

ABSTRACT

Title of dissertation: **THREE ESSAYS ON VOLATILITY
ISSUES IN FINANCIAL MARKETS**

George Panayotov, Doctor of Philosophy, 2005

Dissertation directed by: Professor Dilip Madan
 Department of Finance

Studies of asset returns time-series provide strong evidence that at least two stochastic factors drive volatility. The first essay investigates whether two volatility risks are priced in the stock option market and estimates volatility risk prices in a cross-section of stock option returns. The essay finds that the risk of changes in short-term volatility is significantly negatively priced, which agrees with previous studies of the pricing of a single volatility risk. The essay finds also that a second volatility risk, embedded in longer-term volatility is significantly positively priced. The difference in the pricing of short- and long-term volatility risks is economically significant - option combinations allowing investors to sell short-term volatility and buy long-term volatility offer average profits up to 20% per month.

Value-at-Risk measures only the risk of loss at the end of an investment horizon. An alternative measure (MaxVaR) has been proposed recently, which quantifies the risk of loss at or before the end of an investment horizon. The second essay studies such a risk measure for several jump processes (diffusions with one- and two-sided jumps and two-sided pure-jump processes with different structures of jump arrivals). The main tool of analysis is the first passage probability. MaxVaR for jump processes is compared to standard VaR using returns to five major stock indexes over investment horizons up to one month. Typically MaxVaR is 1.5 - 2 times higher than standard VaR, whereby the excess tends to be higher for longer investment horizons and for lower quantiles of the returns distributions. The results of the essay provide one possible justification for the multipliers applied by the Basle Committee to standard VaR for regulatory purposes.

Several continuous-time versions of the GARCH model have been proposed in the literature, which typically involve two distinct driving stochastic processes. An interesting alternative is the COGARCH model of Kluppelberg, Lindner and Maller (2004), which is driven by a single Levy process. The third essay derives a backward PIDE for the COGARCH model, in the case when the driving process is Variance-Gamma. The PIDE is applied for the calculation of option prices under the COGARCH model.

THREE ESSAYS ON VOLATILITY ISSUES IN FINANCIAL MARKETS

by

George Panayotov

Dissertation submitted to the Faculty of the Graduate School of the
University of Maryland, College Park in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
2005

Advisory Committee:

Professor Dilip Madan, Chair
Professor Haluk Unal
Professor Gurdip Bakshi
Professor Steven Heston
Professor Konstantina Trivisa

ACKNOWLEDGEMENTS

I am deeply grateful to Prof. Dilip Madan for all the time that he generously spent, sharing his ideas. I am indebted to Prof. Haluk Unal for his support over the years, and in particular during the last year of my studies at the University of Maryland.

Contents

Contents	iii
List of tables	v
List of graphs	v
1 Understanding the structure of volatility risks	1
1.1 Introduction	1
1.2 Volatility risk prices in models with one and two volatility factors	4
1.2.1 The price of a single volatility risk	4
1.2.2 Models with multiple volatility factors	8
1.3 Design of the empirical tests	10
1.3.1 Construction of option returns	10
1.3.2 Data and option-pricing models	14
1.3.3 Volatility risk factors	20
1.4 Estimation of volatility risk prices	23
1.4.1 Estimation results	23
1.4.2 Evidence from calendar spreads	37
1.5 Conclusion	40
2 MaxVar for processes with jumps	42
2.1 Introduction	42
2.2 Models and first-passage probabilities	45
2.2.1 CMYD	45
2.2.2 Double exponential jump-diffusion	48
2.2.3 CGMY	51

2.2.4	Finite-Moment Log-Stable (FMLS)	58
2.3	Empirical results	60
2.4	Conclusion	69
3	The COGARCH model and option pricing	70
3.1	Introduction	70
3.2	The COGARCH process	71
3.3	Backward PIDE for European options	74
3.4	Conclusion	88
	Appendix A. Proof of Proposition 1	89
	Appendix B. Volatility risk prices in a two-factor model	90
	Appendix C. Option pricing models	96
	References	101

List of tables

Table 1.1. Fixed moneyness and maturity levels

Table 1.2. Option data

Table 1.3. Estimation errors

Table 1.4. VIX and risk-neutral standard deviation

Table 1.5. Average option excess returns

Table 1.6. Volatility risk prices - one vs. two volatility factors

Table 1.7. Volatility risk prices - raw vs. orthogonal volatility risks

Table 1.8. Average excess returns on delta-hedged option portfolios

Table 1.9. Volatility risk prices - two volatility factors

Table 1.10. Volatility risk prices - one volatility factor

Table 1.11. Volatility risk prices - abs. market returns and one volatility factor

Table 1.12. Average returns to calendar spreads

Table 2.1. Goodness-of-fit tests

Table 2.2. Levy VaR multiples over Normal VaR - 10 days

Table 2.3. Levy VaR multiples over Normal VaR - 20 days

Table 2.4. MaxVaR multiples over Normal VaR - 10 days

Table 2.5. MaxVaR multiples over Normal VaR - 20 days

Table 2.6. Frequency of excessive losses

Table 3.1. Put prices - put-call parity vs. PIDE

List of graphs

Graph 1. Discretization errors - calls

Graph 2. Discretization errors - calls

Graph 3. COGARCH call option prices

Graph 4. COGARCH put option prices

1 Understanding the structure of volatility risks

1.1 Introduction

Recent studies provide evidence that market volatility¹ risk is priced in the stock option market (e.g. Chernov and Ghysels (2000), Benzoni (2001), Coval and Shumway (2001), Bakshi and Kapadia (2003), Carr and Wu (2004)). These studies typically find a negative volatility risk price, suggesting that investors are ready to pay a premium for exposure to the risk of changes in volatility. All these studies consider the price of risk, embedded in a *single* volatility factor.

In contrast, time-series studies find that *more than one* stochastic factor drives asset returns volatility. Engle and Lee (1998) find support for a model with two volatility factors - permanent (trend) and transitory (mean-reverting towards the trend). Gallant, Hsu and Tauchen (1999), Alizadeh, Brandt and Diebold (2002) and Chernov, Gallant, Ghysels and Tauchen (2003) estimate models with one highly persistent and one quickly mean-reverting volatility factor and show that they dominate over one-factor specifications for volatility².

Motivated by the results of the time-series studies, this paper investigates whether the risks in two volatility risks are priced in the stock option market. I construct the volatility factors using implied volatilities from index options with different maturities (between one month and one year).

My main finding is that two volatility risks are indeed priced in the stock option market. The risk of changes in short-term volatility is significantly negatively priced. This result is consistent with the previous volatility risk pricing literature, which uses relatively short-term options and finds negative price of volatility risk.

¹This paper only considers stock market volatility and does not touch upon the volatility of individual stocks. For brevity I will refer to market volatility as "volatility".

²See also Andersen and Bollerslev (1997), Liesenfeld (2001), Jones (2003).

In addition, the paper reports a novel finding - I find that another risk, embedded in longer-term volatility is significantly positively priced. The positive risk price indicates that investors require positive compensation for exposure to long-term volatility risk.

This finding complements previous results on the pricing of two volatility factors in the stock market: Engle and Lee (1998) find that the permanent (or persistent) factor in volatility is significantly positively correlated with the market risk premium, while the transitory factor is not. MacKinlay and Park (2004) confirm the positive correlation of the permanent volatility factor with the risk premium and also find a time-varying and typically negative correlation of the transitory volatility factor with the risk premium.

The differential pricing of volatility risks in the stock option market, found in this paper, is also economically significant, as evidenced by returns on long calendar spreads³. Expected returns on a calendar spread reflect mostly the compensations for volatility risks embedded in the two components of the spread. A short position in a negatively priced volatility risk (short-term) combined with a long position in a positively priced volatility risk (long-term) should then have a positive expected return. I calculate returns on calendar spreads written on a number of index and individual options and find that, in full support of the statistical estimations, spreads on puts gain an impressive 20% monthly on average, while spreads on calls gain about 12% on average. Transaction costs would reduce these numbers, but still, a pronounced difference in volatility risk prices can be captured using calendar spreads.

To perform the empirical tests I construct time-series of daily returns to options

³A long calendar spread is a combination of a short position in an option with short maturity and a long position in an option on the same name, of the same type and with the same strike, but of a longer maturity.

of several fixed levels of moneyness and maturity⁴. For this construction I use options on six stock indexes and twenty two individual stocks. I estimate volatility risk prices in the cross-section of expected option returns using the Fama-MacBeth approach and Generalized Method of Moments (GMM). These methods have not been applied to a cross-section of option returns in previous studies of volatility risk pricing. The cross-sectional analysis allows easily to decompose implied volatility and to estimate separately the prices of the risks in different components. This decomposition turns out to be essential for the disentangling of the risks, embedded in long-term implied volatility.

What is the economic interpretation of the different volatility risk prices found in the paper? There is still little theoretical work on the pricing of more than one volatility risks. Tauchen (2004) studies a model with two consumption-related stochastic volatility factors, which generates endogenously two-factor volatility of stock returns. One feature of Tauchen's model is that the risk prices of the two volatility factors are necessarily of the same sign. This paper offers an alternative model, which is able to generate volatility risk prices of different signs, consistent with the empirical findings. In this model the representative investor's utility function is concave in one source of risk and convex in a second source of risk (for reasonable levels of risk aversion). Such a utility function is closely related to multiplicative habit formation models (e.g. Abel (1990)). Both risk sources exhibit stochastic volatility. The negative volatility risk price is associated to concavity of the utility function, whereas the positively priced volatility risk is associated to convexity of the utility function.

The rest of the paper is organized as follows. Section 1.2 discusses the pricing

⁴Such constructs have been used before on a limited scale - e.g. the CBOE used to derive the price of an at-the-money 30-day option to calculate the Volatility Index (VIX) from 1986 to 1993; Buraschi and Jackwerth (2001) use 45-day index options with fixed moneyness levels close to at-the-money

of volatility risks in models with one and two volatility factors. Section 1.3 describes the construction of option returns and the volatility risk factors. Section 1.4 presents the estimation results and Section 1.5 concludes.

1.2 Volatility risk prices in models with one and two volatility factors

This section considers first the pricing of volatility risk in a model with a single stochastic volatility factor. It complements previous empirical findings of a negative volatility risk price deriving analytically such a negative price in a stochastic volatility model of the type studied in Heston (1993). A negative price of volatility risk is obtained in this model even if no correlation between asset returns and volatility is assumed. This model is later extended to include a second volatility factor. The predictions of the extended model are consistent with the empirical findings of this paper. Next, the section discusses the evidence for two volatility factors and specifies the relation that is tested empirically in the rest of the paper.

1.2.1 The price of a single volatility risk

Consider a standard economy with a single volatility risk. The representative investor in this economy holds the market portfolio and has power utility over the terminal value of this portfolio: $U_T = \frac{(S_T)^{1-\lambda}}{1-\lambda}$. The pricing kernel process in this economy is of the form:

$$\Lambda_t = E_t [S_T^{-\lambda}] \tag{1.2.1}$$

Expectation is taken under the statistical measure, S denotes the value of the market portfolio and λ is the risk aversion coefficient. Assume the following dynamics

for S :

$$\frac{dS_t}{S_t} = D_t^S dt + \sqrt{\sigma_t^S} dW_t^S \quad (1.2.2)$$

where W_t^S is standard Brownian motion. The drift D_t^S is not modeled explicitly, since it does not affect the pricing kernel in the economy. The volatility $\sqrt{\sigma_t^S}$ is stochastic. Assume further that σ_t^S follows a CIR process, which is solution to the following stochastic differential equation:

$$d\sigma_t^S = k(\theta - \sigma_t^S)dt + \eta\sqrt{\sigma_t^S}dW_t \quad (1.2.3)$$

The model (2)-(3) is Heston's (1993) stochastic volatility model (with possibly time-varying drift D_t^S). The Brownian motions W_t^S and W_t are assumed here to be uncorrelated. I discuss below the implication for volatility risk pricing of the correlation between W_t^S and W_t .

Appendix A contains the proof of the following:

Proposition 1 *The stochastic discount factor ζ_t for the economy with one stochastic volatility factor described in (1.2.1)-(1.2.3) is given by*

$$\zeta_t = -\frac{d\Lambda_t}{\Lambda_t} = \lambda \frac{dS_t}{S_t} + \lambda^\sigma d\sigma_t \quad (1.2.4)$$

where the price of market risk λ is strictly positive and the price of volatility risk λ^σ is strictly negative.

The economic intuition for this negative risk price can be provided by the concavity of the utility function. Higher volatility results in lower expected utility (due to concavity). Then any asset, which is positively correlated with volatility has high payoff precisely when expected utility is low. Hence, such an asset acts as

insurance and investors are ready to pay a premium for having it in their portfolio. This argument is well known, e.g. from the risk management literature.

Given the linear form of ζ_t , it is easy to test the model's prediction for the sign of λ^σ in cross-sectional regressions of the type

$$E[R^i] = \lambda^M \beta_i^M + \lambda^\sigma \beta_i^\sigma + \gamma^i \quad (1.2.5)$$

where: $E[R^i]$ is expected excess return on test asset i ; the betas are obtained in time-series regressions of asset i 's returns on proxies for market risk $\left(\frac{dS_t}{S_t}\right)$ and for volatility risk $(d\sigma_t)$; λ^M and λ^σ are the prices of market risk and volatility risk respectively and γ^i is the pricing error. If the two risk factors are normalized to unit variance, this beta-pricing representation yields risk prices equivalent to those in the stochastic discount factor model (1.2.4) (e.g. Cochrane (2001), Ch. 6).

Ang et al. (2004) test a model similar to (1.2.5) and find significantly negative price of volatility risk in the cross-section of expected stock returns. Empirical tests of (1.2.5) involving option returns are reported in Section 1.4 in this paper and also support a negative price of volatility risk.

These tests of the simple relation (1.2.5) are consistent with most of the previous studies of volatility risk pricing in the options markets. Benzoni (2001), Pan (2002), Doran and Ronn (2003), among others, use parametric option-pricing models to estimate a negative volatility risk price from option prices and time series of stock market returns. Coval and Shumway (2001) argue that if volatility risk were not priced, then short delta-neutral at-the-money straddles should earn minus the risk free-rate. In contrast, they find 3% average gain per week, which is (tentatively) interpreted as evidence that market volatility risk is negatively priced. Within a general two-dimensional diffusion model for asset returns, Bakshi and Kapadia

(2003) derive that expected returns to delta-hedged options are positive (negative) exactly when the price of volatility risk is positive (negative) and find significantly negative returns. Carr and Wu (2004) construct synthetic variance swap rates from option prices and compare them to realized variance - the variance risk premium obtained in this way is significantly negative.

The derivation of (1.2.4) assumes that the two Brownian motions are uncorrelated. If ρ is a non-zero correlation between the Brownian motions, then the volatility risk price has two components: $\rho\lambda$ and $\lambda^\sigma\sqrt{1-\rho^2}$. Previous empirical studies have often argued that the negative price of volatility risk they find is due to the negative correlation between changes in volatility and stock returns. So, they focus on the negative $\rho\lambda$ term. This is a powerful argument, given that this negative correlation is among the best-established stylized facts in empirical finance (e.g. Black (1976)). Exposure to volatility risk is thus seen as hedging against market downturns and the negative volatility risk price is seen as the premium investors pay for this hedge. However, focusing on the negative correlation leaves aside the second term in volatility risk price ($\lambda^\sigma\sqrt{1-\rho^2}$). That such a term can be important is indicated e.g. in Carr and Wu (2004) - they find that even after accounting for the correlation between market and volatility risks, there still remains a large unexplained negative component in expected returns to variance swaps. A negative λ^σ as derived above is consistent with this finding. To explore the relative significance of the different components of volatility risk price, this paper reports empirical tests which include a market risk factor and uncorrelated volatility risks.

1.2.2 Models with multiple volatility factors

The study of models with multiple volatility factors has been provoked partly by the observation that the volatility of lower frequency returns is more persistent than the volatility of higher frequency returns. This pattern can be explained by the presence of more than one volatility factors, each with a different level of persistence. Such factors have been interpreted in several ways in the literature:

Andersen and Bollerslev (1997) argue that volatility is driven by heterogeneous information arrival processes with different persistence. Sudden bursts of volatility are typically dominated by the less persistent processes, which die out as time passes to make the more persistent processes influential. Muller et al. (1997) focus on heterogeneous agents, rather than on heterogeneous information processes. They argue that different market agents have different time horizons. The short-term investors evaluate the market more often and perceive the long-term persistent changes in volatility as changes in the average level of volatility at their time scale; in turn, long-term traders perceive short-term changes as random fluctuations around a trend. Liesenfeld (2001) argues that investors' sensitivity to new information is not constant but time-varying and is thus a separate source of randomness in the economy. He finds that the short-term movements of volatility are primarily driven by the information arrival process, while the long-term movements are driven by the sensitivity to news.

MacKinlay and Park (2004) study the correlations between the expected market risk premium and two components of volatility - permanent and transitory. The permanent component is highly persistent and is significantly positively priced in the risk premium, suggesting a positive risk-returns relation. The transitory component is highly volatile and tends to be negatively priced in the risk premium. This component is related to extreme market movements, transitory market regu-

lations, etc. which can be dominating volatility dynamics over certain periods of time.

Tauchen (2004) is the first study to incorporate two stochastic volatility factors in a general equilibrium framework⁵. The two-factor volatility structure is introduced by assuming consumption growth with stochastic volatility, whereby the volatility process itself exhibits stochastic volatility (this is the second source of randomness in volatility). The model generates endogenously a two-factor conditional volatility of the stock return process. It also generates a negative correlation between stock returns and their conditional volatility as observed empirically in data. One feature of his model is that the risk premia on the two volatility factors are both multiples of the same stochastic process (the volatility of consumption volatility) and are necessarily of the same sign. The model thus imposes a restriction on the possible values of volatility risk prices.

The model with two volatility factors which is tested in this paper does not impose *a priori* restrictions on volatility risk prices. In analogy with (1.2.5) I estimate cross-sectional regressions of the type:

$$E [R^i] = \lambda^M \beta_i^M + \lambda^L \beta_i^S + \lambda^L \beta_i^S + \gamma^i. \quad (1.2.6)$$

where:

- the test assets are options (unhedged and delta-hedged) on a number of stock indexes and individual stocks and $E [R^i]$ denotes expected excess returns on option i ;

- the betas are obtained in time-series regressions of option i 's returns on proxies

⁵In a related work Bansal and Yaron (2004) model consumption growth as containing a persistent predictable component plus noise. Stochastic volatility is incorporated both in the persistent component and the noise. However, only one source of randomness in volatility is assumed in their model, common to both components of consumption growth.

for market risk $\left(\frac{dS_t}{S_t}\right)$ and for two volatility risks $(d\sigma_t^S$ and $d\sigma_t^L)$; in particular, I consider a short- and a long-term volatility, denoted by superscripts S and L resp.

- λ^M , λ^S and λ^L are the prices of market risk and the two volatility risks respectively, and γ^i is the pricing error.

Equation (1.2.6) presents a three-factor model with a linear stochastic discount factor. Similar specifications have been widely and successfully employed in empirical asset-pricing tests. It would be interesting, though, to search for an economic justification for equation (1.2.6). One possibility would be to formally extend the model in (1.2.1)-(1.2.3) by adding a second source of randomness with stochastic volatility to (1.2.2). However, this approach would come up with the prediction that both volatility risk prices are negative. An alternative economic model consistent with (1.2.6) is presented in Appendix B. This model is related to habit-formation models and is less restrictive - it predicts that one volatility risk is always negatively priced, while the price of the second risk can have both signs. Such a model is consistent with the empirical finding of this paper that one volatility risk is significantly positively priced.

1.3 Design of the empirical tests

This section describes the construction of the option returns and volatility risk factors, used in testing (1.2.6).

1.3.1 Construction of option returns

I construct daily returns on hypothetical options with fixed levels of moneyness and maturity. In particular, the fixed maturity levels allow to focus on possible

maturity effects in option returns, which would be blurred if, for example, only options held until expiration are considered. The convenience of working with such constructs is well known and has been exploited in different contexts. From 1993 till 2004 the Chicago Board of Options Exchange was calculating the Volatility Index (VIX) as the implied volatility of a 30-day option, struck at-the money forward. In a research context, Buraschi and Jackwerth (2001) construct 45-day options with fixed moneyness levels close to at-the-money. I follow this approach, and extend it to a number of fixed moneyness levels and maturities, ranging from one month to one year.

Options with predetermined strikes and maturities, most likely, were not actually traded on the exchange on any day in the sample. To find their prices, I apply the following two-step procedure. First, I calibrate an option-pricing model to extract the information contained in the available option prices. Next, this estimated model is used to obtain the prices of the specific options I need. This approach has only recently been made feasible by the advent of models, which are capable of accurately calibrating options in the strike and the maturity dimensions together. Section 3.2 presents three models of this type. Extracting information from available options to price other options is a standard procedure. This is how prices are quoted in over-the-counter option markets - if not currently observed, the option price is derived from other available prices by interpolation or a similar procedure. Also, options that have not been traded on a given day are marked-to-market in traders' books in a similar way ⁶.

Once the model is estimated, any option price can be obtained off it. I consider

⁶It is possible to avoid the use of an option-pricing model and apply instead some polynomial smoothing on the implied volatilities of observed options. While an obvious advantage of this approach is that observed prices are fitted exactly, the downside is that when we need to extrapolate to strikes beyond the range of the observed strikes, this procedure is known to be very inaccurate.

three levels of moneyness for puts and for calls, and for each one of five maturities, as given in the Table 1.1. (The table shows the ratios between the actual strikes employed and the at-the-money forward price.) The maturities are one, three, six, nine and twelve months and all strikes are at- or out-of-the-money. To capture the fact that variance increases with maturity, the range of strikes increases accordingly. In this way, for each name I construct 30 time-series of option returns. I estimate volatility risk prices both using all separate time-series and using portfolios, constructed from these series.

Returns to unhedged options are constructed as follows. On each day I calculate option prices on the grid of fixed strikes and maturities. Then I calculate the prices of the same options on the following day - i.e. I keep the strikes but use the following day's parameters and spot price and decrease the time to maturity accordingly. I also take into account the cost of carrying the hedge position to the next day. The daily return on a long zero-cost position in the option is the difference between the second day's option price and the first day's option price with interest:

$$R = O(\theta_2, S_2, K_1, T_2) - O(\theta_1, S_1, K_1, T_1)(1 + r\Delta) \quad (1.3.1)$$

The indexes 1 and 2 refer to the first and second day and O denotes an option price. S , K , r and T are spot, strike, interest rate and time to maturity respectively, θ is an estimated set of parameters and $\Delta = T_1 - T_2$. Note that as the spot price changes from day to day, the strikes used also change, since the grid of moneyness levels is kept constant. Finally, to make the dollar returns obtained in this way comparable across maturities and names I scale these returns by the option price in the first day.

To calculate returns to delta-hedged options, the delta-hedge ratio is needed

Table 1.1. Fixed moneyness and maturity levels

Moneyness levels for five fixed maturities at which option returns are calculated as for each name. Moneyness is the ratio between option strike and spot.

Maturity	Puts			Calls		
1 m.	0.90	0.95	1	1	1.05	1.10
3 m.	0.85	0.90	1	1	1.10	1.15
6 m.	0.80	0.90	1	1	1.10	1.20
9 m.	0.80	0.90	1	1	1.10	1.20
12m.	0.75	0.85	1	1	1.15	1.25

as well. I obtain it numerically in the following way: on each day I move up the spot price by a small amount epsilon, calculate the option price at the new spot (everything else kept the same) and divide the difference between the new and old option prices by epsilon. The daily return on a long zero-cost position in the delta-hedged option is the difference between the second day's option price and the first day's option price with interest less delta times the difference between the second day's spot price and the first day's spot price with interest:

$$R = O(\theta_2, S_2, K_1, T_2) - O(\theta_1, S_1, K_1, T_1)(1 + r\Delta) - \delta(S_2 - S_1(1 + r\Delta)) \quad (1.3.2)$$

where δ is the delta-hedge ratio and all other parameters are as before.

To make dollar returns comparable across maturities and names, I scale them by the price of the underlying asset. Scaling by the option price is possible, but it disregards the hedging component, which can be much higher than the option one.

Of course, the convenient collection of time-series of option returns comes at the price of daily rebalancing - closing the option position at the previous actual strike and maturity and entering into a new position at a new actual strike (but at the fixed moneyness) and same maturity. By constructing option returns in this way, I assume away the thorny issue of transaction costs. However, I provide an alternative check for the statistical results by considering monthly returns to calendar spreads.

1.3.2 Data and option-pricing models

All the data I use come from OptionMetrics, a financial research firm specializing in the analysis of option markets. The "Ivy DB" data set from OptionMetrics

contains daily closing option prices (bid and ask) for all US listed index and equity options, starting in 1996 and updated quarterly. Besides option prices, it also contains daily time-series of the underlying spot prices, dividend payments and projections, stock splits, historical daily interest rate curves and option volumes. Implied volatilities and sensitivities (delta, gamma, vega and theta) for each option are calculated as well. The comprehensive nature of the database makes it most suitable for empirical work on option markets.

The data sample includes daily option prices of six stock indexes and twenty-two major individual stocks for six full years: 1997 - 2002⁷. Table 1.2 displays their names and ticker symbols. The 1997-2002 period offers the additional benefit that it can be split roughly in half to obtain a period of steeply rising stock prices (from January 1997 till mid-2000), and a subsequent period of mostly declining stock prices. As a robustness check, results are presented both for the entire period and for the two sub-periods. Table 1.2 presents also the proportion of three maturity groups in the average daily open interest for at- and out-of-the-money options for each name. It is clear that the longer maturities are well represented, even though the short-maturity group (up to two months) has a somewhat higher proportion in total open interest.

To obtain the option prices needed, I fit a model to the available option prices on each day. The choice among the numerous models that can accurately fit the whole set of options on a given day is of secondary importance in this study. I perform below a limited comparison between three candidate models and pick the one, which is slightly more suitable for my purpose. My main consideration is accuracy of the fit, and I avoid any arguments involving the specifics of the modeled price process. So, both a diffusion-based and a pure-jump model are acceptable,

⁷1996 was dropped, since there were much fewer option prices available for this year.

Table 1.2. Option data

The table displays the names and ticker symbols of the options, used in the estimations. The third column shows average implied at-the-money volatility over 1997 - 2002 for each name. Names are later sorted according to implied volatility in forming portfolios of option returns. The last three columns show the proportion of three maturity groups in the average daily open interest for at- and out-of-the-money options over 1997 - 2002 for each name.

Company name	Ticker	Average implied vol.	Maturity		
			< 2 m.	2-7 m.	> 7m.
Amgen	AMGN	0.43	0.37	0.50	0.14
American Express	AXP	0.37	0.37	0.48	0.14
AOL	AOL	0.56	0.35	0.49	0.16
Boeing	BA	0.34	0.24	0.52	0.24
Bank Index	BKX	0.30	0.70	0.29	0.01
Citibank	C	0.36	0.34	0.49	0.17
Cisco Systems	CSCO	0.52	0.38	0.49	0.13
Pharmaceutical Index	DRG	0.26	0.78	0.22	0.00
General Electric	GE	0.32	0.33	0.50	0.16
Hewlett-Packard	HWP	0.45	0.38	0.51	0.11
IBM	IBM	0.35	0.39	0.46	0.16
Intel	INTC	0.45	0.35	0.48	0.17
Lehman Brothers	LEH	0.48	0.35	0.53	0.12
Merryll Lynch	MER	0.43	0.35	0.49	0.17
Phillip Morris	MO	0.36	0.28	0.55	0.17
Merck	MRK	0.29	0.31	0.54	0.16
Microsoft	MSFT	0.39	0.30	0.48	0.22
National Semicond.	NSM	0.65	0.38	0.48	0.14
Nextel Communic.	NXTL	0.63	0.32	0.56	0.11
Oracle	ORCL	0.57	0.38	0.50	0.12
Pfizer	PFE	0.34	0.31	0.54	0.15
Russel 2000	RUT	0.24	0.68	0.32	0.01
S&P 500	SPX	0.23	0.35	0.47	0.18
Sun Microsystems	SUNW	0.55	0.37	0.48	0.15
Texas Instruments	TXN	0.53	0.31	0.45	0.24
Wal-mart Stores	WMT	0.35	0.38	0.48	0.14
Gold Index	XAU	0.45	0.57	0.40	0.04
Oil Index	XOI	0.24	0.84	0.16	0.00

both models with jumps in volatility and in the price process can be used, etc. It turns out that models, which are conceptually quite different, perform equally well for my purpose.

I focus on the following three models: A stochastic volatility with jumps (SVJ) model studied by Bates (1996) and Bakshi et al. (1997), a double-jump (DJ) model, developed in Duffie et al. (2000) and a pure-jump model with stochastic arrival rate of the jumps (VGSA), as in Carr et al. (2003). Appendix C presents some details on the three models. A full-scale comparison between the models is not my purpose here (see Bakshi and Cao (2003) for a recent detailed study). I only compare their pricing accuracy. To do this, I estimate the three models on each day in the sample of S&P500 options (1509 days for the six years). I employ all out-of-the-money options with strike to spot ratio down to 65% for puts and up to 135% for calls, and maturity between one month and one year (140 options per day on average). The main tool for estimation is the characteristic function of the risk-neutral return density, which is available in closed form for all three models (see Appendix B). Following Carr and Madan (1998), I obtain call prices for any parameter set, by inverting the generalized characteristic function of the call price, using the Fast Fourier Transform⁸. I obtain put prices by put-call parity. I then search for the set of parameters, which minimizes the sum of squared differences between model prices and actual prices. The estimation results are as follows: For the DJ model - 47 days with average % error (A.P.E.) above 5% and average A.P.E. of 2.28% in the remaining days. For the SVJ model - 55 days and 2.42% resp. For the VGSA model - 72 days and 2.48% resp. It is reasonable to apply some filter, when working with estimated, not actual option prices. The above

⁸This procedure is strictly valid only for European-style options. Using only at- and out-of-the-money options mitigates the bias introduced from applying this procedure to options on individual stocks which are traded American-style.

estimations show that by discarding days when average A.P.E. is above 5%, fewer than 5% of the data for S&P500 options are lost. I apply this cut-off in all future estimations.

As expected, the richest model (DJ) performs best, but the differences are modest. So, when choosing among the models the advantages of using a more parsimonious model should also be considered. As discussed in Bakshi and Cao (2003), the estimation of large option pricing models on individual names is hindered by data limitations. In their sample, the majority of the 100 most actively traded names on the CBOE have on average less than ten out-of-the-money options per day. That's why they need to pool together options from all days in a week to perform estimations. Such an approach is not feasible in this study, since pooling across the days in a week may hinder the construction of returns to options with fixed moneyness. So I employ only the names with largest number of options per day and further discard days with insufficient number of options. On average, for all names in the sample, fewer than 12 out-of-the-money options are available on 2.4% of the days, fewer than 15 - on 6.8% of the days and fewer than 18 - on 11.3% of the days. Obviously, increasing the degrees of freedom would come at the price of giving up an increasing amount of the data. That is why, mostly a data-related consideration leads me to choose VGSA (which has the fewest parameters), while sacrificing some accuracy. An additional benefit of such a choice is gain in computational speed. I discard all days with fewer than 12 out-of-the-money options.

Table 1.3 presents the average number of options, used in the estimations for each name, and the percentage errors achieved. I discard days with errors above 5% (usually not more than 4-5% of all days). The errors in the remaining days are around 3% and often less, which is quite satisfactory. This is often within

Table 1.3. Estimation errors

The table displays, for each name, the average number of options per day, used in the estimations (after discarding days with less than 12 options), the proportion of estimations with average percentage error (A.P.E.) greater than 5% (also discarded), and the average A.P.E.. in the remaining days.

Ticker	Aver. daily options	Days with A.P.E.>5%	Remaining A.P.E.
AMGN	28	0.041	0.028
AXP	30	0.020	0.029
AOL	35	0.066	0.026
BA	26	0.042	0.033
BKX	161	0.015	0.024
C	30	0.044	0.032
CSCO	28	0.019	0.025
DRG	58	0.056	0.035
GE	35	0.040	0.031
HWP	27	0.039	0.030
IBM	42	0.019	0.025
INTC	36	0.009	0.025
LEH	26	0.045	0.027
MER	31	0.017	0.027
MO	33	0.079	0.035
MRK	28	0.050	0.030
MSFT	39	0.017	0.025
NSM	19	0.036	0.028
NXTL	23	0.031	0.025
ORCL	23	0.016	0.026
PFE	34	0.071	0.030
RUT	73	0.040	0.031
SPX	121	0.029	0.025
SUNW	34	0.007	0.024
TXN	31	0.004	0.025
WMT	30	0.080	0.030
XAU	35	0.021	0.033
XOI	53	0.092	0.035

the bid-ask spread, in particular for out-of-the-money options. Armed with the estimated parameters for each day, it is easy to generate model prices and returns at the required strikes and maturities.

1.3.3 Volatility risk factors

Cross-sectional estimations of risk prices involve regressions of excess returns on measures of respective risks. I construct these risk measures in two steps. First, I calculate, on each day in the sample, proxies for the market's best estimate of market volatility, realized over different subsequent periods (from one month to one year). Second, I calculate the daily changes in these volatility factors to obtain the volatility risk measures (or "volatility risk factors").

The market's estimate of volatility, realized over a given future period is taken to be the price of the volatility swap with the respective maturity. When "the market" is defined to be the S&P500 index and the length of the period is one month, this best estimate is precisely the CBOE's Volatility Index (VIX). VIX is currently calculated via a non-parametric procedure⁹ employing all current at- and out-of-the-money short-term options on S&P500. One way to obtain the market's volatility estimates for longer future periods would be to extend this procedure, using options with longer maturities. Alternatively, one can use the estimated model parameters for S&P500 and calculate the standard deviations of the S&P500 risk-neutral distribution at the respective horizons¹⁰. I apply the second alternative, mostly for computational convenience.

I verify that the two approaches produce very similar results. First, the corre-

⁹See e.g. Carr and Madan (2001).

¹⁰The risk-neutral variance is obtained by evaluating at zero the first and second derivatives of the characteristic function of the risk-neutral distribution at different horizons. The exact form of the second derivative is quite lengthy, but is readily given by any package, implementing symbolic calculations. The volatility proxy is then the square root of this risk-neutral variance.

lation between VIX and the one-month risk-neutral standard deviation over 1997 - 2002 is 99.1%. Next, I compare the predictive power of the two volatility time-series for realized S&P500 volatility. Following Christensen and Prabhala (1998), I regress realized daily return volatility over 30-day non-overlapping intervals on the two implied volatilities at the beginning of the respective 30-day intervals. Table 1.4 shows the results for 1997 - 2002 and two sub-periods. The estimates involving the VIX and the risk-neutral standard deviation are almost identical. The similarity is observed both in the entire period and the two sub-periods. This comparison justifies the use of the risk-neutral standard deviation. It also provides an indirect check of the quality of the model-based option prices used in constructing option returns.

I define the volatility risk factors to be the daily changes of the volatility proxies at the respective horizons. I also normalize the volatility risk factors to unit standard deviation, which helps to avoid scaling problems and allows for comparing the prices of different volatility risks.

The calculation of volatility risk factors from option prices is motivated by previous findings that option-implied volatilities at different horizons exhibit quite different behavior, indicating that possibly different risks are embedded in these volatilities: Poterba and Summers (1986) find that the changes in forward short-term implied volatility (which is approximately the difference between short-and long-term volatility) are of the same sign but of consistently smaller absolute value than the changes in current short-term implied volatility. Engle and Mustafa (1992) and Xu and Taylor (1994) find that the volatility of short-term implied volatility is larger and mean-reverts faster than that of long-term implied volatility.

Table 1.4. VIX and risk-neutral standard deviation

Panel A shows regression output for $\ln R_{t+30} = \alpha + \beta \ln VIX_t + \varepsilon_t$. R_{t+30} is the realized daily return volatility of S&P500 over a 30-day period starting at time t . Only non-overlapping intervals are involved. VIX_t is the CBOE's Volatility Index calculated at the beginning of each 30-day interval. t-statistics are in parentheses. The two sub-periods are 1/1/1997 - 6/30/2000 and 7/1/2000 - 12/31/2002. Panel B shows regression output for $\ln R_{t+1} = \alpha + \beta \ln SD_t + \varepsilon_t$. SD_t is risk-neutral standard deviation of S&P500 at 30-day horizon, calculated at the beginning of each 30-day interval.

Panel A. VIX

	1997 - 2002		1997 - 2000		2000 - 2002	
α	-0.52	(-1.17)	-0.62	(-0.99)	-0.52	(-0.76)
β	0.82	(4.94)	0.79	(3.42)	0.81	(3.12)
R^2	0.24		0.21		0.24	

Panel B. Risk-neutral standard-deviation

	1997 - 2002		1997 - 2000		2000 - 2002	
α	-0.54	(-1.23)	-0.62	(-0.97)	-0.57	(-0.86)
β	0.80	(4.92)	0.77	(3.33)	0.78	(3.14)
R^2	0.24		0.21		0.24	

1.4 Estimation of volatility risk prices

This section presents results on the estimation of market volatility risk prices in the cross-section of expected option returns, as in equation (1.2.6). It also presents evidence on the economic significance of the difference between the estimated risk prices. For this purpose I employ returns on calendar spreads.

Table 1.5 shows summary statistics of the excess option returns used in estimations. Panel A shows average excess returns to unhedged options across the twenty-eight names, for each strike level and each maturity. Observe that there is a wide variation in these average returns to be explained. There is also a clear pattern across maturities - returns to puts invariably increase with maturity, while those to calls decrease with maturity. Overall, average returns to calls are significantly positive, while returns to puts are mostly negative and sometimes not significantly different from zero. Panel B shows average returns to delta-hedged options. The variation in these returns is still considerable. The maturity pattern for puts is preserved and is even more significant than for unhedged options. Interestingly, returns to longer-term calls now tend to be higher than to short-term ones, in contrast to the unhedged case.

1.4.1 Estimation results

I first estimate the prices of two volatility risk factors with two-step cross-sectional regressions on all individual time-series of excess returns to unhedged options (total of 840 series). I apply the standard procedure of finding the betas on the risk factors at the first step, then regressing, for each day in the sample, excess returns on betas, and finally averaging the second-step regression coefficients and calculating their standard errors.

All regressions involve the market risk, the one-month volatility risk factor

Table 1.5. Average option excess returns

Panel A shows the average of expected excess returns to unhedged option across all names, in each of the strike and maturity groups. Daily returns are multiplied by 30 (monthly basis). E.g. -0.21 stands for -21% of the option price monthly. Each row refers to one of the five maturity groups. O-T-M columns refer to the most out-of-the-money puts / calls; A-T-M columns refer to at-the-money puts / calls. MID columns refer to puts / calls with intermediate moneyness (as in Table 1). Panel B shows the average of expected excess returns to delta-hedged options in the same moneyness and maturity groups. Daily returns are multiplied by 30 (monthly basis) and are now given in % of the spot price. E.g. -0.38 stands for -0.38% of spot monthly. The bottom part of each panel shows the respective t-statistics (average returns divided by standard deviation square root of the number of names). Averages are given for the entire 1997 - 2002 period.

Panel A. Average returns to unhedged options							Panel B. Average returns to delta-hedged options						
Maturity	Puts			Calls			Maturity	Puts			Calls		
	O-T-M	MID	A-T-M	A-T-M	MID	O-T-M		O-T-M	MID	A-T-M	A-T-M	MID	O-T-M
1 m.	-0.21	-0.16	-0.04	0.26	0.58	0.90	-0.38	-0.37	-0.16	-0.04	0.19	-0.01	
3 m.	-0.05	-0.05	-0.02	0.15	0.28	0.46	-0.09	-0.09	-0.06	0.04	0.09	0.06	
6 m.	-0.01	-0.01	-0.02	0.11	0.17	0.40	0.01	-0.01	-0.01	0.08	0.07	0.07	
9 m.	0.02	-0.00	-0.01	0.10	0.16	0.38	0.07	0.08	0.07	0.13	0.12	0.14	
12 m.	0.15	0.03	0.00	0.10	0.29	0.67	0.11	0.12	0.14	0.19	0.15	0.15	
t-statistics (unhedged options)							t-statistics (delta-hedged options)						
Maturity	O-T-M	MID	A-T-M	A-T-M	MID	O-T-M	O-T-M	MID	A-T-M	A-T-M	MID	O-T-M	
1 m.	-4.29	-3.70	-1.38	5.88	11.43	9.28	-7.83	-5.25	-1.68	-0.48	2.48	-0.15	
3 m.	-2.44	-2.86	-1.61	6.91	10.66	6.99	-1.82	-1.49	-0.83	0.52	1.56	1.15	
6 m.	-0.71	-0.80	-1.93	7.79	8.92	4.67	0.18	-0.08	-0.16	1.47	1.29	1.56	
9 m.	1.20	-0.04	-1.28	8.36	7.42	4.00	1.75	1.65	1.32	2.21	1.91	2.25	
12 m.	2.65	1.90	0.33	8.20	4.39	3.70	2.37	2.16	2.20	2.81	2.17	2.35	

and one of the longer term volatility risk factors. In this set-up the one-month risk factor is proxy for short-term volatility risk. This choice is consistent with MacKinlay and Park (2004) who find that monthly volatility exhibits features typical for short-term (transitory) volatility, while three- to six-month volatility represents well long-term (permanent) volatility.

Panel A in Table 1.6 shows the results of these regressions. Market risk is always positively priced (significant at 5%). One-month volatility risk is always significantly negatively priced, while longer-term volatility risks are always significantly positively priced. These results are to a large extent supported by the cross-sectional regressions involving a single volatility factor. Panel B in Table 1.6 presents the results for such single-volatility regressions and show that only the one-month volatility risk has a negative price (insignificant), while all other volatility risk prices are positive and mostly significant. Table 1.6 also shows the importance of a second volatility factor for the explanatory power of the regressions. The numbers in parentheses show, for each combination of risk factors, the proportion of time-series regressions (first pass) with significant alphas. While practically all regressions with one volatility factor have significant alphas (96% in all cases), this proportion dramatically falls to about 15% when a second volatility factor is included. Table 1.6 also shows the adjusted R^2 in regressing *average* excess returns on betas. For any combination of two volatility factors, R^2 increases by 4-5%. These results strongly indicate that, first, two volatility risks are indeed priced in the option market and second, that these risks are of different nature, as evidenced by the different sign of the risk prices. All previous studies of volatility risk pricing in the option market use relatively short-term options (maturity about one month) and mostly find a negative risk price. So, my finding of a negatively priced short-term volatility risk is consistent with previous empirical studies. However, a

Table 1.6. Volatility risk prices - one vs. two volatility factors
(all unhedged option returns time-series)

The table shows volatility risk prices estimated with two-step cross-sectional regressions on all 840 time-series of unhedged daily option returns for 1997 - 2002. The estimated relations are:

$$R^i = \alpha_i + \beta_i^M MKT + \beta_i^{1m} VOL^{1m} + \beta_i^L VOL^L + \varepsilon^i$$

$$E[R^i] = \lambda^M \beta_i^M + \lambda^{1m} \beta_i^{1m} + \lambda^L \beta_i^L + \gamma^i$$

At the second step regressions are run separately for each day and the estimates are then averaged. MKT denotes daily returns on S&P 500, VOL^{1m} denotes daily changes in one-month S&P 500 volatility and VOL^L denotes daily changes in one of the 3, 6, 9 or 12-month volatilities (i.e risk-neutral standard deviations). The λ -s are estimated risk prices for each of the risk factors. Shanken corrected t-statistics are shown for each risk price estimate. In parenthesis is the proportion of alphas in the first-pass regression, estimated to be significant at 5%. In square brackets is the adj. R^2 in regressing average returns on betas. The two panels show results for two and one volatility factors resp.

	Panel A. Two volatility factors		Panel B. One volatility factor		
	Risk price λ	t-stat.	Risk price λ	t-stat.	
MKT	0.07	2.41	MKT	0.07	2.44
VOL 1m	-0.09	-1.99	VOL 1m	-0.03	-0.62
VOL 3m	0.09	1.83		(0.96)	[0.34]
	(0.17)	[0.39]			
MKT	0.06	2.17	MKT	0.07	2.46
VOL 1m	-0.12	-2.61	VOL 3m	0.05	1.04
VOL 6m	0.17	3.20		(0.96)	[0.35]
	(0.16)	[0.40]			
MKT	0.06	1.99	MKT	0.07	2.34
VOL 1m	-0.14	-2.88	VOL 6m	0.12	2.28
VOL 9m	0.24	3.85		(0.96)	[0.36]
	(0.15)	[0.42]			
MKT	0.05	1.90	MKT	0.07	2.22
VOL 1m	-0.14	-2.90	VOL 9m	0.16	2.84
VOL 12m	0.29	4.22		(0.96)	[0.36]
	(0.16)	[0.42]			
			MKT	0.06	2.14
			VOL 12m	0.21	3.26
				(0.95)	[0.38]

positively priced long-term volatility risk has not been identified before. Table 6 also demonstrates the ability of the three-factor model to capture the variance in option excess returns.

One possible concern with the results in Table 1.6 relates to the correlations between the market and the volatility factors. This correlation is well known to be negative and typically high in magnitude and is discussed in Section 1.2.1. For the five volatility risk factors used in the estimations the correlation is -0.60% or less. Besides, the volatility risk factors are highly positively correlated (50-90% in the sample). To eliminate the possible effect of these correlations, I run the regressions also with orthogonal volatility risk factors - I use the component of the one-month volatility risk which is orthogonal to market returns and the component of the longer-term volatility risk which is orthogonal to both the market and the one-month orthogonal volatility risk. Table 1.7 compares the results involving the original (raw) volatility risks with these "orthogonal" volatility risks.

When orthogonal volatility risk factors are used (Panel B), the results are qualitatively the same. The significance of the negative price of the one-month factor is sometimes lower. Note that the price of risk in the orthogonal long-term factor is much higher in magnitude compared to the raw factor. The estimations with volatility risk factors, orthogonal to market returns also indicate that specific volatility risks are priced in the option market. Investors recognize volatility risks beyond those, due to the negative correlation between changes in volatility and market returns.

The above test can be subject to several concerns. First, a more precise evaluation of the explanatory power of the model can be performed using the joint distribution of the errors. However, given the large amount of time series involved, this task is computationally quite demanding, as it involves the computation and

Table 1.7. Volatility risk prices - raw vs. orthogonal volatility risks
(all unhedged option returns time-series)

The table shows volatility risk prices estimated with two-step cross-sectional regressions on all 840 time-series of unhedged daily option returns for 1997 - 2002. The estimated relations are:

$$R^i = \alpha_i + \beta_i^M MKT + \beta_i^{1m} VOL^{1m} + \beta_i^L VOL^L + \varepsilon^i$$

$$E[R^i] = \lambda^M \beta_i^M + \lambda^{1m} \beta_i^{1m} + \lambda^L \beta_i^L + \gamma^i$$

At the second step regressions are run separately for each day and the estimates are then averaged. MKT denotes daily returns on S&P 500, VOL^{1m} denotes daily changes in one-month S&P 500 volatility and VOL^L denotes daily changes in one of the 3, 6, 9 or 12-month volatilities (volatilities here are risk-neutral standard deviations). The λ -s are estimated risk prices for each of the three risk factors. Shanken corrected t-statistics are shown for each risk price estimate. Panel A shows regressions with the volatility risk factors VOL^{1m} and VOL^L . Regressions in panel B use the component of VOL^{1m} orthogonal to MKT , and the component of VOL^L orthogonal to each of the other two factors.

	Panel A. Raw vol.		Panel B. Orthogonal vol.	
	Risk price λ	t-stat.	Risk price λ	t-stat.
MKT	0.07	2.41	0.07	2.42
VOL 1m	-0.09	-1.99	-0.06	-0.95
VOL 3m	0.09	1.83	0.35	3.74
MKT	0.06	2.17	0.06	2.18
VOL 1m	-0.12	-2.61	-0.10	-1.84
VOL 6m	0.17	3.20	0.39	4.57
MKT	0.06	1.99	0.06	2.00
VOL 1m	-0.14	-2.88	-0.13	-2.02
VOL 9m	0.24	3.85	0.42	4.88
MKT	0.05	1.90	0.05	1.92
VOL 1m	-0.14	-2.90	-0.13	-2.09
VOL 12m	0.29	4.22	0.43	4.92

inversion of the covariance matrix of the residuals in the time-series regressions. Second, the standard errors of the estimated parameters are not corrected for heteroskedasticity. This may be an issue when dealing with option data and volatility risks, given the persistence in volatility. Third, unhedged options mix the exposure to the risk in the underlying asset and to volatility risk. It can be argued that investors, seeking exposure to volatility risk will hedge away the risk in the underlyer. So, the price of volatility risk may be better reflected in returns to hedged option. Fourth, the different names are given equal weight, even though their relative importance is highly unequal - for example options on the S&P500 amount to almost half the value of all options in the sample.

To address these concerns I apply a second test where I consider delta-hedged instead of unhedged options, apply GMM for the estimation and reduce the number of asset-return series by forming option portfolios. Delta-hedging allows for more precise exposure to volatility risk. GMM handles the heteroskedasticity of errors and allows to test for all errors being jointly equal to zero. The portfolios allow for an efficient implementation of GMM and account to an extent for the relative importance of different options. Forming portfolios addresses one deficiency of the returns data as well. Because of insufficient out-of-the-money options on certain days and because sometimes the error of estimation has been too high, there are missing observations for certain days for each returns time-series. Since the omissions are relatively few and they come at different days for different names, having several names in a portfolio leaves no missing data in the aggregated returns series.

To form the portfolios I sort the names in the sample according to their average implied volatility (see Table 1.2). Each portfolio belongs to one of the five maturity groups and one of five volatility quintiles (a total of twenty-five portfolios). The different strikes for each name are weighted by the average open interest for the

closest available strikes in the data. For example, I find the proportion of the closest to at-the-money puts on each date and assign the average of these proportions across all days to be the weight of at-the-money puts. I proceed in the same way with the other moneyness levels, both for puts and for calls. The different names within a portfolio are weighted by the average option value for the name, where the average is taken again across all days in the six-year period.

Table 1.8 shows the average excess returns on the portfolios (a total of twenty five) over 1997 - 2002. As in Table 1.3 the numbers are in percent and on a monthly basis. The columns show portfolios arranged from the lowest to the highest average implied volatility of the components. Compared to Table 1.5, the weighted returns (i.e. the portfolio returns) tend to be much lower. There are still portfolios with positive returns, but fewer and with smaller absolute returns. Obviously the larger and less volatile names (in particular the indexes) tend to have lower returns. What is preserved however is the maturity effect - returns to longer-maturity options tend to be higher.

Tables 1.9 and 1.10 contain the main result of this paper. Table 1.9 shows volatility risk prices from estimations with two volatility risk factor. GMM estimations with ten Newey-West lags are reported¹¹. Results are reported both for the entire six-year period and separately for 1997 to mid-2000 and from mid-2000 to 2002.

The one-month volatility factor is always included; the longer-term factors are included both in their raw form, and only with their component orthogonal to the one-month volatility risk factor. Using the orthogonal component does not change the remaining estimates. In all cases the market risk is not priced. This can be expected given delta-hedging. The price of one-month volatility risk is negative and

¹¹Five and twenty lags were also used, producing very similar results which are not reported for brevity.

Table 1.8. Average excess returns on delta-hedged option portfolios

Five portfolios are formed at each maturity by sorting names according to average implied volatility (see Table 1). Volatility quintiles are numbered from 1 (lowest volatility) to 5 (highest volatility). Average daily excess returns for 1997 - 2002 are multiplied by 30 (monthly basis) and given in %; e.g. -0.29 stands for -0.29% of spot monthly.

Maturity	1 Low vol.	2	3	4	5 High vol.
1 m	-0.29	-0.04	-0.03	0.00	0.02
3 m	-0.17	-0.03	-0.03	0.00	0.06
6 m	-0.10	-0.03	0.00	0.02	0.07
9 m	-0.07	-0.02	0.02	0.06	0.05
12 m	-0.03	-0.02	0.04	0.07	0.07

Table 1.9. Volatility risk prices - two volatility factors
(twenty five portfolios)

The table shows volatility risk prices estimated with GMM on twenty five portfolios of delta-hedged options. The moment conditions are:

$$g = \begin{bmatrix} E(R - \alpha - \beta^M MKT - \beta^{1m} VOL^{1m} - \beta^L VOL^L) \\ E \left[(R - \alpha - \beta^M MKT - \beta^{1m} VOL^{1m} - \beta^L VOL^L) MKT \right] \\ E \left[(R - \alpha - \beta^M MKT - \beta^{1m} VOL^{1m} - \beta^L VOL^L) VOL^{1m} \right] \\ E \left[(R - \alpha - \beta^M MKT - \beta^{1m} VOL^{1m} - \beta^L VOL^L) VOL^L \right] \\ E(R - \lambda^M \beta^M - \lambda^{1m} \beta^{1m} - \lambda^L \beta^L) \end{bmatrix} = 0$$

MKT denotes daily returns on S&P 500, VOL^{1m} denotes daily changes in one-month volatility, VOL^L denotes daily changes in one of 3, 6 or 12 month volatility (volatilities here are risk-neutral standard deviations). The λ -s are estimated risk prices for each of the three risk factors. z-statistics are distributed standard normal. Tilded factors (e.g. \widetilde{VOL} 3m.) are the components of the respective raw factors, orthogonal to VOL^{1m} . p-values for the chi-squared test for pricing errors jointly equal to zero are in parenthesis. In square brackets is the adj. R^2 in regressing average returns on betas.

	1997 - 2002		1997 - 2000		2000 - 2002	
	Risk price λ	z-stat.	Risk price λ	z-stat	Risk price λ	z-stat
MKT	0.00	0.06	0.03	0.47	-0.01	-0.07
VOL 1m.	-0.27	-3.77	-0.33	-3.51	-0.11	-1.44
VOL 3m.	0.01	0.16	-0.02	-0.25	0.07	0.78
\widetilde{VOL} 3m.	0.42	2.11	0.46	2.36	0.22	1.31
	(0.42)	[0.69]	(0.72)	[0.68]	(0.85)	[0.20]
MKT	0.02	0.32	0.08	1.22	-0.01	-0.15
VOL 1m.	-0.31	-4.07	-0.46	-3.67	-0.12	-1.60
VOL 6m.	0.12	1.57	0.19	1.40	0.09	1.39
\widetilde{VOL} 6m.	0.46	3.15	0.68	2.79	0.26	1.96
	(0.85)	[0.86]	(0.97)	[0.91]	(0.90)	[0.23]
MKT	0.00	0.00	0.03	0.47	-0.02	-0.26
VOL 1m.	-0.29	-4.37	-0.43	-3.83	-0.12	-1.73
VOL 9m.	0.10	1.81	1.15	1.53	0.08	1.53
\widetilde{VOL} 9m.	0.35	3.59	0.47	2.91	0.22	2.30
	(0.90)	[0.86]	(0.88)	[0.92]	(0.92)	[0.19]
MKT	-0.01	-0.21	0.01	0.12	-0.03	-0.34
VOL 1m.	-0.28	-4.43	-0.40	-3.88	-0.12	-1.81
VOL 12m.	0.11	2.01	0.16	1.79	0.08	1.46
\widetilde{VOL} 12m.	0.30	3.53	0.39	2.92	0.19	2.40
	(0.85)	[0.85]	(0.80)	[0.90]	(0.92)	[0.16]

Table 10. Volatility risk prices - one volatility factor
(twenty five portfolios)

The table shows volatility risk prices estimated with GMM on twenty five portfolios of delta-hedged options. The moment conditions are:

$$g = \begin{bmatrix} E(R - \alpha - \beta^M MKT - \beta^V VOL) \\ E \left[(R - \alpha - \beta^M MKT - \beta^V VOL) MKT \right] \\ E \left[(R - \alpha - \beta^M MKT - \beta^V VOL) VOL \right] \\ E(R - \lambda^M \beta^M - \lambda^V \beta^V) \end{bmatrix} = 0$$

MKT denotes daily returns on S&P 500, *VOL* denotes daily changes in volatility (volatilities are risk-neutral standard deviations at 1, 3, 6, 9 or 12-month horizons). The λ -s are estimated risk prices for each of the two risk factors. z-statistics are distributed standard normal. p-values for the chi-squared test for the pricing errors jointly equal to zero are in parenthesis. In square brackets is the adj. R^2 in regressing average returns on betas. The sub-periods are 1/1/1997 - 6/30/2000 and 7/1/2000 - 12/31/2002

	1997 - 2002		1997 - 2000		2000 - 2002	
	Risk price λ	z-stat.	Risk price λ	z-stat.	Risk price λ	z-stat.
MKT	-0.04	-0.72	-0.03	-0.39	-0.03	-0.43
VOL 1m.	-0.21 (0.24)	-4.31 [0.63]	-0.26 (0.30)	-4.09 [0.65]	-0.09 (0.88)	-1.54 [0.10]
MKT	-0.06	-0.95	-0.06	-0.70	-0.01	-0.12
VOL 3m.	-0.20 (0.11)	-3.87 [0.52]	-0.26 (0.05)	-3.84 [0.63]	-0.03 (0.79)	-0.53 [0.01]
MKT	-0.02	-0.35	0.01	0.16	0.01	0.11
VOL 6m.	-0.12 (0.00)	-2.23 [0.42]	-0.13 (0.05)	-1.87 [0.57]	0.01 (0.63)	0.19 [0.01]
MKT	0.01	0.24	0.06	0.95	0.01	0.10
VOL 9m.	-0.06 (0.00)	-1.29 [0.45]	-0.05 (0.02)	-0.81 [0.60]	0.01 (0.67)	0.16 [0.00]
MKT	0.02	0.51	0.07	1.24	0.00	0.06
VOL 12m.	-0.03 (0.00)	-0.75 [0.47]	-0.01 (0.02)	-0.25 [0.61]	0.00 (0.78)	0.02 [0.00]

highly significant, except for the second sub-period. However, all other volatility risk prices are all positive. The prices of the raw factors are marginally significant (z -statistics about 1.5 - 1.6). Note that the significance is higher for the longer maturity raw factors. However, the risk prices of the orthogonal components are all significant (with two exceptions in the second sub-period). The p -values in all cases are very high (typically 80% or more). The adjusted R^2 in regressing average returns on betas is typically high (0.80 or more), except for the second sub-period. We have thus strong indication that two volatility factors explain most of the variation in expected option returns in the data sample.

Table 1.10 shows volatility risk prices from estimations with one volatility risk factors, which markedly contrast with the one-factor case. The price of market risk is again insignificant. The one-month and three-month volatility risks for the whole period are significantly negatively priced. The prices of longer-term volatility risks are all negative, but not significant. Note that both the significance levels and the absolute magnitude of the volatility risk prices steadily decrease as maturity increases. The same pattern is exactly repeated in the first sub-period. The second sub-period presents mostly insignificant estimates. The table also shows p -values for the chi-squared test for all pricing errors being jointly zero. The entire period and the first sub-period have high p -values for the estimation with the one-month factor (24% and 30% resp.). For longer maturities the p -values decrease sharply. The second sub-period is again different, showing very high p -values (above 60%) for all maturities. The relation between average returns and betas is now weaker (adjusted R^2 about 0.60 or less, and even negligible in the second sub-period).

The results in Tables 1.9 and 1.10 clearly show that long-term volatilities contain two separate risk components with different prices. Including both short- and long-term volatility in the regressions helps disentangle these two risk components.

The second risk, additional to short-term volatility risk is positively priced. The magnitude of the price of this second risk is typically equal or higher than that of the short-term risk price¹².

What is the actual maturity of the short-term volatility risk? So far, the one-month volatility factor is assumed to be the short-term one. However, previous studies have considered various frequencies, sometimes much shorter than a month. It is then possible that the true short-term factor is of much lower maturity, and the one-month factor also mixes two risks. I address this issue by assuming that the absolute value of market returns is proxy for very short-term (one-day) volatility. I run regressions with absolute market returns and each of the former volatility risk factors. If the true maturity of short-term risk is well below one month, this will be reflected in significant estimates of the price of one-day volatility risk. In this case the one-day volatility risk can also be expected to help disentangle the risks in the longer term volatilities. The high explanatory power of two volatility factors for option returns (Table 1.9) should also be preserved. On the other hand, if the true maturity of short-term risk is well above one day, the regressions results should resemble those of the one-volatility factor case (Table 1.10).

Table 1.11 presents the results of regressions involving absolute market returns. These regressions do not support the hypothesis of a very short maturity of the short term-risk. The estimated risk prices for one- to twelve-month volatilities and their significance levels are very close to the one-factor case (Table 1.10). The explanatory power of the regressions is also similar to that in the one-factor case. The absolute value of market returns indeed has a negative and often significant price, but given the other estimates this is not likely to reflect a separate very-short

¹²The tests reported in Tables 8 and 9 were also run with the components of short- and long-term volatilities, orthogonal to market returns. The results were very close, which can be expected, given the insignificant estimates for market risk price.

Table 1.11. Volatility risk prices - absolute market returns and one volatility factor
(twenty five portfolios)

The table shows volatility risk prices estimated with GMM on twenty five portfolios of delta-hedged options. The moment conditions are analogous to those in Table 9. MKT denotes daily returns on S&P 500, $|MKT|$ is the absolute value of MKT , VOL denotes daily changes in one of 1, 3, 6, 9 or 12 month volatility. The λ -s are estimated risk prices for each of the three risk factors. z-statistics are distributed standard normal. Tilded factors (e.g. \widetilde{VOL} 1m.) are the components of the respective raw factors, orthogonal to $|MKT|$. p-values for the chi-square test for pricing errors jointly equal to zero are in parenthesis. In square brackets is the adj. R^2 in regressing average returns on betas. The two sub-periods are as in Table 1.9.

	1997 - 2002		1997 - 2000		2000 - 2002	
	Risk price λ	z-stat.	Risk price λ	z-stat.	Risk price λ	z-stat.
MKT	-0.08	-1.30	-0.03	-0.48	-0.05	-0.74
$ MKT $	-0.16	-1.78	-0.03	-0.25	-0.16	-1.72
VOL 1m.	-0.20	-3.80	-0.26	-4.40	-0.04	-0.68
\widetilde{VOL} 1m.	-0.32	-3.27	-0.38	-3.71	-0.07	-0.56
	(0.36)	[0.74]	(0.37)	[0.71]	(0.88)	[0.21]
MKT	-0.12	-1.68	-0.04	-0.53	-0.05	-0.58
$ MKT $	-0.19	-2.25	0.04	0.39	-0.18	-1.94
VOL 3m.	-0.19	-3.16	-0.25	-4.27	0.01	0.12
\widetilde{VOL} 3m.	-0.39	-2.84	-0.39	-3.35	-0.03	-0.18
	(0.31)	[0.69]	(0.07)	[0.66]	(0.83)	[0.21]
MKT	-0.09	-1.36	0.11	1.83	-0.04	-0.51
$ MKT $	-0.20	-2.38	0.22	1.68	-0.17	-1.90
VOL 6m.	-0.13	-2.07	-0.09	-1.45	0.03	0.47
\widetilde{VOL} 6m.	-0.26	-2.00	-0.04	-0.43	0.00	0.00
	(0.10)	[0.66]	(0.01)	[0.60]	(0.82)	[0.23]
MKT	-0.04	-0.74	0.16	2.80	-0.04	-0.54
$ MKT $	-0.15	-1.98	0.28	1.99	-0.17	-1.89
VOL 9m.	-0.07	-1.28	-0.01	-0.20	0.03	0.55
\widetilde{VOL} 9m.	-0.11	-1.28	0.07	0.89	0.01	0.04
	(0.04)	[0.65]	(0.01)	[0.61]	(0.85)	[0.23]
MKT	-0.02	-0.35	0.17	2.94	-0.04	-0.58
$ MKT $	-0.13	-1.71	0.28	2.01	-0.17	-1.90
VOL 12m.	-0.04	-0.76	0.02	0.35	0.03	0.51
\widetilde{VOL} 12m.	-0.05	-0.75	0.08	1.26	0.01	0.05
	(0.03)	[0.66]	(0.01)	[0.61]	(0.86)	[0.21]

term source of volatility risk.

In summary, I find that two market volatility risks are significantly priced in the sample of excess option returns. The model with two volatility risk factors has little pricing error. While the short-term factor, embedded in close to one-month implied volatility is negatively priced, the long-term factor which can be extracted from longer term-volatilities is positively priced. Next I consider returns to calendar spreads to investigate whether the difference between the prices of short- and long-term volatility risks is also economically significant.

1.4.2 Evidence from calendar spreads

A long calendar spread is a combination of a short position in an option with short maturity and a long position in an option on the same name, of the same type and with the same strike, but of a longer maturity. These spreads are similar to options, in the sense that the possible loss is limited to the amount of the initial net outlay. The results from cross-sectional regressions on unhedged options reported in Table 1.6 showed a positively priced market risk and long-term volatility risk and a negatively priced short-term risk. Expected returns to calendar spreads then have two components - one reflecting the market risks in the two options in the spread and another related to the two volatility risks. The sign and magnitude of the first component should differ across types of options (calls or puts) and moneyness (see e.g. Coval and Shumway (2001)). Results in the previous section imply that the sign of the second component should be unambiguously positive (the position is short the short-term risk and long the long-term one). I do not derive here a formal relation between the two components, but verify that for all moneyness ranges and for both option types (put and calls) the expected returns to calendar spreads are positive. This demonstrates that the component related

to volatility risk is always positive and dominates the market risk component.

Since the results so far only concern market volatility risks, I should strictly focus here only on calendar spreads written on the market. However, given the high explanatory power of market volatility risk factors for option returns (Table 1.9), it can also be expected that calendar spreads written on individual names reflect the difference in the pricing of market volatility risks. So, I consider spreads written on all names in the sample as well.

I use all options in the data set, which allow to calculate the gain of a position in calendar spread. For each name I record, at the beginning of each month all couples of options of the same type and strike and with different maturities, for which prices are available at the end of the month as well. In each spread I use a short-term option of the first available maturity above 50 days. Possible liquidity problems when reversing the position at the end of the month are thus avoided. In this way I replicate a strategy, which trades only twice every month - on opening and closing the spread position. While transaction costs are still present, such trades are definitely feasible. I calculate returns on spreads where the long term option is of the second or third available maturities above 50 days. The results are very similar and I only report results for the second maturity.

Table 1.12 shows average returns to calendar spreads written on the market (S&P500) and on all names in the sample. Separately are shown average returns for different ranges of moneyness for puts and calls, and for different periods. On average, spreads on puts gain an impressive 20% monthly, while those on calls - about 12%. Average returns on individual names are slightly lower than those on the market alone. Spread returns in different moneyness ranges vary considerably, but are all positive. Table 1.12 also show that the Sharpe ratios of calendar spreads are typically about 30-40%, going as high as 100% in one case. Transaction costs

Table 1.12. Average returns to calendar spreads

Panel A shows average returns to calendar spreads formed from short options with maturity at least 50 days and long options with the next available maturity and of the same type and strike. The positions are held for non-overlapping 30-day periods. The strikes of A-T-M options are within $\pm 5\%$ of the spot at the beginning of each 30-day period. O-T-M (I-T-M) options are at least 5% out-of-the-money (in-the-money) at the beginning of each 30-day period. Average spread returns for S&P500 alone and for all 28 names in the sample are shown. The sub-periods are 1/1/1997 - 6/30/2000 and 7/1/2000 - 12/31/2002. Panel B shows Sharpe ratios for calendar spreads in the same moneyness groups and periods.

Panel A. Average one-month returns to calendar spreads

	1997 - 2002		1997 - 2000		2000 - 2002	
	S&P500	All names	S&P500	All names	S&P500	All names
All puts	0.20	0.20	0.16	0.21	0.26	0.17
All calls	0.13	0.12	0.13	0.12	0.14	0.10
A-T-M puts	0.15	0.09	0.20	0.10	0.06	0.06
A-T-M calls	0.05	0.05	0.09	0.07	-0.03	0.03
O-T-M puts	0.16	0.15	0.02	0.13	0.33	0.17
O-T-M calls	0.27	0.20	0.37	0.25	0.15	0.14
I-T-M puts	0.56	0.38	0.70	0.46	0.41	0.26
I-T-M calls	0.13	0.09	0.03	0.07	0.27	0.11

Panel B. Sharpe ratios for calendar spreads

	1997 - 2002		1997 - 2000		2000 - 2002	
	S&P500	All names	S&P500	All names	S&P500	All names
All puts	0.41	0.34	0.37	0.36	0.46	0.37
All calls	0.25	0.31	0.26	0.35	0.24	0.29
A-T-M puts	0.49	0.29	0.64	0.32	0.21	0.24
A-T-M calls	0.20	0.24	0.44	0.33	-0.12	0.12
O-T-M puts	0.31	0.30	0.05	0.25	0.53	0.43
O-T-M calls	0.31	0.43	0.41	0.57	0.19	0.31
I-T-M puts	0.89	0.50	1.15	0.63	0.68	0.43
I-T-M calls	0.31	0.27	0.13	0.22	0.51	0.38

would reduce these numbers, but still the differential pricing of volatility risks is very pronounced. Spread returns thus show that the different prices of short- and long-term volatility risks are not only statistically significant, but economically significant as well.

1.5 Conclusion

A number of volatility-related financial products have been introduced in recent years. Derivatives on realized variance and volatility have been actively traded over the counter. In 2004 the CBOE Futures Exchange introduced futures on the VIX and on the realized three-month variance of the S&P500 index. The practitioners' interest in volatility products has been paralleled by academic research of volatility risk, mostly focused on the risk embedded in a single volatility factor.

This paper complements previous studies of volatility risk by presenting evidence that two implied-volatility risks are priced in a cross-section of expected option returns. I find that the risk in short-term volatility is significantly negatively priced, while another source of risk, orthogonal to the short-term one and embedded in longer-term volatility is significantly positively priced. I show further that the difference in the pricing of short- and long-term volatility risks is also economically significant: I examine returns on long calendar spreads and find that, on average, spreads gain up to 20% monthly.

The estimations for the two sub-periods reveal considerable differences in the parameters, indicating that an extension to time-varying betas and risk prices is justified. The robustness of the findings in this paper to the choice of an option data set and an option-pricing model in constructing option returns can also be examined.

The differential pricing of volatility risks has implication for the modeling of investors' utility. Previous research has found evidence for utility functions over wealth which have both concave and convex sections (Jackwerth (2000), Carr et al. (2002)). It is interesting to explore their results by employing utility functions with more than one arguments and possibly different volatility risks.

Another implication of the findings in this paper relates to the use of options in risk management. It has been argued that firms sometimes face risks which are bundled together in a single asset or liability (e.g. Schrand and Unal (1998)). In this case they can use derivatives to allocate their total risk exposure among multiple sources of risk. This paper suggests that derivatives themselves reflect multiple risks. How do firms chose among derivatives incorporating multiple risks is still to be studied.

2 MaxVar for processes with jumps

2.1 Introduction

The Basle Capital Accord was amended in 1996 to include capital charge for market risk. The Amendment gave the banks the option to use their own internal models for measuring market risk in calculating the capital charge. Among other quantitative requirements, the Amendment required the banks to multiply their VaR estimate by a factor of at least three in calculating the charge¹³. Many market participants expressed the view that the multiplication factors are too high and will possibly undermine the internal models approach.

In "Overview of the Amendment" (BIS (1996)) the Basle Committee recognized the controversy, but still defended the multipliers by arguing that they accounts for potential weaknesses in the modelling process. Among the weaknesses mentioned explicitly in the document were the following:

- the distributions of asset returns often display fatter tails than the normal distribution;
- VaR estimates are typically based on end-of-period positions and generally do not take account of intra-period trading risk¹⁴.

This paper addresses the following question: How much of the multipliers can be explained, first, by non-normality of returns distributions and, second, by the risk

¹³More precisely, the banks should compute VaR using a horizon of 10 days and a 99% confidence interval. At least one year of historical data should be used, updated at least once every quarter. The capital charge is the product of 1.) the higher of the previous days VaR and the average VaR over the preceding 60 days and 2.) a multiplicative factor not smaller than three. This multiplicative factor can be increased to up to four if backtesting reveals that the banks internal model underpredicts losses too often.

¹⁴Other acknowledged weaknesses were that volatilities and correlations can change abruptly, thus rendering the past unreliable approximation to the future, that models cannot adequately capture event risk arising from exceptional market circumstances, and that many models rely on simplifying assumptions, particularly in the case of complex instruments such as options.

of loss *within* the trading period (or interim risk of loss)? To answer this question the paper considers several Levy processes for the underlying assets, all involving jumps. The interim risk of loss is calculated using first passage probabilities for these processes.

In particular, the paper considers diffusions with one- and two-sided jumps and one- and two-sided pure-jump processes. The one-sided pure-jump process employed is the Finite-moment log-stable process (Carr and Wu (2003)). The left tail of the distribution of asset returns in this case declines as a power law and can potentially account for the largest multipliers over normal VaR.

The paper has two main findings: First, Levy models can account for multipliers of 1.05 to 1.5. Second, when the interim risk of loss is also taken into account, the multipliers can increase further - 1.5 to 2.1. Typically, the multipliers for longer periods and for lower loss quantiles are slightly higher. Multipliers higher than 2.1 are not observed during the 5-year period 1998 - 2002 for any of the models employed and for any of the underlying time-series.

While the multipliers obtained in this paper remain well below the factors of three to four, stipulated in the Amendment, the results are still not conclusive. First, the VaR estimates for Levy processes with interim risk are still violated in some cases more often than the respective quantile levels. Second, stochastic changes in volatility are not taken into account. Third, the calculations are based on time-series of daily returns for several major indexes and are thus only illustrative - typical trading books may exhibit profit and loss patterns which greatly differ from index returns. In view of these concerns, the paper contributes mostly to our understanding of the importance of employing Levy processes with jumps and of interim risk consideration in VaR estimations, and should not be construed as an evaluation of whether the Basle multipliers are set at appropriate levels.

Previous studies have demonstrated the ability of VaR estimates based on Levy models to predict more accurately trading losses (e.g. Eberlein et al. (2003) and references therein). The relation to the multipliers, however has not been considered explicitly. VaR with interim risk has also been studied previously. Kritzman and Rich (2002) introduce "Continuous VaR" as a measure of interim risk and show that over long horizons (up to 10 years) hedging based on this measure improves dramatically the performance of portfolio returns, compared to standard VaR. Bodoukh, Richardson, Stanton and Whitelaw (2004) denote the new risk measure as "MaxVar" and calculate ratios between this measure and standard VaR (analogous to the multipliers calculated in this paper)¹⁵. They show that ratios of up to 1.75 can be obtained for certain model parameters and confidence levels, whereby the ratios are increasing in the drift of returns and the length of horizon and are decreasing in the volatility of returns. Both studies employ Brownian motion as the model for asset returns, which is a special case since the first-passage probability in this case is well-known in closed form. The Brownian motion case is also special in that the ratio of MaxVaR to VaR is mainly driven by the drift of returns (Bodoukh et al. (2004)). If we suppose that the drift is zero, which is a reasonable assumption over the short 10-day regulatory period, then the ratio does not depend on the length of horizon or volatility, but only on the VaR confidence level. (An easy application of the reflection principle for Brownian motion with no drift shows, for example, that the ratio is 1.19 for 5% VaR's and 1.11 for 1% VaR's). The contribution of this paper to the study of Continuous VaR / MaxVaR is in applying the concept to a number of models for asset returns involving jumps. For these, more realistic models, significantly higher multipliers are obtained.

¹⁵Interim risk in VaR estimations has been considered earlier in Bodoukh et al. (1995), Stulz (1996), among others.

The rest of the paper is organized as follows: Section 2.2 presents the Levy models and the numerical procedures for calculating first-passage probabilities. Section 2.3 presents the empirical results and Section 2.4 concludes.

2.2 Models and first-passage probabilities

This section presents the models of asset returns considered in the paper. The characteristic functions of log-returns for all models are given in closed form - these functions are used in estimating the model parameters on time-series of returns. The section presents also the numerical procedures for calculating first-passage probability for each model.

2.2.1 CMYD

The CMYD process combines standard Brownian motion and negative jumps. It belongs to the class of spectrally negative processes. The name stands for "CMY plus Diffusion", whereby the "CMY" part comes from the parameters describing the jump component. This process is closely related to the CGMY process studied in Carr, Geman, Madan and Yor (2002). The CMYD process is given by:

$$X_t = \sigma W_t - Z_t \tag{2.2.1}$$

where W_t is a standard Brownian motion, σ is volatility and the jump component Z_t has a Levy measure

$$k(x) = C \frac{\exp(-Mx)}{x^{1+Y}} \quad \text{for } x > 0. \tag{2.2.2}$$

Note that Z_t has only positive jumps, and so X_t has only negative jumps. The CMYD characteristic function is:

$$\begin{aligned}\varphi_{X_t}(u) &\equiv E [e^{iuX_t}] = \exp \left(t \int_0^\infty (e^{-iux} - 1) C \frac{\exp(-Mx)}{x^{1+Y}} dx - \frac{\sigma^2 u^2 t}{2} \right) \\ &= \exp \left\{ t C \Gamma(-Y) [(M - iu)^Y - M^Y] - \frac{\sigma^2 u^2 t}{2} \right\}\end{aligned}\quad (2.2.3)$$

When the uncertainty in asset returns is described by the CMYD process X_t , the asset price dynamics is given by:

$$S_t = S_0 \frac{\exp(\mu t + X_t)}{E [\exp(X_t)]}\quad (2.2.4)$$

The mean rate of return for the stock (under the statistical measure) is μ and the denominator ensures that $E [S_t] = S_0 \exp(\mu t)$. The characteristic function of the log price is:

$$\begin{aligned}E [e^{iu \log(S_t)}] &= \exp \left\{ iu (\log(S_t) + t(\mu - C \Gamma(-Y) [(M + 1)^Y - M^Y])) - \frac{\sigma^2}{2} \right\} \\ &\times \exp \left\{ t C \Gamma(-Y) [(M - iu)^Y - M^Y] - \frac{\sigma^2 u^2 t}{2} \right\}\end{aligned}\quad (2.2.5)$$

Having the characteristic function in closed form allows for efficient estimation of the parameters of the model using FFT (see Carr and Madan (1999)).

Models of stock returns with one-sided jumps have been considered previously, for example in Heston (1993), Carr and Wu (2003). We employ the CMYD model for two main reasons - first, it allows to evaluate the relative performance of models with one- and two-sided jumps in the context of measuring down-jump risk and, second, it offers a significant computational advantage - a technique developed recently by Rogers (2000) provides an efficient method to calculate first-passage

probabilities for Levy processes with one-sided negative jumps¹⁶.

Denote by $f(t, x)$ the probability that a process X_t with only down-sided jumps and starting at zero does not reach the level $x < 0$ before time t . Rogers (2000) suggests the following procedure for calculating $f(t, x)$. First, the double Laplace transform of $f(t, x)$ is shown to be:

$$\tilde{f}(\lambda, z) \equiv \int_0^\infty \int_0^\infty e^{-\lambda t - zx} f(t, x) dt dx = \frac{\beta^*(\lambda) - z}{(\lambda - \psi(z))\beta^*(\lambda)z} \quad (2.2.6)$$

where $\psi(z)$ is the characteristic exponent of X_t :

$$E [\exp(zX_t)] = \exp(t\psi(z)) \quad (2.2.7)$$

and $\beta^*(\lambda)$ is its inverse, i.e. it is the solution of :

$$\psi(\beta(\lambda)) = \lambda \quad (2.2.8)$$

Then the first-passage probability $f(t, x)$ can be found by standard Fourier inversion, but for the difficulty in evaluating $\beta^*(\lambda)$ during the inversion. This difficulty can be avoided by suitably transforming the contour of integration. The transformation is given by $g \equiv \psi \circ \psi_0^{-1}$, where

$$\psi_0^{-1}(z) = \frac{\sqrt{b^2 + 2\sigma^2 z} - b}{\sigma^2} \quad (2.2.9)$$

Here b and σ^2 are the mean rate and variance of the diffusion component of X_t .

¹⁶See Khanna and Madan (2003) for an application of the technique to option pricing.

Now the inversion formula is:

$$f(t, x) = \int_{\Gamma_1} \frac{d\xi}{2\pi i (\sigma^2 \psi_0^{-1}(\xi) + b)} \int_{\Gamma_2} \frac{dz}{2\pi i} \psi'(\psi_0^{-1}(\xi)) \exp(tg(\xi) + xz) \tilde{f}(g(\xi), z) \quad (2.2.10)$$

There is no longer a problem in evaluating the integrand and techniques for two-dimensional Laplace inversion (e.g. Choudhury, Lucantoni, Whitt (1994)) can be efficiently applied to obtain f . In the particular case of the CMYD model :

$$\psi(z) = bz + C\Gamma(-Y) [(M + z)^Y - M^Y] + \frac{\sigma^2 z^2}{2}$$

$$b = \mu - C\Gamma(-Y) [(M + 1)^Y - M^Y] - \frac{\sigma^2}{2}$$

2.2.2 Double exponential jump-diffusion

The double-exponential jump-diffusion process (DEJD) process is studied in Kou and Wang (2002) and differs from CMYD in its structure of jumps. The DEJD process is given by:

$$X_t = \sigma W_t + \sum_{i=1}^{N_t} Y_i, \quad (2.2.11)$$

where σ and W_t are as before, and N_t is a Poisson process. N_t models the arrival of jumps, has intensity λ and is independent of the Brownian motion. Y_i are random variables defining the sizes of the jumps. Y_i have a common two-sided exponential distribution:

$$f_Y(y) = p\eta_1 \exp(-\eta_1 y) 1_{\{y \geq 0\}} + (1 - p)\eta_2 \exp(\eta_2 y) 1_{\{y < 0\}}. \quad (2.2.12)$$

p is the probability of an up-jump given that a jump occurs, and $1/\eta_1$ and $1/\eta_2$ are the means of the exponential distributions for the up- and down-jumps respectively. Since the diffusion and jump components of X_t are independent, the characteristic function of X_t is easily given by:

$$\begin{aligned}\varphi_{X_t}(u) &= E[e^{iuX_t}] \\ &= \exp\left\{-\frac{u^2\sigma^2 t}{2} + \lambda t \left(\frac{p\eta_1}{\eta_1 - iu} + \frac{(1-p)\eta_2}{\eta_2 + iu} - 1\right)\right\},\end{aligned}\quad (2.2.13)$$

and the characteristic exponent is:

$$G(u) = -\frac{u^2\sigma^2}{2} + \lambda \left(\frac{p\eta_1}{\eta_1 - iu} + \frac{(1-p)\eta_2}{\eta_2 + iu} - 1\right).\quad (2.2.14)$$

As before, we model the dynamics of asset prices as

$$S_t = S_0 \frac{\exp(\mu t + X_t)}{E[\exp(X_t)]}.\quad (2.2.15)$$

The characteristic function of the log price is:

$$\begin{aligned}&E[e^{iu \log(S_t)}] \\ &= \exp\left\{iu \left(\log(S_t) + t\mu + t \left[-\frac{u^2\sigma^2}{2} + \lambda \left(\frac{p\eta_1}{\eta_1 - u} + \frac{(1-p)\eta_2}{\eta_2 + u} - 1\right)\right]\right)\right\} \\ &\times \exp\left\{-\frac{u^2\sigma^2 t}{2} + \lambda t \left(\frac{p\eta_1}{\eta_1 - iu} + \frac{(1-p)\eta_2}{\eta_2 + iu} - 1\right)\right\}\end{aligned}\quad (2.2.16)$$

DEJD has a rare advantage when first passage probabilities are concerned - the Laplace transform of the first passage time to a fixed level can be calculated analytically. Such explicit solutions are possible for processes whose jumps are of the phase type (Assmussen et al. (2002)). The double exponential jumps turn out to

be the simplest of this type, hence the calculations are relatively easy to perform. The following theorem (Kou and Wang (2002)) gives the Laplace transform:

Theorem 2 *Let X_t be the DEJD process and τ_b be the first passage time to level $b < 0$ for X_t started from 0. For any $\alpha \in (0, \infty)$ let $\beta_{3,\alpha}$ and $\beta_{4,\alpha}$ be the only negative roots of the Cramer - Lindberg equation $G(\beta) = \alpha$, such that $0 < -\beta_{3,\alpha} < \eta_2 < -\beta_{4,\alpha}$. Then*

$$\begin{aligned} & E[\exp(-\alpha\tau_b)] \\ &= \frac{\eta_2 + \beta_{3,\alpha}}{\eta_2} \frac{\beta_{4,\alpha}}{\beta_{4,\alpha} - \beta_{3,\alpha}} \exp(b\beta_{3,\alpha}) - \frac{\beta_{4,\alpha} - \eta_2}{\eta_2} \frac{\beta_{3,\alpha}}{\beta_{4,\alpha} - \beta_{3,\alpha}} \exp(b\beta_{4,\alpha}) \end{aligned} \quad (2.2.17)$$

For numerical Laplace inversion it is convenient to use the following:

$$\int_0^\infty e^{-\alpha t} P(\tau_b < t) dt = \frac{1}{\alpha} \int_0^\infty e^{-\alpha t} dP(\tau_b < t) = \frac{1}{\alpha} E[\exp(-\alpha\tau_b)] \quad (2.2.18)$$

Kou and Wang (2002) suggest further the use of the Gaver-Stehfest algorithm for inversion.

Given the Laplace transform \hat{f} of f , we have the approximation $f(t) = \lim_{n \rightarrow \infty} \widetilde{f}_n(t)$ where

$$\widetilde{f}_n(t) = \frac{\ln(2)}{t} \frac{(2n)!}{n!(n-1)!} \sum_{k=0}^n (-1)^k \binom{n}{k} \hat{f}\left((n+k) \frac{\ln(2)}{t}\right) \quad (2.2.19)$$

To speed up the convergence Richardson extrapolation can be used, whereby $f(t)$ is approximated by $f_n^*(t)$ for large t , where

$$f_n^*(t) = \sum_{k=1}^n w(k, n) \widetilde{f}_k(t) \quad \text{for } w(k, n) = (-1)^{n-k} \frac{k^n}{k!(n-k)!} \quad (2.2.20)$$

2.2.3 CGMY

The CGMY process was introduced in Carr et al. (2002). The CGMY process is a pure-jump Levy process with the following Levy density:

$$k_X(x) = \begin{cases} C \frac{\exp(-G|x|)}{|x|^{1+Y}} & \text{for } x < 0 \\ C \frac{\exp(-Mx)}{x^{1+Y}} & \text{for } x > 0 \end{cases} \quad (2.2.21)$$

This is obviously a process with two-sided jumps. The C parameter can be considered as a measure of the arrival rate of jumps - both positive and negative. M and G control the rate of exponential decay of the probability of up- and down- jumps of different sizes. The Y parameter allows for a fine distinction between different classes of processes: depending on the value of Y , the process may or may not be completely monotone, and may exhibit finite or infinite activity. The CGMY characteristic function is:

$$\varphi_{X_t}(u) \equiv E [e^{iuX_t}] = \exp \{tC\Gamma(-Y) [(M - iu)^Y - M^Y + (G + iu)^Y - G^Y]\} \quad (2.2.22)$$

The characteristic function of the log price, when the uncertainty is given by the CGMY process is:

$$\begin{aligned} & E [e^{-iu \log(S_t)}] \\ &= \exp \{iu (\log(S_t) + t(\mu - C\Gamma(-Y) [(M + 1)^Y - M^Y + (G - 1)^Y - G^Y])\} \\ &\times \exp \{tC\Gamma(-Y) [(M - iu)^Y - M^Y + (G + iu)^Y - G^Y]\} \end{aligned} \quad (2.2.23)$$

Carr et al. (2002) show that CGMY performs well in calibrating both time-series of stock returns and option prices.

Since no closed-form expression is available for the first-passage probability

of the CGMY process, we apply here a numerical procedure. We note the close relation between the calculation of first-passage probability and the valuation of certain exotic options. In particular, the value of an option paying \$1 if the price of the underlying assets hits certain level within a given time period is equal to the first-passage probability of the price to this level (under the risk-neutral measure). Such an option has been considered within the pure-jump context for example in Hirta (1999). The value function for this option is solution to a partial integro-differential equation (PIDE). Hirta provides an efficient numerical solution to this equation and we follow closely his approach. (See also Madan and Hirta (2003)). Note however, that Value-at-Risk calculations are typically performed under the statistical measure, so we also employ this measure throughout.

Denote by

$$G(S, t, T) = E[1_{(S(u) \leq H}, \quad 0 < u < T] \quad (2.2.24)$$

the conditional expectation at time $t < T$ of the stock price hitting the level $H < S_0$ within the time interval $[0, T]$. If at any $u < t$ we have $S(u) < H$, then $G(S, t, T) = 1$. Otherwise $0 < G(S, t, T) < 1$. By construction $G(S, t, T)$ is a martingale, since it is conditional expectation of a terminal random variable. Then $G(S, t, T)$ is solution to the following PIDE¹⁷:

$$G_t + \mu S G_S + \int_{-\infty}^{+\infty} (G(S e^x, t, T) - G(S, t, T) - S G_S(S, t, T)(e^x - 1)) k(x) dx = 0 \quad (2.2.25)$$

with boundary conditions

$$G(S, t, T) = 1 \quad \text{if } S(u) < H \quad \text{for some } u < t$$

¹⁷See Essay 3 in this thesis for a detailed derivation of the PIDE in a more general, two-dimensional case.

$$G(S, T, T) = 0 \text{ if } S(u) \geq H \text{ for all } u \quad (2.2.26)$$

After changing variables to $s = \ln(S)$ and $g(s, t, T) = G(S, t, T)$ the PIDE (2.2.25) becomes:

$$g_t + \mu g_s + \int_{-\infty}^{+\infty} (g(s+x, t, T) - g(s, t, T) - g_s(s, t, T)(e^x - 1))k(x)dx = 0 \quad (2.2.27)$$

Using the Levy measure for the CGMY process and writing $g(s)$ for $g(s, t, T)$:

$$\begin{aligned} 0 = g_t + \mu g_s + C \int_{-\infty}^0 (g(s+x) - g(s) - g_s(s)(e^x - 1)) \frac{e^{-G|x|}}{|x|^{1+Y}} dx \quad (2.2.28) \\ + C \int_0^{\infty} (g(s+x) - g(s) - g_s(e^x - 1)) \frac{e^{-Mx}}{x^{1+Y}} dx \end{aligned}$$

Note that, because of the boundary condition, for any down-jump which brings the price below the level H , g takes the value of 1. For any given s , these are the jumps such that $x \leq \ln(H) - s$.

Now (2.2.28) is solved using finite differences on the mesh $[s_{\min}, s_{\max}] \times [0, T]$

$$s_i = s_{\min} + ih, \quad i = 0, 1, \dots, N,$$

$$t_j = j\Delta, \quad j = 0, 1, \dots, M,$$

$$h = (s_{\max} - s_{\min})/N \quad \text{and} \quad \Delta = T/M$$

where s denotes the log of the stock price, s_{\min} is the log of the first-passage level $H < S_0$, h is the step in the log-price direction and Δ is the step in the time direction.

The sum of the two integrals in (2.2.28) is approximated by:

$$I \approx I_1 + I_2 + I_3 + I_4 + I_5 + I_6 \quad (2.2.29)$$

where

$$\begin{aligned}
I_1 &= C \int_{-\infty}^{\ln(H)-s} [1 - g(s) + g_s(s)] \frac{e^{-G|x|}}{|x|^{1+Y}} dx, \\
I_2 &= C \int_{\ln(H)-s}^{-\Delta} [g(s+x) - g(s) + g_s(s)] \frac{e^{-G|x|}}{|x|^{1+Y}} dx, \\
I_3 &= -C \int_{-\infty}^{-\Delta} g_s(s) \frac{e^x e^{-G|x|}}{|x|^{1+Y}} dx, \\
I_4 &= C \int_{-\Delta}^0 \left[g(s) + g_s(s)x + \frac{1}{2}g_{ss}x^2 - g(s) - g_s(s)(1+x + \frac{1}{2}x^2 - 1) \right] \frac{e^{-G|x|}}{|x|^{1+Y}} dx, \\
I_5 &= C \int_0^{\Delta} \left[g(s) + g_s(s)x + \frac{1}{2}g_{ss}x^2 - g(s) - g_s(s)(1+x + \frac{1}{2}x^2 - 1) \right] \frac{e^{-Mx}}{|x|^{1+Y}} dx, \\
I_6 &= \int_{\Delta}^U [g(s+x) - g(s) - g_s(s)(e^x - 1)] \frac{e^{-Mx}}{x^{1+Y}} dx.
\end{aligned}$$

The discretization for each of the above six integrals is next written down separately. Here s takes the discrete values s_j for $j = 2 : N - 1$, $s_{j+1} - s_j = \Delta$ and we g_j denotes $g(s_j, t_i, T)$.

Case 1: $0 < Y < 1$.

The following integral is essential in what follows:

$$\begin{aligned}
& C \int_{\beta}^{\alpha} \frac{e^{-\gamma x}}{x^{1+Y}} dx \\
&= -\frac{C\gamma^Y}{Y} \left[\frac{e^{-\gamma\alpha}}{(\gamma\alpha)^Y} - \frac{e^{-\gamma\beta}}{(\gamma\beta)^Y} + \Gamma(1-Y) [\Gamma_{inc}(\gamma\alpha, 1-Y) - \Gamma_{inc}(\gamma\beta, 1-Y)] \right],
\end{aligned} \tag{2.2.30}$$

where $0 \leq \beta < \alpha$, and Γ_{inc} is the lower incomplete gamma function. Using (2.2.30),

for any t_i and any $j = 2 : N - 1$:

$$I_1 = -\left[1 - g_j + \frac{g_{j+1} - g_j}{\Delta}\right] \frac{CG^Y}{Y} \\ \times \left[-\frac{e^{-G\Delta(j-1)}}{(G\Delta(j-1))^Y} + \Gamma(1-Y) [1 - \Gamma_{inc}(G\Delta(j-1), 1-Y)] \right],$$

$$I_3 = \frac{g_{j+1} - g_j}{\Delta} C \int_{-\infty}^{-\Delta} \frac{e^{-(G+1)|x|}}{|x|^{1+Y}} dx = \frac{g_{j+1} - g_j}{\Delta} C \int_{\Delta}^{\infty} \frac{e^{-(G+1)x}}{x^{1+Y}} dx \\ = -\frac{g_{j+1} - g_j}{\Delta} \frac{C(G+1)^Y}{Y} \\ \times \left[-\frac{e^{-(G+1)\Delta}}{[(G+1)\Delta]^Y} + \Gamma(1-Y) [1 - \Gamma_{inc}((G+1)\Delta, 1-Y)] \right],$$

$$I_4 = \frac{C}{2} \left[\frac{g_{j+1} - 2g_j + g_{j-1}}{\Delta^2} - \frac{g_{j+1} - g_j}{\Delta} \right] \int_{-\Delta}^0 \frac{x^2 e^{-G|x|}}{|x|^{1+Y}} dx \\ = \frac{CG^{Y-2}}{2} \left[\frac{g_{j+1} - 2g_j + g_{j-1}}{\Delta^2} - \frac{g_{j+1} - g_j}{\Delta} \right] \Gamma(2-Y) \Gamma_{inc}(G\Delta, 2-Y),$$

$$I_5 = \frac{C}{2} \left[\frac{g_{j+1} - 2g_j + g_{j-1}}{\Delta^2} - \frac{g_{j+1} - g_j}{\Delta} \right] \int_0^{\Delta} \frac{x^2 e^{-Mx}}{x^{1+Y}} dx \\ = \frac{CM^{Y-2}}{2} \left[\frac{g_{j+1} - 2g_j + g_{j-1}}{\Delta^2} - \frac{g_{j+1} - g_j}{\Delta} \right] \Gamma(2-Y) \Gamma_{inc}(M\Delta, 2-Y).$$

For the remaining two integrals I_2 and I_6 note that there is an I_2 term only for $j > 2$, since for $j = 2$ down-jumps are small in absolute value and are given only

by I_4 . Equivalently, there is an I_6 term only for $j < N - 1$.

$$\begin{aligned}
I_2 &\approx C \sum_{k=1}^{j-2} \int_{s_k - s_j}^{s_{k+1} - s_j} \left[g_{k+1} + \frac{g_{k+1} - g_k}{\Delta} (x - s_{k+1} + s_j) - g_j + \frac{g_{j+1} - g_j}{\Delta} \right] \frac{e^{-G|x|}}{|x|^{1+Y}} dx \\
&= C \sum_{k=1}^{j-2} \left[g_{k+1} + \frac{g_{k+1} - g_k}{\Delta} (s_j - s_{k+1}) - g_j + \frac{g_{j+1} - g_j}{\Delta} \right] \int_{s_k - s_j}^{s_{k+1} - s_j} \frac{e^{-G|x|}}{|x|^{1+Y}} dx \\
&+ C \sum_{k=1}^{j-2} \frac{g_{k+1} - g_k}{\Delta} \int_{s_k - s_j}^{s_{k+1} - s_j} \frac{e^{-G|x|}}{|x|^Y} dx. \tag{2.2.31}
\end{aligned}$$

The integral in the first summand in (2.2.31) is equal to:

$$\begin{aligned}
I_2^1 &= \frac{G^Y}{Y} \left[\frac{e^{-G\Delta(j-k)}}{(G\Delta(j-k))^Y} - \frac{e^{-G\Delta(j-k-1)}}{(G\Delta(j-k-1))^Y} \right] \\
&+ \frac{G^Y}{Y} \Gamma(1-Y) [\Gamma_{inc}(G\Delta(j-k), 1-Y) - \Gamma_{inc}(G\Delta(j-k-1), 1-Y)]
\end{aligned}$$

and the integral in the second summand in (2.2.31) is equal to:

$$I_2^2 = G^{1-Y} [\Gamma_{inc}(G\Delta(j-k), 1-Y) - \Gamma_{inc}(G\Delta(j-k-1), 1-Y)].$$

In a similar way:

$$\begin{aligned}
I_6 &\approx C \sum_{k=j+1}^{N-1} \int_{s_k - s_j}^{s_{k+1} - s_j} \left[g_k + \frac{g_{k+1} - g_k}{\Delta} (x - s_k + s_j) - g_j - \frac{g_{j+1} - g_j}{\Delta} (e^x - 1) \right] \frac{e^{-Mx}}{x^{1+Y}} dx \\
&= C \sum_{k=j+1}^{N-1} \left[g_k + \frac{g_{k+1} - g_k}{\Delta} (s_j - s_k) - g_j + \frac{g_{j+1} - g_j}{\Delta} \right] \int_{s_k - s_j}^{s_{k+1} - s_j} \frac{e^{-Mx}}{x^{1+Y}} dx \\
&- C \sum_{k=j+1}^{N-1} \frac{g_{j+1} - g_j}{\Delta} \int_{s_k - s_j}^{s_{k+1} - s_j} \frac{e^{-(M-1)x}}{x^{1+Y}} dx \\
&+ C \sum_{k=j+1}^{N-1} \frac{g_{k+1} - g_k}{\Delta} \int_{s_k - s_j}^{s_{k+1} - s_j} \frac{e^{-Mx}}{x^Y} dx \tag{2.2.32}
\end{aligned}$$

The integral in the first summand in (2.2.32) is equal to:

$$I_6^1 = \frac{e^{-M\Delta(k+1-j)}}{(M\Delta(k+1-j))^Y} - \frac{e^{-M\Delta(k-j)}}{(M\Delta(k-j))^Y} \\ + \Gamma(1-Y) [\Gamma_{inc}(M\Delta(k+1-j), 1-Y) - \Gamma_{inc}(M\Delta(k-j), 1-Y)],$$

the integral in the second summand in (2.2.32) is equal to:

$$I_6^2 = \frac{e^{-(M-1)\Delta(k+1-j)}}{((M-1)\Delta(k+1-j))^Y} - \frac{e^{-(M-1)\Delta(k-j)}}{((M-1)\Delta(k-j))^Y} \\ + \Gamma(1-Y) [\Gamma_{inc}((M-1)\Delta(k+1-j), 1-Y) - \Gamma_{inc}((M-1)\Delta(k-j), 1-Y)],$$

and the integral in the third summand is equal to:

$$I_6^3 = M^{1-Y} [\Gamma_{inc}(M\Delta(k-j), 1-Y) - \Gamma_{inc}(M\Delta(k+1-j), 1-Y)].$$

Case 2: $1 < Y < 2$.

The only difference in this case is that, for $0 \leq \beta < \alpha$, the integral of the Levy density in (2.2.30) becomes:

$$C \int_{\beta}^{\alpha} \frac{e^{-\gamma x}}{x^{1+Y}} dx \\ = \frac{C\gamma^Y}{Y(1-Y)} \left[\frac{e^{-\gamma\alpha}}{(\gamma\alpha)^Y} (1-Y+\gamma\alpha) - \frac{e^{-\gamma\beta}}{(\gamma\beta)^Y} (1-Y+\gamma\beta) \right] \\ + \frac{C\gamma^Y}{Y(1-Y)} \Gamma(2-Y) [\Gamma_{inc}(\gamma\alpha, 2-Y) - \Gamma_{inc}(\gamma\beta, 2-Y)]$$

All other approximations and calculations follow exactly the previous case.

Finally, with this approximation to the summands in (2.2.28) an explicit scheme is applied for solving the PIDE. Assuming that at time t_i the values of g are known,

the values of g at an earlier time t_{i-1} are found by using standard finite difference approximations for g_t and g_s and solving a tri-diagonal linear system.

2.2.4 Finite-Moment Log-Stable (FMLS)

The FMLS process, introduced in Carr and Wu (2003) can be considered as a special case of the CGMY process. It only has down-sided jumps and its Levy density is:

$$k_X(x) = \frac{C}{|x|^{1+Y}} \quad \text{for } x < 0 \quad (2.2.33)$$

For this model $1 < Y < 2$. Since the numerator in (2.2.33) lacks the exponential damping factor of the CGMY of the Levy density, the left tail is fatter. Actually, the left tail is so fat, that FMLS is only made a feasible model for asset returns by disallowing *any* up-jumps. Only this restriction ensures that all moments of asset returns are finite (unlike the stable processes with two-sided jumps). The FMLS characteristic function is:

$$\varphi_{X_t}(u) \equiv E [e^{iuX_t}] = \exp \left[t\sigma^Y u^Y \left(1 - i \times \text{sign}(u) \tan\left(\frac{\pi Y}{2}\right) \right) \right] \quad (2.2.34)$$

where

$$\sigma = \left[\frac{C\Gamma\left(\frac{Y}{2}\right)\Gamma\left(1 - \frac{Y}{2}\right)}{2\Gamma(1+Y)} \right]^{\frac{1}{Y}} \quad (2.2.35)$$

and $\Gamma(x)$ is the gamma function. The PIDE for the value of a claim in the FMLS case is analogous to (2.2.28), but only has an integral corresponding to down-jumps:

$$0 = g_t + \mu g_s + C \int_{-\infty}^0 (g(s+x) - g(s) - g_s(s)(e^x - 1)) \frac{1}{|x|^{1+Y}} dx. \quad (2.2.36)$$

The integral is approximated by:

$$I \approx I_1 + I_2 + I_3 + I_4 \quad (2.2.37)$$

where

$$I_1 = C \int_{-\infty}^{\ln(H)-s} [1 - g(s) + g_s(s)] \frac{1}{|x|^{1+Y}} dx$$

$$I_2 = C \int_{\ln(H)-s}^{-\Delta} [g(s+x) - g(s) + g_s(s)] \frac{1}{|x|^{1+Y}} dx$$

$$I_3 = -C \int_{-\infty}^{-\Delta} g_s(s) \frac{e^x}{|x|^{1+Y}} dx$$

$$I_4 = C \int_{-\Delta}^0 \left[g(s) + g_s(s)x + \frac{1}{2}g_{ss}x^2 - g(s) - g_s(s)(1+x + \frac{1}{2}x^2 - 1) \right] \frac{1}{|x|^{1+Y}} dx$$

The integral of the Levy measure is:

$$C \int_{-\alpha}^{-\beta} \frac{1}{|x|^{1+Y}} dx = \frac{C}{Y} \left(\frac{1}{\beta^Y} - \frac{1}{\alpha^Y} \right) \quad (2.2.38)$$

where $0 < \beta < \alpha$. Using (2.2.38), for any t_i and any $j = 2 : N - 1$:

$$I_1 = [1 - g_j + \frac{g_{j+1} - g_j}{\Delta}] \frac{C}{Y(\Delta(j-1))^Y}$$

$$\begin{aligned} I_3 &= -\frac{g_{j+1} - g_j}{\Delta} C \int_{-\infty}^{-\Delta} \frac{e^{-|x|}}{|x|^{1+Y}} dx = -\frac{g_{j+1} - g_j}{\Delta} C \int_{\Delta}^{\infty} \frac{e^{-x}}{x^{1+Y}} dx \\ &= \frac{g_{j+1} - g_j}{\Delta} \frac{C}{Y(Y-1)} \left[-\frac{e^{-\Delta}}{\Delta^Y} (1 - Y + \Delta) + \Gamma(2 - Y) [1 - \Gamma_{inc}(\Delta, 2 - Y)] \right] \end{aligned}$$

$$\begin{aligned}
I_4 &= \frac{C}{2} \left[\frac{g_{j+1} - 2g_j + g_{j-1}}{\Delta^2} - \frac{g_{j+1} - g_j}{\Delta} \right] \int_{-\Delta}^0 \frac{x^2}{|x|^{1+Y}} dx \\
&= \frac{C}{2} \left[\frac{g_{j+1} - 2g_j + g_{j-1}}{\Delta^2} - \frac{g_{j+1} - g_j}{\Delta} \right] \frac{\Delta^{2-Y}}{2-Y}
\end{aligned}$$

For $j = 3 : N - 1$:

$$\begin{aligned}
I_2 &\approx C \sum_{k=1}^{j-2} \int_{s_k - s_j}^{s_{k+1} - s_j} \left[g_{k+1} + \frac{g_{k+1} - g_k}{\Delta} (x - s_{k+1} + s_j) - g_j + \frac{g_{j+1} - g_j}{\Delta} \right] \frac{1}{|x|^{1+Y}} dx \\
&= C \sum_{k=1}^{j-2} \left[g_{k+1} + \frac{g_{k+1} - g_k}{\Delta} (s_j - s_{k+1}) - g_j + \frac{g_{j+1} - g_j}{\Delta} \right] \int_{s_k - s_j}^{s_{k+1} - s_j} \frac{1}{|x|^{1+Y}} dx \\
&\quad + C \sum_{k=1}^{j-2} \frac{g_{k+1} - g_k}{\Delta} \int_{s_k - s_j}^{s_{k+1} - s_j} \frac{1}{|x|^Y} dx
\end{aligned}$$

The integral in the first summand is equal to

$$I_2^1 = \frac{1}{Y} \left[\frac{1}{(\Delta(j-k-1))^Y} - \frac{1}{(\Delta(j-k))^Y} \right]$$

and the integral in the second summand in (44) is equal to

$$I_2^2 = \frac{1}{Y-1} \left[\frac{1}{(\Delta(j-k-1))^{Y-1}} - \frac{1}{(\Delta(j-k))^{Y-1}} \right].$$

2.3 Empirical results

This section presents the results of the time-series estimation of the Levy models and the Levy and MaxVaR multipliers obtained for each of the models.

Results are reported for the original four models plus an additional modification of the CGMY model. The modification (named CGMYLIM) has the G and M parameters fixed, providing an essentially two-parameter model. By fixing G and

M at very low values (around one¹⁸, when typical estimates in the full CGMY model are from 50 to 150), this approximates the stable model with two-sided jumps, while still preserving the finite moments. So, CGMYLIM is an interesting two-parameter alternative to FMLS.

The estimations are performed for five international stock indexes on daily returns over the five-year period January 1998 - December 2002. The indexes are FTSE, DAX, NIKKEI, Hang Seng and S&P500. Once every week each model is estimated using the 1000 preceding days (about 4 years of mean-adjusted daily returns). The characteristic function, available in closed form for each model, is the main tool of the estimation. Using the fast Fourier transform the characteristic function is inverted once for each parameter setting to obtain the density at a pre-specified set of values for returns (see Carr et al. (2002)). With the density evaluated at these values, the return series are binned by counting the number of observations at each pre-specified return value, assigning data observations to the closest return value. Then a maximization algorithm searches for the parameter estimates that maximize the likelihood of the binned data.

Table 2.1 shows one measure of the performance of the models in these estimations - the proportion of estimations for each model where the chi-squared goodness-of-fit test cannot reject the model at 5% (1%). As expected, the normal model performs the worst and can be rejected in more than half of the cases for most underlyers. The models with one-sided jumps are only slightly better, which can also be expected given the approximate symmetry of the density of daily returns. The two-sided models are much better and typically can be rejected at most in 1-2% of the cases. It is interesting to note that even with only two parameters,

¹⁸The value of 1.01 for both parameters was chosen to ensure stability of the particular discretization scheme for the PIDE. Other schemes may allow for even lower fixed values of the G and M parameters.

Table 2.1 Goodness-of-fit tests
(time-series estimations)

Proportion of estimations for each model where the chi-squared goodness-of-fit test cannot reject the model at 5% (1%).

	5%	1%	5%	1%
	Normal		FMLS	
FTSE	0.36	0.55	0.56	0.68
DAX	0.37	0.46	0.48	0.58
HANG SENG	0.08	0.16	0.16	0.24
NIKKEI	0.26	0.46	0.41	0.59
S&P500	0.19	0.39	0.47	0.62
	CMYD		DEJD	
FTSE	0.60	0.72	0.99	0.99
DAX	0.49	0.63	0.98	1.00
HANG SENG	0.17	0.26	0.82	0.98
NIKKEI	0.41	0.58	1.00	1.00
S&P500	0.42	0.68	0.91	0.99
	CGMY		CGMYLIM	
FTSE	0.98	1.00	0.94	1.00
DAX	0.97	1.00	0.68	1.00
HANG SENG	0.77	0.97	0.52	0.62
NIKKEI	1.00	1.00	1.00	1.00
S&P500	0.89	1.00	0.89	1.00

CGMYLIM performs as well as the larger models on four out of the five indexes.

Tables 2.2 and 2.3 show the maximum and mean ratios between 5% (resp 1%) VaR's obtained from the Levy models with jumps and standard VaR, obtained from a Normal (Gaussian) model for returns. The ratios (or multiples) are modest (rarely exceeding 1.1) for CMYD, DEJD and CGMY. The two models with heaviest left tails - FMLS and CGMYLIM exhibit slightly higher multiples at the 5% confidence level, but significantly higher multiples (up to 1.5 - 1.6) at the 1% level. The multiples for the two horizons - 10 and 20 days are very similar.

Tables 2.4 and 2.5 show the maximum and mean multiples of MaxVaR for the Levy models over standard VaR. The multiples for CMYD and CGMY are the lowest - typically below 1.3. The remaining three models show multiples in the range of 1.4 - 1.7, and going slightly above 2 in a few cases. No multiple above 2.1 is recorded. The multiples for the two horizons are very similar. Since all estimations are performed on demeaned returns, the effect of the drift as observed in Bodoukh et al. (2004) does not appear in our data.

The above results show clearly that taking jumps and interim risk into account can produce significantly higher VaR values compared to Normal VaR. How well do these higher VaR's perform in predicting future losses? To answer this question we calculate the proportion of non-overlapping 10- and 20-day periods during which the actual maximum loss exceeds the MaxVaR estimated at the beginning of the period. For comparison, the proportion of cases when end-of-period loss exceeds Normal VaR is calculated as well. For a good VaR measure the frequency of excessive losses should correspond to the VaR confidence level. Table 2.6 reports the frequencies p-values of the chi-squared test for these frequencies to be significantly different from the respective confidence levels.

Normal VaR performs poorly in this test - in half the cases the test does not

Table 2.2 Levy VaR multiples over Normal VaR - 10 days

10-day VaRs are calculated once every 10 days (between 1/1/1998 and 12/31/2002). Each VaR calculation uses daily returns over the preceding 1000 days. For each Levy model the table shows the maximum (mean) ratio of the 5% (1%) VaRs to the respective normal VaRs.

	5% VaRs		1% VaRs		5% VaRs		1% VaRs	
	max	mean	max	mean	max	mean	max	mean
	CMYD				FMLS			
FTSE	1.06	1.04	1.12	1.06	1.06	1.04	1.31	1.28
DAX	1.08	1.05	1.17	1.10	1.07	1.05	1.31	1.29
HANG SENG	1.04	0.99	1.09	1.04	1.01	0.94	1.21	1.13
NIKKEI	1.05	1.04	1.11	1.07	1.03	0.98	1.25	1.18
S&P500	1.06	1.04	1.19	1.10	1.03	0.99	1.25	1.20
	DEJD							
FTSE	1.04	1.03	1.08	1.05				
DAX	1.05	1.04	1.11	1.07				
HANG SENG	1.01	0.96	1.02	0.99				
NIKKEI	1.02	1.02	1.05	1.03				
S&P500	1.03	1.02	1.09	1.05				
	CGMY				CGMYLIM			
FTSE	1.04	1.03	1.07	1.05	1.18	1.09	1.56	1.22
DAX	1.06	1.04	1.14	1.07	1.17	1.09	1.52	1.26
HANG SENG	1.01	0.96	1.03	1.00	1.18	1.10	1.61	1.38
NIKKEI	1.02	1.02	1.06	1.04	1.13	1.10	1.40	1.28
S&P500	1.03	1.02	1.13	1.07	1.13	1.08	1.49	1.27

Table 2.3. Levy VaR multiples over Normal VaR - 20 days

20-day VaRs are calculated once every 20 days (between 1/1/1998 and 12/31/2002). Each VaR calculation uses daily returns over the preceding 1000 days. For each Levy model the table shows the maximum (mean) ratio of the 5% (1%) VaRs to the respective normal VaRs.

	5% VaRs		1% VaRs		5% VaRs		1% VaRs	
	max	mean	max	mean	max	mean	max	mean
	CMYD				FMLS			
FTSE	1.06	1.04	1.09	1.05	1.08	1.06	1.33	1.29
DAX	1.08	1.05	1.13	1.08	1.09	1.08	1.32	1.29
HANG SENG	1.05	0.99	1.08	1.02	1.03	0.96	1.21	1.12
NIKKEI	1.05	1.04	1.09	1.06	1.05	1.00	1.26	1.17
S&P500	1.06	1.04	1.14	1.08	1.05	1.01	1.27	1.20
	DEJD							
FTSE	1.04	1.03	1.06	1.04				
DAX	1.05	1.04	1.08	1.05				
HANG SENG	1.02	0.97	1.02	0.98				
NIKKEI	1.03	1.02	1.03	1.03				
S&P500	1.03	1.02	1.06	1.03				
	CGMY				CGMYLIM			
FTSE	1.04	1.03	1.05	1.04	1.25	1.12	1.60	1.24
DAX	1.06	1.04	1.11	1.06	1.24	1.13	1.53	1.28
HANG SENG	1.02	0.97	1.02	0.98	1.26	1.16	1.54	1.34
NIKKEI	1.03	1.02	1.04	1.03	1.19	1.14	1.41	1.29
S&P500	1.03	1.03	1.09	1.05	1.19	1.12	1.52	1.29

Table 2.4 MaxVaR multiples over Normal VaR - 10 days

10-day MaxVaRs are calculated once every 10 days (between 1/1/1998 and 12/31/2002). Each MaxVaR calculation uses daily returns over the preceding 1000 days. For each Levy model the table shows the maximum (mean) ratio of the 5% (1%) MaxVaRs to the respective normal VaRs.

	5% VaRs		1% VaRs		5% VaRs		1% VaRs	
	max	mean	max	mean	max	mean	max	mean
	CMYD				FMLS			
FTSE	1.27	1.23	1.25	1.18	1.76	1.59	1.63	1.46
DAX	1.28	1.25	1.31	1.22	1.75	1.64	1.59	1.50
HANG SENG	1.23	1.17	1.21	1.15	1.51	1.40	1.37	1.27
NIKKEI	1.24	1.23	1.23	1.18	1.54	1.46	1.40	1.34
S&P500	1.26	1.23	1.32	1.23	1.70	1.53	1.57	1.41
	DEJD							
FTSE	1.66	1.43	1.52	1.34				
DAX	1.64	1.52	1.59	1.45				
HANG SENG	1.41	1.26	1.35	1.22				
NIKKEI	1.48	1.35	1.41	1.29				
S&P500	1.53	1.46	1.49	1.40				
	CGMY				CGMYLIM			
FTSE	1.35	1.30	1.29	1.24	1.51	1.36	1.82	1.43
DAX	1.33	1.28	1.33	1.23	1.44	1.34	1.76	1.46
HANG SENG	1.23	1.16	1.17	1.13	1.41	1.32	1.94	1.62
NIKKEI	1.27	1.25	1.23	1.19	1.39	1.34	1.62	1.47
S&P500	1.35	1.29	1.36	1.26	1.50	1.36	1.76	1.49

Table 2.5. MaxVaR multiples over Normal VaR - 20 days

20-day MaxVaRs are calculated once every 20 days (between 1/1/1998 and 12/31/2002). Each MaxVaR calculation uses daily returns over the preceding 1000 days. For each Levy model the table shows the maximum (mean) ratio of the 5% (1%) MaxVaRs to the respective normal VaRs.

	5% VaRs		1% VaRs		5% VaRs		1% VaRs	
	max	mean	max	mean	max	mean	max	mean
	CMYD				FMLS			
FTSE	1.25	1.23	1.19	1.15	2.09	1.88	1.90	1.71
DAX	1.27	1.24	1.22	1.18	2.02	1.87	1.83	1.70
HANG SENG	1.25	1.17	1.18	1.11	1.65	1.57	1.50	1.44
NIKKEI	1.25	1.23	1.18	1.16	1.80	1.71	1.64	1.56
S&P500	1.26	1.23	1.23	1.18	2.01	1.81	1.84	1.65
	DEJD							
FTSE	1.76	1.47	1.58	1.36				
DAX	1.66	1.58	1.57	1.47				
HANG SENG	1.40	1.28	1.33	1.21				
NIKKEI	1.48	1.36	1.40	1.28				
S&P500	1.54	1.47	1.47	1.39				
	CGMY				CGMYLIM			
FTSE	1.29	1.26	1.22	1.18	1.53	1.36	1.85	1.42
DAX	1.30	1.25	1.26	1.19	1.49	1.36	1.78	1.46
HANG SENG	1.22	1.15	1.14	1.09	1.51	1.38	1.98	1.64
NIKKEI	1.24	1.23	1.18	1.16	1.43	1.36	1.64	1.48
S&P500	1.29	1.25	1.27	1.20	1.49	1.36	1.77	1.49

Table 2.6 Frequency of excessive losses

Proportion of 10- and 20-day periods when actual loss exceeds the 5% Normal VaR and MaxVaRs (resp. the 1% VaRs) calculated at the beginning of the period. The periods are non-overlapping, from 1/1/1998 to 12/31/2002. * (**) denote p-values of the chi-squared test for the observed loss to exceed VaR / MaxVaR significant at 1% (5%).

	10 days		20 days		10 days		20 days	
	5%	1%	5%	1%	5%	1%	5%	1%
	Normal				FMLS			
FTSE	0.14**	0.03	0.11	0.05*	0.02	0.02	0.02	0.02
DAX	0.09**	0.08**	0.13**	0.09**	0.05	0.00	0.05	0.03
HANG SENG	0.00	0.00	0.08*	0.03	0.00	0.00	0.02	0.00
NIKKEI	0.03	0.00	0.05	0.00	0.00	0.00	0.00	0.00
S&P500	0.08*	0.03	0.11**	0.03	0.02	0.00	0.02	0.00
	CMYD				DEJD			
FTSE	0.05	0.02	0.02	0.00	0.05	0.02	0.03	0.00
DAX	0.08*	0.03	0.06	0.02	0.08*	0.03	0.08*	0.02
HANG SENG	0.00	0.00	0.00	0.00	0.00	0.00	0.05	0.00
NIKKEI	0.00	0.00	0.00	0.00	0.00	0.00	0.02	0.00
S&P500	0.05	0.00	0.02	0.00	0.03	0.00	0.03	0.00
	CGMY				CGMYLIM			
FTSE	0.05	0.02	0.06	0.02	0.05	0.02	0.06	0.02
DAX	0.08*	0.02	0.09*	0.06*	0.08*	0.00	0.09*	0.02
HANG SENG	0.00	0.00	0.05	0.02	0.00	0.00	0.03	0.00
NIKKEI	0.00	0.00	0.03	0.00	0.00	0.00	0.02	0.00
S&P500	0.05	0.00	0.03	0.02	0.03	0.00	0.03	0.00

reject (at 5 % confidence level) the hypothesis that Normal VaR underpredicts future losses. MaxVar performs much better - typically only in one or two cases out of twenty we see significant underprediction of losses. From this perspective the best model is FMLS, for which all p-values are above 5%.

2.4 Conclusion

It is widely recognized that Normal VaR estimates do not capture properly the magnitude of potential losses in banks' portfolios. The 1996 Amendment to the Basle Capital Accord specifies that these estimates should be multiplied by a factor of three or more, as one improvement on the risk measure. This paper examines the multiples over normal VaR which can be obtained when VaR is estimated using processes with jumps as models of asset returns and when taking the risk of interim losses into account. Typical multiples around 1.5 - 1.7 are found (and as high as 2 in a few cases).

One limitation of the paper is that it only considers time-homogeneous (Levy) models. However, it has been shown that incorporating stochastic volatility in models for asset returns improves significantly VaR measures. Second, banks' trading books typically differ from the stock indexes used in the estimations in this paper. It is left for future research to examine the question whether richer models (with stochastic volatility) and more realistic portfolios (e.g. including derivatives) can produce even higher VaR multiples.

3 The COGARCH model and option pricing

3.1 Introduction

The success of GARCH models in capturing important features of asset returns has provoked interest in GARCH option pricing. Duan (1995) studies the measure change when returns of the underlying asset are modeled by a GARCH process. Garcia and Renault (1998) consider hedging in a GARCH option pricing model and Rithcken and Trevor (1999) construct trinomial trees to price American options under GARCH. Different GARCH option pricing models have been suggested and tested empirically in Duan (1996b), Heston and Nandi (2000), Hardle and Hafner (2000), Christoffersen and Jacobs (2004). Duan, Rithcken and Sun (2004) develop a GARCH option pricing model with jumps.

This paper begins the study of option pricing for a new type of a GARCH model - the COGARCH model of Kluppelberg, Lindner and Maller (2004). Like most stochastic volatility models, COGARCH is a continuous-time model. Yet, like GARCH models, it is driven by a single random process and volatility under COGARCH exhibits the feedback and autoregressive properties of volatility in GARCH models. One distinctive feature of COGARCH is that the driving random process is a Levy process with jumps, as in a number of recent models of asset returns¹⁹. A second distinctive feature of COGARCH is that volatility is a pure-jump process, as in Barndorff-Nielsen and Shephard (2001). COGARCH also provides an alternative to the continuous-time limits of certain GARCH processes, which have been previously derived (e.g. Nelson (1990), Duan (1996a)). In con-

¹⁹For example Madan, Carr and Chang (1998) employ the variance-gamma process as a model of asset returns, Shoutens and Teugels (1998) use the Mexiner process, Barndorff-Nielsen and Shephard (2001) use the Normal inverse Gaussian process and Eberlein (2001) applies the hyperbolic process.

trast to these two-dimensional limits, COGARCH preserves the essential GARCH feature that changes in volatility are only caused by moves in the underlying asset and not by changes in a second, latent random variable²⁰.

In the best case options are priced using a closed-form formula or an approximation. Another widely used approach is based on the inversion of the characteristic function of the log-price (e.g. Carr and Madan (1999)). This approach is computationally very efficient, but requires a closed-form expression for the characteristic function. An alternative approach, which is less efficient but is widely used, for example, in pricing exotics, employs the numerical solution of partial differential / integro-differential equations (PDEs / PIDEs). Since for the COGARCH model there is no closed-form option pricing formula or characteristic function currently available, this paper takes the PIDE approach. We exploit the fact that asset returns and volatility under COGARCH are jointly Markov. This Markovian property allows to derive a PIDE for the value of a claim under COGARCH. We employ the PIDE for pricing European calls and puts, but the approach can be easily extended to pricing American options and exotics.

3.2 The COGARCH process

The COGARCH process, introduced in Kluppelberg et al. (2004) is a continuous-time version of the GARCH(1,1) process. The innovations of the COGARCH process are given by the jumps of a Levy process.

Let $(X_t)_{t \geq 0}$ be a Levy process with jumps $\Delta X_t = X_t - X_{t-}$, Levy measure

²⁰For other problems with continuous-time limits of GARCH processes see Corradi (2000), Wang (2002), Brown, Wang and Zhao (2002). But see also Heston and Nandi (2000) and Jeanthau (2004) for discrete time GARCH models which have diffusion limits driven by a single Brownian motion.

k_X and let $0 < \delta < 1$ and $\lambda \geq 0$. Define a process $(Y_t)_{t \geq 0}$ by:

$$Y_t = -t \log(\delta) - \sum_{0 < s < t} \log \left[1 + \frac{\lambda}{\delta} (\Delta X_s)^2 \right]. \quad (3.2.1)$$

Proposition 3 (Kluppelberg et al. (2004)) $(Y_t)_{t \geq 0}$ is a spectrally negative Levy process of bounded variation with drift $\log(\delta)$, with no Gaussian component and with Levy measure k_Y given by

$$k_Y([0, \infty]) = 0 \quad \text{and} \quad k_Y([-\infty, -y]) = k_X \left\{ z \in R : |z| \geq \sqrt{(e^y - 1) \frac{\lambda}{\delta}} \right\}, \quad \text{for } y > 0.$$

Define further a variance process $(\sigma_t^2)_{t \geq 0}$

$$\sigma_t^2 = \left(\beta \int_0^t \exp(Y_s) ds + \sigma_0^2 \right) \exp(-Y_{t-}) \quad (3.2.2)$$

where $\beta > 0$ and σ_0^2 is a finite random variable, independent of $(X_t)_{t \geq 0}$. Then the COGARCH process $(G_t)_{t \geq 0}$ is given by

$$dG_t = \sigma_t dX_t. \quad (3.2.3)$$

The logarithmic asset returns over a time period r are modeled as $G_t^r = G_{t+r} - G_t$. It follows from (3.2.1)-(3.2.3) that both G and σ jump only when X jumps. The jump sizes of G are $\Delta G_t = \sigma_t \Delta X_t$

Proposition 4 (Kluppelberg et al. (2004)) The process $(\sigma_t^2)_{t \geq 0}$ satisfies the following stochastic differential equation:

$$d\sigma_{t+}^2 = \beta dt + \sigma_t^2 \exp(Y_{t-}) d(\exp(-Y_t)) \quad (3.2.4)$$

and it follows that

$$\sigma_t^2 = \beta t + \log(\delta) \int_0^t \sigma_s^2 ds + \frac{\lambda}{\delta} \sum_{0 < s < t} \sigma_s^2 (\Delta X_s)^2 + \sigma_0^2 \quad (3.2.5)$$

and that

$$\sigma_{t+}^2 - \sigma_t^2 = \frac{\lambda}{\delta} \sigma_t^2 (\Delta X_t)^2. \quad (3.2.6)$$

Proposition 5 (*Kluppelberg et al. (2004)*) *The bivariate process (σ_t^2, G_t) is Markovian.*

In what follows we assume that the Levy process X_t is the Variance-Gamma (VG) process (Madan, Carr and Chang (1998)). VG is obtained by evaluating Brownian motion with drift at a random time, given by a gamma process. Let $b(t; \theta, \sigma) = \theta t + \sigma W(t)$ be a Brownian motion with drift θ and volatility σ and let $\gamma(t; \mu, \nu)$ be a gamma process with mean rate μ and variance ν . The VG process X_t is defined in terms of the Brownian motion and the gamma process as:

$$X_t = b(\gamma(t; 1, \nu); \theta, \sigma).$$

Due to the gamma time change VG is a pure-jump process. Its Levy density $k(x)$ is :

$$k(x) = \begin{cases} \frac{\mu_n^2}{\nu_n} \frac{\exp(-\frac{\mu_n}{\nu_n} |x|)}{|x|} & \text{for } x < 0 \\ \frac{\mu_p^2}{\nu_p} \frac{\exp(-\frac{\mu_p}{\nu_p} x)}{x} & \text{for } x > 0 \end{cases} \quad (3.2.7)$$

where

$$\mu_n = \frac{1}{2} \sqrt{\theta^2 + \frac{2\sigma^2}{\nu}} - \frac{\theta}{2}$$

$$\mu_p = \frac{1}{2} \sqrt{\theta^2 + \frac{2\sigma^2}{\nu}} + \frac{\theta}{2}$$

$$\nu_n = \mu_n^2 \nu \quad \text{and} \quad \nu_p = \mu_p^2 \nu.$$

We further change the variables in the Levy density to:

$$C = \frac{\mu_p^2}{\nu_p} = \frac{\mu_n^2}{\nu_n}, \quad M = -\frac{\mu_p}{\nu_p} \quad \text{and} \quad G = \frac{\mu_n}{\nu_n}$$

and the Levy density becomes

$$k(x) = \begin{cases} C \frac{\exp(-G|x|)}{|x|} & \text{for } x < 0 \\ C \frac{\exp(-Mx)}{x} & \text{for } x > 0 \end{cases}. \quad (3.2.8)$$

3.3 Backward PIDE for European options

Assume that the risk-neutral dynamics of the log price is given by the COGARCH process. Then

$$S_t = S_0 e^{\int_0^t a_u du + \int_0^t \int_{-\infty}^{\infty} \sigma_u x \mu(dx, du)} \quad (3.3.1)$$

where $\mu(dx, du)$ is a random jump measure and $\int_0^t a_u du$ is the convexity correction, guaranteeing that the stock price has the proper risk-neutral expectation. To find a_u we use the fact that the discounted stock price is a positive martingale under the risk neutral measure, and therefore it can be represented as a stochastic exponential of a compensated martingale. This compensated martingale has the form:

$$Z_t = \int_0^t \int_{-\infty}^{\infty} (e^{\sigma_u x} - 1) (\mu(dx, du) - \nu(dx, du)) \quad (3.3.2)$$

where $\nu(dx, du)$ is the compensator for $\mu(dx, du)$ (see e.g. Shiryaev (1999)). For the discounted stock price we obtain:

$$\begin{aligned} e^{-rt}S_t &= S_0 e^{\sum_{s \leq t} \Delta Z_s} \prod_{s \leq t} (1 + \Delta Z_s) \\ &= S_0 e^{-\int_0^t \int_{-\infty}^{\infty} (e^{\sigma_u x} - 1) \nu(dx, du) + \int_0^t \int_{-\infty}^{\infty} \sigma_u x \mu(dx, du)}. \end{aligned} \quad (3.3.3)$$

By comparing (3.3.1) and (3.3.3) we obtain

$$-rt + \int_0^t a_u du = - \int_0^t \int_{-\infty}^{\infty} (e^{\sigma_u x} - 1) \nu(dx, du)$$

and so

$$a_t = r - \int_{-\infty}^{\infty} (e^{\sigma_t x} - 1) \nu(dx). \quad (3.3.4)$$

Let s_t denote the log price at time t and $\bar{\omega}(\sigma_t) = - \int_{-\infty}^{\infty} (e^{\sigma_t x} - 1 - \sigma_t x) \nu(dx)$. From (3.3.1) and (3.3.4) we obtain:

$$\begin{aligned} s_t &= s_0 + rt - \int_0^t \int_{-\infty}^{\infty} (e^{\sigma_u x} - 1) \nu(dx, du) + \int_0^t \int_{-\infty}^{\infty} \sigma_u x \mu(dx, du) \\ &= s_0 + \int_0^t (r + \bar{\omega}(\sigma_u)) du + \int_0^t \int_{-\infty}^{\infty} \sigma_u x (\mu(dx, du) - \nu(dx, du)). \end{aligned} \quad (3.3.5)$$

Denoting $\bar{\omega} = \int_{-\infty}^{\infty} x^2 \nu(dx)$ we rewrite (3.2.5) as:

$$\begin{aligned}
\sigma_t^2 &= \sigma_0^2 + \beta t + \log(\delta) \int_0^t \sigma_u^2 du + \frac{\lambda}{\delta} \int_0^t \int_{-\infty}^{\infty} \sigma_u^2 x^2 \nu(dx, du) \\
&\quad + \frac{\lambda}{\delta} \int_0^t \int_{-\infty}^{\infty} \sigma_u^2 x^2 (\mu(dx, du) - \nu(dx, du)) \\
&= \sigma_0^2 + \beta t + (\log(\delta) + \frac{\lambda \bar{\omega}}{\delta}) \int_0^t \sigma_u^2 du + \frac{\lambda}{\delta} \int_0^t \int_{-\infty}^{\infty} \sigma_u^2 x^2 (\mu(dx, du) - \nu(dx, du)).
\end{aligned} \tag{3.3.6}$$

We can now derive the PIDE for any test function f of the two variables s and σ^2 . For this purpose we use (3.3.5) and (3.3.6), and the fact that the bivariate process (s_t, σ_t^2) is jointly Markov (Proposition 3). For this Markov process the infinitesimal generator $I_{(s, \sigma^2)}$ satisfies the relation

$$f(s_t, \sigma_t^2) = f(s_0, \sigma_0^2) + \int_0^t I_{(s, \sigma^2)} f(s_u, \sigma_u^2) du + \text{martingale}. \tag{3.3.7}$$

To obtain the generator we apply Ito's formula for semi-martingales (see e.g. Shiryaev (1999)) and represent $f(s_t, \sigma_t^2)$ as the sum of drift terms and compensated martingales:

$$\begin{aligned}
f(s_t, \sigma_t^2) &= f(s_0, \sigma_0^2) + \int_0^t (r + \bar{\omega}(\sigma_u)) f_s(s_u, \sigma_u^2) du \\
&\quad + \int_0^t \left[\beta + \left(\log(\delta) + \frac{\lambda \bar{\omega}}{\delta} \right) \sigma_u^2 \right] f_{\sigma^2}(s_u, \sigma_u^2) du \\
&\quad + \int_0^t \int_{-\infty}^{\infty} \left[\begin{aligned} & f(s_u + \sigma_u x, \sigma_u^2 (1 + \frac{\lambda}{\delta} x^2)) - f(s_u, \sigma_u^2) \\ & - f_s(s_u, \sigma_u^2) \sigma_u x - f_{\sigma^2}(s_u, \sigma_u^2) \frac{\lambda}{\delta} \sigma_u^2 x^2 \end{aligned} \right] k(x) dx du \\
&\quad + \text{martingale}.
\end{aligned} \tag{3.3.8}$$

It follows from (3.3.7) and (3.3.8) that:

$$\begin{aligned}
I_{(s,\sigma^2)}f(s_t,\sigma_t^2) &= (r + \bar{\omega}(\sigma_t))f_s(s_t,\sigma_t^2) + \left[\beta + (\log(\delta) + \frac{\lambda}{\delta}\bar{\omega})\sigma_t^2 \right] f_{\sigma^2}(s_t,\sigma_t^2) \\
&+ \int_{-\infty}^{\infty} \left[\begin{array}{l} f(s_t + \sigma_t x, \sigma_t^2(1 + \frac{\lambda}{\delta}x^2)) - f(s_t, \sigma_t^2) \\ -f_s(s_t, \sigma_t^2)\sigma_t x - f_{\sigma^2}(s_t, \sigma_t^2)\frac{\lambda}{\delta}\sigma_t^2 x^2 \end{array} \right] k(x) dx. \quad (3.3.9)
\end{aligned}$$

By taking expectation and differentiating in (3.3.7) with respect to t we also obtain that for any test function:

$$E [f_t(s_t, \sigma_t^2)] = E [I_{(s,\sigma^2)}f(s_t, \sigma_t^2)]$$

which gives the PIDE:

$$f_t(s_t, \sigma_t^2) = I_{(s,\sigma^2)}f(s_t, \sigma_t^2). \quad (3.3.10)$$

We now consider the PIDE for a claim in the specific case where the COGARCH process is driven by VG and so $k(x)$ is given by (3.2.7).

Let $g := g(s_t, \sigma_t^2)$ denote the value of this claim. We further change t to $-\tau$ and for brevity omit the time subscript τ . It follows from (3.3.10) that:

$$\begin{aligned}
rg &= g_\tau + rg_s + (\beta + (\log(\delta) + \frac{\lambda}{\delta}\bar{\omega})\sigma^2)g_{\sigma^2} \\
&+ C \int_{-\infty}^0 \left[g(s + \sigma x, \sigma^2(1 + \frac{\lambda}{\delta}x^2)) - g - (e^{\sigma x} - 1)g_s - \frac{\lambda}{\delta}\sigma^2 x^2 g_{\sigma^2} \right] \frac{e^{-G|x|}}{|x|} dx \\
&+ C \int_0^{\infty} \left[g(s + \sigma x, \sigma^2(1 + \frac{\lambda}{\delta}x^2)) - g - (e^{\sigma x} - 1)g_s - \frac{\lambda}{\delta}\sigma^2 x^2 g_{\sigma^2} \right] \frac{e^{-Mx}}{x} dx.
\end{aligned}$$

We also change x to $\frac{x}{\sigma}$ to obtain:

$$\begin{aligned}
rg &= g_\tau + rg_s + (\beta + (\log(\delta) + \frac{\lambda \overline{\omega}}{\delta})\sigma^2)g_{\sigma^2} \\
&+ C \int_{-\infty}^0 \left[g(s+x, \sigma^2 + \frac{\lambda}{\delta}x^2) - g - (e^x - 1)g_s - \frac{\lambda}{\delta}x^2 g_{\sigma^2} \right] \frac{e^{-\frac{G}{\sigma}|x|}}{|x|} dx \\
&+ C \int_0^{+\infty} \left[g(s+x, \sigma^2 + \frac{\lambda}{\delta}x^2) - g - (e^x - 1)g_s - \frac{\lambda}{\delta}x^2 g_{\sigma^2} \right] \frac{e^{-\frac{M}{\sigma}x}}{x} dx. \quad (3.3.11)
\end{aligned}$$

To solve numerically this PIDE we use finite differences. We break down the integrals in (3.3.11) into four parts:

$$I \approx \int_{-D}^{-\Delta_s} (\cdot) + \int_{-\Delta_s}^0 (\cdot) + \int_0^{\Delta_s} (\cdot) + \int_{\Delta_s}^U (\cdot) = I_1 + I_2 + I_3 + I_4 \quad (3.3.12)$$

where D and U are the lower and upper bounds for the discretization in the log-price dimension. For the integrals I_1 and I_4 we apply the approximation developed in Madan and Hirta (2003). We index the discretization grid by subscripts (k and j) in the log-price direction and by superscripts (n) in the volatility direction. At

the node (j, n) on the grid the approximations are:

$$\begin{aligned}
I_1 &\approx C \sum_{k=1}^{j-2} \int_{s_k-s_j}^{s_{k+1}-s_j} \left[g_{k+1}^n + \frac{g_{k+1}^n - g_k^n}{\Delta_s} (x - s_{k+1} + s_j) \right] \frac{e^{-\frac{G}{\sigma}|x|}}{|x|} dx \\
&+ C \sum_{k=1}^{j-2} \int_{s_k-s_j}^{s_{k+1}-s_j} \left[\frac{g_{k+1}^{n+1} - g_{k+1}^n}{\Delta_{\sigma^2}} \frac{\lambda}{\delta} (x^2 - (s_{k+1} - s_j)^2) \right] \frac{e^{-\frac{G}{\sigma}|x|}}{|x|} dx \\
&- C \sum_{k=1}^{j-2} \int_{s_k-s_j}^{s_{k+1}-s_j} \left[g_j^n + \frac{g_{j+1}^n - g_j^n}{\Delta_s} (e^x - 1) + \frac{g_j^{n+1} - g_j^n}{\Delta_{\sigma^2}} \frac{\lambda}{\delta} x^2 \right] \frac{e^{-\frac{G}{\sigma}|x|}}{|x|} dx \\
&= C \sum_{k=1}^{j-2} \left[A_k^1 \int_{s_k-s_j}^{s_{k+1}-s_j} \frac{e^{-\frac{G}{\sigma}|x|}}{|x|} dx - A_k^2 \int_{s_k-s_j}^{s_{k+1}-s_j} e^{-\frac{G}{\sigma}|x|} dx \right] \\
&+ C \sum_{k=1}^{j-2} \left[A_k^3 \int_{s_k-s_j}^{s_{k+1}-s_j} |x| e^{-\frac{G}{\sigma}|x|} dx - A_k^4 \int_{s_k-s_j}^{s_{k+1}-s_j} \frac{e^{-(\frac{G}{\sigma}+1)|x|}}{|x|} dx \right] \quad (3.3.13)
\end{aligned}$$

where

$$\begin{aligned}
A_k^1 &= \left[g_{k+1}^n + \frac{g_{k+1}^n - g_k^n}{\Delta_s} (s_j - s_{k+1}) - \frac{g_{k+1}^{n+1} - g_{k+1}^n}{\Delta_{\sigma^2}} \frac{\lambda}{\delta} (s_j - s_{k+1})^2 - g_j^n + \frac{g_{j+1}^n - g_j^n}{\Delta_s} \right] \\
A_k^2 &= \frac{g_{k+1}^n - g_k^n}{\Delta_s} \\
A_k^3 &= \frac{\lambda}{\delta \Delta_{\sigma^2}} [g_{k+1}^{n+1} - g_{k+1}^n - g_j^{n+1} + g_j^n] \\
A_k^4 &= \frac{g_{j+1}^n - g_j^n}{\Delta_s}.
\end{aligned}$$

For the integral I_4 involving up-jumps we obtain in a similar way:

$$\begin{aligned}
I_4 &\approx C \sum_{k=j+1}^{N-1} \int_{s_k-s_j}^{s_{k+1}-s_j} \left[g_k^n + \frac{g_{k+1}^n - g_k^n}{\Delta_s} (x - s_k + s_j) \right] \frac{e^{-\frac{M}{\sigma}x}}{x} dx \\
&+ C \sum_{k=j+1}^{N-1} \int_{s_k-s_j}^{s_{k+1}-s_j} \left[\frac{g_k^{n+1} - g_k^n}{\Delta_{\sigma^2}} \frac{\lambda}{\delta} (x^2 - (s_k - s_j)^2) \right] \frac{e^{-\frac{M}{\sigma}x}}{x} dx \\
&- C \sum_{k=j+1}^{N-1} \int_{s_k-s_j}^{s_{k+1}-s_j} \left[g_j^n + \frac{g_{j+1}^n - g_j^n}{\Delta_s} (e^x - 1) + \frac{g_j^{n+1} - g_j^n}{\Delta_{\sigma^2}} \frac{\lambda}{\delta} x^2 \right] \frac{e^{-\frac{M}{\sigma}x}}{x} dx \\
&= C \sum_{k=j+1}^{N-1} \left[B_k^1 \int_{s_k-s_j}^{s_{k+1}-s_j} \frac{e^{-\frac{M}{\sigma}x}}{x} dx + B_k^2 \int_{s_k-s_j}^{s_{k+1}-s_j} e^{-\frac{M}{\sigma}x} dx + B_k^3 \int_{s_k-s_j}^{s_{k+1}-s_j} x e^{-\frac{M}{\sigma}x} dx \right] \\
&+ C \sum_{k=j+1}^{N-1} \left[B_k^3 \int_{s_k-s_j}^{s_{k+1}-s_j} x e^{-\frac{M}{\sigma}x} dx - B_k^4 \int_{s_k-s_j}^{s_{k+1}-s_j} \frac{e^{-(\frac{M}{\sigma}-1)x}}{x} dx \right] \tag{3.3.14}
\end{aligned}$$

where

$$\begin{aligned}
B_k^1 &= \left[g_k^n + \frac{g_{k+1}^n - g_k^n}{\Delta_s} (s_j - s_k) - \frac{g_k^{n+1} - g_k^n}{\Delta_{\sigma^2}} \frac{\lambda}{\delta} (s_j - s_k)^2 - g_j^n + \frac{g_{j+1}^n - g_j^n}{\Delta_s} \right] \\
B_k^2 &= \frac{g_{k+1}^n - g_k^n}{\Delta_s} \\
B_k^3 &= \frac{2\lambda}{\delta \Delta_{\sigma^2}} [g_k^{n+1} - g_k^n - g_j^{n+1} + g_j^n] \\
B_k^4 &= \frac{g_{j+1}^n - g_j^n}{\Delta_s}.
\end{aligned}$$

The integrals in (3.3.13) and (3.3.14) are given by exponentials and in terms of the exponential integral.

For I_2 we have:

$$\begin{aligned} I_2 &\approx C \int_{-\Delta_s}^0 \left[g + xg_s + \frac{\lambda}{\delta} x^2 g_{\sigma^2} - g - \left(1 + x + \frac{x^2}{2} - 1\right)g_s - \frac{\lambda}{\delta} x^2 g_{\sigma^2} \right] \frac{e^{-\frac{C}{\sigma}|x|}}{|x|} dx \\ &= -g_s \frac{C}{2} \int_{-\Delta_s}^0 x^2 \frac{e^{-\frac{C}{\sigma}|x|}}{|x|} dx \end{aligned}$$

so the approximation at the (j, n) node is

$$I_2 \approx -\frac{g_{j+1}^n - g_j^n}{\Delta_s} \frac{C}{2} \int_{-\Delta_s}^0 |x| e^{-\frac{C}{\sigma}|x|} dx. \quad (3.3.15)$$

In the same way

$$I_3 \approx -\frac{g_{j+1}^n - g_j^n}{\Delta_s} \frac{C}{2} \int_0^{\Delta_s} x e^{-\frac{M}{\sigma}x} dx. \quad (3.3.16)$$

When the claim g is European call option we solve (3.3.11) with with initial conditions:

$$g(s, \sigma^2, 0) = (e^s - K)^+$$

and boundary conditions in the log-price direction:

$$\begin{aligned} g(s, \sigma^2, \tau) &= 0 \quad \text{for } s \rightarrow 0 \\ g(s, \sigma^2, \tau) &= (e^{s-q\tau} - Ke^{-r\tau})^+ \quad \text{for } s \rightarrow \infty \end{aligned}$$

where K is the strike price and r and q are the risk-free rate and the dividend yield respectively. (Similar boundary conditions apply for an European put option). In the volatility direction we impose no boundary conditions. Instead, we use one-sided derivatives of g at the volatility end-points of the grid. Having the required derivatives at each grid point at time τ (starting at $\tau = 0$), we apply an explicit scheme to solve the PIDE and find the value of g at each each grid point at time

$\tau + 1$ (including all points with boundary volatility values).

One issue remaining is the discretization error. We do not derive here analytically an expression for the error, but perform a small numerical experiment - we calculate option prices at several different values of the grid steps and try to infer from them the order of convergence of the scheme. First, it turns out that the solution depends very little on the step in the volatility direction - the option prices change at most by fraction of a percent when doubling or quadrupling the number of volatility grid points. In what follows we use 30 volatility grid points. However, the sensitivity to the step in the log-price direction is much larger.

Graphs 1 and 2 show, for a number of steps, the prices (large dots) of options with different strikes. The spot price is 100, the grid is limited between 20 and 300 in the log-price direction, and we use 50 to 300 uniform log-price steps. The lines on the graphs are quadratic functions fitted through the prices with the three largest steps (i.e. the three right-most dots). In most cases these lines pass exactly through the prices calculated with smaller steps. We extrapolate these quadratics and take the value at zero to be the option price. (Fitting linear functions through any two of the three right-most dots proved to be less precise.)

As a check on the option prices obtained in the above way, we compare put prices calculated directly with the PIDE and put prices obtained from put-call parity using the PIDE to calculate the respective call prices. Table 3.1 shows the prices of puts with strikes from 70 to 130 (where the spot is 100, time to maturity is 0.5, interest and dividends are zero and the COGARCH parameters are: $\sigma = 0.25$, $\nu = 0.1$, $\theta = -0.3$, $\delta = 0.9$, $\lambda = 0.1$ and $\beta = 0.1$).

Table 3.1. Put prices - put-call parity vs. PIDE








strike	put price via put-call parity	put price via PIDE	% difference
70	0.3262	0.3274	0.3639
80	1.1689	1.1701	0.1023
90	3.2939	3.2950	0.0361
100	7.2444	7.2456	0.0163
110	13.3758	13.3770	0.0089
120	21.3859	21.3871	0.0056
130	30.4951	30.4963	0.0039

The differences are very small - about 0.1 cent and not more than a third of a percent even for the most out-of-the-money option.

Graphs 3 and 4 show COGARCH option prices for a range of strikes, maturities and for several parameter sets. The parameters are chosen to vary widely across parameter sets. At the lowest maturity (0.25 years) there is little difference between the prices obtained at different parameter sets. For longer maturities (up to one year) the differences become larger with one notable exception - it appears that at all maturities option prices are extremely insensitive to the value of the parameter λ . The prices obtained for $\lambda = 10$ (the large dots) hardly differ from those obtained with $\lambda = 0.1$ (the thicker solid line). Of course, only the calibration of the model to actual option prices will reveal what are the appropriate values of the parameters








Graph 1. Discretization errors – Calls

Call prices are calculated on grids with log-price steps from 0.054 to 0.009 and denoted by large dots. The lines are quadratics through the three right-most dots. Spot is 100, maturity is 0.5, strikes are between 70 and 130.





Graph 2. Discretization errors – Puts

Put prices are calculated on grids with log-price steps from 0.054 to 0.009 and denoted by large dots. The lines are quadratics through the three right-most dots. Spot is 100, maturity is 0.5, strikes are between 70 and 130.



Graph 3. COGARCH call option prices



spot=100, strikes from 70 to 130



VG parameters: sigma = 0.25, nu = 0.1, theta = -0.3

Graph 4. COGARCH put option prices

spot=100, strikes from 70 to 130



VG parameters: sigma = 0.25, nu = 0.1, theta = -0.3

and the price sensitivity.

3.4 Conclusion

This paper studies option pricing under the COGARCH model of Kluppelberg et al. (2004). The paper derives a backward PIDE for the value of a claim written on an asset, following COGARCH. Some properties of European option prices under COGARCH are demonstrated.

The backward PIDE derived, however, has one obvious drawback if one intends to use it for calibration of the model to option prices: it needs to be solved separately for each option (or at least for each strike), which may be very costly in computation time. An alternative is the forward PIDE. By solving a single forward equation one can calculate simultaneously the prices of options of all strikes and maturities. Forward PIDEs for models of the log price involving jumps have been derived in Andersen and Andreasen (1999), Carr and Andreasen (2002), Carr and Hirsu (2002), Madan (2005). The derivation of a forward equation for the COGARCH process and its use for calibration of the model is left for future research.

Appendix A. Proof of Proposition 1

Expectations are taken under the statistical measure. The pricing kernel process Λ_t is obtained by conditioning on the realizations of the volatility σ^S . Let $\overline{E}_t[\cdot]$ denote expectation, conditional on the volatility path.

$$S_T = S_t \exp \left(\int_t^T D_s^S ds + \int_t^T \sqrt{\sigma_s^S} dW_s^C - \frac{1}{2} \int_t^T \sigma_s^S ds \right) \quad (1)$$

$$\begin{aligned} \Lambda_t &= E_t \left[S_t^{-\lambda} \exp \left(-\lambda \int_t^T D_s^S ds - \lambda \int_t^T \sqrt{\sigma_s^S} dW_s^S + \frac{1}{2} \lambda \int_t^T \sigma_s^S ds \right) \right] \\ &= S_t^{-\lambda} \exp \left(-\lambda \int_t^T D_s^S ds \right) E_t \left[\exp \left(\frac{1}{2} \lambda \int_t^T \sigma_s^S ds \right) \overline{E}_t \left[\exp \left(-\lambda \int_t^T \sqrt{\sigma_s^S} dW_s^S \right) \right] \right] \\ &= S_t^{-\lambda} \exp \left(-\lambda \int_t^T D_s^S ds \right) E_t \left[\exp \left(\frac{\lambda(\lambda+1)}{2} \int_t^T \sigma_s^S ds \right) \right] \end{aligned} \quad (2)$$

To evaluate the expectation in (2) we use the fact that the characteristic function of the integral of a CIR variable is known in closed form. Let $\tau = T - t$ and $Y_\tau = \int_t^T \sigma_s^S ds$. Then the characteristic function of Y_τ is:

$$\phi_{Y_\tau}(u) = E_t [e^{iuY_\tau}] = A(\tau, u) \exp [B(\tau, u)\sigma_t^S] \quad (3)$$

$$\begin{aligned} A(\tau, u) &= \frac{\exp(\frac{\kappa^2 \theta \tau}{\eta^2})}{\left(\cosh(\frac{\gamma \tau}{2}) + \frac{\kappa}{\gamma} \sinh(\frac{\gamma \tau}{2}) \right)^{\frac{2\kappa \theta}{\eta^2}}} \\ B(\tau, u) &= \frac{2iu}{\kappa + \gamma \coth(\frac{\gamma \tau}{2})} \quad \text{and} \quad \gamma = \sqrt{\kappa^2 - 2\eta^2 iu} \end{aligned}$$

On evaluating this characteristic function at $u = -i\frac{\lambda(\lambda+1)}{2}$ it follows from (2) that the stochastic discount factor is:

$$\zeta_t = -\frac{d\Lambda_t}{\Lambda_t} = \lambda \frac{dS_t}{S_t} - \frac{\lambda(\lambda+1)}{\kappa + \gamma \coth(\frac{\gamma\tau}{2})} d\sigma_t^S \quad (4)$$

whereby terms of order dt in the differentiation of Λ_t are ignored.

Appendix B. Volatility risk prices in a two-factor model

This appendix presents a stylized model which is consistent with the empirical findings of the paper. The model is an extension of the basic model in (1)-(3), which includes a second stochastic volatility. The stochastic discount factor derived for the extended model corresponds closely to the two-factor model (1.2.6) tested empirically in this paper. For the extended model I show that one volatility risk has always a negative price, while the price of the second volatility risk can be of both signs. In particular, this risk price is positive (as found empirically above) when the utility function is convex in the second volatility factor.

In the spirit of habit-formation models, I assume an economy where utility is of the form:

$$U_T = \frac{\left(\frac{C_T}{H_T}\right)^{1-\lambda}}{1-\lambda} \quad (5)$$

and the pricing kernel process is:

$$\Lambda_t = E_t \left[\left(\frac{C_T}{H_T}\right)^{-\lambda} \right] \quad (6)$$

where C denotes an aggregate consumption good, and λ is the risk aversion coefficient. H is interpreted as a time-varying habit in a multiplicative form, as introduced by Abel (1990) and Gali (1994). In models of this type utility does not depend on the absolute level of consumption, but on the level of consumption relative to a benchmark (habit). The habit is usually related to past consumption. Here I also allow for randomness in habit. Assume the following dynamics for H and C and their respective volatility processes:

$$\frac{dH_t}{H_t} = D_t^H dt + \sqrt{\sigma_t^H} dW_t^H \quad (7)$$

$$\frac{dC_t}{C_t} = D_t^C dt + \sqrt{\sigma_t^C} dW_t^C + \beta \sqrt{\sigma_t^H} dW_t^H \quad (8)$$

$$d\sigma_t^C = k^C (\theta^C - \sigma_t^C) dt + \eta^C \sqrt{\sigma_t^C} dW_t^1 \quad (9)$$

$$d\sigma_t^H = k^H (\theta^H - \sigma_t^H) dt + \eta^H \sqrt{\sigma_t^H} dW_t^2 \quad (10)$$

All four Brownian motions W_t^H , W_t^C , W_t^1 and W_t^2 are assumed to be independent. The drift of habit growth (D_t^H) can be a function of past consumption as in the "catching up with the Joneses" versions of habit models; it is not modeled explicitly as before. What is essential is the separate source of randomness in habit - W_t^H , with stochastic volatility $\sqrt{\sigma_t^H}$ ("habit volatility"). Such a source of randomness reflects the notion that habit also depends on some current variables, similar to the "keeping up with the Joneses" versions of habit-formation models. The drift D_t^C is not modeled explicitly as before. A key feature of the the model is that it allows for random changes in habit to affect the dynamics of the consumption good (β is a sensitivity parameter, so it should take values between zero and one). We now prove the following:

Proposition 1 *The stochastic discount factor ζ_t in the economy with two volatility factors described in (5)-(10) is given by*

$$\zeta_t = -\frac{d\Lambda_t}{\Lambda_t} = \lambda \frac{dC_t}{C_t} - \lambda \frac{dH_t}{H_t} - B^C(\tau, u^C) d\sigma_t^C - B^H(\tau, u^H) d\sigma_t^H \quad (11)$$

where

$$B^C > 0 \quad \text{and}$$

$$B^H > 0 \quad \text{if} \quad \frac{\lambda - 1}{\lambda + 1} < \beta < 1 \quad (12)$$

Proof. Expectations are taken under the statistical measure. For brevity in notation let:

$$\overline{C}_\tau = \exp \left(\int_t^T D_s^C ds \right) \quad \text{and} \quad \overline{H}_\tau = \exp \left(\int_t^T D_s^H ds \right)$$

Condition on the realizations of σ^H and σ^C , and obtain for H_T and C_T the following:

$$H_T = H_t \overline{H}_\tau \exp \left(\int_t^T \sqrt{\sigma_s^H} dW_s^H - \frac{1}{2} \int_t^T \sigma_s^H ds \right) \quad (13)$$

$$C_T = C_t \overline{C}_\tau \exp \left(\int_t^T \sqrt{\sigma_s^C} dW_s^C + \beta \int_t^T \sqrt{\sigma_s^H} dW_s^H \right) \\ \times \exp \left(-\frac{1}{2} \int_t^T \sigma_s^C ds - \frac{1}{2} \beta^2 \int_t^T \sigma_s^H ds \right) \quad (14)$$

The pricing kernel is:

$$\Lambda_t = C_t^{-\lambda} \overline{C}_\tau^{-\lambda} H_t^\lambda \overline{H}_\tau^{-\lambda} E_t [L \overline{E}_t [M]] \quad (15)$$

where

$$L = \exp \left(\frac{1}{2} \lambda \left(\int_t^T \sigma_s^C ds + (\beta^2 - 1) \int_t^T \sigma_s^H ds \right) \right)$$

$$M = \exp \left(-\lambda \left(\int_t^T \sqrt{\sigma_s^C} dW_s^C + (\beta - 1) \int_t^T \sqrt{\sigma_s^H} dW_s^H \right) \right)$$

Take first the conditional expectation in (29):

$$\begin{aligned} \Lambda_t &= C_t^{-\lambda} \overline{C}_\tau^{-\lambda} H_t^\lambda \overline{H}_\tau^{-\lambda} E_t \left[L \exp \left(\frac{1}{2} \lambda^2 \left(\int_t^T \sigma_s^C ds + (\beta - 1)^2 \int_t^T \sigma_s^H ds \right) \right) \right] \\ &= C_t^{-\lambda} \overline{C}_\tau^{-\lambda} H_t^\lambda \overline{H}_\tau^{-\lambda} E_t [N] \end{aligned} \quad (16)$$

where

$$N = \exp \left(\frac{1}{2} \lambda(\lambda + 1) \int_t^T \sigma_s^C ds - \frac{1}{2} \lambda(1 - \beta) [1 + \beta - \lambda(1 - \beta)] \int_t^T \sigma_s^H ds \right)$$

Using the independence of the two volatility processes and (3), the pricing kernel process is:

$$\Lambda_t = C_t^{-\lambda} \overline{C}_\tau^{-\lambda} H_t^\lambda \overline{H}_\tau^{-\lambda} A^C(\tau, u^C) \exp(B^C(\tau, u^C) \sigma_t^C) A^H(\tau, u^H) \exp(B^H(\tau, u^H) \sigma_t^H) \quad (17)$$

where superscripts C and H denote parameters related to consumption and habit respectively, and

$$u^C = -i \frac{\lambda(\lambda + 1)}{2}$$

$$u^H = i \frac{1}{2} \lambda(1 - \beta) [1 + \beta - \lambda(1 - \beta)]$$

$$B^C(\tau, u^C) = \frac{\lambda(\lambda + 1)}{\kappa^C + \gamma^C \coth(\frac{\gamma^C \tau}{2})} > 0 \quad (18)$$

$$B^H(\tau, u^H) = -\frac{\lambda(1-\beta)[1+\beta-\lambda(1-\beta)]}{\kappa^H + \gamma^H \coth(\frac{\gamma^H \tau}{2})}$$

The stochastic discount factor is then:

$$\zeta_t = -\frac{d\Lambda_t}{\Lambda_t} = \lambda \frac{dC_t}{C_t} - B^C(\tau, u^C) d\sigma_t^C - \lambda \frac{dH_t}{H_t} - B^H(\tau, u^H) d\sigma_t^H \quad (19)$$

and $B^H(\tau, u^H) > 0$ if $\frac{\lambda-1}{\lambda+1} < \beta < 1$ ■

The first two terms are exactly as in the case with one volatility factor and without habit. The third term reflects a negative habit risk price. Given that an increase in habit decreases utility, any asset positively correlated with the change in habit has high payoffs in low-utility states, which provides the intuition for the negative price of habit risk. The last term reflects the price of "habit volatility" risk. Assume further that consumption equals dividends and that the market price-to-dividend ratio is constant. So, the market risk price is positive and the prices of the two market volatility risks have the signs of the risk prices of σ_t^C and σ_t^H respectively. Then the stochastic discount factor can also be written as:

$$\zeta_t = -\frac{d\Lambda_t}{\Lambda_t} = \lambda \frac{dS_t}{S_t} - \lambda \frac{dH_t}{H_t} + \lambda^\sigma d\sigma_t + \lambda^{\sigma^H} d\sigma_t^H \quad (20)$$

$\lambda^\sigma < 0$ as follows from (18) and (19) and the sign of λ^{σ^H} is determined from (12).

What is the intuition for condition (12)? We now show that this condition is closely related to the concavity or convexity of the utility function (5) in habit.

Denote for brevity

$$\widetilde{C}_T = C_t \overline{C}_\tau \exp \left(\int_t^T \sqrt{\sigma_s^C} dW_s^C - \frac{1}{2} \int_t^T \sigma_s^C ds \right)$$

Re-write (14) as:

$$\begin{aligned}
C_T &= \tilde{C}_T \exp \left(\beta \int_t^T \sqrt{\sigma_s^H} dW_s^H - \frac{1}{2} \beta^2 \int_t^T \sigma_s^H ds \pm \frac{1}{2} \beta \int_t^T \sigma_s^H ds \right) \\
&= \tilde{C}_T H_t^{-\beta} \overline{H}_\tau^{-\beta} \exp \left(\frac{1}{2} \beta (1 - \beta) \int_t^T \sigma_s^H ds \right) H_T^\beta
\end{aligned} \tag{21}$$

Then utility at time T is:

$$U_T = \frac{1}{1 - \lambda} \left[\tilde{C}_T H_t^{-\beta} \overline{H}_\tau^{-\beta} \exp \left(\frac{1}{2} \beta (1 - \beta) \int_t^T \sigma_s^H ds \right) \right]^{1-\lambda} H_T^{(\beta-1)(1-\lambda)} \tag{22}$$

This function is convex in H_T when $\frac{\lambda-2}{\lambda-1} < \beta < 1$. While this condition is not equivalent to (12), it can be easily checked that for $\lambda > 3$ the positive price of habit volatility risk follows from the convexity in habit of the utility function. Empirical studies show that $\lambda > 3$ is a plausible range for the values of the risk aversion parameter. So, intuition for the the price of "habit volatility" risk can be provided, as in the single volatility risk case by the shape of the utility function - concavity in habit reflects a decrease in expected utility when "habit volatility" increases, hence the desirability of assets correlated with changes in "habit volatility" and the negative price of "habit volatility" risk. Exactly the opposite argument applies when the utility function is convex in habit, so the price of "habit volatility" risk is positive in this case.

The model presented above is illustrative in nature. It leaves unspecified some important components, in particular the form of the two drifts D_t^H and D_t^C . It also assumes a constant price-dividend ratio. Still, it points to the essential role that concavity / convexity of utility functions can have in modeling volatility risks.

Utility functions which exhibit both concavity and convexity have been found

empirically in different contexts. From results in Jackwerth (2000) it follows that investors' utility from wealth (proxied by a market index) has been changing over time. In particular after 1987 it has exhibited a convex shape over certain ranges of market moves. In a similar spirit, Carr et al. (2002) estimate marginal utility over instantaneous market moves (jumps) of different size and show that it has both decreasing and increasing sections over different ranges of jumps. The increasing sections correspond to convex utility. It is interesting to explore their results in the context of utility functions with two arguments and possibly two volatility risks.

Now compare (20) with (1.2.6), estimated above. Note that (1.2.6) lacks a term corresponding to $\frac{dH_t}{H_t}$. However, this lack is reflected only in the pricing errors, but not in the prices of volatility risks (for uncorrelated risk factors). Besides, with excess returns and normalized risk factors the risk prices in the two models are equivalent, as discussed above. It follows that the extended model is consistent with the empirical findings of this paper

Appendix C. Option pricing models

This appendix presents details on the three option pricing models compared in Section 3:

1. Stochastic volatility and jumps model

Bates (1996) and Bakshi et al. (1997) consider an eight-parameter model for the log price with stochastic volatility and jumps (SVJ):

$$\begin{aligned} dS_t &= (r - \lambda)S_t dt + \sqrt{V_t}S_t dW_t^1 + J_y S_t dq_t^y \\ dV_t &= (\theta - kV_t)dt + \sigma\sqrt{V_t}dW_t^2 \end{aligned} \tag{23}$$

W_t^1 and W_t^2 are standard Brownian motions with correlation ρ ,

J_y is log-normal with mean μ_y and variance σ_y^2

q_t^y is a Poisson process with arrival rate λ_y

r is interest rate and λ is the jump-compensator

Let: $a = u^2 + iu + 2iu\mu_y\lambda_y - 2\lambda_y(e^{iu\mu_y - \frac{u^2\sigma_y^2}{2}} - 1)$

$$b = iu\sigma\rho - k$$

$$\gamma = \sqrt{b^2 + a\sigma^2}$$

$$c = \theta \left[(\gamma + b)t + 2\log\left(1 - \frac{\gamma+b}{2\gamma}(1 - e^{-\gamma t})\right) \right]$$

$$d = 2\gamma - (\gamma + b)(1 - e^{-\gamma t})$$

$$B = -\frac{a(1 - e^{-\gamma t})}{d}$$

The characteristic function of the log-price at horizon t under the SVJ process is:

$$E^Q \left[e^{iu \log(S_t)} \right] = e^{iu \log(S_0) + iurt - \frac{c}{\sigma^2} + BV_0} \quad (24)$$

where V_0 is instantaneous variance. The expectation is taken here under the risk-neutral measure Q .

2. Double jumps model

Duffie et al. (2000) develop a flexible model specification, which also adds jumps in volatility. I use here one of their double-jump (DJ) models:

$$\begin{aligned} dS_t &= (r - \lambda)S_t dt + \sqrt{V_t}S_t dW_t^1 + J_y S_t dq_t^y \\ dV_t &= (\theta - kV_t)dt + \sigma\sqrt{V_t}dW_t^2 + J_v dq_t^v \end{aligned} \quad (25)$$

All variables, except for J_v and q_t^v have the same meaning as above. q_t^v is a second Poisson process with arrival rate λ_v . The sizes of the jumps are exponentially

distributed with mean μ_v

$$\text{Let: } a = u^2 + iu$$

$$b = iu\sigma\rho - k$$

$$\gamma = \sqrt{b^2 + a\sigma^2}$$

$$c = \theta \left[(\gamma + b)t + 2 \log(1 - \frac{\gamma+b}{2\gamma}(1 - e^{-\gamma t})) \right]$$

$$d = 2\gamma - (\gamma + b)(1 - e^{-\gamma t})$$

$$f = \frac{\lambda_y e^{\mu_y + \frac{\sigma_y^2}{2}} + \lambda_v}{\lambda_y + \lambda_v} - 1$$

$$g = t e^{iu\mu_y - \frac{u^2\sigma_y^2}{2}}$$

$$h = \frac{\gamma-b}{\gamma-b+a}t - \frac{2\mu_v a}{\gamma^2 - (b-\mu_v a)^2} \log(1 - \frac{\gamma+b-\mu_v a}{2\gamma}(1 - e^{-\gamma t}))$$

$$A = -(\lambda_y + \lambda_v)(1 + iuf)t + \lambda_y g + \lambda_v h$$

$$B = -\frac{a(1-e^{-\gamma t})}{d}$$

The characteristic function of the log-price at horizon t under the DJ process is:

$$E^Q [e^{iu \log(S_t)}] = e^{iu \log(S_0) + iurt - \frac{c}{\sigma^2} + A + BV_0} \quad (26)$$

3. VGSA model

Carr and Wu (2003) provide a general study of the financial applications of time-changed Levy processes. They show that most of the stochastic processes employed as models of asset returns, including SVJ and DJ, belong to this class. The third model I consider here is also based on a process of this class. The VGSA process (Carr et al. (2003)) is a six-parameter pure-jump process with stochastic arrival rate of jumps. (Stochastic arrival is an analogue of stochastic volatility for pure-jump processes.) VGSA is introduced in two steps, each step being an explicit time-change of a Levy process.

At the first step a Brownian motion with drift, denoted as

$$b(t; \theta, \sigma) = \theta t + \sigma W_t \quad (27)$$

where W_t is standard Brownian motion, is evaluated at a time given by an independent gamma process $\gamma(t; 1, \nu)$ with unit mean rate and variance rate ν . The process, obtained in this way is the Variance Gamma (VG) process (Carr and Madan (1998)):

$$VG(t; \sigma, \nu, \theta) = b(\gamma(t; 1, \nu); \theta, \sigma) \quad (28)$$

VG is a Levy process, like the Brownian motion, but unlike it, is a pure-jump process, due to the gamma time-change. Its characteristic function is:

$$\phi_{VG_t}(u) = E [e^{iuVG_t}] = \left(\frac{1}{1 - i\theta\nu u + (\sigma^2\nu/2)u^2} \right)^{t/\nu} \quad (29)$$

and its characteristic exponent is:

$$\psi_{VG}(u) = \frac{1}{\nu} \ln (1 - i\theta\nu u + (\sigma^2\nu/2)u^2) \quad (30)$$

At the second step, the VG process itself is time-changed. The time-change is independent of the VG process and is given by the integral of the mean-reverting CIR process. The CIR process is solution to the following stochastic differential equation:

$$dy_t = k(\eta - y_t)dt + \lambda\sqrt{y_t}dW_t \quad (31)$$

Denote $Y_t = \int_0^t y_s ds$, then

$$VGSA(t; \sigma, \nu, \theta, k, \eta, \lambda) = VG(Y_t; \sigma, \nu, \theta) \quad (32)$$

The characteristic function of Y_t is:

$$\phi_{Y_t}(u) = E [e^{iuY_t}] = A(t, u) \exp B(t, u)y_0$$

where $A(t, u)$, $B(t, u)$ and γ are as given in Section 2.2. Since Y_t is independent of the VG process, the characteristic function of the VGSA process can be obtained via conditional expectation:

$$\phi_{VGSA_t}(u) = E [e^{iuVGSA_t}] = \phi_{Y_t}(i\psi_{VG}(u)) \quad (33)$$

Then the characteristic function of the log-price at horizon t under the VGSA process is:

$$E^Q [e^{iu \log(S_t)}] = e^{iu \log(S_0) + iurt} \frac{\phi_{VGSA_t}(u)}{\phi_{VGSA_t}(-i)} \quad (34)$$

References

- [1] Abel, A. (1990), "Asset prices under habit formation and catching up with the Joneses", *American Economic Review* 80, 38-42.
- [2] Alizadeh, S., M. Brandt and F. Diebold (2002), "Range-based estimation of stochastic volatility models", *Journal of Finance* 57, 1047-1091.
- [3] Andersen, T. and T. Bollerslev (1997), "Heterogeneous information arrivals and return volatility dynamics: Uncovering the long-run in high frequency returns", *Journal of Finance* 52, 975-1005.
- [4] Andersen, L. and J. Andreasen (1999), "Jumping smiles", *Risk* 65-68
- [5] Andreasen, J. and P. Carr (2002), "Put-call reversal", Working paper, New York University
- [6] Ang, A., R. Hodrick, Y. Xing and X. Zhang (2004), "The cross-section of volatility and expected returns", Working paper, Columbia University.
- [7] Asmussen, S., F. Avram and M. Pistorius (2002), "Russian and American put options under exponential phase-type Levy models". *Stochastic Processes and their Application* (under revision).
- [8] Bansal, R. and A. Yaron (2004), "Risks for the long run", *Journal of Finance* 59, 1481-1509.
- [9] Bakshi, G. and C. Cao (2003), "Risk-neutral kurtosis, jumps and option pricing: Evidence from 100 most actively traded firms on the CBOE", Working paper, University of Maryland.

- [10] Bakshi, G. and N. Kapadia (2003), "Delta-hedged gains and the negative volatility risk premium", *Review of Financial Studies* 16, 101-143.
- [11] Barndorff-Nielsen, O. and N. Shephard (2001), "Non-Gaussian Ornstein-Uhlenbeck-based models and some of their uses in financial econometrics" *Journal of the Royal Statistical Society Ser. B* 63, 167-241
- [12] Bekaert, G. and G. Wu (2000), "Asymmetric volatilities and risk in equity markets", *Review of Financial Studies* 13, 1-42.
- [13] Benzoni, L. (2001) "Pricing options under stochastic volatility", Working paper, University of Minnesota.
- [14] BIS (1996), "Overview of the Amendment to the Capital Accord to incorporate market risks", Basle Committee on Banking Supervision, January
- [15] Black, F. (1976), "Studies of stock price volatility changes", *Proceedings of the 1976's meeting of the Business and Economic Statistics Section*, American Statistical Association, 177-181.
- [16] Bodoukh, J., M. Richardson, R. Stanton and R. Whitelaw (1995), "Expect the worst - rethinking the Value-at-Risk concept using worst case scenario analysis and its implications for risk management", *Risk*, September
- [17] Bodoukh, J., M. Richardson, R. Stanton and R. Whitelaw (2004), "MaxVaR: long-horizon Value-at-Risk in a mark-to-market environment", *Journal of Investment Management* 2, 1-6
- [18] Buraschi, A. and J. Jackwerth (2001), "The price of a smile: hedging and spanning in option markets", *Review of Financial Studies* 14, 495-527.

- [19] Carr, P. and D. Madan (1999), "Option pricing and the Fast Fourier transform", *Journal of Computational Finance* 4, 61-73.
- [20] Carr, P. and D. Madan (2001), "Towards a theory of volatility trading", in *Volatility* R. Jarrow (ed.) Risk publications.
- [21] Carr, P., H. Geman, D. Madan and M. Yor (2002), "The fine structure of asset returns", *Journal of Business* 75, 305-332.
- [22] Carr, P., H. Geman, D. Madan and M. Yor (2003), "Stochastic volatility for Levy processes", *Mathematical Finance* 13, 345-382.
- [23] Carr, P. and L. Wu (2003), "The finite-moment log-stable process and option pricing", *Journal of Finance* 58, 753-778
- [24] Carr, P. and L. Wu (2004), "Time-changed Levy processes and option pricing", *Journal of Financial Economics* 71, 113-141.
- [25] Carr, P. and L. Wu (2004), "Variance risk premia", Working paper, Baruch College.
- [26] Carr, P and A. Hirsu (2002), "Why be backward? Forward equations for American options", Working paper, New York University
- [27] Chernov, M. and E. Ghysels (2000), "A study towards a unified approach to the joint estimation of objective and risk-neutral measures for the purpose of option valuation", *Journal of Financial Economics* 56, 407-458.
- [28] Chernov, M., A. Gallant, E. Ghysels and G. Tauchen (2003), "Alternative models for stock price dynamics", *Journal of Econometrics* 116, 225-257.

- [29] Choudhury, G., D. Lucantoni and W. Whitt (1994), "Multidimensional transform inversion with applications to the transient M/G/1 queue", *Annals of Applied Probability* 4, 115-123
- [30] Cochrane, J. (2001), *Asset Pricing*, Princeton University Press.
- [31] Corradi, V. (2000), "Reconsidering the continuous-time limit of the GARCH(1,1) process", *Journal of Econometrics* 96, 145-153
- [32] Coval, J. and T. Shumway (2001), "Expected option returns," *Journal of Finance* 56, 983-1009.
- [33] Christensen, B. and N. Prabhala (1998), "The relation between implied and realized volatility", *Journal of Financial Economics* 50, 125-150.
- [34] Christoffersen, P. and K. Jacobs (2004), "Which GARCH model for option valuation?", Working paper, McGill University
- [35] Doran, J. and E. Ronn (2003), "On the market price of volatility risk", Working paper, University of Texas.
- [36] Duan, J.-C. (1995), "The GARCH option pricing model", *Mathematical Finance* 8, 13-32
- [37] Duan, J.-C. (1996a), "A unified theory of option pricing under stochastic volatility - from GARCH to diffusion", Working paper, University of Toronto
- [38] Duan, J.-C. (1996b), "Cracking the smile", *Risk* 9, 55-59
- [39] Duffie, D., J. Pan and K. Singleton (2000) "Transform analysis and asset pricing for affine jump-diffusions", *Econometrica* 68, 1343-1376.

- [40] Eberlein, E. (2001), "Application of generalized hyperbolic Levy motions to finance", in *Levy Processes, Theory and Applications*, Barndorff-Nielsen, O., T. Mikosch and S. Resnick eds., Birkhauser, Boston 319-336
- [41] Eberlein, E, J. Kallsen and J. Kristen (2002), "Risk management based on stochastic volatility", *Journal of Risk* 5, 19-44
- [42] Engle, R. and G. Lee (1998), "A long-run and short-run component model of stock return volatility", in R. Engle and H. White (eds.), *Co-integration, causality and forecasting: A Festschrift in Honor of Clive W. J. Granger*, Oxford University Press, 475-497.
- [43] Engle, R. and C. Mustafa (1992), "Implied ARCH models from option prices", *Journal of Econometrics* 52, 289-311.
- [44] Gali, J. (1994), "Keeping up with the Joneses: consumption externalities, portfolio choice and asset prices", *Journal of Money, Credit and Banking* 26, 1-8.
- [45] Gallant, A., C. Hsu, and G. Tauchen (1999), "Using daily range data to calibrate volatility diffusions and extract the forward integrated variance", Working paper, Duke University.
- [46] Garcia, R and E. Renault (1998), "A note on hedging in ARCH and stochastic volatility option pricing models", *Mathematical Finance* 8, 153-161
- [47] Hardle, W. and C. Hafner (2000), "Discrete-time option pricing with flexible volatility estimation", *Finance and Stochastics* 4, 189-207

- [48] Heston, S. (1993), "A closed-form solution for options with stochastic volatility and application to bond and currency options", *Review of Financial Studies* 6, 327-343.
- [49] Heston, S. (1993), "Invisible parameters in option prices", *Journal of Finance* 48, 933-947
- [50] Heston, S. and S. Nandi (2000), "A closed-form GARCH option valuation model", *Review of Financial Studies* 13, 585-626
- [51] Hirsa, A. (1999), Ph.D. Thesis, University of Maryland
- [52] Hirsa, A. and D. Madan (2003), "Pricing American options under variance gamma", *Computational Finance* 7
- [53] Jeanthau, T. (2004), "A link between complete models with stochastic volatility and ARCH models", *Finance and Stochastics* 8, 111-131
- [54] Jones, C. (2003), "The dynamics of stochastic volatility", *Journal of Econometrics* 116, 181-224.
- [55] Jackwerth, J. (2000), "Recovering risk aversion from option prices and realized returns", *Review of Financial Studies* 13, 433-451.
- [56] Khanna, A. and D. Madan (2004), "Understanding option prices", *Quantitative Finance* 4, 55-63
- [57] Kluppelberg, C., A. Lindner and R. Maller (2004), "A continuous-time GARCH process driven by a Levy process: stationarity and second-order behavior", *Journal of Applied Probability* 41

- [58] Kou, S. and H. Wang (2003), "First-passage times for a jump diffusion process", *Advances in Applied Probability* 35, 504-531
- [59] Kritzman, M and D. Rich (2002), "The mismeasurement of risk", *Financial Analysts Journal* 58, 91-99
- [60] Liesenfeld, R. (2001) "A generalized bivariate mixture model for stock price volatility and trading volume", *Journal of Econometrics* 104, 141-178.
- [61] MacKinlay, C. and J. Park (2004), "The relationship between expected risk premium and conditional permanent and transitory volatility", Working paper, Wharton School.
- [62] Madan, D., P. Carr and E. Chang (1998), "The Variance Gamma process and option pricing", *European Finance Review* 1, 79-105.
- [63] Muller, U., M. Dacorogna, R. Dave, R. Olsen and O. Pictet (1997), "Volatilities of different time resolutions - analyzing the dynamics of market components", *Journal of Empirical Finance* 4, 213-239.
- [64] Nelson, D. (1990), "Conditional heteroskedasticity in asset returns: a new approach", *Econometrica* 59, 347-370
- [65] Pan, J. (2002), "The jump-risk premia implicit in options: Evidence from an integrated time-series study", *Journal of Financial Economics* 63, 3-50.
- [66] Poterba, J. and L. Summers (1986), "The persistence of volatility and stock-market fluctuations", *American Economic Review* 76, 1142-1151.
- [67] Ritchken, P. and R. Trevor (1999), "Pricing options under generalized GARCH and stochastic volatility processes", *Journal of Finance* 54, 377-402

- [68] Rogers, L. C. G. (2000), Evaluating first-passage probabilities for spectrally one-sided Levy processes", *Journal of Applied Probability* 37, 1173-1180
- [69] Shoutens, W. and J. Teugels (1998), Levy processes, polynomials and martingales", *Communications of Statistics and Stochastic Models* 14, 335-349
- [70] Schrand, C and H. Unal (1998), "Hedging and coordinated risk management", *Journal of Finance* 53, 979-1013.
- [71] Stulz, R. (1996), "Rethinking risk management", *Journal of Applied Corporate Finance* 9, 8-24
- [72] Tauchen, G. (2004), "Stochastic volatility in general equilibrium", Working paper, Duke University.
- [73] Xu, X. and S. Taylor (1994),. "The term structure of volatility implied by foreign exchange options", *Journal of Financial and Quantitative Analysis* 29, 57-74.
- [74] Wang, Y. (2002), "Asymptotic non-equivalence of GARCH models and diffusions", *Annals of Statistics* 30, 754-783