{\kappa_v^Q - \lambda_1 \mu_v^Q} \right), \quad (6.71)
\end{aligned}$$

for details of the computation of Equation (6.71) see Appendix 6.8. The term structure estimators' accuracy depends on how variable Equation (6.71) is with respect to $T_i - t$. Also from above, the parameters which could impact the estimator are: $m_t - \theta_m^Q$, κ_v^Q , κ_m^Q , λ_0 , λ_1 , λ_2 . Unless there are large risk premiums $\theta_m^Q \approx \theta_m^P$. Bardgett et al. (2019) find evidence for modest risk premium, suggesting θ_m^Q is larger under the historical measure, when using returns and options on the S&P 500. Implying on average $m_t \approx \theta_m^Q$. However, for minimal bias the term structure must also be flat to ensure even in periods of high volatility that V_t is close to m_t thus concluding V_t is also close to θ_m^Q . Table 6.1 reports the implied variance difference between the interpolated 15 and 45 day ATM options, on trading days which are not affected by FOMC announcements. We choose these maturities as they closely represent the front two maturities. We interpolate each IV cross section to get the 15 and 45 day values as each cross section does not certainly contain these exact maturities. This process is repeated for every 30 minute observation in each trading day and then averaged over each year. From Table 6.1 all mean slope estimates using all data are negative, with the exception of 2008. This is likely due to the impact of the financial crisis where short term volatility was abnormally high. Even with 2008 the slope estimates are quite flat, especially for years 2012 on-wards, which is in agreement with Broadie et al. (2007) who find the gradient of the IV term structure is small ($< 1\%$) for S&P 500 options. As the IV term structure is flat this implies V_t will be close to θ_m^Q , on average. This is also supported by evidence from

parameter estimations of the mean reversion. Bardgett et al. (2019) finds $\kappa_m^Q = 0.383$ as well as evidence for a lower rate of mean reversion under the historical measure, implying a very small rate of mean reversion in the stochastic central tendency. We find in Section 3.2.8 κ_m^Q is between 0.003 and 0.320, in support for a low rate of mean reversion. Furthermore, as we only use short maturity options for our estimates, the conclusion is any impact from stochastic volatility and shocks to the term structure estimator is small.

Next, consider the time series estimator. The estimators accuracy depends on how variable Equation (6.71) is as a function of t . The time series estimator is given as:

$$(\sigma_{t,T_i-t}^{\text{BS}})^2 - (\sigma_{t+\delta t,T_i-t-\delta t}^{\text{BS}})^2 = EIV_{t,T_i-t} - EIV_{t+\delta t,T_i-t-\delta t} + (T_i - t)^{-1}(\sigma_j^Q)^2, \quad (6.72)$$

with EIV_{t,T_i-t} being a function of V_t and $EIV_{t+\delta t,T_i-t-\delta t}$ a function of $V_{t+\delta t}$. If $V_t \approx V_{t+\delta t}$ then the time series estimator will be accurate as the effect of mean reversion over δt ($\delta t = \text{one-hour}$) is negligible. However, if volatility changes substantially over δt then the performance rapidly diminishes. Changes in V_t , from the specification above, are given by: $\sigma_v, \sigma_m, \lambda$, i.e., the Brownian paths and the jump intensity. Table 3.9 displays summary statistics on the 1-hour variability in implied variance (IV). We choose a 1-hour interval as when calculating the time series estimator we use information in the observations before and after the announcement which is a 45 or 60 minute interval, depending on the announcement time. Unsurprisingly, the highest average variability is in 2008 due to the financial crisis. However, even with this, the 1-hour variability in IV is between 0.45% and 0.78%. Dubinsky et al. (2018), comment in a daily frequency setting for individual equity options the daily variability is between three to five percent, which implies normal variation could cause significant movements in IV which would cause substantial errors in the time series estimator. However, as our IV variability between observations is significantly lower errors will be negligible. Next, consider the effect of maturity. Intuitively diffusive volatility is more important for longer tenor options, magnifying the impact of shocks. As a result bias will increase with maturity, however we use short term options, thus mitigating this damage.

Our conclusion is as follows: both term structure and time series estimates are reliable and are robust to stochastic volatility and jumps. The term structure estimator depends on

Table 6.1: Term structure slope estimates

This table provides the average term structure slope calculated as the difference between 15 and 45 days at-the-money implied volatilities on trading days that are not within effect of FOMC announcements (we remove all data from 20 days prior to and 5 days after an FOMC announcement). The columns High Vol use only trading dates on which the short-term ATM implied volatility is at least 50% above its average. The columns Low Vol use only trading dates on which the short-term ATM implied volatility is more than 25% below its average. The columns 5% and 95% provide the 5% and 95% percentiles, respectively.

Year	All			High Vol			Low Vol		
	Mean	5%	95%	Mean	5%	95%	Mean	5%	95%
2008	1.65	-1.74	11.99	6.75	-1.22	15.79	-0.61	-1.79	0.74
2010	-1.06	-2.73	1.23	1.93	-0.14	4.31	-2.00	-3.25	-0.89
2012	-0.68	-2.05	1.37	1.51	-0.31	2.58	-1.81	-2.81	-1.39
2014	-0.72	-2.06	0.81	1.93	-0.08	4.53	-1.67	-2.17	-1.27
2016	-0.71	-2.18	1.86	0.77	-0.76	3.39	-1.71	-2.62	-1.08
Avg	-0.31	-2.15	-2.15	2.58	-0.50	-0.50	-1.56	-2.53	-2.53

$m_t - \theta_m^Q, \kappa_v^Q, \kappa_m^Q, \lambda$ and the maturity, for reasonable parameters any bias will be small. The time series estimator depends on $\sigma_v, \sigma_m, \lambda$ and the volatility shocks. In theory, as the term structure estimator is subject to one additional parameter than the time series estimator, it is a noisier estimator.

6.7 Pricing Method

The models are priced using the Fourier Cosine Transform method of Fang and Oosterlee (2009). In this part of the appendix we detail the method and its adaptation to price the two-factor models. For a probability distribution $f(Y)$, the associated characteristic function, $\psi(u)$ is

$$\psi(u) = \int_{-\infty}^{\infty} e^{iuY} f(Y) dY, \quad (6.73)$$

To improve computational efficiency, we truncate the integration bounds to constants $[a, b]$, such that $\hat{\psi}(u) \approx \psi(u)$, where $\hat{\psi}(u)$ denotes the truncated characteristic function. The details of how to calculate these constants are explained below. Choosing to evaluate the

characteristic function at $u = \frac{k\pi}{b-a}$ and multiplying by $e^{-i\frac{k\pi a}{b-a}}$, yields:

$$\hat{\psi}\left(\frac{k\pi}{b-a}\right)e^{-i\frac{k\pi a}{b-a}} = e^{-i\frac{k\pi a}{b-a}} \int_a^b e^{i\frac{k\pi}{b-a}Y} f(Y) dY, \quad (6.74)$$

$$= \int_a^b e^{ik\pi \frac{Y-a}{b-a}} f(Y) dY, \quad (6.75)$$

$$= \int_a^b \left(\cos(k\pi(\frac{Y-a}{b-a})) + i\sin(k\pi(\frac{Y-a}{b-a})) \right) f(Y) dY, \quad (6.76)$$

where in the last step the Euler formula has been used. Taking the real part yields:

$$\operatorname{Re}\left\{\hat{\psi}\left(\frac{k\pi}{b-a}\right)e^{-i\frac{k\pi a}{b-a}}\right\} = \int_a^b \cos(k\pi(\frac{Y-a}{b-a})) f(Y) dY. \quad (6.77)$$

Where the right-hand-side of Equation (6.77) is of the form of the cosine coefficient from the Fourier series expansion. Thus, we may write the Fourier series expansion of the probability density $f(Y)$ as $\tilde{f}(Y)$, with form:

$$\tilde{f}(Y) = \sum_{k=0}^{N-1} \frac{2}{b-a} \operatorname{Re}\left\{\hat{\psi}\left(\frac{k\pi}{b-a}\right)e^{-i\frac{k\pi a}{b-a}}\right\} \cos(k\pi \frac{Y-a}{b-a}), \quad (6.78)$$

after further truncation of the summation. Where $\overline{\sum}$ denotes the first term in the summation is weighted by one-half. It is now possible to use this formalism to express path independent European option prices. The value up to a constant C at time t of an option may be written as

$$v(Y, t) = C \int_a^b v(Y_T, T) f(Y_T|Y) dY_T, \quad (6.79)$$

where v denotes the value of the option and Y_T is the natural logarithm of the forward price at maturity, T . Substituting for the probability density function from the cosine expansion, yielding

$$v(Y, t) = C \int_a^b v(Y_T, T) \overline{\sum}_{k=0}^{\infty} A_k \cos(k\pi \frac{Y_T - a}{b-a}) dY_T. \quad (6.80)$$

Substituting for the Fourier cosine coefficient A_k and truncating the summation yields

$$v(Y, t) \approx C \sum_{k=0}^{N-1} \operatorname{Re} \left\{ \hat{\psi} \left(\frac{k\pi}{b-a} \right) e^{-i \frac{k\pi a}{b-a}} \right\} V_k, \quad (6.81)$$

Where we have defined

$$V_k = \frac{2}{b-a} \int_a^b v(Y_T, T) \cos(k\pi \frac{Y_T - a}{b-a}). \quad (6.82)$$

As we only consider vanilla European options the payoff $v(Y_T, T)$ is expressed as

$$v(Y_T, T) = \max(\iota K(e^{Y_T} - 1), 0), \quad (6.83)$$

where $\iota = 1$ for calls and $\iota = -1$ for puts. After integration of Equation (6.80)

$$V_k^{\text{call}} = \frac{2}{b-a} K(\Xi_k(0, b) - \Phi_k(0, b)), \quad (6.84)$$

$$V_k^{\text{put}} = \frac{2}{b-a} K(\Phi_k(0, a) - \Xi_k(0, a)), \quad (6.85)$$

where the functions Ξ_k and Φ_k have come from the integration of the payoff function with the cosine and are given as:

$$\begin{aligned} \Xi_k(c, d) &= \int_c^d e^{Y_T} \cos(k\pi \frac{Y_T - a}{b-a}) dY_T, \\ &= \frac{1}{1 + (\frac{k\pi}{b-a})^2} [\cos(k\pi \frac{d-a}{b-a}) e^d - \cos(k\pi \frac{c-a}{b-a}) e^c \\ &\quad + \frac{k\pi}{b-a} \sin(k\pi \frac{d-a}{b-a}) e^d - \frac{k\pi}{b-a} \sin(k\pi \frac{c-a}{b-a}) e^c]. \end{aligned} \quad (6.86)$$

The Φ_k for $k \neq 0$, is given as

$$\begin{aligned} \Phi_k(c, d) &= \int_c^d \cos(k\pi \frac{Y_T - a}{b-a}) dY_T, \\ &= \frac{b-a}{k\pi} \left(\sin(k\pi \frac{d-a}{b-a}) - \sin(k\pi \frac{c-a}{b-a}) \right), \end{aligned} \quad (6.87)$$

for the case $k = 0$

$$\Phi_k(c, d) = d - c. \quad (6.88)$$

In terms of calculating the truncation range $[a, b]$ we follow Fang and Oosterlee who quote, without proof, the formula:

$$[a, b] = \left[c_1 - L\sqrt{c_2 + \sqrt{c_4}}, c_1 + L\sqrt{c_2 + \sqrt{c_4}} \right], \quad (6.89)$$

with $L = 10$ and c_n denoting the n -th cumulant. The cumulants are calculated as follows. For a random variable X its corresponding cumulant generating function is given by

$$G(\omega) = \log(\mathbb{E}[\exp(\omega X)]), \quad (6.90)$$

$$= \log(\psi(-i\omega)), \quad (6.91)$$

then the cumulants can be calculated using:

$$c_n = G^{(n)}(0). \quad (6.92)$$

Addressing the application to two-factor models. The complexity comes from the characteristic function only being available via numerical solution of the Equations in (6.61)-(6.63), in contrast to all the one-factor models where the characteristic function is available in closed-form (even with the inclusion of the FOMC jumps as these simply modify the constant term). In order to calculate the cumulants for the two-factor models numerical derivative methods must be relied upon. We solve the differential Equations in (6.61)-(6.63) via numerical integration over a very fine grid in close proximity to zero. We use a cubic spline interpolation to return query values from the numerical derivation procedure.²

On the subject of calculating the characteristic function we use the following procedure. The differential Equations (6.61)-(6.63) are solved overall unique maturities using a logarithmic grid with 4,500 points, ranging from 0 to 5,000. The reason for this choice of point spacing

²The algorithm for the numerical differentiation is the *derivest* suite in Matlab.

is the characteristic function behaves similarly to a damped sinusoidal motion. Thus for small u large sample point density is required, however, for large u the function oscillates closely around zero, requiring sparser sample point density. We use a cubic spline to interpolate the query points $\frac{k\pi}{b-a}$, which are the desired points of evaluation.

6.8 Extracting Latent Volatility

This appendix provides a derivation of the expression used to calculate the risk-neutral expected cumulative variance, an extension of the work of Duan and Yeh (2010). Once again, for completeness we will consider a two-factor stochastic volatility model with non-constant intensity, which depends on the level of the factors. Using Equation (3.1) to obtain:

$$\mathbb{E}_t^{\mathbb{Q}}\left(\ln\frac{F_T}{F_t}\right) = \int_t^T \mathbb{E}_t^{\mathbb{Q}}\left(-\frac{1}{2}V_s - \lambda\bar{\mu}\right)ds + \sum_{j=N_t^d+1}^{N_{t+\tau}^d} \mathbb{E}_t^{\mathbb{Q}}[\bar{Z}_j], \quad (6.93)$$

$$\begin{aligned} &= (-\lambda_0\bar{\mu})(T-t) - (\lambda_1\bar{\mu} + \frac{1}{2}) \int_t^T \mathbb{E}_t^{\mathbb{Q}}(V_s)ds \\ &\quad - \lambda_2\bar{\mu} \int_t^T \mathbb{E}_t^{\mathbb{Q}}(m_s)ds + \frac{1}{2} \sum_{j=N_t^d+1}^{N_{t+\tau}^d} (\sigma_j^{\mathbb{Q}})^2. \end{aligned} \quad (6.94)$$

In order to calculate the expected cumulative variance we start by considering:

$$d\left(e^{(\kappa_v^{\mathbb{Q}} - \lambda_1\mu_v^{\mathbb{Q}})t}V_t\right) = (\kappa_v^{\mathbb{Q}} - \lambda_1\mu_v^{\mathbb{Q}})e^{(\kappa_v^{\mathbb{Q}} - \lambda_1\mu_v^{\mathbb{Q}})t}V_tdt + e^{(\kappa_v^{\mathbb{Q}} - \lambda_1\mu_v^{\mathbb{Q}})t}dV_t. \quad (6.95)$$

Using the result for dV_t from Equation (3.2), with a non-constant intensity, this produces:

$$\begin{aligned} d\left(e^{(\kappa_v^{\mathbb{Q}} - \lambda_1\mu_v^{\mathbb{Q}})t}V_t\right) &= e^{(\kappa_v^{\mathbb{Q}} - \lambda_1\mu_v^{\mathbb{Q}})t}[(m_t(\kappa_v^{\mathbb{Q}} + \lambda_2\mu_v^{\mathbb{Q}}) + \lambda_0\mu_v^{\mathbb{Q}})dt \\ &\quad + \sigma_v\sqrt{V_t}dW_t^v + \xi_t^v dZ_t - \lambda\mu_v^{\mathbb{Q}}dt], \end{aligned} \quad (6.96)$$

integration and taking the expected value yields:

$$\begin{aligned} \mathbb{E}_t^{\mathbb{Q}}(V_T) &= e^{-(\kappa_v^{\mathbb{Q}} - \lambda_1\mu_v^{\mathbb{Q}})(T-t)}V_t + \frac{\lambda_0\mu_v^{\mathbb{Q}}}{\kappa_v^{\mathbb{Q}} - \lambda_1\mu_v^{\mathbb{Q}}}\left(1 - e^{-(\kappa_v^{\mathbb{Q}} - \lambda_1\mu_v^{\mathbb{Q}})(T-t)}\right) \\ &\quad + (\kappa_v^{\mathbb{Q}} + \lambda_2\mu_v^{\mathbb{Q}}) \int_t^T \mathbb{E}_t^{\mathbb{Q}}(m_s)e^{-(\kappa_v^{\mathbb{Q}} - \lambda_1\mu_v^{\mathbb{Q}})(T-s)}ds. \end{aligned} \quad (6.97)$$

In order to calculate the integral in Equation (6.97) use the dynamics outlined out in Equation (3.3). Wishing to evaluate:

$$\begin{aligned} \int_t^T \mathbb{E}_t^{\mathbb{Q}}(m_s) e^{-(\kappa_v^{\mathbb{Q}} - \lambda_1 \mu_v^{\mathbb{Q}})(T-s)} ds &= \int_t^T e^{-(\kappa_v^{\mathbb{Q}} - \lambda_1 \mu_v^{\mathbb{Q}})(T-s)} \left(e^{-\kappa_m^{\mathbb{Q}}(s-t)} m_t + \theta_m^{\mathbb{Q}} \left(1 - e^{-\kappa_m^{\mathbb{Q}}(s-t)} \right) \right) ds, \\ &= \frac{\theta_m^{\mathbb{Q}}}{\kappa_v^{\mathbb{Q}} - \lambda_1 \mu_v^{\mathbb{Q}}} \left(1 - e^{-(\kappa_v^{\mathbb{Q}} - \lambda_1 \mu_v^{\mathbb{Q}})(T-t)} \right) + \frac{e^{-(\kappa_v^{\mathbb{Q}} - \lambda_1 \mu_v^{\mathbb{Q}})(T-t)} - e^{-\kappa_m^{\mathbb{Q}}(T-t)}}{\kappa_m^{\mathbb{Q}} - (\kappa_v^{\mathbb{Q}} - \lambda_1 \mu_v^{\mathbb{Q}})} (m_t - \theta_m^{\mathbb{Q}}), \end{aligned} \quad (6.98)$$

substitution into Equation (6.97) yields:

$$\begin{aligned} \mathbb{E}_t^{\mathbb{Q}}(V_T) &= e^{-(\kappa_v^{\mathbb{Q}} - \lambda_1 \mu_v^{\mathbb{Q}})(T-t)} V_t + \frac{\lambda_0 \mu_v^{\mathbb{Q}}}{\kappa_v^{\mathbb{Q}} - \lambda_1 \mu_v^{\mathbb{Q}}} \left(1 - e^{-(\kappa_v^{\mathbb{Q}} - \lambda_1 \mu_v^{\mathbb{Q}})(T-t)} \right) \\ &\quad + (\kappa_v^{\mathbb{Q}} + \lambda_2 \mu_v^{\mathbb{Q}}) \left(\frac{\theta_m^{\mathbb{Q}}}{\kappa_v^{\mathbb{Q}} - \lambda_1 \mu_v^{\mathbb{Q}}} \left(1 - e^{-(\kappa_v^{\mathbb{Q}} - \lambda_1 \mu_v^{\mathbb{Q}})(T-t)} \right) \right. \\ &\quad \left. + \frac{e^{-(\kappa_v^{\mathbb{Q}} - \lambda_1 \mu_v^{\mathbb{Q}})(T-t)} - e^{-\kappa_m^{\mathbb{Q}}(T-t)}}{\kappa_m^{\mathbb{Q}} - (\kappa_v^{\mathbb{Q}} - \lambda_1 \mu_v^{\mathbb{Q}})} (m_t - \theta_m^{\mathbb{Q}}) \right). \end{aligned} \quad (6.99)$$

Final integration yields:

$$\begin{aligned} \int_t^T \mathbb{E}_t^{\mathbb{Q}}(V_s) ds &= \frac{1 - e^{-(\kappa_v^{\mathbb{Q}} - \lambda_1 \mu_v^{\mathbb{Q}})(T-t)}}{\kappa_v^{\mathbb{Q}} - \lambda_1 \mu_v^{\mathbb{Q}}} V_t \\ &\quad + (\kappa_v^{\mathbb{Q}} + \lambda_2 \mu_v^{\mathbb{Q}}) \frac{\kappa_m^{\mathbb{Q}} (1 - e^{-(\kappa_v^{\mathbb{Q}} - \lambda_1 \mu_v^{\mathbb{Q}})(T-t)}) - (\kappa_v^{\mathbb{Q}} - \lambda_1 \mu_v^{\mathbb{Q}}) (1 - e^{-\kappa_m^{\mathbb{Q}}(T-t)})}{\kappa_m^{\mathbb{Q}} (\kappa_v^{\mathbb{Q}} - \lambda_1 \mu_v^{\mathbb{Q}}) (\kappa_m^{\mathbb{Q}} - (\kappa_v^{\mathbb{Q}} - \lambda_1 \mu_v^{\mathbb{Q}}))} (m_t - \theta_m^{\mathbb{Q}}) \\ &\quad + \left(\frac{(\lambda_0 \mu_v^{\mathbb{Q}} + \theta_m^{\mathbb{Q}} (\kappa_v^{\mathbb{Q}} + \lambda_2 \mu_v^{\mathbb{Q}}))}{\kappa_v^{\mathbb{Q}} - \lambda_1 \mu_v^{\mathbb{Q}}} \right) \left((T-t) - \frac{1 - e^{-(\kappa_v^{\mathbb{Q}} - \lambda_1 \mu_v^{\mathbb{Q}})(T-t)}}{\kappa_v^{\mathbb{Q}} - \lambda_1 \mu_v^{\mathbb{Q}}} \right). \end{aligned} \quad (6.100)$$

Note if we set $\lambda_1 = \lambda_2 = 0$, we achieve the result for the expected cumulative variance of Kaeck and Alexander (2012).

The central tendency calculation is as follows. As we do not model for jumps in the central tendency we start from:

$$d(e^{\kappa_m^{\mathbb{Q}} t} m_t) = \kappa_m^{\mathbb{Q}} e^{\kappa_m^{\mathbb{Q}} t} m_t + e^{\kappa_m^{\mathbb{Q}} t} dm_t. \quad (6.101)$$

Using Equation (3.3) to substitute for dm_t , after cancellation:

$$d(e^{\kappa_m^{\mathbb{Q}}t}m_t) = e^{\kappa_m^{\mathbb{Q}}t}(\kappa_m^{\mathbb{Q}}\theta_m^{\mathbb{Q}}dt + \sigma_m\sqrt{m_t}dW_t^m), \quad (6.102)$$

integration yields:

$$m_T = e^{-\kappa_m^{\mathbb{Q}}(T-t)}m_t + \theta_m^{\mathbb{Q}}\left(1 - e^{-\kappa_m^{\mathbb{Q}}(T-t)}\right) + e^{-\kappa_m^{\mathbb{Q}}} \int_t^T \sigma_m\sqrt{m_s}dW_s^m, \quad (6.103)$$

taking expectation values and final integration yields the end result:

$$\int_t^T \mathbb{E}_t^{\mathbb{Q}}(m_s)ds = \frac{1 - e^{-\kappa_m^{\mathbb{Q}}(T-t)}}{\kappa_m^{\mathbb{Q}}}m_t + \theta_m^{\mathbb{Q}}\left((T-t) - \frac{1 - e^{-\kappa_m^{\mathbb{Q}}(T-t)}}{\kappa_m^{\mathbb{Q}}}\right). \quad (6.104)$$

By the definition of the VIX over the interval $(T-t)$:

$$VIX_t^2(T-t) = \frac{2}{T-t}e^{r(T-t)} \int_0^x \frac{P_t(K)}{K^2}dK + \frac{2}{T-t}e^{r(T-t)} \int_x^\infty \frac{C_t(K)}{K^2}dK, \quad (6.105)$$

using the generic payoff decomposition theorem

$$\begin{aligned} e^{r(T-t)} \int_0^x \frac{P_t(K)}{K^2}dK + e^{r(T-t)} \int_x^\infty \frac{C_t(K)}{K^2}dK &= \frac{F_t(T-t) - x}{x} \\ &\quad - \ln(F_t/x) - \mathbb{E}_t^{\mathbb{Q}}[\ln(F_T/F_t)]. \end{aligned} \quad (6.106)$$

Choosing the threshold $x = F_t(T-t)$ yields:

$$VIX_t^2(T-t) = \frac{-2}{T-t}\mathbb{E}_t^{\mathbb{Q}}[\ln(F_T/F_t)], \quad (6.107)$$

yielding:

$$\begin{aligned} \text{VIX}_t^2(T-t) = & \frac{2}{T-t} \left[\phi(T-t) + V_t \gamma_v \left(\lambda_1 \bar{\mu} + \frac{1}{2} \right) \right. \\ & + m_t \left(\lambda_2 \bar{\mu} \gamma_m + (\kappa_v^{\mathbb{Q}} + \lambda_2 \mu_v^{\mathbb{Q}}) (\lambda_1 \bar{\mu} + 1/2) \right) \end{aligned} \quad (6.108)$$

$$\begin{aligned} & \times \frac{\kappa_m^{\mathbb{Q}} (1 - e^{-(\kappa_v^{\mathbb{Q}} - \lambda_1 \mu_v^{\mathbb{Q}})(T-t)}) - (\kappa_v^{\mathbb{Q}} - \lambda_1 \mu_v^{\mathbb{Q}}) (1 - e^{-\kappa_m^{\mathbb{Q}}(T-t)})}{\kappa_m^{\mathbb{Q}} (\kappa_v^{\mathbb{Q}} - \lambda_1 \mu_v^{\mathbb{Q}}) (\kappa_m^{\mathbb{Q}} - (\kappa_v^{\mathbb{Q}} - \lambda_1 \mu_v^{\mathbb{Q}}))} \Big) \\ & + \lambda_2 \bar{\mu} \theta_m^{\mathbb{Q}} ((T-t) - \gamma_m) + ((T-t) - \gamma_v) \left(\lambda_1 \bar{\mu} + \frac{1}{2} \right) \\ & \times \frac{\theta_m^{\mathbb{Q}} (\kappa_v^{\mathbb{Q}} + \lambda_2 \mu_v^{\mathbb{Q}}) + \lambda_0 \mu_v^{\mathbb{Q}}}{\kappa_v^{\mathbb{Q}} - \lambda_1 \mu_v^{\mathbb{Q}}} - \theta_m^{\mathbb{Q}} (\lambda_1 \bar{\mu} + 1/2) (\kappa_v^{\mathbb{Q}} + \lambda_2 \mu_v^{\mathbb{Q}}) \\ & \times \frac{\kappa_m^{\mathbb{Q}} (1 - e^{-(\kappa_v^{\mathbb{Q}} - \lambda_1 \mu_v^{\mathbb{Q}})(T-t)}) - (\kappa_v^{\mathbb{Q}} - \lambda_1 \mu_v^{\mathbb{Q}}) (1 - e^{-\kappa_m^{\mathbb{Q}}(T-t)})}{\kappa_m^{\mathbb{Q}} (\kappa_v^{\mathbb{Q}} - \lambda_1 \mu_v^{\mathbb{Q}}) (\kappa_m^{\mathbb{Q}} - (\kappa_v^{\mathbb{Q}} - \lambda_1 \mu_v^{\mathbb{Q}}))} - \frac{1}{2} \sum_{j=N_t^d+1}^{N_{t+\tau}^d} \sigma_j^2 \Big]. \end{aligned} \quad (6.109)$$

Where:

$$\phi = \lambda_0 \bar{\mu}, \quad (6.110)$$

$$\gamma_m = \frac{1 - e^{-\kappa_m^{\mathbb{Q}}(T-t)}}{\kappa_m^{\mathbb{Q}}}, \quad (6.111)$$

$$\gamma_v = \frac{1 - e^{-(\kappa_v^{\mathbb{Q}} - \lambda_1 \mu_v^{\mathbb{Q}})(T-t)}}{\kappa_v^{\mathbb{Q}} - \lambda_1 \mu_v^{\mathbb{Q}}}. \quad (6.112)$$

6.9 Time Series Simulation Multi-Factor Models

The purpose of this appendix is to detail the checking procedures we use to ensure the numerical methodology to price the two-factor models is correct and errors produced are acceptable. We compare the numerical ordinary differential equation (ODE) framework solution to the model (as detailed in Appendix 6.7) with that of a Monte Carlo (MC) framework.

We use an Euler discretization scheme for the asset price and variance processes, with a full truncation to tackle the issue of negative variances (see Lorde et al. (2010)). Conditional

on a time s for $t > s$ the discretization scheme for the asset price and variance processes are:

$$Y_t = Y_s - (\lambda\bar{\mu} + \frac{1}{2}V_s^+)\Delta t + \sqrt{V_s^+\Delta t}z_S + \left(e^{\xi_t^y} - 1\right)b_t, \quad (6.113)$$

$$V_t = V_s^+ - \kappa_v(m_s^+ - V_s^+)\Delta t + \sigma_v\sqrt{V_s^+\Delta t}z_V + \xi_t^v b_t, \quad (6.114)$$

$$m_t = m_s^+ + \kappa_m(\theta_m - m_s^+)\Delta t + \sigma_m\sqrt{m_s^+}\Delta t z_m. \quad (6.115)$$

where Y_t denotes $\log(F_t)$, the log forward price. Denote $\Delta t = t - s$, $z_V, z_m \sim \mathcal{N}(0, 1)$ and $z_S = \rho z_V + \sqrt{1 - \rho^2}z$ with $z \sim \mathcal{N}(0, 1)$ and $x^+ = \max(x, 0)$, refers to the full truncation scheme. To simulate the Poisson process we use a Bernoulli process, denoted by b_t , which takes the value either 0 or 1 and is identically distributed and independent. This scheme is used to generate 20 different sample paths of weekly returns and latent variances over one year, i.e., 52 observations with $\Delta t = 7$ days.³ At each observation time we simulate seven unique maturities between 4 and 550 days to maturity. Following a maturity profile of the first three monthly's, then followed by three quarterly option expiries and finally one yearly expiry. We simulate strikes which create a moneyness range between 0.8 and 1.1 across one-hundred strikes. Each option price is computed using the MC framework⁴ using 500,000 simulations and a time-step of $1/40^{\text{th}}$ of a trading day for maturities less than 30-days and 50,000 simulations and a time-step of $1/10^{\text{th}}$ of a trading day for maturities greater than or equal to 30-days.⁵ This choice of simulation paths and time-step based upon maturity is somewhat ad-hoc, but reasonable. Table 6.2 documents static VWRMSE values between the MC procedure and the numerical ODE framework solution, compared for five maturities in the maturity range. This static exercise is done using an ad-hoc but reasonable parameter vector. From column two of Table 6.2 the SV model has errors ranging from 0.0437% to 0.1185%, column three documents a range for the SV2 model of 0.0204% to 0.1643% and column four presents an error range for the SVCJ2 model of 0.0187% to 0.1897%.

While all model errors from Table 6.2 across the maturity range are negligible, it is

³We choose a time step of one week so that the variance and central tendency demonstrate reasonable change.

⁴In order to reduce the variance of the estimated prices we use antithetic variates when simulating our standard normal random numbers used in z_V, z_m, z_S .

⁵The reason we use two different time-step and number of simulations regimes is due to computational burden.

Table 6.2: Static VWRMSE of Monte Carlo against numerical ODE solution

Static VWRMSE between MC solution and ODE framework solution. This is taken across five different maturities ranging between 4 and 550 days to expiry, displayed in the first column under DTM. The other three columns display the VWRMSE, as a percentage, for each model: SV, SV2 and SVCJ2. The parameter vectors are as follows: $\Theta^{SVCJ2} = \{\kappa_v, \theta_m, \sigma_v, \rho, \kappa_m, \sigma_m, \lambda, \mu_y, \sigma_y, \mu_v\}$ and $\Theta^{SVCJ2} = (1.1768, 0.082259, 0.43746, -0.54589, 2, 0.2, 5, -0.005, 0.014, 0.013)$. Where $\Theta^{SV}, \Theta^{SV2}$ are the first four and six parameters of Θ^{SVCJ2} respectively.

DTM	SV	SV2	SVCJ2
4	0.0701	0.0701	0.0484
30	0.0437	0.0284	0.0397
100	0.1185	0.0204	0.0187
200	0.0535	0.164	0.1430
550	0.0535	0.0291	0.1897

necessary to undertake a calibration exercise against the simulated time-series data in order to establish there are no errors or inconsistencies in the ODE framework. The calibration process is done using the objective function defined in Equation (3.15). As initial conditions we start with the true parameter vector, i.e., Θ^{true} , we also pass the variance paths at each cross-section. The results of the simulated calibration exercise for models SV2 and SVCJ2 are reported in Table 6.3, we also report results for the SV model in order to benchmark the errors expected from the MC on the first four parameters. Comparison between the numerically related models (SV2 and SVCJ2) with the SV Bias and RMSE show negligible errors come from the numerical framework as the Bias and RMSE across the three models are comparable. In examining the RMSE and Bias columns of the SV2 and SVCJ2 models, it appears the following parameters: $\kappa_V, \kappa_m, \sigma_m$ might prove a challenge to estimate in a full calibration setup, as they have the highest Bias and RMSE values, consistent across models.

Table 6.3: Simulated calibration exercise

Simulated results for calibration exercise between the SV, SV2 and SVCJ2 models generated from the MC framework and the proposed numerical framework (only relevant for the SV2, SVCJ2 models). The models are estimated with the VWRMSE objective function defined in Equation (3.15). True parameters used in the MC simulation are given by the True row. Mean bias between the true parameters and the calibrated parameters are given in the second row denoted Bias. Third row RMSE reports the standard deviation of calibrated parameter results.

	SV			SV2			SVCJ2		
	True	Bias	RMSE	True	Bias	RMSE	True	Bias	RMSE
κ_v	1.177	0.004	0.005	1.177	0.005	0.009	1.177	0.002	0.008
θ	0.080	0.000	0.000	0.080	0.000	0.000	0.080	0.000	0.000
σ_v	0.400	0.004	0.004	0.400	0.005	0.003	0.400	0.000	0.000
ρ	-0.546	0.019	0.007	-0.546	0.013	0.004	-0.546	0.002	0.004
κ_m				2.000	0.005	0.040	2.000	0.019	0.060
σ_m				0.200	0.071	0.055	0.200	0.074	0.087
λ							5.000	0.000	0.000
μ_y							-0.005	0.001	0.001
σ_y							0.014	0.000	0.001
μ_v							0.013	0.000	0.000

6.10 Short Maturity Data Set Calibrations

This appendix reports the results of the short maturity data set calibrations we use to produce the results in Table 3.13. We calibrate the original model (SV or SVJ) to each the short and ultra-short maturity categories. The results from the calibration exercise is displayed in Table 6.4. In comparison between the reduced maturity data sets and the full maturity profile the SV model changes to parameters do not qualitatively change the probability density function, except for the σ_v estimates, which are consistently and significantly higher. This increased value causes weight to shift out of the tails and into a sharp peak. Regarding the SVJ model the change in parameters does not qualitatively change the probability distribution significantly. The final step to produce Table 3.13 is to fix the original structural parameters and allow the time adjusted component to vary. In this way we ensure any error reduction is completely down to the effect of accounting for a particular volatility seasonality and not some non-trivial interplay of the time component with one or more of the original model's parameters.

6.11 Maturity Related Effects

Here we demonstrate the effects of using too many long term options in the sample data and its effect on time adjusted parameter estimates from the calibration exercise. Figure 6.1 displays the objective function (Vega-Weighted RMSE (VWRMSE)) as a function of the trading ratio R_{tr} for the SV_{tr} model in 2010 using the short maturity data sample. Clearly for R_{tr} less than approximately 0.25 the objective function becomes incredibly flat, with little variation for large moves in value of R_{tr} . As a result the optimiser⁶ will struggle to find a value and often ends up choosing a fairly extreme value in these cases.

⁶We use a derivative based optimiser called *patternsearch* from the Global Optimizer Toolbox in Matlab.

Figure 6.1: Trading ratio VWRMSE

Vega-Weighted RMSE (VWRMSE) plot for the SV_{tr} model calibrated to the short maturity data set (between 8 and 20, inclusive, days to maturity) in 2010. The horizontal axis displays the trading ratio between the model and the calendar time convention R_{tr} and the vertical axis displays the VWRMSE in percentage.

