Three Essays in Theoretical and Empirical Derivative Pricing

Ali Boloorforoosh

A Thesis
in
The John Molson School of Business
Department of Finance

Presented in Partial Fulfillment of the Requirements
for the Degree of
Doctor of Philosophy at
Concordia University
Montreal, Quebec, Canada

August 2014

CONCORDIA UNIVERSITY

School of Graduate Studies

This is to certify that the thesis prepared

By: Ali Boloorforoosh

Entitled: Three Essays in Theoretical and Empirical Derivative Pricing

and submitted in partial fulfillment of the requirements for the degree of

Doctor of Philosophy (Business Administration)

complies with the regulations of the University and meets the accepted standards with respect to originality and quality.

Signed by the final examining committee:

Dr. Linda Dyer	Chair
Dr. P. Christoffersen	External Examiner
Dr. Prosper Dovonon	External to Program
Dr. L. Kryzanowski	Examiner
Dr. B. Croitoru	Examiner
Dr. D. Newton	Examiner
Dr. S. Perrakis	Supervisor

Approved by ______________________________

Dr. H. Bhabra, Graduate Program Director

August 26, 2014

____

Dr. S. Harvey, Dean

John Molson School of Business

ABSTRACT

Three Essays in Theoretical and Empirical Derivative Pricing

Ali Boloorforoosh, Ph.D.
Condordia University, 2014

In this thesis we present three papers in empirical and theoretical derivative pricing. In the first essay, which is an empirical study on equity options, we use simultaneous data from equity, index and option markets in order to estimate a single factor market model in which idiosyncratic volatility is allowed to be priced. We model the index dynamics' P-distribution as a mean-reverting stochastic volatility model as in Heston (1993), and the equity returns as single factor models with stochastic idiosyncratic volatility terms. We derive theoretically the underlying assets' Q-distributions and estimate the parameters of both P- and Q-distributions using a joint likelihood function. We document the existence of a common factor structure in option implied idiosyncratic variances. We show that the average idiosyncratic variance, which proxies for the common factor, is priced in the cross section of equity returns, and that it reduces the pricing error when added to the Fama-French model. We find that the idiosyncratic volatilities differ under P- and Q-measures, and we estimate the price of this idiosyncratic volatility risk, which turns out to be always significantly different from zero for all the stocks in our sample. Further, we show that the idiosyncratic volatility risk premiums are not explained by the usuall equity risk factors. Finally, we explore the implications of our results for the estimates of the conditional equity betas.

In the next two essays, we present a theoretical methodology for the pricing of catastrophe derivatives. In the second essay, we present a new approach to the pricing of catastrophe event derivatives that does not assume a fully diversifiable event risk. Instead, we assume that the event occurrence and intensity affect the return of the market portfolio of an agent that

trades in the event derivatives. Based on this approach, we derive values for a CAT option and a reinsurance contract on an insurers assets using recent results from the option pricing literature. We show that the assumption of unsystematic event risk seriously underprices the CAT option. Last, we present numerical results for our derivatives using real data from hurricane landings in Florida.

In the third and final essay, we extend the methodology developed in the second essay by relaxing several of its assumptions, and apply it to the valuation of a reinsurance contract given the value of a futures contract indexed on the CAT event. Since the payoff of the reinsurance contract has the form of a vertical spread, our methodology is also applicable to the valuation of derivatives with non-convex payoffs in other markets. Our approach recognizes the fundamental incompleteness of financial markets arising from the occurrence of rare events. Moreover, our approach does not rely on the existence of a representative agent and his/her risk preferences, and we do not assume knowledge of the martingale probability measure beyond the futures price. Using stochastic dominance methodology, we derive bounds for the value of the reinsurance contract. The derived bounds represent reservation write and purchase prices for the reinsurance contract, violation of which result in second degree stochastically dominating strategies that increase the expected utility of any risk averse investor. Our method is general in nature and is independent of distributional assumptions on the CAT event amplitude. It can be generalized without reformulation to any Markovian process that may include dependence between the amplitude distribution and the frequency of occurrence of the event.

To my grandmother, maman mehri.

ACKNOWLEDGMENTS

I would like to extend my gratitude to my supervisor Dr. Stylianos Perrakis for his continual guidance and support. He has been a mentor and a role model for me, and his attitude of hard work and excellence has always been an inspiration to me. I'm also grateful to my committee members, Dr. Christoffersen, Dr. Croitoru, Dr. Dovonon, Dr. Kryzanowski, Dr. Isaenko, and Dr. Newton for their valuable comments and constructive suggestions. I also acknowledge the help and support from the staff at the John Molson School of Business, Edite Almeida, Karen Fada, Karen Fiddler, Sandra Romanini, and Norma Paradis. Last, I would like to thank my fellow Ph.D. students, in particular Hamed Ghanbari and Kaveh Moradi, for the encouragement and especially for their friendship.

My long-term debt is to my wife and my best friend, Sussan. She has always made me believe in myself and selflessly supported me throughout my studies, and she has sacrificed a lot for my success. I would not be able to finish my dissertation without her love and support.

I am eternally indebted to my parents, Sepideh and Mohsen, for all they have done for me. Their teachings and support have influenced every aspect of my life. My mother has taught me how to be considerate and care for others, and my father has been my intellectual coach for the past thirty four years.

Special thanks go to my long time office mate Ash Hossain. He has always been a valuable friend, both socially and intellectually, and he has always had my interest at heart.

I would like to thank my brother, Amir, for always being there for me. I also thank my long time friends Hessam Babaee and Kaveh Laksari, who have taught me many things, and whose friendship has always been an asset in my life.

Finally, I am grateful to my son, Arman, for being the utmost joy in my life, and for the sleepless nights!

Contents

List of Tables

List of Figures

Chapter 1

Is Idiosyncratic Volatility Risk Priced? Evidence from the Physical and Risk-Neutral Distributions

1.1 Introduction

The most important result of the capital asset pricing model (CAPM) states that only the systematic risk is priced in equilibrium and idiosyncratic risk is not. Some earlier studies such as Levy (1978), Merton (1987), and Xu and Malkiel (2003) challenged this finding, by suggesting that investors may not be able to diversify properly. In such a case idiosyncratic risk should be positively related to the expected stock returns to compensate for this imperfect diversification. Although the definition of idiosyncratic risk has changed over time because of various redefinitions of the CAPM[1], the very definition of idiosyncratic risk implies that it should be either uncorrelated or, if the non-diversification argument is accepted, positively correlated to expected stock returns.

In the light of this theory, the results of the influential study of Ang et al. (2006), constitute a puzzle, since they document the underperformance of the stocks with high idiosyncratic return volatilities. Several subsequent studies have tried to explain this puzzle. Chen and Petkova (2012) propose that idiosyncratic volatility is priced because it correlates with changes in the average equity return variance, which is part of the aggregate variance. Duarte et al. (2012) introduce the predictive idiosyncratic variance component that correlates with macro economic factors, and argue that the puzzling findings of Ang et al. (2006) do not hold when portfolios are sorted based on the unpredicted idiosyncratic volatilities. Other papers have suggested that the puzzling findings are due to the choice of frequency, weighting, illiquidity, or the specific measurement of volatility.[2]

Another set of relevant CAPM studies examined beta estimation and systematic risk through the option market. Buss and Vilkov (2012, BV) and Chang et al. (2011, CCJV) show that information extracted from that market is forward-looking, while the usual CAPM estimations rely on historical data. Hence, CAPM results extracted from the option market may be better proxies for future estimates of systematic risk than those stemming from the conventional approach. Christoffersen et al. (2013, CFJ) adopt a similar reasoning and estimate betas from a cross section of index and equity options, assuming that the

[1] In particular, the conditional CAPM and the Fama-French (1993) model.

[2] See Huang et al. (2010), Bali and Cakici (2008), Han and Lesmond (2011), and Fu (2009).

idiosyncratic volatility is not priced.

A key issue in using the option market for the CAPM estimations is the change in probability measure between underlying and option markets. Once the complete markets assumption is abandoned, the underlying asset return distribution extracted from the option market, the risk neutral or Q-distribution, differs from the one observed in the underlying asset market, the physical or P-distribution. In particular, the P-distribution market volatility differs from the risk neutral volatility by the price of volatility risk. The CAPM studies relying on option market data have addressed the issue by correcting the volatility through an ad hoc modeling of the correlation matrix of the returns (as in BV) or by adopting several assumptions about the structure of idiosyncratic risk (as in CCJV). All these studies rely only on option market data for their empirical work and assume explicitly that idiosyncratic risk is not priced, or that idiosyncratic volatility is the same under both P- and Q-distributions.

In this paper we combine the two strands of literature, by investigating the pricing of idiosyncratic risk using option-implied volatilities. We estimate the parameters of the P- and Q-distributions from both underlying and option market data, and we assume in our underlying market model that idiosyncratic volatility is priced. Our empirical results reject decisively the hypothesis that idiosyncratic risk is not priced.[3] To the best of our knowledge this is the first study that links the physical and risk neutral distributions of idiosyncratic volatilities.

Our estimation methodology has also implications for equity betas, which are very important in portfolio selection and corporate finance. A popular estimation method is to use a rolling window of historical returns. Several studies have proposed using option prices to obtain forward looking estimates of stock betas. In our modeling framework beta enters the equity price dynamics and is estimated directly as part of the structural parameters. Moreover, we estimate the stock beta using the information in both returns and options prices, taking into account the market variance and idiosyncratic variance risk premiums. Further, we develop a new procedure to estimate conditional equity betas, which we estimate out-of-sample, using

[3] Note that our idiosyncratic volatility is not the same as in the study of Ang et al. (2006), since it is extracted from a single factor model and not from the Fama-French (1993) factors. The latter cannot generate an option pricing model.

only the option prices observed on a given day.

We use a continuous-time modeling framework that allows for a factor structure in equity returns, where the factor is the market return. The idiosyncratic return volatility of the stock ($IVol$) follows a square-root stochastic process and is allowed to be priced.[4] We estimate the model parameters and idiosyncratic volatility state variables, conditional on market parameters and state variables, using the joint information from the equity option prices and equity returns. Our data set contains historical returns and option prices for the market index and 27 blue chip stocks over the period 1996 to 2011, and contains more than 3.4 million option quotes. Our estimation is based on a likelihood function that has a return component and an option component, while the structural parameters are internally consistent between the P and Q measures. The simultaneous estimation of the P- and Q-distribution parameters allows us to filter the spot idiosyncratic volatilities under both distributions, using an internally consistent set of parameters.

Previous studies document a strong factor structure in implied volatility levels, moneyness slopes and term structure slopes, with the factor being the market implied volatility. In our results we find that even after removing the market return factor from the equity returns there is still a strong factor structure left in the implied idiosyncratic volatility ($IIVol$) levels, slopes and term structure slopes, very similar to that observed in total implied volatilities. The first two principal components of the $IIVol$ levels explain 58% and 23% of the cross-sectional variations, respectively. The first common component has a 99% correlation with the average implied idiosyncratic volatility levels of all firms. Further, the first and the second common components have a correlation of 65% and 55% with the index implied volatility levels, respectively. The first two principal components of $IIVol$ moneyness slopes explain 48% and 6% of the cross-sectional variations, respectively. The first common component has a 99% correlation with the average implied idiosyncratic volatility slopes of all firms. The first and the second common components have a correlation of 42% and 8% with the index implied volatility slope, respectively. Finally, the first two principal components of the term structure slopes explain 61% and 7% of the variations, respectively. The first principal component has

[4] This modeling framework was also used by CFJ, except that these authors assumed that $IVol$ is the same under the P- and Q-distributions.

a 99% correlation with the average implied idiosyncratic volatility term structure slopes of all firms. The first and the second common components have a correlation of 78% and -13% with the index implied volatility term structure slope, respectively.

These findings are consistent with those of Herskovic et al. (2013), who show that there is a strong factor structure in idiosyncratic volatilities, similar to that in total volatilities, by looking at firms fundamentals. Our findings complement the empirical literature that shows there are common factors in the idiosyncratic volatilities of stock returns under the P distribution, by documenting the existence of the same factor structure in distributions extracted from equity options prices.

We use the average idiosyncratic variance, AIV, as the potentially priced risk factor, and we test whether this factor can help explain the cross section of equity returns. We show that the AIV factor can reduce the pricing error of the Fama-French 25 size and value portfolios. Moreover, AIV has a positive risk premium, and its cross sectional explanatory power, in our sample period, is more than that of the HML and SMB factors.

Further, we derive the expected option return, and we form portfolios that contain the equity option, the stock, the index option, and the market index. These portfolios are formed and rebalanced in such a way that they are only exposed to the idiosyncratic variance of the equity. The return on these portfolios can be considered as the risk premium for the equity idiosyncratic variance. Using calls and puts with different moneyness ratios, we present evidence of the existence of the idiosyncratic variance risk premium.

Our estimation results show that the idiosyncratic volatility is priced and it can bear a negative or positive sign. Moreover, we define a measure of the idiosyncratic variance risk premium, defined over a 30-day period as the difference between the expected integrated idiosyncratic variance under the P and Q measures that is only partially driven by the market volatility risk premium. Further, the idiosyncratic volatility risk premium is significantly different from zero for all of the stocks in our sample, and has different signs for different stocks.

We show that the market return and the Fama-French factors, as well as the momentum factor cannot explain the time-series variations in the idiosyncratic volatility risk premiums. These time series variations can, however, be partly explained by the market variance risk premium. Further, we show that the average variance risk premium of all firms, $AIVRP$, together with the component of market variance risk premium orthogonal to $AIVRP$, have a strong explanatory power in the time-series variations of the idiosyncratic variance risk premiums.

The rest of the paper is organized as follows. In section 2 we present the model. Section 3 contains the description of the data and the estimation methodology, as well as the results regarding the common structure in idiosyncratic volatilities. In section 4 we discuss the ability of our proposed factor in explaining the cross section of equity returns. Section 5 presents the measure and the properties of the idiosyncratic volatility risk premiums. In section 6 we discuss the estimation of conditional betas. Section 7 concludes.

1.2 The Model

Here we present an equity option valuation model using a single-factor structure that links the equity return dynamics to the market return dynamics. We model an equity market consisting of N stocks and a market index. The individual stock prices are denoted by $S_{i,t}$ for $i = 1, 2, \dots, N$, and the market index price is denoted by S_t. We assume that investors have access to a risk free bond whose return is r.

We assume the following stochastic volatility dynamics for the market index under the physical distributions (hereafter P):

$$\begin{aligned} dS_t/S_t &= (\mu)dt + \sqrt{v_t}dz_t, \\ dv_t &= \kappa(\theta - v_t)dt + \sigma\sqrt{v_t}dw_t \end{aligned} \tag{1.2.1}$$

As in Heston (1993), θ is the unconditional average variance, κ captures the speed of mean reversion of v_t to its long-run average, and σ measures the volatility of variance. The market equity risk premium is represented by μ. The correlation between the shocks to market return and its variance is represented by ρ, and it captures the market leverage effect.

Further, we assume that the stock return follows a one-factor model, where the factor is the excess return on the market. The volatility of the idiosyncratic part of the stock return, referred to as idiosyncratic volatility ($IVol$), is assumed to be stochastic, and to follow a square-root type process. The following describes the stock price dynamics for firm i:

$$\begin{aligned} dS_{i,t}/S_{i,t} &= (\mu_i)dt + \beta_i(dS_t/S_t - rdt) + \sqrt{\xi_{i,t}}dz_{i,t}, \\ d\xi_{i,t} &= \kappa_i(\theta_i - \xi_{i,t})dt + \sigma_i\sqrt{\xi_{i,t}}dw_{i,t} \end{aligned} \tag{1.2.2}$$

where, $dS/S - rdt$ is the instantaneous excess return on the market, β_i is the market beta, μ_i is the idiosyncratic return, ξ_i is the variance of the idiosyncratic return[5] , σ_i is the volatility of the idiosyncratic variance, κ_i is the speed of mean reversion for idiosyncratic volatility, θ_i is the long-run average of the idiosyncratic volatility, and ρ_i is the correlation between the shocks to idiosyncratic return and its variance.

Proposition 1.1. *The market index has the following dynamics under the risk-neutral measure (hereafter Q):*

$$\begin{aligned} dS_t/S_t &= rdt + \sqrt{v_t}d\tilde{z}_t, \\ dv_t &= \tilde{\kappa}(\tilde{\theta} - v_t)dt + \sigma\sqrt{v_t}d\tilde{w}_t \end{aligned} \tag{1.2.3}$$

where, $\tilde{\kappa} = \kappa + \lambda$, and $\tilde{\theta} = \frac{\kappa\theta}{\kappa+\lambda}$, and where λ is the price of market volatility risk as in Heston (1993). Moreover, the equity dynamics under Q is as follows

$$\begin{aligned} dS_{i,t}/S_{i,t} &= rdt + \beta_i(dS_t/S_t - rdt) + \sqrt{\xi_{i,t}}d\tilde{z}_{i,t}, \\ d\xi_{i,t} &= \tilde{\kappa}_i(\tilde{\theta}_i - \xi_{i,t})dt + \sigma_i\sqrt{\xi_{i,t}}d\tilde{w}_{i,t} \end{aligned} \tag{1.2.4}$$

where, $\tilde{\kappa}_i = \kappa_i + \lambda_i$, and $\tilde{\theta}_i = \frac{\kappa_i\theta_i}{\kappa_i+\lambda_i}$, and where λ_i is the price of idiosyncratic volatility risk.

Proof. See Appendix A. □

[5] Here the idiosyncratic return is defined as the excess stock return in a one-factor model framework.

The above model has been used by Christoffersen et al. (2013) who, assume that the idiosyncratic volatility is not priced and that the idiosyncratic variance follows the same dynamics under the P and Q distributions. They discuss the consistency of their model with some of the empirical evidence in the equity option literature such as Duan and Wei (2009) and Dennis and Mayhew (2002). Moreover, they derive a close form solution for the equity option price, and present estimation results based on equity options.

The assumption that the idiosyncratic volatility is not priced is equivalent to implying that the market excess return is the only priced factor. There is, however, significant evidence that there are other priced factors in the economy. If the CAPM is misspecified and there are other priced factors, then the idiosyncratic variance would consist of exposure to those missing factors, and the price of the idiosyncratic variance would reflect the linear combination of the prices of the variance of the missing factors. This is a verifiable hypothesis, which we test by relaxing the assumption of non-priced $IVol$, and letting the idiosyncratic volatility dynamics be different under the P and Q distributions.

Following Heston (1993), we assume that the price of the $IVol$ risk is proportional to the level of the idiosyncratic variance. Based on this assumption, the same closed form solution for the European equity options that Christoffersen et al. (2013) derive holds in our framework.[6] In our empirical work we test the hypothesis of priced idiosyncratic volatility by using information from both equity returns and equity options and verifying whether the idiosyncratic volatility dynamics are different under the physical and risk-neutral distributions.

1.3 Estimation and Results

There are several approaches to estimating stochastic volatility models. The main challenge in estimating stochastic volatility models is the filtering of the unobserved volatility. One approach is to treat the unobserved volatility as a parameter, and estimate all parameters using a single cross section of option prices. This is done by Bakshi et al. (1997). Another approach is to use multiple cross sections of option prices. However, for every cross section,

[6] For the derivation of the option price formula please refer to Christoffersen et al. (2013).

a different initial volatility estimate is required. Bates (2000) and Huang and Wu (2004) use this approach. A third approach provides a likelihood-based estimation that can combine the information from the option data and the underlying returns, and imposes consistency between the P and Q distributions. Ait-Sahalia and Kimmel (2007), Eraker (2004), Jones (2003), and Bates (2006) provide an MCMC analysis within this framework. A last group of papers takes a frequentist approach that can also combine the information from the option prices and the underlying returns. Chernov and Ghysels (2000) use the efficient method of moments, while Pan (2002) uses a method of moments technique. Santa-Clara and Yan (2010) and Christoffersen et al. (2013) use likelihood functions that contain a return component and an option component. Our empirical setup is most closely related to this last group of papers.

1.3.1 Data

We collect option data for the S&P 500 and for 27 equities, all components of the Dow Jones index. We did not include in our sample Bank of America, Kraft Food, and Travellers because of data unavailability. The option data that we use comes from the OptionMetrics volatility surface, which is based on the bid-ask midpoint. Our data spans the period from January 4, 1996, to December 29, 2011. We focus on options with maturity of up to six months. Since our estimation is computationally very demanding, we excluded options with longer maturities to keep the estimation manageable. Moreover, our data contains out-of-the-money options with moneyness[7] less than 1.1 for calls and greater than 0.9 for puts. We filter out options with implied volatility less than 5% and greater than 150%, and options that violate the apparent arbitrage conditions as described in Bakshi et al. (2003).

We also collect data for the index levels, daily returns, and dividend yield, as well as stock prices, returns, and cash dividends from CRSP. The implied volatility surface data is calculated using binomial trees. When evaluating the option model price, for every option on every stock on every day we deduct the present value of dividends, which is assumed to be known during the life of the option, from the stock price on that day, and we treat the

[7] Moneyness is defined as the strike price divided by the underlying asset's price.

option as European. The discounting is done using the appropriate interest rate estimated by linear interpolation of the Zero Coupon Yield Curve available from OptionMetrics. We do the same discounting for the index using the index dividend yield.

Table 1.1 presents the names of the companies in our sample as well as the number of calls, puts, and total options for each firm, and for the market index. The number of option contracts is highest for the S&P500 index. On average there are $120,811$ options available for each firm, with Cisco and Johnson & Johnson having the lowest and highest number of contracts among all firms, respectively. Our estimation is based on a total of $3,430,176$ option quotes for the market and all equities.

In Tables 1.2 and 1.3 we report the sample average implied volatility, minimum and maximum implied volatilities, along with average option delta, option vega[8], and average days-to-maturity of all calls and puts, separately. The average implied volatility of the market in our sample is 19.2%. Cisco and Johnson & Johnson have the maximum and minimum average implied volatility in our sample. Moreover, the average days-to-maturity is close to 80 days for all firms and the market index.

1.3.2 Joint Estimation

In order to capture the difference between the physical and risk neutral distributions of the equity idiosyncratic volatilities, it is required to fit both distributions using the same internally consistent set of structural parameters. We do so by using a joint likelihood function that has two components, one based on returns and one based on options. Since the market variance and equity idiosyncratic variance are unobserved, we filter these state variables by using the Particle Filter (PF) method. The PF methodology offers a convenient filter for nonlinear models such as the stochastic volatility models and is used extensively in engineering, with some applications in finance.[9]

Our estimation consists of two steps. First, we estimate the market's parameters and the

[8] These are Black-Scholes vegas evaluated at the implied volatilities.

[9] For other studies that use the PF method see Christoffersen et al. (2010), Johannes et al. (2009) and Gordon et al. (1993).

filtered spot market variances. Then conditional on the market model's parameters and the spot market variances, we estimate each equity's parameters and the spot idiosyncratic variances. In what follows we describe the detailed estimation procedure.

Market Model

Here we describe the estimation of the market model, presented in (1.2.1) and (1.2.3). We describe how the return-based and the option-based likelihood functions are calculated, and finally how the parameters and the spot variances are estimated. Applying Ito's lemma to (1.2.1) we get:

$$\begin{aligned} dln(S_t) &= (\mu - \frac{1}{2}v_t)dt + \sqrt{v_t}dz_t, \\ dv_t &= \kappa(\theta - v_t)dt + \sigma\sqrt{v_t}dw_t \end{aligned} \tag{1.3.1}$$

The above equation shows how the unobserved volatility states are related to the observed index prices. This relationship allows the filtering of the market spot volatilities from the returns. First we discretize the model in (1.3.1). We use the Euler scheme to get:

$$\begin{aligned} ln(S_{t+\Delta t}) - ln(S_t) &= (\mu - \frac{1}{2}v_t)\Delta t + \sqrt{v_t}z_{t+\Delta t}, \\ v_{t+\Delta t} &= v_t + \kappa(\theta - v_t)\Delta t + \sigma\sqrt{v_t}w_{t+\Delta t} \end{aligned} \tag{1.3.2}$$

where, $z_{t+\Delta t}$ and $w_{t+\Delta t}$ are normal random variables with mean zero and variance Δt. We implement the discretized model in (1.3.2) using daily index log-returns, but all the results are expressed in annual terms.

The PF method approximates the true density of the variance state $v_{t+\Delta t}$ by a set of "particles" that are updated through the equations in (1.3.2). At any time $t + \Delta t$ we generate N particles $\{v^j_{t+\Delta t}\}^N_{j=1}$ from the empirical distribution of $v_{t+\Delta t}$, conditional on N particles $\{v^j_t\}^N_{j=1}$ from the empirical distributions of v_t. This particular implementation of the PF is referred to as the sampling-importance-resampling (SIR) PF and follows Pitt (2002).[10]

[10] We refer to Christoffersen et al. (2010) and Pitt (2002) for a more detailed description of the PF algorithm.

Starting from a vector of particles $v_1^j = \theta \, \forall j$, on every day t we simulate a new set of particles,$\{\tilde{v}_{t+\Delta t}^j\}$ from the set of smoothly resampled particles $\{v_t^j\}$, according to (1.3.2), as follows:

$$\begin{aligned} z_{t+\Delta t}^j &= (ln(S_{t+\Delta t}/S_t) - (\mu - \frac{1}{2}v_t^j))/\sqrt{v_t^j} \\ w_{t+\Delta t}^j &= \rho z_{t+\Delta t}^j + \sqrt{1-\rho^2}\epsilon_{t+\Delta t}^j \end{aligned} \tag{1.3.3}$$

where $\epsilon_{t+\Delta t}^j$ are independent normal random variables with mean zero and variance Δt. Replacing (1.3.3) into the variance dynamics in (1.3.2), we get a simulated set of particles.

$$\tilde{v}_{t+\Delta t}^j = v_t^j + \kappa(\theta - v_t^j)\Delta t + \sigma\sqrt{v_t^j} w_{t+\Delta t}^j \tag{1.3.4}$$

So far we have N possible values for $v_{t+\Delta t}$. Now we want to give weights to the simulated particles. The weight for every particle, $\tilde{W}_{t+\Delta t}^j$, is the likelihood that the next day return at $t + 2\Delta t$ is generated by this particle.

$$\tilde{W}_{t+\Delta t}^j = \frac{1}{\sqrt{2\pi \tilde{v}_{t+\Delta t}^j \Delta t}} \cdot exp\left(-\frac{1}{2} \frac{\left(ln(\frac{S_{t+2\Delta t}}{S_{t+\Delta t}}) - (\mu - \frac{1}{2}\tilde{v}_{t+\Delta t}^j)\Delta t\right)^2}{\tilde{v}_{t+\Delta t}^j \Delta t} \right) \tag{1.3.5}$$

We can then get the probability of each particle by normalizing the weights:

$$\breve{W}_{t+\Delta t}^j = \frac{\tilde{W}_{t+\Delta t}^j}{\sum_{j=1}^N \tilde{W}_{t+\Delta t}^j} \tag{1.3.6}$$

This procedure provides us with the set of the raw particles and the associated probabilities, from which we can apply Pitt (2002) algorithm to get the empirical distribution of smoothly resampled particles. These particles can be used to simulate the next period particles until we have the empirical distributions of variances for each day in the sample.

The return-based likelihood function, which is a function of the P-distribution parameters $\Theta \equiv \{\mu, \kappa, \theta, \sigma, \rho\}$, can be defined as follows:

$$\ln L^R \propto \sum_{t=1}^{T} ln\Big(\frac{1}{N}\sum_{j=1}^{N}\tilde{W}_t^j(\Theta)\Big) \tag{1.3.7}$$

The P-measure filtered spot variance v_t^P would be the average of the smoothly resampled particles.

$$\hat{v}_t^P = \frac{1}{N}\sum_{j}^{N} v_t^j \tag{1.3.8}$$

Moreover, the filtered shocks to the index return, conditional on the filtered spot variance would be:

$$\hat{z}_{t+\Delta t}^P = (ln(S_{t+\Delta t}/S_t) - (\mu - \frac{1}{2}\hat{v}_t^P))/\sqrt{\hat{v}_t^P} \tag{1.3.9}$$

For the market model under the Q-distribution we need to estimate the vector of spot variances $\{v_t\}$, and a set of structural parameters $\tilde{\Theta} \equiv \{\kappa, \theta, \lambda, \sigma, \rho\}$. These parameters completely identify the data generating process under the risk-neutral measure. The unobserved spot variance under the Q-measure is filtered from the returns using the PF method as described before, but this time based on the mapped structural parameters $\{\tilde{\kappa}, \tilde{\theta}, \sigma, \rho\}$, where $\tilde{\kappa} = \kappa + \lambda$, and $\tilde{\theta} = \frac{\kappa\theta}{\kappa+\lambda}$. After repeating the same procedure as described before, the Q-measure spot variance on every day can be estimated as the average of the smoothly resampled particles.

$$\hat{v}_t^Q = \frac{1}{N}\sum_{j}^{N} v_t^j \tag{1.3.10}$$

Similarly, the Q-measure filtered shocks to the index return, conditional on the filtered spot variance would be:

$$\hat{z}_{t+\Delta t}^Q = (ln(S_{t+\Delta t}/S_t) - (\mu - \frac{1}{2}\hat{v}_t^Q))/\sqrt{\hat{v}_t^Q} \tag{1.3.11}$$

Now define the vega weighted option pricing error of an option n as:

$$\eta_n = (C_n^O - C_n^M(\tilde{\Theta}, \hat{v}^Q))/Vega_n, \quad n = 1, \ldots, M \tag{1.3.12}$$

where, C_n^O is the observed price of index option n, $C_n^M(\tilde{\Theta}, \hat{v}^Q))$ is the model price for the same option,[11] M is the total number of index options, and $Vega_n$ is the Black-Scholes option vega evaluated at the implied volatility. These vega weighted option pricing errors serve as an approximation to the implied volatility errors, and since they do not require a numerical inversion of the Black-Scholes model, they are very helpful in large scale optimization problems such as ours. Assuming that these disturbances are i.i.d. normal, the option-based likelihood can be obtained as follows:

$$\ln L^O \propto -\frac{1}{2}\Big(M \ln(2\pi) + \sum_{n=1}^{M} \big(\ln(s^2) + \eta_n^2/s^2\big)\Big) \tag{1.3.13}$$

where we can replace s^2 by its sample analog $\hat{s}^2 = \frac{1}{M}\sum_{n=1}^{N} \eta_n^2$. The set of structural parameters $\hat{\Theta}$ and $\hat{\tilde{\Theta}}$ can be found as the solution to the following optimization problem:

$$\max_{\Theta, \tilde{\Theta}} \ln L^R + \ln L^O \tag{1.3.14}$$

Equity Model

We estimate the equity model parameters and the spot idiosyncratic variances for every stock, conditional on the filtered market spot variances and the filtered shocks to the index returns. We want to estimate the set of structural parameters $\Theta_i \equiv \{\mu_i, \kappa_i, \theta_i, \sigma_i, \rho_i, \beta_i\}$ of the model in (1.2.2), as well as the vector of spot idiosyncratic variances $\{\xi_{i,t}\}$. The Euler discretization of (1.2.2) yields:

[11] The time subscript is dropped for compactness.

$$\begin{aligned}\frac{(S^i_{t+\Delta t}-S^i_t)}{S^i_t} &= (\mu_i)\Delta t+\beta_i\Big((\mu-r)\Delta t+\sqrt{v_t}z_{t+\Delta t}\Big)+\sqrt{\xi_{i,t}}z_{i,t+\Delta t}\\ \xi_{i,t+\Delta t} &= \xi_{i,t}+\kappa_i(\theta_i-\xi_{i,t})\Delta t+\sigma_i\sqrt{\xi_{i,t}}w_{i,t+\Delta t}\end{aligned} \tag{1.3.15}$$

As in the case of the market model, for a set of smoothly resampled particles $\{\xi^j_{i,t}\}$ at time t, the P-measure shocks to stock returns $\{z^j_{i,t+\Delta t}\}$ can be obtained conditional on the filtered shocks to market return, $\{\hat{z}^P_{t+\Delta t}\}$.

We then generate correlated shocks to idiosyncratic variance dynamics:

$$w^j_{i,t+\Delta t}=\rho_i z^j_{i,t+\Delta t}+\sqrt{1-\rho_i^2}\epsilon^j_{i,t+\Delta t} \tag{1.3.16}$$

where, $\epsilon^j_{i,t+\Delta t}$ are independent random variables with mean zero and variance Δt. We can now simulate a raw set of particles, $\{\tilde{\xi}^j_{i,t+\Delta t}\}$ according to equations (1.3.10), given the set of smoothly resampled particles $\{\xi^j_{i,t}\}$.

$$\tilde{\xi}^j_{i,t+\Delta t}=\xi^j_{i,t}+\kappa_i(\theta_i-\xi^j_{i,t})\Delta t+\sigma_i\sqrt{\xi^j_{i,t}}w^j_{i,t+\Delta t} \tag{1.3.17}$$

We then have a set of N possible values for $\xi_{i,t+\Delta t}$ to which we want to assign weights. The weight for every particle would be the likelihood that the next day stock return is generated by this particle, given that the next day's index return shock is revealed first.

$$\tilde{W}^j_{i,t+\Delta t}=\frac{1}{\sqrt{2\pi M^j_1}}\cdot exp\left(-\frac{1}{2}\frac{\Big((\frac{S_{i,t+2\Delta t}-S_{i,t+\Delta t}}{S_{i,t+\Delta t}})-M^j_2\Big)^2}{M^j_2}\right) \tag{1.3.18}$$

where M^i_1 and M^j_2 are the conditional mean and variance of the stock return at $t+2\Delta t$.

$$
\begin{aligned}
M^j_{1,t+\Delta t} &= E\left[\Big(\frac{S_{i,t+2\Delta t}-S_{i,t+\Delta t}}{S_{i,t+\Delta t}}\Big)\Big|S_{i,t+\Delta t},\xi^j_{i,t+\Delta t},\hat{v}_{t+\Delta t};\hat{z}_{t+2\Delta t}\right]\\
&= \mu_i\Delta t+\beta_i(\mu-r)\Delta t\\
M^j_{2,t+\Delta t} &= var\left[\Big(\frac{S_{i,t+2\Delta t}-S_{i,t+\Delta t}}{S_{i,t+\Delta t}}\Big)\Big|S_{i,t+\Delta t},\xi^j_{i,t+\Delta t},\hat{v}_{t+\Delta t};\hat{z}_{t+2\Delta t}\right]\\
&= (\beta_i^2\hat{v}_{t+\Delta t}+\xi^j_{i,t+\Delta t})\Delta t
\end{aligned}
\tag{1.3.19}
$$

After normalizing the weights $\tilde{W}^j_{i,t+\Delta t}$ we would have the empirical distribution of the $\xi_{i,t+\Delta t}$, from which we can smoothly resample the next period's particles. We start from $\xi^j_{i,1}=\theta_i\,\forall j$, and repeat this procedure for every day in the sample.

The return-based likelihood functions, which is a function of the P-distribution parameters $\Theta_i\equiv\{\mu_i,\kappa_i,\theta_i,\sigma_i,\rho_i,\beta_i\}$, can be defined as follows:

$$
\ln L_i^R \propto \sum_{t=1}^{T} ln\Big(\frac{1}{N}\sum_{j=1}^{N}\tilde{W}^j_{i,t}(\Theta_i)\Big) \tag{1.3.20}
$$

Moreover, the vector of P-measure filtered spot idiosyncratic variances can be obtained as follows:

$$
\hat{\xi}^P_{i,t}=\frac{1}{N}\sum_{j}^{N}\xi^j_{i,t} \tag{1.3.21}
$$

For the equity model under the Q-distribution (1.2.4) we need to estimate the vector of spot variances $\{\xi_{i,t}\}$, and a set of structural parameters $\tilde{\Theta}_i\equiv\{\kappa_i,\theta_i,\lambda_i,\sigma_i,\rho_i,\beta_i\}$. The unobserved spot idiosyncratic variance under the Q-measure is filtered from the returns using the PF method as described before, based on the mapped structural parameters $\{\tilde{\kappa}_i,\tilde{\theta}_i,\sigma_i,\rho_i,\beta_i\}$, where $\tilde{\kappa}_i=\kappa_i+\lambda_i$, and $\tilde{\theta}_i=\frac{\kappa_i\theta_i}{\kappa_i+\lambda_i}$. After repeating the same procedure as described before, the Q-measure spot idiosyncratic variance on every day can be estimated as the average of the smoothly resampled particles.

$$\hat{\xi}_{i,t}^{Q} = \frac{1}{N}\sum_{j}^{N} \xi_{i,t}^{j} \tag{1.3.22}$$

Given now the market structural parameters $\hat{\tilde{\Theta}}$, estimated market spot variances $\{\hat{v}_t\}$, filtered shocks to the market returns $\{\hat{z}_t^Q\}$, and the estimated Q-measure spot idiosyncratic variances $\hat{\xi}_{i,t}^{Q}$, we can find the option pricing error for every option as a function of the Q-measure equity structural parameters $\tilde{\Theta}_i$. As with the market model, for an option n written on stock i, we define the vega weighted option pricing error as:

$$\eta_{i,n} = (C_{i,n}^{O} - C_{i,n}^{M}(\tilde{\Theta}_i, \hat{\tilde{\Theta}}, \hat{v}^{Q}, \hat{\xi}_i^{Q}))/Vega_{i,n}, \quad n = 1, \ldots, M_i \tag{1.3.23}$$

where, $C_{i,n}^{O}$ is the observed price of equity option n for stock i, $C_{i,n}^{M}$ is the model price for the same option, M_i is the total number of options available for stock i, and $Vega_{i,n}$ is the Black-Scholes option vega evaluated at the implied volatility. Assuming that these disturbances are i.i.d. normal, the option-based likelihood can be obtained as follows:

$$\ln L_i^{O} \propto -\frac{1}{2}\Big(M_i \ln(2\pi) + \sum_{n=1}^{M_i}\big(\ln(s_i^2) + \eta_{i,n}^2/s_i^2\big)\Big) \tag{1.3.24}$$

where we can replace the s_i^2 by its sample analog $\hat{s}_i^2 = \frac{1}{M_i}\sum_{n=1}^{M_i}\eta_{i,n}^2$. The set of equity structural parameters $\hat{\Theta}_i$ and $\hat{\tilde{\Theta}}_i$ can be found as the solution to the following optimization problem:

$$\max_{\Theta_i, \tilde{\Theta}_i} \ln L_i^{R} + \ln L_i^{O} \tag{1.3.25}$$

1.3.3 Estimation Results

In this section we report the parameters estimates for the S&P 500 index and the 27 equities in our sample for the period $1996 - 2011$. As stated previously, in our estimation we use information from the returns as well as the option prices. For both the index and the equities we use option contracts on each trading day. The index options are European, but the equity options are American. Since the closed form equity option pricing formula in our framework is only available for European options, we eliminate in-the-money options from the sample to avoid biases due to the early premium exercise of American options.[12]

Parameter Estimates

Table 1.4 presents the estimated structural parameters for the market and the equity models. In our joint estimation we restrict the P- and Q-measure parameters to be consistent. Hence only the unconditional mean and the speed of mean reversion of the volatility dynamics would be different between the two measures due to the prices of the market and the idiosyncratic volatility risks. Consistent with the previous studies of the market index, we find the price of the volatility risk to be negative, $\lambda = -1.21$. The unconditional mean of the market variance under P and Q are $\theta = 0.037$ and $\tilde{\theta} = 0.061$. Moreover, the speeds of mean reversion of the market volatility dynamics are $\kappa = 3.157$ and $\tilde{\kappa} = 1.94$ under the P and Q dynamics, respectively. These parameter estimates are consistent with those of other studies such as CFJ.

In our estimations we set the market equity risk premium equal to the annual sample average $\mu = 0.078$. Moreover, instead of estimating μ_i for each firm, we run a regression of the equity return on the market excess return, and we set μ_i equal to the OLS alpha of the stock.

The unconditional idiosyncratic variance means are mostly larger than that of the market. Under the physical distribution they range from $\theta_i = 0.034$ for MMM to $\theta_i = 0.146$ for HPQ. The average θ_i for the firms in our sample is 0.068. On the other hand, the speed of mean reversion of equity idiosyncratic variances is much lower than that of the index for all stocks in the sample. It ranges from $\kappa_i = 0.138$ for IBM to $\kappa_i = 1.678$ for AA, with an average of 0.701 for all stocks. The price of idiosyncratic variance risk varies substantially among the firms in our sample, and is of different signs for different stocks. Alcoa has the lowest price of idiosyncratic variance risk $\lambda_i = -0.872$, and XOM has the largest with $\lambda_i = 1.145$. The

[12] Bakshi et al. (2003) show that the difference between Black-Scholes implied volatilities and American option implied volatilities are negligible for out-of-the-money calls and puts.

average absolute value of the price of idiosyncratic variance is 0.213.

Consistent with the literature, $\rho = -0.494$ is large and negative, capturing the so-called leverage effect. The correlation between the shocks to equity return and the shocks to idiosyncratic variance is negative for all stocks except for CSCO, HPQ, IBM, MSFT, and UTX. It ranges from $\rho = -0.649$ for JPM to $\rho = 0.1$ for CSCO. Moreover, the beta estimates seems reasonable, ranging from $\beta_i = 0.55$ for MCD to $\beta_i = 1.23$ for AXP. The average beta of the firms in our sample is 0.91.

Filtered Spot Idiosyncratic Variances

As described before, in our estimation we filter the unobserved market variance and the equity idiosyncratic variances from the returns, under both P and Q measures. Table 1.5 presents the average, standard deviation, minimum, and maximum spot idiosyncratic variance of all firms during the time period in our sample. In the top row we also report the same statistics for the spot index variance.

In Table 1.6 we present the correlation matrix of the spot idiosyncratic variances, as well as the market spot variance. We can see that there is high degree of correlation between the spot idiosyncratic variances, and all pairwise correlations are positive. The average pairwise correlation between the equity spot idiosyncratic variances is 58%, and the average pairwise correlation of equity spot idiosyncratic variance with market spot variance is 43%. These results show that there is a common structure in idiosyncratic variance levels of equities, and the common factors might be priced. In the following sections, we conduct a systematic analysis of the common structure in equity idiosyncratic variances.

1.3.4 The Common Structure in Idiosyncratic Volatilities

Several studies have found a common structure in idiosyncratic volatilities[13]. All of these studies, however, have focused on the idiosyncratic volatilities estimated from the returns

[13] See Duarte et al. (2012) and the references within.

under the physical distribution. Here we impose consistency in the parameters of the P- and Q-distributions, we verify whether the same common structure exists in the idiosyncratic volatilities estimated from the equity option prices under the Q distribution, and we compare the common factors under both P- and Q-measures.

Implied Idiosyncratic Volatilities

Similar to the implied volatility of an option, we define the implied idiosyncratic volatility ($IIVol$) of an option as the idiosyncratic volatility that would make the model option price equal to the observed price, given the estimated parameters and the estimated spot market volatility. The $IIVol$ can be found as the solution to the following equation.

$$C^O_{i,t,n} - C^M_{i,t,n}(\hat{\tilde{\Theta}}_i, \hat{\tilde{\Theta}}, \hat{v}^Q_t, IIVol_{i,t,n}) = 0 \tag{1.3.26}$$

where, as before, $\hat{\tilde{\Theta}}$ and $\hat{\tilde{\Theta}}_i$ are the estimated Q-measure parameters of the market and equity models, respectively, and $\hat{v}^Q$ is the estimated spot market variance under the risk-neutral measure. For every option written on every stock in our sample we find the $IIVol$ as described above, and we run the following two regressions on every day, one for the implied idiosyncratic volatilities and one for the total implied volatilities. [14]

$$IIVol_{i,n,t} = a^{IVol}_{i,t} + b^{IVol}_{i,t} \cdot (\frac{S_{i,t}}{K_{i,n}}) + c^{IVol}_{i,t} \cdot (DTM_{i,n}) + \epsilon_{i,n,t} \tag{1.3.27}$$

$$IV_{i,n,t} = a^{IV}_{i,t} + b^{IV}_{i,t} \cdot (\frac{S_{i,t}}{K_{i,n}}) + c^{IV}_{i,t} \cdot (DTM_{i,n}) + \epsilon_{i,n,t} \tag{1.3.28}$$

Where, i denotes the stock, and n denotes the option contract available for that stock on day t, with strike price $K_{i,n}$ and time to maturity $DTM_{i,n}$. The coefficients, $a^{IVol}_{i,t}, b^{IVol}_{i,t}, c^{IVol}_{i,t}$ in the regression (1.3.27) represent measures of idiosyncratic volatility level, moneyness slope,

[14] Duan and Wei (2009) and Christoffersen et al. (2013) use similar regressions for Black-Scholes implied volatilities.

and term structure slope, respectively. Moreover, the coefficients, $a_{i,t}^{IV}, b_{i,t}^{IV}, c_{i,t}^{IV}$ in the regression (1.3.28) represent measures of implied volatility level, moneyness slope, and term structure slope, respectively. After estimating the regression coefficients for every stock-day we have three matrices, $\{a_{i,t}^{IVol}\}$, $\{b_{i,t}^{IVol}\}$, and $\{c_{i,t}^{IVol}\}$, for the implied idiosyncratic volatilities, and three matrices, $\{a_{i,t}^{IV}\}$, $\{b_{i,t}^{IV}\}$, and $\{c_{i,t}^{IV}\}$, for the implied volatilities.

We run a similar regression as in (1.3.28) for the index options to get the level, slope, and the term structure slope of the market implied volatilities, represented by $\{a_t\}$, $\{b_t\}$, and $\{c_t\}$, respectively.

$$IV_{n,t} = a_t + b_t \cdot (\frac{S_t}{K_n}) + c_t \cdot (DTM_n) + \epsilon_{n,t} \tag{1.3.29}$$

In what follows, we present a principal component analysis (PCA) of the regression coefficients estimated in (1.3.27) - (1.3.29).

Common Structure in Levels

As expected, the PCA of the implied volatility levels indicates a strong factor structure. The first two principal components explain 79% and 11% of the variations across all stocks, respectively, and the first component has a 90% correlation with the market's implied volatility levels. What is surprising is that after accounting for the common market factor in equity returns, there is still a strong factor structure remaining in the implied idiosyncratic volatility levels. Principal component analysis of the $IIVol$ levels shows that the first two principal components explain 58% and 23% of the cross-sectional variations, respectively. Table 1.7 presents the loadings on the first two components, as well as the percentage of variation captured by each component. The loadings on the first factor are positive for all stocks, while the loadings on the second factor are positive and negative for different stocks. Moreover, both factors are sizeable in terms of explaining the variation, and the fact that the loadings on the second factor take on different signs for different stocks, suggest that it may be related to firm specific characteristics.

The first common component of the $IIVol$ levels has a 99% correlation with the average implied idiosyncratic volatility level of all firms. Further, the first and the second common components have correlations of 65% and 55% with the index implied volatility levels, respectively.

Common Structure in Moneyness Slope

The first two principal components of the implied volatility slopes explain 32% and 8% of the cross-sectional variation, and the first common component has a 38% correlation with the market's implied volatility slope. After removing the common market factor from the returns, there are still commonalities in the moneyness slopes of the implied idiosyncratic volatilities. PCA of the $IIVol$ slopes shows that the first two principal components explain 48% and 6% of the cross-sectional variations, respectively. So the percentage of the variation explained by the first common component of slopes is higher for $IIVol$ than for IV. Table 1.8 presents the loadings on the first two components, as well as the percentage of variation captured by each component. Similar to those of the levels, the loadings on the first factor are positive for all stocks, while the loadings on the second factor are positive and negative for different stocks, although the signs are not consistent with those of the levels.

The first common component of the $IIVol$ moneyness slopes has a 99% correlation with the average implied idiosyncratic volatility slopes of all firms. Further, the first and the second common components have a correlation of 42% and 8% with the index implied volatility slope, respectively.

Common Structure in Term Structure Slope

The first two principal components of term slope of the implied volatilities explain 57% and 9% of the variations, respectively, and the first principal component has a 78% correlation with the market's term slope. Moreover, the two first principal components of the $IIVol$ term sloes explain 61% and 7% of the variations, respectively. As with the results for the moneyness slope, the proportion of variation explained by the first two principal components of the term slopes are higher for $IIVol$ than IV. Table 1.9 presents the loadings on the first two components, as well as the percentage of variation captured by each component. As before, the loadings on the first common component are all positive, while the loadings on the second components have different signs for different stocks.

Moreover, the first principal component has a 99% correlation with the average implied idiosyncratic volatility term slope of all firms. Further, the first and the second common components have a correlation of 78% and -13% with the index implied volatility term slope, respectively.

These results show that after removing the common market factor from the returns, there is still a strong common factor structure in implied idiosyncratic volatilities. This is consistent with the findings of Herskovic et al. (2013) who study the common structure in idiosyncratic volatilities by looking at the firms' fundamentals. Our findings complement the empirical literature that shows there are common factors in the idiosyncratic volatilities of stock returns under the P distribution, by showing that the same factor structure is evident in equity options prices.

Filtered Spot Idiosyncratic Volatilities

Our results indicate the existence of a common structure in the equity idiosyncratic volatilities obtained from the equity option prices. While there appears to be a market volatility factor in the cross-section of equity idiosyncratic volatilities, the average idiosyncratic volatility of all firms shows stronger correlation with the common components of the implied idiosyncratic volatilities. Moreover, the existence of the common structure is strongest in the levels of the implied idiosyncratic volatilities. Here, we further investigate the common structure in idiosyncratic volatilities by analyzing the spot idiosyncratic volatility levels filtered from the returns.

In our framework, idiosyncratic variance of a stock is a state variable. In section 1.3.2 we discussed the filtration of this unobserved state variable from the equity returns using the physical and risk-neutral structural parameters, and we presented the properties of these estimated spot idiosyncratic volatilities. Here we perform principal component analysis of the filtered idiosyncratic volatilities under the P and Q distributions. These idiosyncratic variances are estimated from the returns, based on the P- and Q-measure parameters. The spot idiosyncratic volatilities are theoretically the same under the two distributions; however over any discrete interval, such as a day, they would be different due to the price of the idiosyncratic volatility risk.

Principal component analysis of the spot idiosyncratic volatilities under the physical distribution indicates that the first two principal components explain 57% and 30% of the cross-sectional variations, respectively. Table 1.10 presents the loadings on the first two components for all stocks. The loadings on the first principal component are positive for all stocks, while the loadings on the second principal component have different signs for different stocks. The first principal component has a 98% correlation with the average idiosyncratic volatility. Moreover, the first and the second principal components have correlations of 68% and 48% with the spot market volatilities, respectively.

Principal component analysis of the return-based idiosyncratic volatilities under the Q-distribution yields qualitatively similar results to those under P. The first two principal components explain 55% and 31% of the variations, respectively. The first principal component has a 98% correlation with the average idiosyncratic volatility. Moreover, the first and the second components have a correlation of 67% and -50% with the market spot volatility under the risk-neutral distribution. Table 1.10 presents the loading on the first two components. The loadings on the first principal component are very close for all stocks under the two distributions. On the one hand, the loadings on the second principal component, while very close in absolute value, are of the exact opposite sign under the P- and Q-distributions. In other words, the second principal component under the P measure is almost perfectly and negatively correlated with that under the Q measure. This estimation is very similar to the one done in the previous section, and it is not surprising that the results are very similar. Indeed, the data used for the PCA here is the same as was used to generate the model parameters, which in turn were used to generate the data for the PCA of the previous section. The interest of the last section is that it shows that the commonality exists across slope and term structure.

The average idiosyncratic volatility seems to explain the cross-sectional variation of the idiosyncratic volatilities very well, and it is highly correlated with the market volatility, which is also highly correlated with the first two components of the idiosyncratic volatilities. We regress the market spot variance on the average idiosyncratic variances to find the component of the market variance that is orthogonal to the average idiosyncratic variance. The R^2 of the regression is 30% and 26%, under the P and Q distributions, respectively. Moreover, the vector of the residuals, which is the component of the market variance that is orthogonal to the average idiosyncratic variance, denoted by F_P^{orth}, has a 69% correlation with the second principal component of the idiosyncratic variances. Under the Q distribution, the orthogonal component, denoted by F_Q^{orth}, has a -71% correlation with the second principal component. Our results suggest that the average idiosyncratic variance and the component of the market spot variance that is orthogonal to it are good proxies for the common principal components of the equity idiosyncratic volatility levels.

1.4 Cross Section of Equity Returns

The previous sections show that there is a common factor structure in the idiosyncratic volatilities. The existence of common factors among equity idiosyncratic volatilities does not necessarily mean that idiosyncratic volatility is priced. If however, there are one or more systematic risk factors missing from the model, the exposure to those factors would be captured by the idiosyncratic volatilities, and the common components of idiosyncratic volatilities are related to the variances of the missing systematic factors. In our one factor model, the equity return is represented as follows:

$$r_i = \mu_i + \beta_i(r_m - r_f) + \epsilon_i \tag{1.4.1}$$

If the market excess return is not the only systematic risk factor and there are K factors, $F_1, \ldots, F_K$, missing from the model, then the residuals in (1.4.1) are in fact equal to:

$$\epsilon_i = \beta_{i,1}F_1 + \ldots + \beta_{i,K}F_K + u_i \tag{1.4.2}$$

where, u_i is the true idiosyncratic residual. So, the idiosyncratic variance, as defined in a one factor model, is related to the variance of the missing factors and their corresponding loadings, as follows:

$$var(\epsilon_i) = \beta_{i,1}^2 \cdot var(F_1) + \ldots + \beta_{i,K}^2 \cdot var(F_K) + var(u_i) \qquad (1.4.3)$$

Moreover, if idiosyncratic volatility is priced due to exposure to the missing factors, we would expect the common components of the equity idiosyncratic volatilities to help explain the cross-section of equity returns. This is what we investigate in this section. In particular, we test whether the average idiosyncratic variance, AIV, that we found to proxy for the first principal component of the idiosyncratic variances has any explanatory power in explaining the cross section of the 25 Fama-French portfolios, which are formed on size and book-to-market. We follow the standard two-step procedure for cross-sectional asset pricing. First, we run a time series regression of the excess portfolio returns on the Fama-French factors, as well as the proposed AIV factor for each portfolio.

$$r_{p,t}^e = a_p + b_p^m \cdot (r_t^m - r_t) + b_p^{smb} \cdot r_t^{SMB} + b_p^{hml} \cdot r_t^{HML} + b_p^{AIV} \cdot AIV_t + \varepsilon_{p,t} \qquad (1.4.4)$$

where, $r_{p,t}^e$ is the excess return of portfolio p at time t, $(r^m - r)$ is the excess market return, and r^{SMB} and r^{HML} are the returns on the size (small-minus-big) and value (high-minus-low) factors, respectively.[15] Moreover, AIV is the average idiosyncratic variance of the firms in our sample, and is estimated as described in the previous sections. Figure 1.1 presents the coefficients of the AIV factor in the time series regressions (1.4.4). The coefficients of AIV are all positive, and even in the presence of the Fama-Franch factors, are significantly different from zero for 20 out of the 25 portfolios.

In the second step, we regress the average portfolio excess returns on the time series coefficient estimates:

[15] Data on these portfolios and factors are downloaded from Kenneth French's website.

$$\bar{r}_p^e = \gamma_0 + \gamma_m \cdot \hat{b}_p^m + \gamma_{smb} \cdot \hat{b}_p^{smb} + \gamma_{hml} \cdot \hat{b}_p^{hml} + \gamma_{AIV} \cdot \hat{b}_p^{AIV} + \varepsilon_p \tag{1.4.5}$$

In order to compare the explanatory power of AIV to that of the Fama-French factors, we use different linear combinations of the factors in (1.4.5). In Table 1.11 we present the regression results. As we see in the table, the AIV factor premium is significantly and economically different from zero. Moreover, when added to the Fama-French model, the AIV factor reduces the pricing error from 0.137% per day to 0.11%, and it increases the adjusted R^2 from 82% to 84%. Further, the combination of the market factor and the AIV factor performs better than the combination of market and SMB, as well as the market and HML, in terms of explaining the cross sectional variations of the 25 portfolios' returns. It is worth noting that our sample contains only 27 equities, and our definition of the average idiosyncratic variance is rather narrow. We would expect that as the sample size gets larger, the AIV factor would be able to better explain the cross section of equity returns.

In our sample the risk premium of SMB is not significantly different from zero in any of the regressions. However, HML has a positive and significant risk premium. The AIV also carries a positive risk premium equal to 0.01% per day. Moreover, the fact that AIV does not drive out the HML in the regression where all factors are present, combined with the fact that SMB has an insignificant risk premium in all regressions, suggests that the explanatory power of AIV comes from exposure to missing systematic risk factors other than SMB and HML. The cross sectional asset pricing test above is based on a single cross section. As a robustness check we also perform a Fama-MacBeth cross sectional test on daily returns. The estimated risk premiums and the t-stats are presented in 1.12.

The proposed AIV factor in our study resembles to some extent the residual market factor in the APT literature. The residual market factor is obtained as the residuals in the regression of the market factor on known factors, and it captures the exposure to the missing factors. In our case however, the proposed factor is constructed as the linear combination of the variances of the missing factors, and it is based on the combined information from the returns and option prices.

1.5 Idiosyncratic Variance Risk Premium

So far we have shown that equity idiosyncratic variance is priced, and that there is factor structure among idiosyncratic variances. Moreover, we showed that the common component of idiosyncratic variances is a priced factor in the cross section of equity returns, and that it can reduce the pricing error when added to the Fama-French factors. In this section we discuss how we can create portfolios that are only exposed to the idiosyncratic variance risk of an equity, and we present evidence that for the majority of the equities in our sample, the mean return of these portfolios is significantly different from zero. Moreover, we present a measure of idiosyncratic variance risk premium, and we investigate whether this risk premium is explained by the usual equity risk factors.

1.5.1 Evidence from Portfolio Returns

Equity options are investment assets, and their expected returns are of particular interest to practitioners and academics. In a stochastic volatility framework, the underlying price and its volatility are state variables, and the expected index option return depends on the risk premiums associated with them. In our framework, idiosyncratic volatility is also a state variable, and since we show that it is priced, its risk premium would be a part of the expected equity option return. The following proposition provides an expression for the expected index and equity option returns under the physical distribution.

Proposition 1.2. *For a derivative $f(t, S, v)$ written on the index with price S and variance v at time t, the instantaneous expected excess return on the derivative contract is given by:*

$$\frac{1}{dt}E_t^P[\frac{df}{f} - rdt] = f_s\frac{S_t}{f}(\mu - r) + f_v\frac{1}{f}\lambda v_t \tag{1.5.1}$$

and for a derivative $f^i(t, S_i, v_i)$ written on the equity with price S_i and total variance v_i at time t, the instantaneous expected excess return on the derivative contract is given by:

$$\frac{1}{dt}E_t^P[\frac{df^i}{f^i} - rdt] = f^i_{s_i}\frac{S_{i,t}}{f^i}\big((\mu_i - r) + \beta_i(\mu - r)\big) + f^i_{v_i}\frac{1}{f^i}(\beta_i^2\lambda v_t + \lambda_i\xi_{i,t}) \qquad (1.5.2)$$

where, f_s, f_v, $f^i_{s_i}$, and $f^i_{v_i}$ represent partial derivatives, and the structural parameters and state variables are as defined before.

Proof. See Appendix B. □

Therefore the expected excess return of an equity option depends on the equity and variance risk premiums of the index through the equity beta, as well as on the idiosyncratic return and idiosyncratic variance risk premiums.

Now consider a delta hedged portfolio of an index option and the index, denoted by π and a delta hedged portfolio of an equity option and the underlying stock, denoted by π^i. These portfolios are by construction insensitive with respect to the changes in the underlying asset's price. Using the results of Proposition 2, the instantaneous expected excess return on these delta neutral portfolios can be shown to be the following.

$$\frac{1}{dt}E_t^P[\frac{d\pi_t}{\pi_t} - rdt] = f_v\frac{1}{\pi_t}\lambda v_t \qquad (1.5.3)$$

$$\frac{1}{dt}E_t^P[\frac{d\pi^i_t}{\pi^i_t} - rdt] = f^i_{v_i}\frac{1}{\pi^i_t}(\beta_{i,t}^2\lambda v_t + \lambda_i\xi_{i,t}) \qquad (1.5.4)$$

The portfolio π's return depends only on the market variance risk premium, while the return of portfolio π^i depends on the market variance risk premium through the equity beta, as well as on the idiosyncratic variance risk premium. So we can create a portfolio by taking positions in π and π^i, that is insensitive with respect to the market variance, and is only exposed to the idiosyncratic variance risk. Consider a hedge portfolio, Π, that consists of y units of a delta hedged index option π, and x units of a delta hedged equity options π^i. The portfolio value at any time is the following:

$$\Pi_t = x_t \cdot \pi^i + y_t \cdot \pi_t \tag{1.5.5}$$

The instantaneous expected excess return of this portfolio can be found directly from (1.5.3) and (1.5.4), and it is equal to:

$$\frac{1}{dt}E_t^P[\frac{d\Pi_t}{\Pi_t} - rdt] = \Big((x_t f^i_{v_i}\beta^2_{i,t} + y_t f_v)\lambda v_t + x_t f^i_{v_i}\lambda_i \xi_{i,t}\Big)\frac{1}{\Pi_t} \tag{1.5.6}$$

Our goal is to choose x and y at any time so that the portfolio Π is not exposed to the market variance risk premium. This can be accomplished if x and y at any time have the following relationship.

$$y_t = -x_t \frac{f^i_{v_i}}{f_v}\beta^2_{i,t} \tag{1.5.7}$$

At any time t we choose $x_t = \frac{1}{f^i_{v_i}}$ and $y_t = \frac{-1}{f_v}\beta^2_{i,t}$. So, according to (1.5.6) the instantaneous expected excess return on our hedge portfolio would be:

$$\frac{1}{dt}E_t^P[\frac{d\Pi_t}{\Pi_t} - rdt] = \lambda_i \xi_{i,t}\frac{1}{\Pi_t} \tag{1.5.8}$$

Therefore, a portfolio with a long position in $\frac{1}{f^i_{v_i}}$ units of a delta neutral equity option, and a short position in $\frac{\beta^2_{i,t}}{f_v}$ units of a delta neutral index option is instantaneously insensitive with respect to the changes in the equity price, index price, and the market variance. The only risk that this portfolio bears is the idiosyncratic variance risk of the equity, and the return on this portfolio would be the idiosyncratic volatility risk premium.

In what follows, we describe how we form and rebalance portfolios that only loads on the idiosyncratic variance risk premiums of equities, and we investigate whether the return on these portfolios are statistically and economically different from zero. For every firm in our sample, as well as for the market index, we create hedge portfolios using options with

30 days to maturity, and with three moneyness rations 1, 1.025, and 1.05 (1, 0.975, and 0.95) for calls (puts), respectively. The detail description of the data used and the way the portfolios are formed and rebalanced are presented in Appendix C. If the equity idiosyncratic variance bears a risk premium, we would expect these portfolios to have mean returns that are significantly different from zero.

Portfolio Returns

In tables 1.13 and 1.14 we present the mean annualized returns of the hedge portfolios for calls and puts, respectively. Present in the tables are also the t-statistics for the null hypothesis that mean portfolio return is zero. These portfolios are by construction only exposed to the idiosyncratic variance risk of the equity, and a significant mean portfolio return is an indication of a non-zero idiosyncratic volatility risk premium. There is quite a bit of variation across call and put portfolios, and across different moneyness rations. For call portfolios, the mean returns are of different signs for different equities, and they generally decrease as the moneyness increases. In the last row of the table we also report the number of firms for which the mean portfolio return is significantly different zero. For call portfolios, as we move further away from the money, the number of firms with significant idiosyncratic volatility risk premium decreases.

Put portfolios also indicate positive and negative idiosyncratic volatility risk premiums for different stocks. The number of stocks with significant idiosyncratic volatility risk premium, as evident from the mean put portfolio returns, is large for any moneyness ratio. In general, the results obtained from the put portfolios are more indicative of the significant equity idiosyncratic risk premium.

It should be noted that, in our portfolio rebalancing we do not take into account the transaction costs. We merely provide these results as indication of non-zero idiosyncratic volatility risk premium, and our portfolio formation and rebalancing is a statistical procedure for highlighting this evidence, rather than an implementable trading strategy. Motivated by these results, in the next section, we proceed with introducing a measure of idiosyncratic volatility risk premium, and we analyze the properties of these risk premiums for the firms in our sample.

1.5.2 Measure of Idiosyncratic Volatility Risk Premium

Here we present a measure of idiosyncratic variance risk premium, and investigate whether is it statistically and economically significant, and whether it can be explained by the usual equity risk factors. The instantaneous idiosyncratic variance risk premium in our modeling framework is:

$$E_t^P[d\xi_t] - E_t^Q[d\xi_t] = \lambda_i \xi_t dt \tag{1.5.9}$$

The idiosyncratic variance risk premium (RP) at any time t over a discrete time interval of length $T-t$ can be obtained as the difference between the expected integrated idiosyncratic variance under physical and risk-neutral distributions.

$$\begin{aligned} IVRP_{i,t} &= \frac{1}{T-t} E_t^P \Big[\int_t^T \xi_{i,s} ds \Big] - \frac{1}{T-t} E_t^Q \Big[\int_t^T \xi_{i,s} ds \Big] \\ &= \Big(\theta_i + \frac{1-e^{-\kappa_i (T-t)}}{\kappa_i (T-t)} (\xi_{i,t}^P - \theta_i) \Big) - \Big(\tilde{\theta}_i + \frac{1-e^{-\tilde{\kappa}_i (T-t)}}{\tilde{\kappa}_i (T-t)} (\xi_{i,t}^Q - \tilde{\theta}_i) \Big) \end{aligned} \tag{1.5.10}$$

Where, $IVRP_{i,t}$ is the idiosyncratic variance RP of stock i, ξ_t^P and ξ_t^Q are the estimated spot idiosyncratic variances at time t under P and Q, and the rest of the structural parameters are as defined before. In our calculations we choose T to be 30 days, and we calculate the annualized 30-day idiosyncratic variance RP. We also calculate the 30-day market variance risk premium as:

$$\begin{aligned} MVRP_t &= \frac{1}{T-t} E_t^P \Big[\int_t^T v_s ds \Big] - \frac{1}{T-t} E_t^Q \Big[\int_t^T v_s ds \Big] \\ &= \Big(\theta + \frac{1-e^{-\kappa (T-t)}}{\kappa (T-t)} (v_t^P - \theta) \Big) - \Big(\tilde{\theta} + \frac{1-e^{-\tilde{\kappa} (T-t)}}{\tilde{\kappa} (T-t)} (v_t^Q - \tilde{\theta}) \Big) \end{aligned} \tag{1.5.11}$$

Where, $MVRP_{i,t}$ is the market variance RP and, v_t^P and v_t^Q are the estimated spot market variances at time t under P and Q. The descriptive statistics for the idiosyncratic variance RP of each stock, as well as the market volatility RP, are presented in Table 1.15. We also report in the last column the t-stat of the null hypothesis that the average idiosyncratic variance RP is zero. Consistent with the estimation results for the structural parameters, the average idiosyncratic variance of the stocks in our sample can take positive or negative signs. Moreover, the mean RP is statistically different from zero for all stocks. Earlier studies have found that the market variance RP is negative. Consistent with these previous results[16], the average market variance RP in our sample is -0.48% per annum, and it is strongly significantly different from zero.

Principal component analysis of the idiosyncratic variance RP's shows that the first two principal components explain 70% and 20% of the variation, respectively. These results were expected, since the RP's reflect the difference between the P- and Q-estimates of the levels of idiosyncratic variances, for both of which there is a high commonality as we saw before. The first principal component has a 89% correlation with the average idiosyncratic variance risk premium defined as follows.

$$AIVRP_t = \frac{1}{N}\sum_{i=1}^{N} IVRP_{i,t} \tag{1.5.12}$$

The first principal component of the idiosyncratic variance risk premiums is also highly correlated with the market variance risk premium with the correlation coefficient of -0.78%. We regress the market variance risk premium on to the average idiosyncratic variance risk premium. With $R^2 = 75\%$, the average idiosyncratic variance risk premium is a strong predictor of the market variance risk premium. Moreover, the residuals of the regression, denoted by F_{RP}^{orth} , which is the component of market variance risk premium that is orthogonal to the $AIVRP$ has a correlation of 45% with the second principal component of the idiosyncratic variance risk premiums. These results are qualitatively and quantitatively consistent with those presented in the previous section.

[16]See Carr and Wu (2009), and Driessen et al. (2009).

Are Idiosyncratic Variance Risk Premiums Explained by the Usual Equity Risk Factors?

In our modeling framework, the independent variation in idiosyncratic variance represents an additional source of risk, independent from the equity return risk premium that is due to covariation of the equity return and market return. Under the classical CAPM and its extensions, idiosyncratic variance risk premium cannot come from an independent source of risk, and can only come from the correlation between the idiosyncratic variance and the market return. In this section we investigate whether the Fama-French three factors plus the momentum factor can explain the variation in the idiosyncratic variance risk premiums. For every firm in our sample we run the following regression:

$$IVRP_{i,t} = a_i + b_i^m \cdot (r_t^m - r_t) + b_i^{smb} \cdot r_t^{SMB} + b_i^{hml} \cdot r_t^{HML} + b_i^{mom} \cdot r_t^{mom} + \varepsilon_{i,t} \tag{1.5.13}$$

where, $IVRP_{i,t}$ is the idiosyncratic variance risk premium of stock i at time t, and the rest of the factors are as defined before. Table 1.16 presents the OLS coefficient estimates, the t-stats, and the R^2 of the regressions. As is evident from the R^2, these equity risk factors cannot explain the time-series variations in the idiosyncratic variance risk premiums. Moreover, the coefficients of the market return and SMB factors are not significantly different from zero for almost all of the stocks.

Idiosyncratic variance risk premium can also come from the correlation between the idiosyncratic variance and the market variance. In a stochastic volatility model such as Heston (1993), a negative market variance risk premium is generated because of the negative correlation between the market variance and market return. Christoffersen et al. (2013) present a pricing kernel in which both the equity premium and the variance premium have two distinct components originating in preferences. So, if the idiosyncratic variance of a firm's equity is correlated with aggregate volatility, then the market variance risk premium should be able to explain at least part of the variation in idiosyncratic variance risk premiums. To test this conjecture we run the following regression:

$$IVRP_{i,t} = a_i + b_i \cdot MVRP_t + \varepsilon_{i,t} \tag{1.5.14}$$

The results are presented in Table 1.17. The coefficient of the $MVRP$ is significant for all stocks except for JNJ. The R^2's are much larger compared to the previous regression, ranging from zero for JNJ to 82% for CVX, and with the average R^2 equal to 26%. Moreover, the coefficient of the $MVRP$ is positive (negative) for stocks with negative (positive) average idiosyncratic variance risk premium.

Motivated by our PCA results in the previous sections, we also test the explanatory power of the average idiosyncratic variance risk premium, $AIVRP$, together with the component of the market variance risk premium orthogonal to it (F_{RP}^{orth}) in the following regression:

$$IVRP_{i,t} = a_i + b_{AIVRP,i} \cdot AIVRP_t + b_{orth,i} \cdot F_{RP}^{orth} + \varepsilon_{i,t} \tag{1.5.15}$$

The two factors can significantly improve the R^2 compared to the previous regressions. The R^2 now ranges from 10% for KO to 89% for CVX, with the average R^2 equal to 51%. Moreover, the coefficients as well as the intercept are significantly different from zero for all stocks. The results are presented in Table 1.18.

1.6 Conditional Equity Betas

Stock betas are one of the most important measures of equity risk. The importance of stock betas in corporate finance and asset pricing has motivated researchers to look for better methods of estimating these variables. An accurate measurement of betas is crucial in many applications such as cost of capital estimation and detection of abnormal returns. Stock betas are usually estimated using historical returns on the stock and the market index. There is a consensus that stock betas are time varying, and the popular approach to account for the time variation is to use a rolling window of historical returns. There are other sophisticated estimation methods based on historical betas to capture the time variation.

All these techniques are based on historical information, and the main premise is that the future will be similar to the past.

Other proposed methods of estimating the equity betas use the information inherent in option prices. Option prices contain information about the future distribution of the underlying asset, and incorporating this information can potentially lead to better estimates for any variable, especially when the historical patterns are unstable or when there have been structural breaks.

In estimating stock betas from option prices, a few important issues should be noted. First, the information inherent in option prices is related to the risk neutral distribution of the underlying asset. Since betas are ultimately used as a measure of equity risk under the physical measure, a proper link should be made between the P and Q distributions, and the premiums for the priced risks in the market should be taken into account.

Second, the consistency between the index option prices and equity option prices as well as the consistency between the market's P and Q distributions are very important. Driessen et al. (2009) study the relationship between equity options and index options, and they find structural differences which they explain by the existence of correlation risk. Bates (2000) indicates that any successful option pricing model should be able to reconcile the P and Q distributions of the underlying asset. Constantinides et al. (2011) show widespread evidence of mispricing in the index options resulting from the lack of consistency between the index option and the underlying markets.

Several studies have proposed to use the option implied information in estimating betas. French et al. (1983) compute betas based on historical correlations and the implied volatilities of the market and the equity. Siegel (1995) proposes the creation a derivative from which implied betas can be estimated. Perhaps most related to our study are two recent papers by Chang et al. (2011), hereafter CCVJ, and Buss and Vilkov (2012), hereafter BV.

BV use implied volatilities of the market and equity along with a parametric model for implied correlations, and estimate the betas. CCVJ show that in a one factor model, and under the assumption of zero skewness for the idiosyncratic returns, the implied betas can be estimated as the product of the ratio of equity to market implied skewness, and the ratio of equity to market implied volatilities. They derive option implied moments of the equity risk neutral distribution based on the method of Bakshi et al. (2003) using only one day of observed option prices, and their estimation does not rely on historical returns.

All of these papers estimate the equity beta as the product of the correlation between the market and equity returns and the ratio of equity to market volatility, and the main distinguishing feature among them is how the correlation is estimated. In this paper the equity betas appear in equation (1.2.2) in the model and are estimated as part of the structural parameters in our framework. The advantages of our beta estimates are as follows. First, since the beta enters both the P- and Q-equity return dynamics explicitly, it is estimated directly, without making any assumptions regarding the correlations. Moreover, since our estimation methodology is based on the joint information in the stock returns and equity option prices, our beta estimates take all available information from the P and Q distributions into account. In Table 1.19 we present the estimated betas for the firms in our sample, based on the information in returns alone, based on the information in option prices alone, and based on the joint information. We also report the OLS betas.

Second, since we use an equity option pricing model that links the equity price dynamics to the market price dynamics, the consistency between the index option market and equity option market is taken into account. Third, our option pricing model is based on the assumptions regarding the physical dynamics of the equity price, and the explicit transformation to the risk neutral measure. So, betas are estimated taking into account the price of market variance risk, as well as the price of idiosyncratic variance risk. The methods in the previous studies that estimate the betas from the option markets use exclusively the RN measure; since we are ultimately interested in the application of these betas under the physical measure, these estimates might be biased because of the presence of the market risk premiums.

In the model presented in this paper the betas are assumed to be constant. There is, however, widespread agreement in the literature that equity betas are time varying. The reason that betas change over time is the time variation in market volatility, stock volatility, and the correlation between the market and stock returns. To capture the time variation in equity betas and account for the conditionality, we propose the following procedure to estimate betas.

We fix the risk neutral structural parameters, $\hat{\hat{\Theta}}$ and $\hat{\hat{\Theta}}_i$, that we estimated before for the market dynamics and equity dynamics. Our parameter estimates are based on the full sample of option prices for the index and the equities. Since the data generating process for the stochastic volatility process of the market return and idiosyncratic return are assumed to remain the same over time, and since the parameters are constant, we use the entire sample to estimate these parameters to increase estimation precision.

On every day t in our sample and given the structural parameters of the market and of every equity, excluding the constant beta estimate, and given the estimated spot market variance, we can find the conditional beta and the spot idiosyncratic variance for every equity i using the equity option prices observed at time t:

$$\{\hat{\xi}_{i,t}, \hat{\beta}_{i,t}\} = \underset{\{\xi_{i,t},\beta_{i,t}\}}{\operatorname{argmin}} \sum_{n=1}^{M_{i,t}} (C^O_{i,t,n} - C^M_{i,t,n}(\hat{\hat{\Theta}}_i, \hat{\hat{\Theta}}, \hat{v}^Q_t, \xi_{i,t}))^2 / Vega^2_{i,t,n}, t = 1, 2, \ldots, T, \; i = 1, \ldots, N. \quad (1.6.1)$$

Where, as before, $C^O_{i,t,n}$ is the observed price, $C^M_{i,t,n}$ is the model price, and M_i is the number of option contract with six months to maturity available for stock i at time t. The options used in estimating the conditional betas are not in the sample that we used for our estimation in previous sections. Thus, our estimation of betas is done out-of-sample. The choice of a six-month horizon is to create a balance between the option liquidity that is largest for short maturity, and the relevant horizon for firm risk, which is arguably considerably longer. Moreover, given the estimated market and equity structural parameters, and given the spot market variance, the estimation of beta on any day relies on the observed options on that day alone. This feature is similar to that in CCJV, and allows for more reliable beta estimates when new information is released about the firm.

In Table 1.20 we repost the mean, standard deviation, minimum and maximum of the conditional betas for each stock. In the last column we also present the unconditional betas estimated before for comparison. The mean conditional betas are larger than the unconditional ones for most of the firms. The average is 1.1 across the 27 stocks compared to the average of 0.91 for the unconditional betas. Nonetheless, in almost all cases the mean of the conditional betas lies within one standard deviation of its unconditional value.

1.7 Conclusion

We use a one factor model for equity return dynamics in which the idiosyncratic volatility of the stock follows a stochastic process, and is allowed to be priced. We develop a method to estimate the structural parameters as well as the spot idiosyncratic variances using the return data and the option data. Our estimation is based on a joint likelihood function that has a return component and an option component, while the structural parameters are internally consistent between the physical and risk neutral measures. In a recent study Christoffersen et al (2013) use the same model, but they assume that the idiosyncratic variance is not priced. This assumption implies that the market return is the only priced risk factor, and that the one factor model is the correct model. In our estimation we show that the price of idiosyncratic variance is significantly different from zero for all the 27 firms in our sample.

We calculate the implied idiosyncratic variance of the options in our sample, and we show that even after removing the market factor from equity returns, there is still a strong factor structure in implied idiosyncratic variances. Our principal component analysis shows that there is a strong factor structure in option implied levels, moneyness slopes, and term structure slopes of idiosyncratic variances. We show that the average idiosyncratic variance, AIV, of all stocks is a good proxy for the common factor. These findings complement the literature that documents the existence of a common factors structure among idiosyncratic variances under the physical distribution.

We show that if there are factors missing from our one factor model, the idiosyncratic variance as defined in a one factor model, captures the exposure to the variance of the missing factors. We show that the AIV factor has positive loadings in the time series regression of the 25 Fama-French value and size portfolios. Moreover, we show that the AIV factor reduces the cross sectional pricing error of these portfolios when added to the Fama-French model. The AIV has a significant and positive risk premium, and its cross sectional explanatory power, in our sample period, is distinct from and more than that of the HML and SMB factors.

We derive the expected option return in our framework, and we discuss a trading strategy involving the equity option, the underlying equity, the index option, and the index. These portfolios are constructed and rebalanced in such a way that they are only exposed to the idiosyncratic variance risk of the equity. We show that the mean returns on these portfolios are significantly different from zero for the majority of the equities in our sample, which is an indication of the existence of a premium for the idiosyncratic variance risk.

We propose a measure for the idiosyncratic variance risk premium, defined as the difference between the P and Q expected integrated idiosyncratic variance. We show that the mean idiosyncratic variance risk premium is significantly different from zero for the firms in our sample, and that this risk premium is not explained by the usual risk factors. Moreover, we show that the time series variations of idiosyncratic variance risk premiums are well explained by the average idiosyncratic variance risk premium together with the component of the market variance risk premium orthogonal to the average idiosyncratic variance risk premium.

Finally, we discuss the implications of our model for the estimation of equity betas. The equity beta in our model is estimated as part of the structural parameters, using the simultaneous information from returns and options. Moreover, we propose a method to characterize the time variation in equity betas.

Appendix

1.A Proof of Proposition 1

The market index is assumed to follow the stochastic volatility model below:

$$dS_t/S_t = (\mu)dt + \sqrt{v_t}dz_t,$$
$$dv_t = \kappa(\theta - v_t)dt + \sigma\sqrt{v_t}dw_t \tag{1.A.1}$$

Using Girsanov's theorem we can write the following transformation for the two Brownian motions dz and dw.

$$dz = d\tilde{z} - (\psi_1 + \rho\psi_2)dt,$$
$$dw = d\tilde{w} - (\rho\psi_1 + \psi_2)dt \tag{1.A.2}$$

where, ψ_1 and ψ_2 are the price of risk for dz and dw, respectively, and ρ is the correlations between the two Brownian motions. The drift of the index return dynamics under the risk neutral measure is equal to the risk free rate, so we have the following restriction:

$$\psi_1 + \rho\psi_2 = \frac{\mu - r}{\sqrt{v}} \tag{1.A.3}$$

Moreover, we assume that the price of volatility risk is proportional to volatility. So we have the second restriction as follows:

$$\rho\psi_1 + \psi_2 = \frac{\lambda\sqrt{v}}{\sigma} \tag{1.A.4}$$

The unique prices of risk can be found from (1.A.3) and (1.A.4), and are the following:

$$\psi_1 = \frac{\mu\sigma - \rho\lambda v}{\sigma\sqrt{v}(1-\rho^2)},$$
$$\psi_2 = \frac{\lambda v - \mu\rho\sigma}{\sigma\sqrt{v}(1-\rho^2)} \tag{1.A.5}$$

Replacing (1.A.2)-(1.A.5) into (1.A.1) yields the index return dynamics under the risk neutral measure.

$$dS_t/S_t = rdt + \sqrt{v_t}d\tilde{z}_t,$$
$$dv_t = \tilde{\kappa}(\tilde{\theta} - v_t)dt + \sigma\sqrt{v_t}d\tilde{w}_t \tag{1.A.6}$$

where, $\tilde{\kappa} = \kappa + \lambda$, and $\tilde{\theta} = \frac{\kappa\theta}{\kappa+\lambda}$.

For the equity return dynamics we have the following dynamics under the physical measure:

$$\begin{aligned} dS_{i,t}/S_{i,t} &= (\mu_i)dt + \beta_i(dS_t/S_t - rdt) + \sqrt{\xi_{i,t}}dz_{i,t}, \\ d\xi_{i,t} &= \kappa_i(\theta_i - \xi_{i,t})dt + \sigma_i\sqrt{\xi_{i,t}}dw_{i,t} \end{aligned} \tag{1.A.7}$$

Representing the prices of idiosyncratic shocks dz_i and dw_i by ψ_1^i and ψ_2^i, respectively, we have the following transformation using Girsanov's theorem:

$$\begin{aligned} dz_i &= d\tilde{z}_i - (\psi_1^i + \rho_i\psi_2^i)dt, \\ dw_i &= d\tilde{w}_i - (\rho_i\psi_1^i + \psi_2^i)dt \end{aligned} \tag{1.A.8}$$

Similar to the case of the market return dynamics, we apply the following restrictions to the prices of risk, assuming that the price of idiosyncratic volatility risk is proportional to idiosyncratic volatility.

$$\begin{aligned} \psi_1^i + \rho_i\psi_2^i &= \frac{\mu_i - r}{\sqrt{\xi_i}}, \\ \rho_i\psi_1^i + \psi_2^i &= \frac{\lambda_i\sqrt{\xi_i}}{\sigma_i} \end{aligned} \tag{1.A.9}$$

solving for ψ_1^i and ψ_1^i we have the following prices of idiosyncratic risk.

$$\begin{aligned} \psi_1^i &= \frac{(\mu_i - r)\sigma_i - \rho_i\lambda_i\xi_i}{\sigma_i\sqrt{\xi_i}(1-\rho_i^2)}, \\ \psi_2^i &= \frac{\lambda_i\xi_i - (\mu_i - r)\rho_i\sigma_i}{\sigma_i\sqrt{\xi_i}(1-\rho_i^2)} \end{aligned} \tag{1.A.10}$$

Replacing (1.A.8)-(1.A.10) into (1.A.7) we have the following equity return dynamics under the risk neutral measure:

$$\begin{aligned} dS_{i,t}/S_{i,t} &= rdt + \beta_i(dS_t/S_t - rdt) + \sqrt{\xi_{i,t}}d\tilde{z}_{i,t}, \\ d\xi_{i,t} &= \tilde{\kappa}_i(\tilde{\theta}_i - \xi_{i,t})dt + \sigma_i\sqrt{\xi_{i,t}}d\tilde{w}_{i,t} \end{aligned} \tag{1.A.11}$$

where, $\tilde{\kappa}_i = \kappa_i + \lambda_i$, and $\tilde{\theta}_i = \frac{\kappa_i \theta_i}{\kappa_i + \lambda_i}$. $\square$

1.B Proof of Proposition 2

We derive the instantaneous return for the equity option. The return for index option can be derived similarly. Let $f^i(t, S_i, v_i)$ be the price of a derivative whose price depends on the spot price and spot variance of the equity. Applying Ito's lemma to $f^i(t, S_i, v_i)$, we have:

$$df^i = f^i_t dt + f^i_{s_i} dS_i + \frac{1}{2} f^i_{s_i s_i} dS_i dS_i + f^i_{v_i} dv_i + \frac{1}{2} f^i_{v_i v_i} dv_i dv_i + f^i_{s_i v_i} dS_i dv_i \tag{1.B.1}$$

where, f^i_x and f^i_{xy} denote the first and second partial derivative of f with respect to x and xy, respectively. Moreover, in our one factor model, total variance of the equity return is defined as $v_i = \beta_i^2 v + \xi_i$. Using (1.2.1)-(1.2.4) together with the pricing PDE, we have the following equation for df^i under the physical measure:

$$df^i = \Big(rf^i - rS_i f^i_{s_i} - f^i_{v_i}\Big(\beta_i^2 \tilde{\kappa}(\theta - v) + \tilde{\kappa}_i(\tilde{\theta}_i - \xi_i)\Big)\Big)dt + f^i_{s_i} dS_i + f^i_{v_i} dv_i \tag{1.B.2}$$

Note that our model implies that:

$$\begin{aligned} \frac{1}{dt} E^P_t[dS_i] &= \Big(\mu_i + \beta_i(\mu - r)\Big) S_i, \\ \frac{1}{dt} E^P_t[dv_i] &= \beta_i^2 \kappa(\theta - v) + \kappa_i(\theta_i - \xi_i) \end{aligned} \tag{1.B.3}$$

Relations (1.B.3) together with (1.B.2) yields the following:

$$\frac{1}{dt} E^P_t[\frac{df^i}{f^i} - rdt] = \frac{f^i_{s_i}}{f^i}(\frac{1}{dt} E^P_t[dS_i] - rS_i) + \frac{f^i_{v_i}}{f^i}\left(\frac{1}{dt} E^P_t[dv_i] - \Big(\beta_i^2 \tilde{\kappa}(\tilde{\theta} - v) + \tilde{\kappa}_i(\tilde{\theta}_i - \xi_i)\Big)\right) \tag{1.B.4}$$

which simplifies to:

$$\frac{1}{dt}E_t^P[\frac{df^i}{f^i} - rdt] = f^i_{s_i}\frac{S_i}{f^i}\Big((\mu_i - r) + \beta_i(\mu - r)\Big) + f^i_{v_i}\frac{1}{f^i}(\beta_i^2\lambda v_t + \lambda_i\xi_{i,t}) \tag{1.B.5}$$

□

1.C Portfolio Formation and Rebalancing

The data that we use in this section is the same data described in section 3.1. Since equity option data is relatively scant both in cross-sectional and maturity dimensions, we apply a simple weighting scheme to derive option prices for the target values of moneyness and maturity. First, we calculate the implied volatilities of the options. In the next step, for each of the two maturities bracketing the target days-to-maturity of 30 days, for two values of a given target moneyness, we derive averages of IVs weighted by the reciprocal of the absolute distance from this target moneyness. Last, the two observations for IV weighted by the reciprocal of the absolute distance from the target maturity of 30 days yield the final value for IV that is subsequently inverted into the option price via the Black-Scholes formula.

In cases where we had observations with at most three days from the target maturity and/or at most 0.01 from the target moneyness, we used a single observation. In cases we had no bracketing observations for a given target, we used a nearest neighbourhood value. For the index, where we have richer data, in addition to the procedure above we also apply the practitioner's Black-Scholes via least squares and verify the sensitivity of the results to the relatively crude approach by necessity applied to equities. We choose the target moneyness ratios of 1, 1.025 and 1.05 (0.95, 0.975 and 1) for calls (puts). Standardization of option contracts is necessary to insure that the variability of the portfolio returns is only due to exposure to the risk factors. Note that the use the Black-Scholes formula to derive our target prices leaves them free of the Black-Scholes assumptions since this formula is used merely as a translation device. Once we obtain our daily option prices from cross-sections, we screen them again by rejecting observations with ATM prices below 10 cents and whose ATM IV is outside the range 5-150%. This last set of filters resulted in rejecting less than 0.1% of firm-days.

On every day and for each target moneyness ratio we set up a zero-net-cost portfolio with a long position in $\frac{1}{f^i_{v_i}}$ units of a delta hedged equity call (put) and a short position in $\frac{\beta_t^2}{f_v}$ units of a delta hedged index call (put), with the proceeds invested or borrowed at the risk free rate, r. In our calculations we approximate $f^i_{s_i}(f_s)$ and $f^i_{v_i}(f_v)$ by equity (index) option's Black-Shcoles delta and vega. Moreover, the equity beta is estimated using a rolling window of 250 days historical returns on the equity and the market. The value of the hedge portfolio at time t is:

$$\Pi_t = \frac{1}{vega_{i,t}}(f^i_t - \Delta_{i,t} S_{i,t}) - \frac{1}{vega_t}\beta^2_{i,t}(f_t - \Delta_t S_t) \tag{1.C.1}$$

If Π_t is positive we invest the proceeds at the risk free rate, and if it negative, we borrow this amount at the risk free rate. So after one day the gain (loss) for our zero-net-cost portfolio is:

$$\begin{aligned} G_{t+1} = \Pi_{t+1} - \Pi_t = & \frac{1}{vega_{i,t}}\Big((f^i_{t+1} - f^i_t) - \Delta_{i,t}(S_{i,t+1} - S_{i,t})\Big) \\ & - \frac{1}{vega_t}\beta^2_{i,t}\Big((f_{t+1} - f_t) - \Delta_t(S_{t+1} - S_t)\Big) - \Pi_t(\frac{r}{252}) \end{aligned} \tag{1.C.2}$$

We register the gain G_{t+1}, and repeat this exercise until it is done for every day in our sample. This hedge portfolio is, by construction, insensitive with respect to the changes in the equity price, index price and the market variance, and is only exposed to the idiosyncratic variance risk of the equity. So, the daily gains can be thought of as excess dollar return for bearing idiosyncratic volatility risk. In order to transform the excess dollar returns into percentage return and since the option price is homogenous of first degree with respect to the initial stock price, we scale the dollar returns by the initial stock price. Finally, we compound the daily portfolio returns into monthly returns.

Table 1.1: **Company Names, Tickers, and the Number of Options**

		Number of Options		
Company Name	**Ticker**	**Calls**	**Puts**	**All**
S&P500 Index	**SPX**	97,355	70,934	168,289
Alcoa	**AA**	52,254	45,251	97,505
American Express	**AXP**	58,391	49,264	107,655
Boeing	**BA**	60,897	52,182	113,079
Caterpillar	**CAT**	57,977	50,036	108,013
Cisco	**CSCO**	49,533	42,555	92,088
Chevron	**CVX**	76,709	63,163	139,872
Dupont	**DD**	68,408	57,032	125,440
Disney	**DIS**	60,406	51,482	111,888
General Electric	**GE**	64,687	52,571	117,258
Home Depot	**HD**	59,296	50,778	110,074
Hewlett-Packard	**HPQ**	50,888	44,397	95,285
IBM	**IBM**	69,281	58,509	127,790
Intel	**INTC**	49,624	42,572	92,196
Johnson & Johnson	**JNJ**	84,686	67,461	152,147
JP Morgan	**JPM**	61,264	49,993	111,257
Coca Cola	**KO**	78,875	63,576	142,451
McDonald's	**MCD**	70,192	59,008	129,200
3M	**MMM**	75,277	62,499	137,776
Merck	**MRK**	67,354	55,852	123,206
Microsoft	**MSFT**	56,630	47,837	104,467
Pfizer	**PFE**	60,483	51,830	112,313
Procter & Gamble	**PG**	84,061	67,733	151,794
AT&T	**T**	74,071	58,506	132,577
United Technologies	**UTX**	70,371	58,781	129,152
V erizon	**VZ**	71,543	56,265	127,808
Walmart	**WMT**	70,668	59,539	130,207
Exxon Mobil	**XOM**	76,187	63,202	139,389

For each firm we present the name of the company, and its ticker symbol. We also report the number of Calls, Puts, and the total number of options available in our sample for each firm.

Table 1.2: **Option Statistics: Calls**

Ticker	Avg. IV	Min (IV)	Max (IV)	Avg. Delta	Avg. vega	Avg. DTM
SPX	0.180	0.073	0.750	0.417	198.015	80.608
AA	0.342	0.169	1.496	0.476	6.910	79.574
AXP	0.296	0.127	1.482	0.460	10.144	79.509
BA	0.293	0.161	0.896	0.461	10.666	78.889
CAT	0.309	0.160	1.033	0.466	11.417	78.849
CSCO	0.355	0.159	1.071	0.471	5.841	78.540
CVX	0.232	0.128	0.944	0.447	13.470	79.645
DD	0.259	0.123	0.923	0.459	8.852	80.528
DIS	0.283	0.069	0.959	0.466	6.720	81.597
GE	0.257	0.069	1.489	0.467	8.622	84.010
HD	0.294	0.153	1.009	0.464	7.310	80.183
HPQ	0.342	0.153	0.979	0.473	7.531	79.760
IBM	0.253	0.119	0.868	0.443	19.123	79.475
INTC	0.349	0.173	0.909	0.481	7.443	81.488
JNJ	0.199	0.097	0.708	0.438	11.388	82.155
JPM	0.293	0.112	1.489	0.461	8.993	79.048
KO	0.211	0.083	0.693	0.445	9.410	83.176
MCD	0.244	0.116	0.789	0.455	8.349	80.995
MMM	0.233	0.125	0.796	0.443	15.130	79.134
MRK	0.263	0.143	0.852	0.463	9.902	80.410
MSFT	0.290	0.122	0.879	0.475	9.099	83.745
PFE	0.272	0.142	1.010	0.471	7.145	84.019
PG	0.201	0.093	0.643	0.436	12.913	81.913
T	0.238	0.102	0.822	0.465	6.177	80.458
UTX	0.249	0.132	0.823	0.446	13.176	79.456
VZ	0.240	0.092	0.870	0.473	7.905	82.906
WMT	0.239	0.112	0.673	0.446	8.903	81.216
XOM	0.227	0.126	0.848	0.445	11.574	80.630

For each firm we report the average, minimum, and maximum implied volatility (IV) for the call options. In the last three columns we also report the average call delta, average vega, and average days-to-maturity (DTM) of the calls. The implied volatilities and the options deltas are provided by OptionMetrics, while the vega is calculated using Black-Scholes formula evaluated at the implied volatility of the options.

Table 1.3: **Option Statistics: Puts**

Ticker	Avg. IV	Min (IV)	Max (IV)	Avg. Delta	Avg. vega	Avg. DTM
SPX	0.208	0.089	0.784	-0.349	188.464	80.669
AA	0.354	0.174	1.484	-0.384	6.607	77.623
AXP	0.312	0.122	1.494	-0.379	9.826	78.367
BA	0.305	0.174	0.924	-0.380	10.172	78.019
CAT	0.324	0.179	1.026	-0.382	10.963	78.630
CSCO	0.365	0.163	1.096	-0.391	5.508	76.969
CVX	0.248	0.117	0.965	-0.366	12.955	79.686
DD	0.276	0.137	0.942	-0.373	8.537	80.319
DIS	0.300	0.143	0.987	-0.381	6.382	80.566
GE	0.277	0.071	1.496	-0.375	8.329	83.703
HD	0.309	0.140	1.020	-0.382	7.030	79.012
HPQ	0.352	0.164	0.920	-0.390	7.304	78.551
IBM	0.270	0.124	0.885	-0.372	18.553	79.914
INTC	0.358	0.164	0.900	-0.392	7.300	79.531
JNJ	0.220	0.096	0.768	-0.362	10.834	81.336
JPM	0.315	0.120	1.492	-0.375	8.848	79.156
KO	0.231	0.095	0.670	-0.367	9.005	82.628
MCD	0.263	0.125	0.699	-0.375	7.902	80.480
MMM	0.250	0.138	0.844	-0.367	14.682	80.075
MRK	0.279	0.091	0.881	-0.375	9.589	79.789
MSFT	0.306	0.112	0.913	-0.389	8.960	82.884
PFE	0.288	0.139	0.709	-0.383	6.984	84.079
PG	0.222	0.096	0.682	-0.361	12.427	81.336
T	0.260	0.103	0.836	-0.373	6.032	80.409
UTX	0.269	0.136	0.856	-0.372	12.667	79.524
VZ	0.262	0.109	0.896	-0.370	7.655	82.216
WMT	0.256	0.114	0.675	-0.373	8.557	80.655
XOM	0.243	0.128	0.953	-0.368	11.056	80.150

For each firm we report the average, minimum, and maximum implied volatility (IV) for the put options. In the last three columns we also report the average delta, average vega, and average days-to-maturity (DTM) of the puts. The implied volatilities and the options deltas are provided by OptionMetrics, while the vega is calculated using Black-Scholes formula evaluated at the implied volatility of the options.

Table 1.4: **Market and Equity Models Parameter Estimates**

Ticker	μ	κ	θ	σ	ρ	β	λ	$\tilde{\kappa}$	$\tilde{\theta}$
SPX	0.078	3.157	0.037	0.318	-0.494		-1.211	1.946	0.061
AA	0.021	1.678	0.041	0.297	-0.277	1.09	-0.872	0.806	0.086
AXP	0.116	0.409	0.049	0.201	-0.355	1.23	0.013	0.422	0.048
BA	0.072	0.480	0.062	0.084	-0.521	1.06	0.040	0.519	0.057
CAT	0.152	0.719	0.063	0.138	-0.275	1.11	-0.012	0.707	0.064
CSCO	0.116	0.417	0.107	0.165	0.100	1.08	-0.013	0.405	0.111
CVX	0.118	1.583	0.042	0.185	-0.289	0.81	0.020	1.603	0.041
DD	0.050	0.378	0.053	0.090	-0.458	0.96	0.080	0.458	0.044
DIS	0.065	0.579	0.060	0.128	-0.269	1.07	0.055	0.634	0.054
GE	0.049	0.844	0.049	0.287	-0.204	0.86	-0.229	0.615	0.067
HD	0.117	0.801	0.101	0.242	-0.121	1.06	0.289	1.090	0.074
HPQ	0.061	0.330	0.146	0.148	0.052	1.14	0.234	0.565	0.086
IBM	0.145	0.138	0.078	0.068	0.046	0.97	0.024	0.162	0.066
INTC	0.111	0.677	0.126	0.185	-0.101	0.99	0.308	0.985	0.087
JNJ	0.087	0.713	0.038	0.086	-0.481	0.60	-0.282	0.431	0.063
JPM	0.084	0.785	0.130	0.326	-0.649	0.85	0.427	1.212	0.084
KO	0.055	0.365	0.049	0.080	-0.151	0.75	-0.100	0.265	0.067
MCD	0.121	1.645	0.055	0.182	-0.511	0.55	-0.543	1.101	0.081
MMM	0.085	0.239	0.034	0.057	-0.122	0.98	0.006	0.245	0.033
MRK	0.061	0.277	0.063	0.077	-0.231	0.89	0.092	0.368	0.047
MSFT	0.110	0.268	0.066	0.077	0.060	1.04	-0.053	0.216	0.082
PFE	0.078	0.760	0.089	0.143	-0.233	0.78	0.173	0.933	0.072
PG	0.102	0.346	0.044	0.080	-0.108	0.75	0.053	0.398	0.038
T	0.066	1.026	0.049	0.173	-0.362	0.81	-0.093	0.933	0.054
UTX	0.131	0.556	0.036	0.087	0.032	1.03	-0.115	0.441	0.046
VZ	0.075	1.228	0.059	0.192	-0.602	0.70	-0.117	1.111	0.065
WMT	0.120	0.991	0.039	0.139	-0.092	0.70	-0.366	0.625	0.061
XOM	0.115	0.699	0.106	0.191	-0.370	0.75	1.145	1.845	0.040
Average	0.092	0.701	0.068	0.152	-0.240	0.91	0.006	0.707	0.064
Min	0.021	0.138	0.034	0.057	-0.649	0.551	-0.872	0.162	0.033
Max	0.152	1.678	0.146	0.326	0.100	1.225	1.145	1.845	0.111

We use OTM options over the period 1996-2011 to estimate the market and equity parameters. The estimation is based on a joint likelihood function that has a return component and an option component. The estimation of the equity model is conditional on the estimates of the market model. For the market model, μ is set to the sample average risk premium. For the equity model, μ is set equal to the intercept of the CAPM regression of the equity returns on market excess returns. Equity beta is a free parameter and is assumed constant.

Table 1.5: **Distributional Properties of Filtered Spot Idiosyncratic Variances**

	P-distribution				Q-disribution			
Ticker	**mean**	**std**	**min**	**max**	**mean**	**std**	**min**	**max**
SPX	0.0376	0.0342	0.0034	0.2608	0.0409	0.0394	0.0035	0.2956
AA	0.0942	0.0814	0.0083	0.4974	0.1078	0.1001	0.0091	0.6166
AXP	0.0750	0.0889	0.0008	0.4934	0.0711	0.0839	0.0008	0.4645
BA	0.0594	0.0280	0.0126	0.1362	0.0564	0.0264	0.0123	0.1313
CAT	0.0649	0.0323	0.0091	0.1564	0.0629	0.0308	0.0086	0.1443
CSCO	0.1311	0.1146	0.0138	0.4957	0.1292	0.1144	0.0134	0.4935
CVX	0.0396	0.0237	0.0079	0.1763	0.0384	0.0226	0.0079	0.1685
DD	0.0502	0.0305	0.0047	0.1286	0.0463	0.0281	0.0043	0.1204
DIS	0.0556	0.0383	0.0077	0.1409	0.0524	0.0360	0.0075	0.1345
GE	0.0638	0.0688	0.0016	0.4675	0.0643	0.0696	0.0016	0.4732
HD	0.0688	0.0535	0.0068	0.2583	0.0624	0.0478	0.0066	0.2350
HPQ	0.1089	0.0840	0.0105	0.3509	0.0968	0.0740	0.0102	0.3085
IBM	0.0512	0.0440	0.0050	0.1763	0.0485	0.0421	0.0049	0.1691
INTC	0.1267	0.0905	0.0168	0.4613	0.1138	0.0794	0.0159	0.4079
JNJ	0.0314	0.0200	0.0033	0.1009	0.0348	0.0230	0.0036	0.1118
JPM	0.1257	0.1356	0.0024	0.8304	0.1125	0.1160	0.0023	0.7225
KO	0.0327	0.0260	0.0036	0.1066	0.0340	0.0277	0.0036	0.1122
MCD	0.0543	0.0314	0.0067	0.1525	0.0596	0.0354	0.0071	0.1700
MMM	0.0272	0.0151	0.0076	0.0661	0.0259	0.0145	0.0074	0.0633
MRK	0.0516	0.0212	0.0127	0.1124	0.0475	0.0190	0.0119	0.1001
MSFT	0.0618	0.0439	0.0061	0.1778	0.0626	0.0450	0.0061	0.1820
PFE	0.0638	0.0340	0.0136	0.1535	0.0598	0.0315	0.0133	0.1442
PG	0.0328	0.0295	0.0037	0.1295	0.0312	0.0282	0.0036	0.1246
T	0.0551	0.0415	0.0023	0.1636	0.0550	0.0416	0.0024	0.1639
UTX	0.0342	0.0235	0.0058	0.1066	0.0348	0.0245	0.0058	0.1100
VZ	0.0555	0.0393	0.0023	0.1842	0.0560	0.0399	0.0021	0.1856
WMT	0.0519	0.0428	0.0060	0.1880	0.0564	0.0483	0.0061	0.2112
XOM	0.0481	0.0299	0.0047	0.2079	0.0381	0.0214	0.0042	0.1571

For every firm we report the mean, standard deviation, minimum, and maximum of time-series of the filtered idiosyncratic variances. The spot idiosyncratic variance are filtered from the returns, based on the optimal parameter estimates under the P and Q measures. In the top row we also report the statistics for the time-series of the market spot variance.

Table 1.6: **Correlation Matrix of Idiosyncratic Variances**

	AA	AXP	BA	CAT	CSCO	CVX	DD	DIS	GE	HD	HPQ	IBM	INTC	JNJ	JPM	KO	MCD	MMM	MRK	MSFT	PFE	PG	T	UTX	VZ	WMT	XOM	SPX
AA	1.00																											
AXP	0.88	1.00																										
BA	0.48	0.59	1.00																									
CAT	0.63	0.66	0.68	1.00																								
CSCO	0.04	0.15	0.65	0.18	1.00																							
CVX	0.66	0.62	0.46	0.59	0.03	1.00																						
DD	0.55	0.63	0.83	0.83	0.51	0.54	1.00																					
DIS	0.34	0.42	0.88	0.55	0.76	0.29	0.74	1.00																				
GE	0.85	0.89	0.65	0.60	0.32	0.61	0.61	0.53	1.00																			
HD	0.46	0.49	0.74	0.47	0.64	0.47	0.75	0.77	0.56	1.00																		
HPQ	0.03	0.11	0.64	0.30	0.85	0.08	0.63	0.80	0.23	0.73	1.00																	
IBM	0.01	0.12	0.61	0.36	0.78	0.14	0.70	0.70	0.20	0.71	0.89	1.00																
INTC	0.21	0.29	0.73	0.33	0.91	0.20	0.66	0.83	0.43	0.81	0.88	0.86	1.00															
JNJ	0.10	0.18	0.63	0.47	0.56	0.37	0.71	0.66	0.29	0.66	0.73	0.83	0.68	1.00														
JPM	0.88	0.91	0.61	0.59	0.19	0.69	0.60	0.49	0.93	0.58	0.15	0.15	0.36	0.26	1.00													
KO	0.16	0.26	0.66	0.54	0.65	0.32	0.82	0.69	0.28	0.74	0.82	0.92	0.77	0.88	0.26	1.00												
MCD	0.26	0.30	0.70	0.44	0.50	0.43	0.64	0.76	0.37	0.78	0.68	0.66	0.68	0.79	0.41	0.71	1.00											
MMM	0.09	0.19	0.66	0.57	0.68	0.30	0.79	0.68	0.21	0.63	0.77	0.85	0.72	0.78	0.17	0.90	0.58	1.00										
MRK	0.56	0.68	0.64	0.63	0.37	0.64	0.76	0.51	0.63	0.66	0.46	0.55	0.53	0.66	0.63	0.68	0.61	0.58	1.00									
MSFT	0.19	0.33	0.80	0.47	0.87	0.22	0.79	0.84	0.40	0.79	0.90	0.93	0.91	0.80	0.33	0.89	0.70	0.85	0.62	1.00								
PFE	0.27	0.37	0.68	0.64	0.47	0.53	0.80	0.64	0.42	0.67	0.60	0.75	0.65	0.85	0.42	0.82	0.68	0.77	0.75	0.74	1.00							
PG	0.11	0.23	0.60	0.58	0.58	0.34	0.82	0.60	0.21	0.66	0.72	0.85	0.66	0.84	0.19	0.94	0.63	0.93	0.67	0.83	0.82	1.00						
T	0.34	0.38	0.75	0.43	0.63	0.42	0.68	0.82	0.49	0.83	0.72	0.73	0.81	0.77	0.50	0.74	0.88	0.60	0.61	0.80	0.69	0.61	1.00					
UTX	0.20	0.28	0.76	0.50	0.78	0.25	0.75	0.85	0.36	0.77	0.88	0.83	0.84	0.78	0.32	0.84	0.77	0.80	0.60	0.89	0.68	0.77	0.80	1.00				
VZ	0.42	0.44	0.73	0.45	0.57	0.49	0.68	0.76	0.55	0.84	0.62	0.68	0.76	0.73	0.57	0.69	0.82	0.56	0.65	0.74	0.72	0.58	0.92	0.73	1.00			
WMT	0.18	0.25	0.62	0.55	0.61	0.38	0.80	0.68	0.28	0.79	0.79	0.89	0.75	0.87	0.26	0.95	0.76	0.88	0.67	0.86	0.82	0.93	0.76	0.84	0.73	1.00		
XOM	0.61	0.59	0.51	0.52	0.17	0.91	0.55	0.40	0.66	0.59	0.21	0.26	0.34	0.46	0.70	0.38	0.54	0.32	0.68	0.33	0.55	0.36	0.55	0.37	0.62	0.44	1.00	
SPX	0.76	0.71	0.54	0.46	0.18	0.77	0.48	0.41	0.81	0.54	0.11	0.10	0.33	0.26	0.84	0.20	0.41	0.11	0.50	0.26	0.37	0.12	0.51	0.25	0.59	0.23	0.80	1.00

We present the sample correlation matrix of the spot idiosyncratic variances under the P distribution. Presented in the last row is the pairwise correlation between the spot idiosyncratic variances and the market spot variance. The results under the Q distribution is virtually the same.

Table 1.7: **PCA of Implied Idiosyncratic Variance Levels**

Ticker	1st Component	2nd Component
AA	0.1407	0.4762
AXP	0.2624	0.4128
BA	0.1339	0.0320
CAT	0.1107	0.1676
CSCO	0.3199	-0.3363
CVX	0.0642	0.1247
DD	0.1494	0.0926
DIS	0.1836	-0.0227
GE	0.2323	0.2998
HD	0.2109	0.0100
HPQ	0.2347	-0.2954
IBM	0.2230	-0.1783
INTC	0.2650	-0.1442
JNJ	0.1507	-0.1072
JPM	0.2771	0.3544
KO	0.1775	-0.1058
MCD	0.1435	-0.0485
MMM	0.1114	-0.0485
MRK	0.1038	0.0509
MSFT	0.2414	-0.1407
PFE	0.1351	-0.0211
PG	0.1608	-0.0988
T	0.2238	-0.0279
UTX	0.1594	-0.0485
VZ	0.1953	-0.0095
WMT	0.2209	-0.1148
XOM	0.0751	0.0656
Average	0.1817	0.0125
Minimum	0.0642	-0.3363
Maximum	0.3199	0.4762
Variation explained	58%	23%
Correlation with the average implied idiosyncratic volatility level	99%	4.30%
Corrleation with market implied volatility level	65%	55%

We report the loadings on the first two principal components of the implied idiosyncratic variance levels obtained from the option prices. We also present the percentage of variance explained by the first two components, as well as their correlations with average implied idiosyncratic variance level of all firms, and with the implied variance levels of the market.

Table 1.8: **PCA of Implied Idiosyncratic Variance Moneyness Slopes**

Ticker	**1st Component**	**2nd Component**
AA	0.0876	0.0617
AXP	0.1936	0.1124
BA	0.2002	-0.0591
CAT	0.2364	0.0158
CSCO	0.1815	-0.0252
CVX	0.1805	-0.0590
DD	0.1921	0.0414
DIS	0.1934	-0.0880
GE	0.1340	0.8307
HD	0.2016	-0.1263
HPQ	0.2016	-0.2163
IBM	0.2304	-0.1086
INTC	0.1591	-0.0009
JNJ	0.1870	0.1110
JPM	0.1520	0.0657
KO	0.1856	-0.0971
MCD	0.1271	-0.0581
MMM	0.2559	-0.1304
MRK	0.1877	0.0315
MSFT	0.2277	0.0553
PFE	0.1110	0.2855
PG	0.2846	0.0278
T	0.2084	-0.0101
UTX	0.2159	-0.2536
VZ	0.1580	0.0500
WMT	0.2057	0.0435
XOM	0.1718	-0.0113
Average	0.1878	0.0181
Minimum	0.0876	-0.2536
Maximum	0.2846	0.8307
Variation explained	48%	6%
Correlation with the average implied idiosyncratic volatility moneyness slope	99%	3.40%
Corrleation with market implied volatility moneyness slope	42%	8%

We report the loadings on the first two principal components of the implied idiosyncratic variance moneyness slopes obtained from the option prices. We also present the percentage of variance explained by the first two components, as well as their correlations with average implied idiosyncratic variance moneyness slopes of all firms, and with the implied variance moneyness slope of the market.

Table 1.9: **PCA of Implied Idiosyncratic Variance Term Structure Slopes**

Ticker	1st Component	2nd Component
AA	0.1978	-0.4083
AXP	0.2834	-0.3145
BA	0.2173	0.0617
CAT	0.2234	-0.1185
CSCO	0.2523	0.4405
CVX	0.1471	0.0145
DD	0.1912	-0.0345
DIS	0.2364	0.0845
GE	0.1579	-0.2983
HD	0.1801	0.1234
HPQ	0.2554	0.3761
IBM	0.2317	0.0521
INTC	0.1682	-0.0250
JNJ	0.1208	0.0078
JPM	0.2145	-0.4184
KO	0.1430	-0.0096
MCD	0.1101	0.0332
MMM	0.2007	-0.0189
MRK	0.1747	0.0396
MSFT	0.2170	0.1780
PFE	0.1215	-0.0006
PG	0.1507	0.1007
T	0.1864	0.1056
UTX	0.2334	-0.1498
VZ	0.1546	0.0610
WMT	0.1282	0.0863
XOM	0.1519	0.0114
Average	0.1870	-0.0007
Minimum	0.1101	-0.4184
Maximum	0.2834	0.4405
Variation explained	61%	7%
Correlation with the average implied idiosyncratic volatility term structure slope	99%	-0.12%
Corrleation with market implied volatility term structure slope	78%	-13%

We report the loadings on the first two principal components of the implied idiosyncratic variance term structure slopes obtained from the option prices. We also present the percentage of variance explained by the first two components, as well as their correlations with average implied idiosyncratic variance term structure slopes of all firms, and with the implied variance term structure slope of the market.

Table 1.10: **PCA of Filtered Spot Idiosyncratic Variances**

	Physicla distribution		Risk-neutral distribution	
Ticker	**1st Component**	**2nd Component**	**1st Component**	**2nd Component**
AA	0.2395	0.3321	0.3087	-0.4391
AXP	0.2930	0.3277	0.2901	-0.3195
BA	0.1079	-0.0171	0.1070	0.0177
CAT	0.0918	0.0502	0.0908	-0.0433
CSCO	0.3541	-0.4436	0.3701	0.4550
CVX	0.0591	0.0642	0.0577	-0.0597
DD	0.1144	-0.0075	0.1113	0.0108
DIS	0.1429	-0.0720	0.1398	0.0754
GE	0.2505	0.2019	0.2658	-0.2126
HD	0.2070	-0.0619	0.1930	0.0648
HPQ	0.2562	-0.3417	0.2363	0.3092
IBM	0.1305	-0.1711	0.1331	0.1671
INTC	0.3324	-0.2861	0.3043	0.2617
JNJ	0.0570	-0.0506	0.0672	0.0614
JPM	0.4806	0.4829	0.4354	-0.4078
KO	0.0818	-0.0735	0.0903	0.0809
MCD	0.0981	-0.0490	0.1170	0.0582
MMM	0.0421	-0.0466	0.0423	0.0466
MRK	0.0721	0.0140	0.0670	-0.0095
MSFT	0.1595	-0.1389	0.1708	0.1508
PFE	0.1062	-0.0428	0.1018	0.0469
PG	0.0805	-0.0792	0.0817	0.0791
T	0.1506	-0.0670	0.1592	0.0721
UTX	0.0807	-0.0692	0.0880	0.0759
VZ	0.1436	-0.0366	0.1544	0.0403
WMT	0.1327	-0.1153	0.1598	0.1338
XOM	0.0859	0.0603	0.0611	-0.0357
Average	0.1611	-0.0236	0.1631	0.0252
Minimum	0.0421	-0.4436	0.0423	-0.4391
Maximum	0.4806	0.4829	0.4354	0.4550
Variation explained	57%	30%	55%	31%
Correlation with the average implied idiosyncratic volatility all firms	98%	-10.50%	98%	11.50%
Corrleation with market spot volatilities	68%	48%	67%	-50%

We report the loadings on the first two principal components of the spot idiosyncratic variance levels obtained from the returns. We also present the percentage of variance explained by the first two components, as well as their correlations with average spot idiosyncratic variances of all firms, and with the spot variances of the market.

Table 1.11: **Cross-Sectional Test of Different Asset Pricing Models**

	1	p-val	2	p-val	3	p-val	4	p-val	5	p-val
intercept	0.00141	(0)	0.00145	(0)	0.00137	(0)	0.0012	(0)	0.0011	(0)
Mrkt	-0.001	(0)	-0.0011	(0)	-0.00105	(0)	-0.0009	(0)	-0.0009	(0)
SMB	0.000076	(0.103)			8.60E-05	(0.079)			0.0001	(0.1248)
HML			0.0001	(0.12)	0.0001	(0.0407)			0.0001	(0.0226)
AIV							0.0001	(0.0326)	0.0001	(0.043)
Adj. R2	0.8233		0.796		0.8207		0.8347			0.8431

The table presents the factor risk premiums in the cross sectional regression of the Fama-French 25 portfolios mean excess returns on the time series coefficient estimates. We run the regression for different combination of the factors. Model 1 is the market and the SMB factor. Model 2 is the market and HML factor. Model 3 is the Fama-French model. Model 4 is the market and average idiosyncratic volatility, AIV, factor. Model 6 includes the Fama-French factors as well as the AIV factor. Presented in the table are also the P-values of the estimated risk premiums, as well as the adjusted R^2 for each model.

Table 1.12: **Fama-MacBeth Cross Sectional Test**

	1	t-stat	2	t-stat	3	t-stat
intercept	0.0014	(9.18)	0.0012	(6.27)	0.0011	(7.33)
Mrkt	-0.0011	(-4.04)	-0.0009	(-3.401)	-0.0009	(-3.48)
SMB	0.0001	(0.801)			0.0001	(0.6594)
HML	0.0001	(0.882)			0.0001	(0.938)
AIV			0.0001	(1.974)	0.0001	(2.126)

The table presents the factor risk premiums in the Fama-MacBeth cross sectional test of the Fama-French 25 portfolios mean excess returns. Model 1 is the Fama-French model. Model 2 is the market and average idiosyncratic volatility, AIV, factor. Model 3 includes the Fama-French factors as well as the AIV factor. Presented in the table are also the t-stats of the estimated risk premiums.

Table 1.13: **Mean Annualized Monthly Hedge Portfolio Returns: Calls**

Ticker	ATM	t-stat	OTM1	t-stat	OTM2	t-stat
SPX	-0.39	-15.00	-0.23	-9.00	0.26	1.23
AA	10.90	5.01	6.57	3.74	-4.14	-0.88
AXP	5.18	4.59	3.41	3.72	-2.16	-0.98
BA	0.32	0.66	-0.39	-0.87	-4.65	-1.87
CAT	1.65	2.50	0.43	0.79	-3.80	-1.47
CSCO	21.89	8.41	12.91	5.54	-13.90	-1.20
CVX	13.18	5.38	6.70	3.18	-11.49	-0.81
DD	-2.57	-4.37	-2.51	-4.90	-6.45	-2.55
DIS	6.21	5.84	4.11	4.56	-2.39	-0.68
GE	0.60	0.40	-0.80	-0.63	-7.10	-2.31
HD	0.85	0.82	-0.85	-0.97	-8.87	-1.81
HPQ	8.14	6.08	5.02	4.16	-3.25	-0.56
IBM	0.89	4.45	0.32	1.55	-2.69	-1.79
INTC	12.62	6.80	6.79	4.13	-13.86	-1.45
JNJ	-1.86	-5.65	-1.32	-5.22	-1.65	-1.12
JPM	5.00	2.37	4.00	2.51	0.51	0.13
KO	-1.92	-4.76	-1.21	-3.60	-2.74	-1.74
MCD	-1.39	-1.56	-1.05	-1.37	-0.58	-0.21
MMM	-1.32	-5.28	-1.23	-5.68	-2.19	-1.95
MRK	-5.78	-5.53	-2.13	-0.98	-5.23	-2.43
MSFT	0.48	0.21	-0.77	-0.45	-6.86	-1.58
PFE	-8.90	-4.31	-7.46	-4.50	-12.89	-2.27
PG	-2.24	-6.12	-1.60	-5.46	-2.58	-2.04
T	-13.83	-2.73	-11.35	-2.44	-20.12	-3.08
UTX	0.48	1.62	0.14	0.53	-1.49	-1.39
VZ	-19.71	-2.72	-15.47	-2.55	-14.33	-2.58
WMT	0.57	1.41	0.04	0.10	-3.98	-1.89
XOM	-0.57	-1.59	-0.74	-2.52	-2.69	-1.95
Mean	1.07		0.06		-5.98	
Max	21.89		12.91		0.51	
Min	-19.71		-15.47		-20.12	
No. significant	19		17		7	

The table presents the mean annualized monthly returns of the call portfolios for each firm in our sample. The portfolios are constructed and rebalanced in such a way that are only exposed to the idiosyncratic variance risk of the equity. For each firm we construct portfolios with target days-to-maturity of 30 days and three different target moneyness rations of 1, 1.025, and 1.05, denoted by ATM, OTM1, and OTM2, respectively. For each portfolio we also report the t-statistic of the null hypothesis that the mean return is zero. For each moneyness ratio we also report the minimum, maximum and mean returns, as well as the number of firms with returns significantly different from zero.

Table 1.14: **Mean Annualized Monthly Hedge Portfolio Returns: Puts**

Ticker	ATM	t-stat	OTM1	t-stat	OTM2	t-stat
SPX	0.25	13.32	-0.81	-20.49	-1.14	-11.98
AA	-2.28	-1.53	40.10	8.11	55.60	5.90
AXP	-2.68	-2.90	12.70	7.97	14.87	6.01
BA	1.35	3.64	3.91	6.32	-0.10	-0.10
CAT	0.08	0.14	5.58	6.32	0.58	0.33
CSCO	-11.86	-9.40	28.01	8.01	33.75	4.49
CVX	-7.09	-4.55	12.53	4.69	5.26	0.85
DD	4.17	9.21	7.46	11.63	1.72	1.52
DIS	-0.70	-0.86	9.08	7.94	4.75	2.33
GE	2.50	1.46	16.65	7.16	15.51	4.99
HD	-0.03	-0.04	3.67	3.45	-8.42	-4.77
HPQ	-1.79	-2.45	13.87	8.27	11.80	3.69
IBM	-0.18	-1.45	0.89	3.06	-1.46	-2.82
INTC	-3.71	-3.30	22.69	13.07	16.63	3.94
JNJ	2.08	6.32	0.69	1.32	-1.88	-2.12
JPM	-0.13	-0.10	18.69	8.74	14.88	3.75
KO	2.49	7.51	0.30	0.60	-4.14	-4.88
MCD	2.37	3.28	-0.40	-0.50	-9.59	-8.10
MMM	1.47	7.63	0.54	1.92	-3.22	-7.09
MRK	6.64	5.57	3.94	4.13	-4.92	-3.73
MSFT	2.39	0.88	11.44	4.29	4.61	1.86
PFE	11.83	5.99	10.79	5.26	-8.20	-3.09
PG	2.37	7.05	-0.63	-0.72	-5.52	-4.04
T	8.93	4.86	7.89	4.20	-1.58	-0.56
UTX	0.64	2.70	2.49	5.92	-1.16	-1.44
VZ	12.11	5.40	5.41	3.34	-9.48	-4.14
WMT	0.18	0.72	1.57	2.09	-3.65	-2.80
XOM	1.89	6.54	2.20	5.26	-3.46	-4.24
Mean	1.22		8.97		4.19	
Max	12.11		40.10		55.60	
Min	-11.86		-0.63		-9.59	
No. significant	18		22		20	

The table presents the mean annualized monthly returns of the put portfolios for each firm in our sample. The portfolios are constructed and rebalanced in such a way that are only exposed to the idiosyncratic variance risk of the equity. For each firm we construct portfolios with target days-to-maturity of 30 days and three different target moneyness rations of 1, 0.975, and 0.95, denoted by ATM, OTM1, and OTM2, respectively. For each portfolio we also report the t-statistic of the null hypothesis that the mean return is zero. For each moneyness ratio we also report the minimum, maximum and mean returns, as well as the number of firms with returns significantly different from zero.

Table 1.15: **Idiosyncratic Variance Risk Premiums**

Ticker	Mean (%)	std (%)	Skewness	Kurtosis	t-stat
SPX	-0.48	0.65	-3.675	18.687	-46.34
AA	-1.63	2.17	-3.204	14.014	-47.03
AXP	0.38	0.52	2.843	12.005	46.17
BA	0.30	0.19	0.513	2.595	97.83
CAT	0.19	0.25	3.064	13.951	48.49
CSCO	0.18	0.25	2.447	10.819	45.06
CVX	0.12	0.14	2.759	11.447	53.82
DD	0.39	0.30	1.088	4.296	81.52
DIS	0.33	0.30	1.334	4.819	67.97
GE	-0.10	0.22	-3.757	21.760	-28.17
HD	0.69	0.68	1.328	3.834	63.26
HPQ	1.29	1.23	1.313	4.173	65.45
IBM	0.27	0.24	1.050	3.996	69.55
INTC	1.39	1.27	1.273	3.861	68.98
JNJ	-0.37	0.42	-2.181	9.387	-55.86
JPM	1.47	2.23	3.252	14.382	41.32
KO	-0.14	0.30	-3.273	16.391	-30.26
MCD	-0.62	0.50	-1.267	4.292	-77.25
MMM	0.13	0.09	0.622	2.168	92.34
MRK	0.42	0.31	1.477	5.385	85.48
MSFT	-0.09	0.31	-2.090	10.644	-17.41
PFE	0.43	0.32	1.071	3.770	83.00
PG	0.17	0.15	0.813	2.552	68.90
T	-0.02	0.10	0.032	7.750	-9.92
UTX	-0.08	0.18	-1.890	8.385	-26.36
VZ	-0.08	0.11	-1.599	6.594	-42.97
WMT	-0.51	0.63	-1.721	5.080	-50.95
XOM	1.15	1.10	2.706	11.110	65.71

For each firm we present the annualized first four moments of the time-series of idiosyncratic variance risk premium. The idiosyncratic variance risk premium is calculated as the difference between the excepted integrated idiosyncratic variance under P and Q distributions. In the last column we report the t-stat for the null hypothesis that the average idiosyncratic variance risk premium is zero.

Table 1.16: **Fama-French Regression of the Idiosyncratic Variance Risk Premiums**

Ticker	a	t-stat	b^m	t-stat	b^{smb}	t-stat	b^{hml}	t-stat	b^{mom}	t-stat	R^2 (%)
AA	-0.016	-47.418	0.049	1.746	-0.059	-1.074	0.072	1.316	0.233	6.330	1.03
AXP	0.004	46.532	-0.007	-1.123	0.019	1.419	-0.024	-1.841	-0.053	-6.053	0.98
BA	0.003	97.938	-0.001	-0.472	0.008	1.677	-0.002	-0.477	-0.012	-3.557	0.39
CAT	0.002	48.857	-0.002	-0.728	0.011	1.684	-0.011	-1.791	-0.026	-6.105	1.03
CSCO	0.002	45.470	-0.004	-1.408	0.010	1.645	-0.017	-2.760	-0.026	-6.282	1.08
CVX	0.001	54.253	-0.004	-2.027	0.002	0.638	-0.012	-3.507	-0.014	-5.815	0.94
DD	0.004	81.669	-0.001	-0.280	0.003	0.411	-0.004	-0.499	-0.020	-3.968	0.44
DIS	0.003	68.209	-0.003	-0.766	0.013	1.687	-0.010	-1.345	-0.024	-4.642	0.61
GE	-0.001	-28.374	0.003	0.945	0.003	0.445	-0.002	-0.363	0.019	4.923	0.75
HD	0.007	63.420	-0.017	-1.916	0.006	0.374	-0.015	-0.854	-0.044	-3.822	0.39
HPQ	0.013	65.314	-0.001	-0.048	0.019	0.614	0.043	1.369	-0.005	-0.223	0.06
IBM	0.003	69.428	0.000	-0.007	0.000	-0.044	0.007	1.114	-0.002	-0.542	0.06
INTC	0.014	68.953	-0.022	-1.345	0.034	1.042	0.020	0.616	-0.037	-1.711	0.15
JNJ	-0.004	-55.748	-0.003	-0.643	0.019	1.746	-0.012	-1.114	-0.005	-0.761	0.12
JPM	0.015	41.687	-0.050	-1.756	0.040	0.708	-0.126	-2.240	-0.231	-6.096	0.95
KO	-0.001	-30.142	-0.003	-0.732	0.014	1.817	-0.016	-2.200	-0.010	-1.958	0.26
MCD	-0.006	-77.139	0.003	0.489	0.003	0.245	-0.001	-0.045	0.006	0.716	0.02
MMM	0.001	92.270	0.000	0.308	0.004	1.675	0.002	0.854	-0.003	-1.855	0.21
MRK	0.004	85.590	-0.002	-0.555	0.000	-0.036	0.001	0.163	-0.019	-3.565	0.38
MSFT	-0.001	-17.280	-0.003	-0.719	0.016	2.039	-0.016	-2.057	-0.017	-3.267	0.41
PFE	0.004	83.118	-0.003	-0.784	-0.004	-0.427	-0.012	-1.419	-0.018	-3.333	0.30
PG	0.002	68.832	0.000	-0.102	-0.003	-0.681	0.004	1.056	-0.003	-1.253	0.12
T	0.000	-9.744	-0.001	-0.646	0.003	1.359	-0.007	-2.869	-0.007	-4.124	0.57
UTX	-0.001	-26.210	-0.001	-0.360	0.005	1.071	-0.013	-2.830	-0.010	-3.174	0.41
VZ	-0.001	-42.881	0.001	0.661	0.004	1.571	-0.005	-1.809	-0.003	-1.842	0.26
WMT	-0.005	-50.812	0.000	-0.031	0.010	0.652	-0.033	-2.092	-0.015	-1.440	0.16
XOM	0.012	65.989	-0.031	-2.214	-0.037	-1.344	-0.075	-2.692	-0.075	-4.014	0.53

For every firm we run a time-series regression of the idiosyncratic variance risk premium on the market excess return ($r_t^m - r_t$), Fama-French factors (r^{SMB} and r^{HML}) as well as the momentum factor (r^{mom}), according to the following equations:

$$IVRP_{i,t} = a_i + b_i^m \cdot (r_t^m - r_t) + b_i^{smb} \cdot r_t^{SMB} + b_i^{hml} \cdot r_t^{HML} + b_i^{mom} \cdot r_t^{mom} + \varepsilon_{i,t}$$

We report the regression coefficients, t-stats, and the R^2 of the regression.

Table 1.17: **Explanatory Power of $MVRP$ for $IVRP$**

Ticker	a	t-stat	b^m	t-stat	R^2 (%)
AA	-0.004	-15.305	2.448	68.301	54.35
AXP	0.001	15.893	-0.543	-58.968	47.01
BA	0.002	69.258	-0.146	-35.452	24.28
CAT	0.001	18.929	-0.252	-56.188	44.62
CSCO	0.000	10.205	-0.305	-87.259	66.02
CVX	0.000	22.477	-0.193	-134.361	82.16
DD	0.003	54.028	-0.238	-38.009	26.93
DIS	0.002	40.447	-0.274	-45.610	34.68
GE	0.000	-9.268	0.128	25.516	14.25
HD	0.004	34.883	-0.698	-56.188	44.62
HPQ	0.011	47.068	-0.313	-10.537	2.76
IBM	0.002	49.928	-0.067	-11.636	3.34
INTC	0.010	44.277	-0.769	-27.142	15.82
JNJ	-0.004	-45.858	-0.016	-1.599	0.07
JPM	0.002	6.571	-2.665	-78.404	61.07
KO	-0.002	-30.189	-0.067	-9.406	2.21
MCD	-0.005	-54.568	0.190	15.964	6.11
MMM	0.001	65.639	-0.043	-21.647	10.68
MRK	0.003	57.821	-0.237	-36.155	25.01
MSFT	-0.002	-25.803	-0.138	-18.814	8.28
PFE	0.003	55.357	-0.268	-40.224	29.22
PG	0.001	45.949	-0.074	-20.875	10.01
T	-0.001	-29.710	-0.075	-34.628	23.43
UTX	-0.001	-30.682	-0.064	-14.963	5.40
VZ	-0.001	-37.475	-0.013	-4.771	0.58
WMT	-0.005	-39.185	0.047	3.046	0.24
XOM	0.005	37.687	-1.273	-73.038	57.65

For each firm we run a time-series regression of the idiosyncratic variance risk premium on the market variance risk premium, according to the following equation:

$$IVRP_{i,t} = a_i + b_i \cdot MVRP_t + \varepsilon_{i,t}$$

We report the regression coefficients and the associated t-stats, along with the regression R^2.

Table 1.18: **Explanatory Power of $AVRP$ and F_{orth} for $IVRP$**

Ticker	a	**t-stat**	b^{AIVRP}	**t-stat**	b^{orth}	**t-stat**	R^2 (%)
AA	0.001	3.049	-8.238	-74.027	1.439	29.341	61.81
AXP	-0.001	-15.524	2.291	102.929	-0.093	-9.436	73.17
BA	0.001	56.727	0.861	113.724	0.095	28.534	77.82
CAT	0.000	-6.436	1.017	83.419	-0.065	-12.107	64.46
CSCO	0.000	-14.455	1.038	103.983	-0.173	-39.429	75.94
CVX	0.000	-7.213	0.608	155.640	-0.133	-77.314	88.52
DD	0.001	31.299	1.288	95.511	0.099	16.657	70.58
DIS	0.000	9.501	1.461	165.524	0.103	26.512	87.76
GE	0.000	11.882	-0.709	-49.298	-0.060	-9.533	39.15
HD	0.000	5.534	3.082	110.011	-0.051	-4.156	75.57
HPQ	0.005	24.568	3.737	51.893	1.132	35.678	50.30
IBM	0.001	28.054	0.720	49.302	0.202	31.427	46.59
INTC	0.003	19.145	5.062	83.300	0.760	28.400	66.41
JNJ	-0.003	-28.704	-0.540	-17.055	-0.301	-21.581	16.19
JPM	-0.006	-27.189	10.080	121.447	-1.021	-27.909	79.85
KO	-0.001	-17.532	-0.130	-5.606	-0.213	-20.959	10.73
MCD	-0.003	-34.296	-1.338	-38.270	-0.229	-14.899	30.09
MMM	0.001	45.379	0.298	60.871	0.051	23.510	52.08
MRK	0.002	35.471	1.232	75.821	0.074	10.380	59.92
MSFT	-0.001	-17.039	0.150	6.130	-0.235	-21.843	11.61
PFE	0.001	32.376	1.338	82.768	0.056	7.918	63.83
PG	0.001	22.912	0.524	59.050	0.092	23.479	50.76
T	0.000	-22.412	0.145	19.914	-0.097	-30.186	25.03
UTX	-0.001	-17.900	-0.025	-1.853	-0.155	-25.583	14.38
VZ	-0.001	-22.442	-0.101	-11.740	-0.079	-20.669	12.60
WMT	-0.003	-22.409	-0.990	-20.644	-0.380	-17.997	16.07
XOM	0.003	18.054	4.141	75.481	-0.818	-33.851	63.59

We first regress the market variance risk premium on the average variance risk premiums of all firms to get the component,F_{RP}^{orth}, of the market variance risk premium that is orthogonal to the average variance risk premium. Then we regress the idiosyncratic variance risk premium of each firm on the average variance risk premium,$AVRP$, and the component of the market variance risk premium orthogonal to it.

$$IVRP_{i,t} = a_i + b_{AIVRP,i} \cdot AIVRP_t + b_{orth,i} \cdot F_{RP}^{orth} + \varepsilon_{i,t}$$

We report the regression coefficients and the associated t-stats, along with the regression R^2.

Table 1.19: **Estimated Unconditional Betas**

Ticker	**Joint**	**OLS**	**P**	**Q**
AA	1.09	1.30	1.41	0.99
AXP	1.23	1.45	1.08	0.99
BA	1.06	0.93	1.11	0.98
CAT	1.11	1.08	1.28	0.99
CSCO	1.08	1.41	1.15	0.97
CVX	0.81	0.81	0.94	0.91
DD	0.96	0.99	1.20	1.00
DIS	1.07	1.05	1.05	1.00
GE	0.86	1.19	1.03	1.00
HD	1.06	1.09	1.17	0.99
HPQ	1.14	1.13	1.01	0.95
IBM	0.97	0.91	0.73	0.91
INTC	0.99	1.32	1.24	0.98
JNJ	0.60	0.57	0.48	0.67
JPM	0.85	1.55	1.12	1.01
KO	0.75	0.58	0.66	0.75
MCD	0.55	0.60	0.71	0.74
MMM	0.98	0.77	0.85	0.90
MRK	0.89	0.77	0.74	0.91
MSFT	1.04	1.09	0.90	1.00
PFE	0.78	0.81	0.93	0.95
PG	0.75	0.57	0.57	0.70
T	0.81	0.79	0.77	0.83
UTX	1.03	0.95	1.04	0.96
VZ	0.70	0.74	0.74	0.81
WMT	0.70	0.75	0.73	0.70
XOM	0.75	0.80	0.94	0.92
Average	0.91	0.96	0.95	0.91
Min	0.55	0.57	0.48	0.67
Max	1.23	1.55	1.41	1.01

The first column reports the beta estimates based on the joint likelihood function. The second column presents the OLS beta estimate over the whole sample. The betas in the third column are estimated by fitting the equity return dynamics (1.2.2) to equity returns under the P measure, while the betas in the last column are estimated by fitting the model option price to the observed option prices under the Q measure.

Table 1.20: **Conditional Beta Estimates**

	Conditional beta				
Ticker	**mean**	**std**	**min**	**max**	**unconditional beta**
AA	1.26	0.34	0.47	2.94	1.09
AXP	1.44	0.37	0.44	3.86	1.23
BA	1.26	0.35	0.24	2.07	1.06
CAT	1.33	0.36	0.22	2.50	1.11
CSCO	1.62	0.52	0.00	3.14	1.08
CVX	0.95	0.27	0.25	1.60	0.81
DD	1.12	0.33	0.22	2.09	0.96
DIS	1.25	0.37	0.20	2.12	1.07
GE	1.07	0.38	0.00	2.34	0.86
HD	1.25	0.32	0.15	2.40	1.06
HPQ	1.42	0.39	0.01	2.78	1.14
IBM	1.29	0.32	0.01	2.22	0.97
INTC	1.36	0.47	0.10	2.72	0.99
JNJ	0.90	0.29	0.23	1.81	0.60
JPM	0.51	0.39	0.00	2.64	0.85
KO	1.01	0.34	0.08	1.99	0.75
MCD	0.49	0.22	0.00	1.70	0.55
MMM	1.10	0.29	0.04	1.83	0.98
MRK	1.08	0.42	0.09	1.96	0.89
MSFT	1.20	0.62	0.00	3.20	1.04
PFE	0.99	0.40	0.17	2.20	0.78
PG	0.98	0.31	0.05	2.00	0.75
T	0.81	0.40	0.00	2.36	0.81
UTX	1.22	0.24	0.12	1.95	1.03
VZ	0.65	0.31	0.00	2.02	0.70
WMT	1.07	0.37	0.09	2.41	0.70
XOM	0.95	0.22	0.02	1.49	0.75

We estimate conditional betas on a daily frequency for every firm. These estimates are based on the estimated market and equity models structural parameters, and the estimated spot market variance, conditional on the option prices observed on each day. We present the mean, standard deviation, minimum, and maximum of the time-series of the conditional betas. In the last column we also report our estimate of the unconditional beta baed on the joint likelihood estimation for comparison.

Figure 1.1: **Time-Series Coefficients of the AIV Factor**

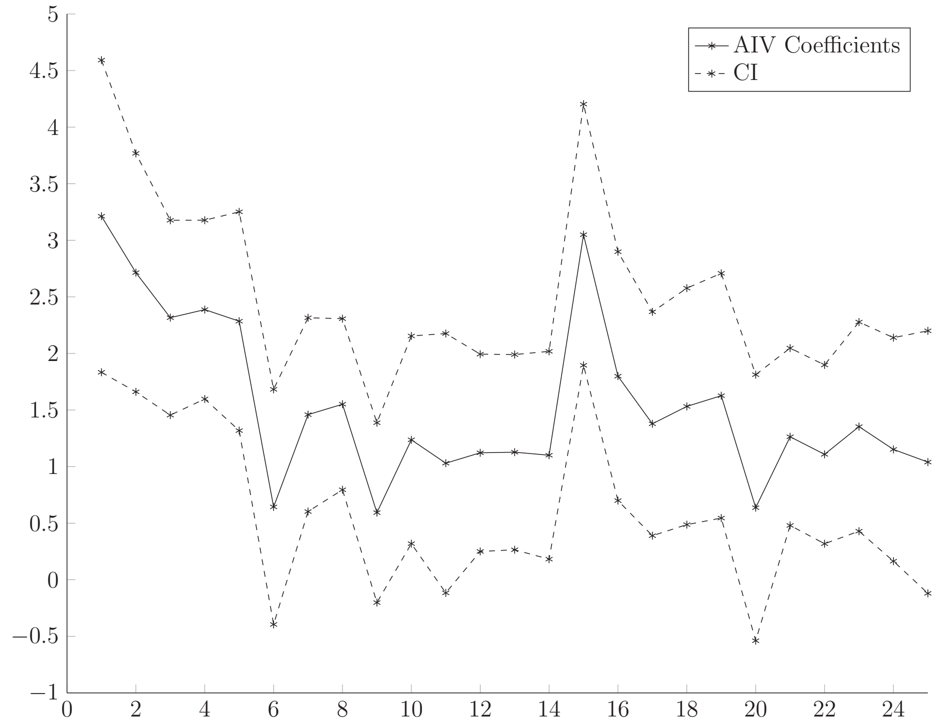

For each of the 25 Fama-French portfolios we run the time series regression (1.4.4) of the portfolio's excess return onto the Fama-French factors and the average idiosyncratic volatility factor (AIV). The solid line plots the coefficient estimates, b^{AIV}, and the dashed lines represent the 95% confidence intervals.

Chapter 2

Valuing Catastrophe Derivatives under Limited Diversification: a Stochastic Dominance Approach

2.1 Introduction

This paper examines the valuation of catastrophe instruments (CAT), financial contracts whose payoffs depend on the occurrence of a rare event that causes significant losses to a category of economic agents. A literature has been developed recently on the valuation of such contracts which, however, leaves many questions unanswered. We demonstrate here that the value of such instruments is crucially dependent on the assumptions made about the agents that would include such instruments in their portfolio. Plausible alternative assumptions may, in turn, bring major changes in the value of the instruments. We introduce the stochastic dominance methodology for valuing the catastrophe instruments and illustrate its application with numerical examples in the case of a European call option and a reinsurance contract.

Catastrophe insurance derivatives (futures and options) were introduced by the Chicago Board of Trade as early as 1992 as hedging instruments for the risk faced by insurers. They have not had much success as traded instruments in organized markets, although there is apparently active over-the-counter trading in them. As of mid-2010 the only such instruments listed in the Chicago Mercantile Exchange (CME) were futures and options contracts on the CHI, the CME Hurricane Index (formerly the Carville HI) for various parts of the US, but there was very little open interest and very few trades in the recent past. This does not eliminate the need for a valuation methodology for financial claims contingent on CAT events, but it does raise some questions about the assumptions on financial market equilibrium adopted during their valuation.

The importance of the valuation assumptions stems from the fact that in the presence of rare events financial markets are incomplete. Merton (1976), who was the first to note this property, suggested a contingent claims valuation method in which the rare event risk would be fully diversifiable and as such could be treated as unsystematic risk and not priced. The Merton assumption has been accepted by several authors valuing CAT financial instruments, who assume that there is an efficient reinsurance market that diversifies the CAT event risk[1], but it is not to be accepted as a panacea, since it is clearly not applicable in many situations. As Duan and Yu (2005, p. 2441) note, catastrophe risk cannot be hedged if it has economy-wide implications.[2] More recently, Ibragimov et al. (2009) note that the efficient reinsurance assumption is not satisfied in real markets, since insurers specialize in geographical regions and particular types of coverage, thus putting individual firms at high risk to specific catastrophe events.[3] These same authors show that catastrophe risk may be non diversifiable even if it does not have economy-wide impact, depending on the characteristics of the probability distribution of the event risk.

In most of the literature the catastrophe event is modeled as a pure jump process, with Poisson arrivals and amplitudes that follow a given unspecified distribution.[4] Unfortunately the arbitrage methodology that is used almost exclusively in financial valuation is not particularly suited to the pricing of cash flows that depend on such rare events. This methodology uses most often continuous time valuation and a language that relies heavily on mathematical formulation, often ignoring troublesome elements of the underlying economic reasoning.[5] The difficulties are compounded by the fact that the underlying catastrophe process does not generally correspond to a traded financial instrument.

If catastrophe risk cannot be hedged then the dominant approach to the valuation of the rare event risk is to assume some kind of general equilibrium model that embodies strong assumptions like weak aggregation or the existence of a representative investor. This approach is used almost exclusively in the option pricing literature, and has only recently been introduced explicitly into the catastrophe instrument valuation.[6] The representative investor is almost always represented by a constant proportional risk aversion (CPRA) utility function, and the valuation results depend on the investor's risk parameter. This parameter enters into

[1]See, for instance, Dassios and Jang (2002), Duan and Yu (2005), Lee and Yu (2007), Lin and Wang (2009), and Chang et al. (2010).

[2]This was already known from the option pricing literature in the presence of event risk. See, for instance, Bates (1991).

[3]Similar remarks were also made by Barrieu and Loubergé (2009), who rely on behavioral considerations to account for the lack of diversification.

[4]See Geman and Yor (1997), Froot (2001), and Muermann (2003)

[5]See, for instance, the discussion on market incompleteness in Geman and Yor (1997, p. 187) and in Bakshi and Madan (2002, pp. 107-108)

[6]See Chang et al. (2010), who use the representative investor in deriving the valuation expressions but assume that the event risk is fully diversifiable in their numerical work.

the transformation of the physical distribution of the occurrence of the catastrophe event (the P-distribution) into the risk neutral distribution used in the valuation of the financial instruments (the Q-distribution).

Apart from the fact that such a parameter is notoriously difficult to estimate and the values appearing in the literature vary from 0 to more than 50,[7] it is also completely unobservable in the case of most catastrophe instruments, which do not trade in organized financial markets. Most applications in finance estimate simultaneously the P- and Q-distributions, from the underlying security and the option market respectively. This presupposes liquid markets and sufficient alternative strike prices to generate the Q- distribution, which are found in stock index options but are unavailable in catastrophe derivatives.[8]

We solve this problem by adopting an alternative approach to the valuation of contingent claims, that of *stochastic dominance* (SD) which uses a much weaker set of assumptions than equilibrium. Unlike equilibrium, SD does not rely on the existence of a representative investor, let alone one with a CPRA utility function. Its only assumption is a pricing kernel that is monotone with respect to the contingent claim's payoffs. A sufficient condition for such monotonicity to be satisfied is the existence of a set of investors that hold portfolios comprised only of the underlying asset and other assets independent of it, as well as the riskless asset. The SD approach, originally introduced by Perrakis and Ryan (1984), Ritchken (1985), Levy (1985), Perrakis (1986, 1988), and Ritchken and Kuo (1988), has recently been extended to incorporate proportional transaction costs.[9]

The major advantage of SD in valuing CAT derivatives is the fact that the only information that it needs in order to value the contingent claim comes from the underlying asset's market, from the P-distribution. Instead of a single value for the contingent claim SD computes an

[7]See the survey article by Kocherlakota (1996).

[8]Muermann (2003) states that the link between the P- and Q-distributions can be found from simultaneously priced insurance derivatives and insurance contracts. Similarly, Chang et al. (2010) claim that it is possible to use observed reinsurance premiums in order to calibrate the prices of the catastrophe derivatives. These claims assume that markets are in equilibrium and that derivatives and underlying assets are "correctly" priced with respect to each other. As Constantinides et al. (2009) and Constantinides et al. (2011) show, this is certainly not the case in S&P 500 index options, which makes it highly unlikely that it would be true in CAT instruments.

[9]See Constantinides and Perrakis (2002, 2007). An empirical application is in Constantinides et al. (2011).

upper and a lower bound on the value, which are reservation-purchase and reservation-write prices for the contingent claim. Violation of either one of these bounds allows the option holder to adopt a trading strategy that would increase the expected utility of *any* risk averse investor satisfying the conditions that guarantee a monotone pricing kernel. The bounds are derived in a discrete time multiperiod context, and are eventually extended to continuous time by a limiting argument.

We apply our methodology to the valuation of two contingent claims, both indexed on hurricane events. The first is a call option on hurricane intensity measured by the CHI, similar to the ones offered by the CME, while the second is a reinsurance contract on an insurer's assets. The two claims are different because the former is contingent on a pure jump while the latter on a jump diffusion process. The claim's value under the Merton (1976) assumption of fully diversifiable CAT event risk lies below the two bounds in the first case and coincides with the lower bound in the reinsurance contract. We also show, using realistic data from the CHI distribution, that adopting the Merton (1976) assumption seriously underestimates the value of the CAT call option.

In the remainder of this section we complete the literature review. As noted, almost all earlier studies adopt the Merton (1976) assumption that the CAT event risk is fully diversifiable. The differences in the valuation expressions come from alternative specifications of the continuous time dynamics of the CAT event and the associated financial claims on it. Geman and Yor (1997) and Muermann (2003) model the claim arrival process as a mixed jump-diffusion, in which the jump component has a fixed amplitude. Dassios and Jang (2002) use the Cox process to represent the amplitude of the CAT event. Duan and Yu (2005) use similarly jump-diffusion dynamics with stochastic interest rates and a lognormally distributed jump amplitude to model the contractual liability of an insurer facing catastrophe events. The stochastic interest rate feature is also present in the mixed jump-diffusion model of Jaimungal and Wang (2006). By contrast, Lee and Yu (2007) use a diffusion process with stochastic interest rates for the insurer's assets and liabilities and model separately the CAT event as a Poisson process with lognormal jump amplitudes, which they value using the Merton (1976) assumption. Lin and Wang (2009) use a jump diffusion model of asset dynamics to represent the aggregate catastrophe losses and apply it to value a perpetual American put option. Last, Chang et al. (2010) use a trinomial discrete time model to value the claim arrival process and a representative investor to evaluate the risk neutral distribution, but state, correctly, that "the introduction of utility functions to resolve the problems generated by the incomplete nature of the market is, in fact, often impractical as they are too much preference-specific";[10] for this reason the authors use the Merton (1976) assumption in their numerical work. As this paper shows, stochastic dominance can solve the incompleteness problem without resorting to a representative investor with an unobservable preference parameter.

2.2 The Model in Discrete Time

Except for the trivial case where the underlying asset return takes only two values in discrete time, the market for rare event risk is incomplete in a discrete time context. The valuation of an option in such a market cannot yield a unique price. Our market equilibrium for a CAT option is derived under the following set of assumptions that are sufficient for our results:

- There exists at least one utility-maximizing risk averse investor (the *trader*) in the economy who holds a portfolio containing an index and the riskless asset, with the index return depending linearly on the intensity of the CAT event.
- This particular investor is marginal in the option market.
- The riskless rate is non-random.[11]

[10] See their footnote 6, p. 28.

[11] Although this assumption may not be justified in practice, its effect on option values is generally recognized as minor in short- and medium-lived options. It has been adopted without any exception in *all* equilibrium based jump-diffusion index and stock option valuation models that have appeared in the literature. See the comments in Bates (1991, p. 1039, note 30) and Amin and Ng (1993, p. 891). In order to evaluate the various features of option pricing models, Bakshi, Cao, and Chen (1997) applied without deriving it a risk-neutral model featuring stochastic interest rate, stochastic volatility and jumps. They found that stochastic interest rates offer no goodness of fit improvement.

Each trader holds a portfolio of in the riskless asset and in the index by maximizing recursively the expected utility of final wealth[12] over the periods $t = 0, 1, \ldots, T'$ of length Δt. Let S_t denote the current value of the index and R the return of the riskless asset per period. Time is initially assumed discrete $t = 0, 1, \ldots, T$, with intervals of length Δt, implying that $R = e^{r\Delta t} = 1 + r\Delta t + o(\Delta t)$, where r denotes the interest rate in continuous time. In each interval the index has returns

$$\frac{S_{t+\Delta t} - S_t}{S_t} \equiv z_{t+\Delta t} = \nu_{t+\Delta t} + \gamma\eta\Delta N \tag{2.2.1}$$

We assume that $E[z_{t+\Delta t}|S_t] \geq R - 1$. The term η represents the amplitude of the CAT event, where N is a Poisson counting process with intensity with intensity λ. The return $z_{t+\Delta t}$ is the convolution of $\eta_{t+\Delta t} = \gamma\eta$ with probability $\lambda\Delta t$ and $\eta_{t+\Delta t} = 0$ with probability $1 - \lambda\Delta t$, and of $\nu_{t+\Delta t}$, which is the index return component that is independent of $\eta_{t+\Delta t}$ and whose distribution may depend onS_t. We assume that the CAT events are independent and identically distributed ($i.i.d.$); the amplitude η may be an earthquake or hurricane intensity measure, and the parameter denotes the impact of the event on the trader's portfolio. If $\gamma = 0$ then the CAT event risk is not part of the index return risk and we have a case similar to Merton (1976). We assume, therefore, that the coefficient γ is negative.

We note that this modeling of the impact of the CAT event on the trader's portfolio is economically different from the more common approach[13] of representing the claims arrival as a mixed jump-diffusion process. As shown further on in this section, it is relatively easy to extend the model to cover the inclusion of a diffusion component to the contribution of the CAT event to the trader's portfolio returns, by replacing the pure jump term in (3.2.1) by a mixed jump-diffusion process. Nonetheless, such a diffusion component should not be included in the derivative instrument's payoff if the latter is indexed to the amplitude of the CAT event given that it occurs as it happens with the hurricane derivatives, the only ones still trading in the financial markets. If, on the other hand, the derivative to be valued

[12]The results are unchanged if the traders are assumed to maximize the expected utility of the consumption stream.

[13]See, for instance, Geman and Yor (1997), Jaimungal and Wang (2006), and most of the references surveyed in the previous section.

is a claims reinsurance contract then the instrument's payoff may also include the diffusion component. These two cases will be examined sequentially in this section.

Case A: Valuation of a call option on the cumulative amplitude of the CAT event

Let $A_t \equiv \kappa \sum_{\tau=0}^{\tau=N_t} \eta_\tau$, $\kappa > 0$ denote the accumulated losses from the catastrophe event, assumed proportional to the CAT amplitude. If the trader also has a marginal open position in a given CAT call option with value $C_t(A_t)$ and with terminal condition $C_T = (A_T - K)^+$ at option expiration time $T < T'$[14] then the following relations characterize market equilibrium in any *single* trading period $(t, t + \Delta t)$, assuming no transaction costs and no taxes:

$$
\begin{aligned}
&E[Y(z_{t+\Delta t})|S_t] = 1 \\
&E[(1 + z_{t+\Delta t})Y(z_{t+\Delta t})|S_t] = R \\
&C_t(A_t) = \frac{1}{R}E[C_{t+\Delta t}(A_t + \kappa\eta_{t+\Delta t})Y(z_{t+\Delta t})|S_t]
\end{aligned}
\tag{2.2.2}
$$

In (3.2.2) $Y(z_{t+\Delta t}) \geq 0$ represents the pricing kernel, the state-contingent discount factor or normalized marginal rate of substitution of the trader evaluated at her optimal portfolio choice. Because of the assumed risk aversion and portfolio composition of our traders it can be easily seen that the pricing kernel would be monotone non-increasing in the stock return $z_{t+\Delta t}$ for every $t = 0, \ldots, T$.

Let now $\hat{Y}_t(\eta_{t+\Delta t}) = E[Y(z_{t+\Delta t})|S_t, \eta_{t+\Delta t}]$, $\bar{\nu}_t = E[Y(z_{t+\Delta t})\nu_{t+\Delta t}|S_t]$, $\hat{\nu}_t = E[\nu_{t+\Delta t}|S_t]$, and $\hat{\eta} = E[\eta_{t+\Delta t}|\eta \neq 0], E[\eta_{t+\Delta t}] = \lambda\hat{\eta}\Delta t$. The following relations are obvious results derived from (3.2.1)-(3.2.2).

$$
\begin{aligned}
&E[\hat{Y}_t(\eta_{t+\Delta t})] = 1 \\
&E[\hat{Y}_t(\eta_{t+\Delta t})\eta_{t+\Delta t}|S_t]] = \frac{1 + \bar{\nu}_t - R}{-\gamma} \equiv \phi = \alpha_t\Delta t + o(\Delta t), \\
&C_t(A_t) = \frac{1}{R}E[C_{t+\Delta t}(A_t + \kappa\eta_{t+\Delta t})\hat{Y}_t(\eta_{t+\Delta t})|S_t]
\end{aligned}
\tag{2.2.3}
$$

Observe that since $\gamma < 0$ and $Y(z_{t+\Delta t})$ is non-increasing, the pricing kernel $\hat{Y}(\eta_{t+\Delta t})$ is now non-decreasing in the intensity $\eta_{t+\Delta t}$ of the CAT event. Similarly, observe that

[14] All results in this paper are derived for call options. They are applicable without reformulation to European put options, either directly or through put-call parity.

$R \leq 1 + \bar{\nu}_t \leq 1 + \hat{\nu}_t$; if $R = 1 + \hat{\nu}_t$ then $\phi = 0$ and we are again in the case of a CAT event risk that is fully diversifiable; this is not possible here, since by assumption $\eta_{t+\Delta t}$ takes only non-negative values, and $E[\hat{Y}_t(\eta_{t+\Delta t})\eta_{t+\Delta t}|S_t] \geq E[\eta_{t+\Delta t}] = \lambda\hat{\eta}\Delta t > 0$. We can thus apply the modified LP approach originally introduced by Ritchken (1985) and extended by Perrakis (1988) in the case of a trinomial distribution in order to derive the following result, whose proof is in the appendix.

Lemma 2.1. *If the option price $C_t(A_t)$ is convex for any t then it lies within the following bounds:*

$$\frac{1}{R}E^{L_t}[C_{t+\Delta t}(A_t + \kappa\eta_{t+\Delta t})|S_t] \leq C_t(A_t) \leq \frac{1}{R}E^{U_t}[C_{t+\Delta t}(A_t + \kappa\eta_{t+\Delta t})|S_t] \tag{2.2.4}$$

where E^{U_t} and E^{L_t} denote respectively expectations taken with respect to the following distributions. If $\eta_{t+\Delta t}$ follows a discrete distribution (q_i, η_i),$i = 0, 1, \ldots, n$,[15] *with the amplitudes ordered in terms of size, then:*

$$\begin{aligned} U_{i,t} &= \frac{\eta_n - \phi}{\eta_n - \lambda\hat{\eta}\Delta t} q_{i,t+\Delta t}, i = 0, 1, \ldots, n-1 \\ U_{n,t} &= \frac{\eta_n - \phi}{\eta_n - \lambda\hat{\eta}\Delta t} q_n + \frac{\phi - \lambda\hat{\eta}\Delta t}{\eta_n - \lambda\hat{\eta}\Delta t} \end{aligned} \tag{2.2.5a}$$

$$\begin{aligned} L_{i,t} &= 0, i = 0, 1, \ldots, h-1 \\ L_{h,t} &= \frac{\hat{\eta}_{h+1} - \phi}{\hat{\eta}_{h+1} - \hat{\eta}_h} \frac{q_{h,t}}{\sum_{k=h}^{k=n} q_k} \\ L_{i,t} &= \frac{\hat{\eta}_{h+1} - \phi}{\hat{\eta}_{h+1} - \hat{\eta}_h} \frac{q_{i,t}}{\sum_{k=h}^{k=n} q_k} + \frac{\phi - \hat{\eta}_h}{\hat{\eta}_{h+1} - \hat{\eta}_h} \frac{q_{i,t}}{\sum_{k=h+1}^{k=n} q_k}, \ i = h+1, \ldots, n \end{aligned} \tag{2.2.5b}$$

Where $\hat{\eta}_i = \frac{\sum_{k=i}^{k=n} q_k \eta_k}{\sum_{k=i}^{k=n} q_k}$, $i = 0, 1, \ldots, n$ and $\hat{\eta}_{h+1} \geq \phi \geq \hat{\eta}_h$.

For a continuous distribution $P(\eta)$ of the CAT event amplitude the expectations are taken with respect to the following distributions:

[15] See the appendix for the definition of the distribution. We assume that the state $\eta_{t+\Delta t} = 0$ implies that the CAT event does not take place.

$$U(\eta) = \begin{cases} P(\eta) & \textit{with probability} \quad \frac{\eta_n - \phi}{\eta_n - \lambda\hat{\eta}\Delta t} \\ 1_{\eta_n} & \textit{with probability} \quad \frac{\phi - \lambda\hat{\eta}\Delta t}{\eta_n - \lambda\hat{\eta}\Delta t} \end{cases} \tag{2.2.6}$$
$$L(\eta) = P(\eta|\eta \geq \eta^*),\ E[\eta|\eta \geq \eta^*] = \phi$$

Proof. See Appendix A. □

The results allows us to prove the following.

Proposition 2.1. *Under the conditions of Lemma 1 all admissible CAT call option prices lie between the upper and lower bounds $\overline{C}_t$ and $\underline{C}_t$, evaluated by the following recursive expressions*

$$\begin{aligned} \overline{C}_T(A_T) = \underline{C}_T(A_T) &= (A_T - K)^+ \\ \overline{C}_t(A_t) &= \frac{1}{R} E^{U_t}[\overline{C}_{t+\Delta t}(A_t + \kappa\eta_{t+\Delta t})|S_t] \\ \underline{C}_t(A_t) &= \frac{1}{R} E^{L_t}[\underline{C}_{t+\Delta t}(A_t + \kappa\eta_{t+\Delta t})|S_t] \end{aligned} \tag{2.2.7}$$

where E^{U_t} and E^{L_t} denote expectations taken with respect to the distributions given in (3.2.5) or (3.2.6).

Proof. We use induction to prove that (3.2.7) yields expressions that form upper and lower bounds on admissible option values. It is clear that (3.2.7) holds at T and that $\overline{C}_T(A_T)$ and $\underline{C}_T(A_T)$ are both convex in A_T. Assume now that $\overline{C}_{t+\Delta t}(A_t+\kappa\eta_{t+\Delta t})$ and $\underline{C}_{t+\Delta t}(A_t+\kappa\eta_{t+\Delta t})$ are respectively upper and lower bounds on the convex function $C_{t+\Delta t}(A_t+\kappa\eta_{t+\Delta t})$, implying that:

$$\frac{1}{R}\underline{C}_{t+\Delta t}(A_t + \kappa\eta_{t+\Delta t}) \leq C_{t+\Delta t}(A_t + \kappa\eta_{t+\Delta t}) \leq \frac{1}{R}\overline{C}_{t+\Delta t}(A_t + \kappa\eta_{t+\Delta t}) \tag{2.2.8}$$

By Lemma 1 we also have:

$$\frac{1}{R}E^{L_t}[C_{t+\Delta t}(A_t + \kappa\eta_{t+\Delta t})] \leq C_t(A_t) \leq \frac{1}{R}E^{U_t}[C_{t+\Delta t}(A_t + \kappa\eta_{t+\Delta t})] \tag{2.2.9}$$

(3.2.8) and (3.2.9), however, imply that:

$$\underline{C}_t(A_t) = \frac{1}{R}E^{L_t}[\underline{C}_{t+\Delta t}(A_t + \kappa\eta_{t+\Delta t})] \leq C_t(A_t) \leq \frac{1}{R}E^{U_t}[\overline{C}_{t+\Delta t}(A_t + \kappa\eta_{t+\Delta t})] = \overline{C}_t(A_t) \tag{2.2.10}$$

□

In Proposition 1 it is shown that under our mild set of assumptions the unique Q-distribution of the traditional arbitrage or equilibrium methods is replaced by two risk-neutral distributions that yield values bracketing the economically correct price of the derivative. Observe that both distributions transform the P-distribution of the catastrophe event in such a way that its expected amplitude becomes equal to ϕ. Further, these two boundary Q-distributions may be easily estimated from data on the CAT event: we note from (3.2.5)-(2.2.10) that with the exception of the parameter ϕ the bounds in Proposition 1 are derived exclusively with data from the P-distribution, in this case the distribution of the jump intensity and amplitude. As for ϕ, its boundary values are $\lambda\Delta t\hat{\eta}$ and $\dfrac{1+\hat{\nu}_t - R}{-\gamma}$.[16]In other words, it is bounded between the expected CAT event amplitude and the risk premium on the non-CAT trader's portfolio per unit CAT amplitude, which would be typically the expected return on the market index. Since the bounds are clearly increasing in ϕ, an upper bound independent of ϕ can be found if we replace ϕ with $\dfrac{1+\hat{\nu}_t - R}{-\gamma}$. In such a case the lower bound is found for $\phi = \lambda\Delta t\hat{\eta}$, yielding the discounted expectation of the option payoff under the P-distribution, consistent with a fully diversifiable CAT risk. Hence, without any further assumptions on ϕ the value of the option lies between this discounted expected payoff and a higher value dependent on the market risk premium. In other words, the derivative is a "negative beta" asset, as it befits an insurance instrument.

Case B: Valuation of a reinsurance contract

The reinsurance contract paying the loss in asset value below a reference point is like a protective put option on the insurer's assets, including the premiums and the accumulated

[16]The parameter ϕ is observable if there are traded future contracts on the CAT event as with the CME Hurricane Index, provided their maturity exceeds the expiration date of the option.

claims for the insured losses incurred in the period $[0, T]$, with the exercise price K playing the role of the reference point on the loss reimbursements. The main difference with the CAT event option is that there are additional randomly occurring cash flows from premium income and from additional insured risks that accrue continuously and are not associated with the CAT event. Without loss of generality we assume that these cash flows can be incorporated into the random term $\nu_{t+\Delta t}$. In such a case it is the entire return in equation (3.2.1) that becomes part of the payoff of the reinsurance contract. Let also I_t denote the accumulated assets of the insurer subject to the reinsurance contract, such that $I_{t+\Delta t} = I_t(1 + wz_{t+\Delta t})$, $w > 0$, with the payoff of the contract at maturity equal to $P_T = (K - I_T)^+$. Hence, the reinsurance contract is a contingent claim on I_t.

As before, we assume $E[z_{t+\Delta t}|S_t] \geq R - 1$. The market equilibrium equations (3.2.2) for the period $(t, t + \Delta t)$ become now

$$\begin{aligned} &E[Y(z_{t+\Delta t})|S_t] = 1 \\ &E[(1 + z_{t+\Delta t})Y(z_{t+\Delta t})|S_t] = R \\ &P_t(I_t) = \frac{1}{R}E[P_{t+\Delta t}(I_t(1 + wz_{t+\Delta t}))Y(z_{t+\Delta t})|S_t] \end{aligned} \tag{2.2'}$$

Note that $Y(z_{t+\Delta t})$ is now non-increasing, and $P_{t+\Delta t}$ in (2.2)' is a decreasing and convex function of $z_{t+\Delta t}$. Hence, the LP approach of Ritchken (1985) can be applied with little reformulation. Now the P-distribution is the distribution of the entire return $z_{t+\Delta t}$, assumed discrete without loss of generality, (z_i, p_i), $i = 1, \ldots, m$, where $m > n$. We define the ordered states $\hat{z}_i$ in the same way as $\hat{\eta}_i$:

$$\hat{z}_i = \frac{\sum_{k=1}^{k=i} p_k z_k}{\sum_{k=1}^{k=i} p_k}, \; i = 1, \ldots, m \tag{2.3'}$$

In the next section we argue that for a sufficiently small Δt the minimum value of $z_{t+\Delta t}$, denoted by z_{min}, is equal to $\nu_{min} + \gamma\eta_n$. Similarly, the fact that $E[z_{t+\Delta t}|S_t] \equiv \hat{z}_m \geq R - 1$ implies that there exist two states indexed by h and $h + 1 \leq m$ such that $\hat{z}_{h+1} \geq R - 1 \geq \hat{z}_h$; for a sufficiently small Δt these states belong to the component $\nu_{t+\Delta t}$ of the return. Instead of Lemma 1 we now have the following results, which yield bounds on $P_t(I_t)$ on the basis of

(2.3)'. The proof is virtually identical to that of Ritchken (1985) and will be omitted.[17]

Lemma 2.2. *If the option price $P_t(I_t)$ is convex for any t then it lies within the following bounds:*

$$\frac{1}{R}E^{L_t}[P_{t+\Delta t}(I_t(1+wz_{t+\Delta t})|S_t] \leq P_t(I_t) \leq \frac{1}{R}E^{U_t}[P_{t+\Delta t}(I_t(1+wz_{t+\Delta t})|S_t], \tag{2.4'}$$

where E^{U_t} and E^{L_t} denote respectively expectations taken with respect to the following distributions:

$$\begin{aligned} U_{1,t} &= \frac{R-1-z_{min}}{\hat{z}-z_{min}}p_1 + \frac{\hat{z}+1-R}{\hat{z}-z_{min}} \\ U_{i,t} &= \frac{R-1-z_{min}}{\hat{z}-z_{min}}p_i, i = 2,\dots,m \end{aligned} \tag{2.5a'}$$

$$\begin{aligned} L_{i,t} &= \frac{\hat{z}_{h+1}+1-R}{\hat{z}_{h+1}-\hat{z}_h}\frac{p_i}{\sum_{k=1}^{k=h}p_k} + \frac{R-1-\hat{z}_h}{\hat{z}_{h+1}-\hat{z}_h}\frac{p_i}{\sum_{k=1}^{k=h+1}p_k}, \ i = 1,\dots,h \\ L_{h,t} &= \frac{R-1-\hat{z}_h}{\hat{z}_{h+1}-\hat{z}_h}\frac{p_{h+1}}{\sum_{k=1}^{k=h+1}p_k} \\ L_{i,t} &= 0, \ i \geq h+1 \end{aligned} \tag{2.5b'}$$

Likewise, for continuous distributions for both $\gamma\eta$ and $\nu_{t+\Delta t}$, with $P(z)$ denoting their convolution,

$$\begin{aligned} U(z) &= \begin{cases} P(z) & \textit{with probability} \quad \frac{R-1-z_{min}}{\hat{z}-z_{min}} \\ 1_{z_{min}} & \textit{with probability} \quad \frac{\hat{z}+1-R}{\hat{z}-z_{min}} \equiv \Theta \end{cases} \\ L(z) &= P(z|z \leq z^*), \ E[z|z \leq z^*] = R-1 \end{aligned} \tag{2.6'}$$

Proof. See Appendix B. □

These expressions are very similar to (3.2.6) and (3.2.5). Note also that (2.6)' becomes identical to (3.2.6) when there is no diffusion component in $z_{t+\Delta t}$. Similarly, Proposition 1 holds with the substitution of I_t instead of A_t. In the next section we examine the value of the

[17]For economy of notation we also omit the time subscript from the distribution.

options in Case A and Case B in continuous time, in which there are significant differences in the two cases.

2.3 The Convergence to the Continuous Time Limits

For the convergence to continuous time we adopt the definition (2.3.1) of the return $z_{t+\Delta t}$ for both cases A and B. It corresponds to a non-CAT return component $\nu_{t+\Delta t}$ that becomes a pure diffusion at the limit, while the accumulated claims become a convolution of a diffusion and a jump components.

$$z_{t+\Delta t} = \nu_{t+\Delta t} + +\gamma\eta\Delta N = (\mu_t - \lambda\mu_J)\Delta t + \sigma_{\nu t}\epsilon\sqrt{\Delta t} + J\Delta N \tag{2.3.1}$$

In this expression ϵ has a bounded distributions of mean zero and variance one, $\epsilon \sim D(0,1)$ and $\epsilon_{min} \leq \epsilon \leq \epsilon_{max}$, but otherwise unrestricted. The instantaneous mean of the insurer's assets is $E[z_{t+\Delta t}] = \mu_t\Delta t$, with $\mu_t > r$ by assumption. Similarly, in (2.3.1) we set $\gamma\eta \equiv J$ and we use the conventional notation of $\mu_J = E[J]$; note that since $\gamma < 0$, $\mu_J < 0$ as well. In Case A we have $\sigma_{\nu t} = 0$. In the appendix we prove the following result.

Lemma 2.3. *For $\Delta t \to 0$ the discrete process for the return $z_{t+\Delta t}$ described by (2.3.1) tends to the following jump-diffusion process.*

$$\frac{dS_t}{S_t} = (\mu_t - \lambda\mu_J)dt + \sigma_{\nu t}dW + \gamma\eta dN, \tag{2.3.2}$$

where dW is a Wiener processes with $E[dW] = 0$, and N is a Poisson counting process with intensity λ.[18]

Reinterpreting Lemma 3 in the context of the discrete time formulation of Case B, we set the diffusion component of the insurer's assets equal to $\nu_{t+\Delta t} = (\mu_t - \lambda\mu_J)\Delta t + \sigma_{\nu t}\epsilon\sqrt{\Delta t}$, with $\hat{\nu}_t = E[\nu_{t+\Delta t}] = (\mu_t - \lambda\mu_J)\Delta t$. Setting for simplicity $w = 1$, the insurer's assets subject to

[18] It is assumed that the amplitude is such that $1 + \gamma\eta > 0$ always.

the reinsurance contract become $I_t(1 + wz_{t+\Delta t}) = I_t(1 + (\mu_t - \lambda\mu_J)\Delta t + \sigma_{\nu t}\epsilon\sqrt{\Delta t} + \gamma\eta\Delta N)$. Since by assumption the returns $\nu_{t+\Delta t}$ are bounded, it follows that for a sufficiently small Δt the return component $\nu_{t+\Delta t}$ is dominated by the CAT event term, and its minimum occurs when $\eta = \eta_n$, as argued earlier. Set also $\phi = \alpha_t\Delta t$, $\alpha_t \in [\lambda\hat{\eta}, \frac{\mu_t - \lambda\mu_J - r}{-\gamma}]$, $(\phi - \lambda\hat{\eta}\Delta t) =\equiv \beta_t\Delta t$. We may now prove the continuous time results of this paper in the following propositions, with both proofs relegated to the appendix.

Proposition 2.2. *When the trader's portfolio returns are given by the jump-diffusion process (3.3.2), the value $C_t(A_t)$ of the CAT call option described in Case A is bounded by the following values:*

$$\begin{aligned} C_t(A_t) &\geq \underline{C}_t(A_t) = e^{-rT}\sum_{N=0}^{N=\infty} e^{-\lambda_t^L T}\cdot\frac{(\lambda_t^L T)^N}{N!}C_N^L(A_t) \\ C_t(A_t) &\leq \overline{C}_t(A_t) = e^{-rT}\sum_{N=0}^{N=\infty} e^{-(\lambda+\lambda_t^U)T}\cdot\frac{((\lambda+\lambda_t^U)T)^N}{N!}C_N^U(A_t) \end{aligned} \tag{2.3.3}$$

where $C_n^L(A_t) = E_t^L[(A_t + \kappa n\eta - K)^+|n]$ and $C_n^U(A_t) = E_t^U[(A_t + \kappa n\eta - K)^+|n]$ denote the conditional expectations of the option payoffs given n catastrophe events till option expiration. The conditional probabilty distribution of CAT amplitudes, along with the jump intensities under U- and L-distributions, can be derived from (3.2.5) and (3.2.6) as bellow.

$$U(\eta|N \neq 0) = \begin{cases} P(\eta) & \text{with probability } \frac{\lambda}{\lambda+\lambda_t^U} \\ 1_{\eta_n} & \text{with probability } \frac{\lambda_t^U}{\lambda+\lambda_t^U} \end{cases}, \ \lambda_t^U = \frac{\beta_t}{\eta_n} \tag{2.3.4a}$$

$$L(\eta|N \neq 0) = P(\eta),\ \lambda_t^L = \frac{\alpha_t}{\hat{\eta}} \tag{2.3.4b}$$

Proof. See Appendix C. □

The intuition of this result becomes clear if we observe that both the U-and L-distributions transform the expectations of the amplitudes of the CAT event, changing it from $\lambda\hat{\eta}\Delta t$ to ϕ and rendering the index return $z_{t+\Delta t}$ risk neutral. These transformations raise the payoff expectation for valuation purposes, as befits an insurance instrument. Hence, the Merton (1976) case of a fully diversifiable CAT event risk corresponds to the case $\phi = \lambda\hat{\eta}\Delta t$ and lies below both values given by (3.3.3)-(3.3.4).

Turning now to the valuation of the reinsurance contract of Case B, we find that the derived solution is no longer simple, and that there are no closed form expressions for the bounds on the value of the contract. By Lemma 2 the limiting P-distribution of the contract payoff is a mixed jump-diffusion process. The following proposition, proven in the appendix, gives similar limits on the U- and L-distributions of Lemma 1 and Proposition 1 under this mixed process.

Proposition 2.3. *When the trader's portfolio returns are given by the jump-diffusion process (3.3.2) the value $P_t(I_t)$ of the reinsurance contract described in Case B is bounded by the discounted expectations of the contract payoff under the following distributions that form the continuous time limits of the U- and L-distributions given in (3.2.5) and (3.2.6). For the upper bound the limiting distribution is:*

$$\frac{dI_t}{I_t} = [r - (\lambda + \lambda_t^U)\mu_{Jt}^U]dt + \sigma_{\nu t}dW + J_t^U dN, \tag{2.3.5}$$

where,

$$\lambda_t^U = \frac{\mu_t - r}{-J_{min}},\ J_{min} = \gamma\eta_n \tag{2.3.6}$$

and J_t^U is a mixture of jumps with intensity $\lambda + \lambda_t^U$ and distribution and mean

$$J_t^U = \begin{cases} J & \text{with probability} \quad \frac{\lambda}{\lambda+\lambda_t^U} \\ J_{min} & \text{with probability} \quad \frac{\lambda_t^U}{\lambda+\lambda_t^U} \end{cases}$$
$$\mu_{Jt}^U = \frac{\lambda}{\lambda+\lambda_t^U}\mu_J + \frac{\lambda_t^U}{\lambda+\lambda_t^U}J_{min} \tag{2.3.7}$$

For the lower bound the limiting distribution is:

$$\frac{dI_t}{I_t} = [r - \lambda\mu_{Jt}]dt + \sigma_{\nu t}dW + JdN \tag{2.3.8}$$

Proof. See Appendix D. □

Although both distributions (3.3.5)-(2.3.7) and (2.3.8) are risk neutral, it is not possible to find closed form solutions for the expectations of the contract payoff. As Merton (1976) first pointed out, closed form expressions for options under jump diffusion asset dynamics are available only in the special case where the jump amplitude is lognormally distributed, which is not a reasonable assumption for CAT events. Nonetheless, estimation by Monte Carlo simulation is clearly feasible, while efficient tree-based numerical methods are also available for such problems[19]

We note that the lower bound distribution (2.3.8) is identical to the case where the CAT event risk is diversifiable. Hence, the Merton (1976) case appears as a lower bound to the reinsurance contract value in our analysis. On the other hand, the width of the bounds in (2.3.7) clearly depends on the risk premium $\mu_t - r$, on the highest possible amplitude η_n of the CAT event, as well as on its impact on the agent portfolio represented by the parameter γ. This dependence will appear very clearly in the numerical results of the next section.

We close this section by pointing out one potential difficulty and one limitation of our analysis. Although, the CAT event under the P-distribution is independent of the diffusion components in both Case A and Case B, this property is not necessarily preserved in the two martingale distributions if the instantaneous mean of the diffusion is state-dependent. As we can see from (3.3.5)-(2.3.8), the arrival intensity and/or the amplitude distribution in these two processes are both dependent on the instantaneous means of the diffusion processes, which may not be constant. This may introduce state dependence into the U- and L-mixed processes, which may complicate the numerical estimation. Second, the closed form expressions for the two boundary distributions are valid only if the payoff and the option price are convex with respect to the underlying asset price. If convexity does not hold[20] the

[19] See, for instance, Amin (1993).

[20] For instance, if the reinsurance contract in Case B has a ceiling as well as a deductible.

general approach is still valid, but the modified LP approach described in Section 2 can only provide numerical solutions.

2.4 Numerical Results

In this section we apply the results derived in the previous sections to price a call option on the catastrophe event, as well as a reinsurance contract on the assets of an insurance company. In both cases the underlying process stems from the cumulative losses associated with hurricane landings. We calibrate this loss process from the available data from CME for hurricane landings in the state of Florida for the period 1998-2007 and the associated changes in the CHI. For Case A the option contract multiplier is set at $\kappa = \$1000$ per unit index change, as in the CME hurricane options. For Case B the multiplier should be proportional to the losses incurred given a landing, and their distribution among the regional insurers according to our limited diversification assumption. For the aggregate losses, assuming the same contract is offered to all eligible insurance companies. We estimate $\kappa = 0.41$ (in billion dollars) from the reported losses associated with the Florida landings during the period of 1975-2005.[21]

Case A

In order to calculate the bounds, we discretize the time and state space in the following manner. We divide the time $T - t$ until the maturity of the CAT event option into N subdivisions of equal length, each equal to $\Delta t = \frac{T-t}{N}$. We then construct a multinomial lattice with $n+1$ branches emanating from each node to represent the cumulative catastrophe loss associated with the hurricane landings. Figure 2.1 depicts the cumulative loss process for two periods.

[Figure 4.1 about here]

If $A_t$denotes the cumulated loss at some time t then at time $t + \Delta t$ we have $n + 1$ possible

[21] The coefficient was estimated from a linear regression of the hurricane losses on the intensity of the landed hurricanes in all states, based on the data provided by CME for the period of 1975-2005. This parameter does not enter in the calculation in Case B explicitly, but its implicit in the value of γ.

outcomes for the next period cumulative loss. The cumulative loss could stay the same, corresponding to the event of no hurricane landing ($\eta_0 = 0$), or it could go up to $A_t + \kappa\eta_i$, $i = 1, \ldots, n$. After N periods we will have the terminal cumulative losses at time T. As shown in the previous section, this discrete multinomial process is known to converge to a compound Poisson process as we increase the number of subdivisions.

We extract the conditional distribution of intensities, (p_i, η_i), from the histogram of CHI values for Florida landings for $n = 2$.[22] So, conditional on landing, the hurricane intensity could be 4.425 with probability 0.444 or 10.475 with probability 0.556, with $\hat{\eta} = 7.114$. We also estimate the intensity of the hurricane landing as the average number of landings per year and equal to 0.9 for the state of Florida, and we assume it to be time independent; this yields $\lambda\hat{\eta} = 6.403$. Hence, the combined distribution of landing and intensities, (q_i, η_i), results in a trinomial lattice for the underlying loss process. We assume that in every period there can be only one hurricane landing. So, when the number of subdivisions is N, the minimum and maximum number of hurricane landings are zero and N, respectively.

The lattice representing the underlying loss process described above is non-recombining for the most parts, and after a few periods it becomes too large to be handled. However, since we assume the combined distribution of hurricane landing and intensities to be *i.i.d.*, and since our derivative of interest is a European option,[23] we only need the terminal distribution of accumulated loss under the U and L distributions. In building up the lattice we use an approximation that reduces the computational intensity significantly. In every time step we first order the state values and their associated probabilities. We aggregate the nodes with the same state value into one node, and associate a probability to that node equal to the sum of the probabilities of the aggregated nodes. Then we truncate the distribution at a point beyond which the sum of the probabilities is less than 0.01. We allocate the entire probability mass of the truncated part to the last node, and we set the state value of the last node such that the contribution of the truncated part to the expected value of the whole distribution does not change. We then use these state values and the associated probabilities

[22] We choose $n = 2$ because the lattice is non-recombining and after t periods we have $(n+1)^t$ states; this makes the computations very intensive and even infeasible for large values of N.

[23] Although the CME hurricane options are American, the lack of dividends on the underlying implies that early exercise is not profitable.

to calculate the accumulated loss distribution at the next period.

We calculate, based on (3.2.5), the two risk neutral transformations of the physical probability distribution, corresponding to the upper and the lower bounds. These probability measures are then used, as in the previous section, to calculate the distribution of the terminal underlying accumulated loss. The upper bound and the lower bound on the option price are then calculated as the discounted expected payoff under the upper and lower bound distributions, respectively.

In what follows we calculate the option price under the U and L distributions given by (3.2.5) for varying time partitions, as well as their continuous time limits given by (3.3.3). We also estimate the CAT event option price under the physical distribution, corresponding to Mertons assumption of the diversifiability of the CAT risk.

Figure 2.2 and Table 2.1 show the two bounds and the Merton price in the discrete time and continuous time for different numbers of subdivisions. The strike prices for the CME options are set at the various successive index level points; we use a strike price of 5 as our base case, or $\$5,000$ given the multiplier of $\$1000$. The option has six months left to maturity. The initial accumulated loss is assumed to be zero. The annual risk free rate is 1%, and the annual expected return on the market is equal to 5%. We set $\gamma = -0.004$, in which case the parameter α lies between 6.403 and 10; we choose $\alpha = 7$.

[Figure 4.2 about here]

[Table 4.1 about here]

As we see in the graph the two bounds and the Merton price increase as we increase the number of subdivisions, and converge to their continuous time values for large values of N. The bounds are very tight, with the gap about 2% at the limit. On the other hand, the Merton price is significantly lower, a bit less than 13% from the midpoint of the bounds; this shows the size of the error in assuming that the CAT event risk is diversifiable.

Figure 2.3 and Table 2.2 show the bounds for the same contract mentioned above against different values of the strike price. The two bounds are decreasing convex functions of the strike price, with the width of the bounds remaining proportionally the same at different strike prices, and at roughly the same proportional distance from the Merton price.

[Figure 4.3 about here]

[Table 4.2 about here]

Figure 2.4 and Table 2.3 show the values of the bounds and the Merton price for the same parameter values and contract specifications as before, but for different values of alpha. As we see, when $\alpha = \lambda\hat{\eta} = 6.403$ the CAT risk is diversifiable and both bounds as well as the Merton price coincide. However, as α increases the two bounds increase and become wider, with the width reaching about 6.5% of the midpoint for $\alpha = 10$, reflecting market incompleteness and the higher CAT event risk premium. On the other hand the Merton price does not change and becomes much lower than the two bounds, essentially irrelevant for the pricing of the CAT event option.

[Figure 4.4 about here]

[Table 4.3 about here]

Case B

In this section we price a put option reinsurance contract on the total accumulated assets of the insurer, where the assets return is assumed to follow a jump-diffusion process. We use the same discretization as in case A, with the exception that here we have an extra source of randomness. We approximate the jump part by the same trinomial as in case A, and we approximate the diffusion part by a Bernoulli random variable that takes the values 1 and -1 with equal probabilities. Since it is assumed that the jump part is independent of the diffusion part, and since the jump process and the Bernoulli process are assumed to be *i.i.d.*, the combined discrete process representing the jump-diffusion process consists of six states in every period that are independently and identically distributed. We first calculate, based on (2.6)', the two risk neutral transformations of the physical probability distribution, corresponding to the upper and the lower bounds. We then use these transformed probability measures to calculate the distribution of the terminal asset value by applying the same lattice construction technique described in case A. Figure 2.5 depicts this lattice for a single period. The upper bound and the lower bound on the option price are then calculated as the discounted expected payoff under the upper and lower bound distributions, respectively.

[Figure 4.5 about here]

The intensity of the jump process and the jump sizes are based on the data for the state of Florida, as described in case A. We set the drift and the volatility of the diffusion equal to 0.05 and 0.1 respectively. We find bounds on the price of a reinsurance contract that is represented by a put option maturing in six months. The initial value of the assets is $\$100M$ and the option is at the money. Unlike in case A, we cannot derive closed form solutions for the bounds in the continuous time limit. Further, the convergence from the discrete to the continuous time is rather slow. Accordingly, we estimate the continuous time limits of the bounds from Proposition 3 in the previous section by Monte Carlo simulation.

Figure 2.6 and Table 2.4 show the bounds for different values of the market return. As demonstrated, the lower bound which coincides with the Merton price does not vary. However, as we increase the market return, the upper bound increases and the bounds get wider. At the upper range, corresponding to a 6% risk premium, the width of the bounds rises to about 14% of its midpoint from a bit more than 9% for the base case of a 4% risk premium.

[Figure 4.6 about here]

[Table 4.4 about here]

We also demonstrate in Figure 2.7 and Table 2.2 the two bounds for different values of the strike price for the base case of a 5% mean (4% risk premium). As expected the bounds are increasing and convex functions of the strike price. Further, the bounds become proportionally much wider as the moneyness of the option decreases, consistent with what happens in the actual option markets.

[Figure 4.7 about here]

[Table 4.5 about here]

Last, Figure 2.8 and Table 2.6 show the bounds for different values of the parameter gamma, the sensitivity of the portfolio return to the CAT event. As this sensitivity decreases in absolute value the bounds get lower and tighter, and when the traders portfolio return does not depend on the CAT event, consistent with the assumption that the CAT event is diversifiable, the two bounds coincide and become equal to the Merton price. The increase in the size of the bounds when γ increases stems from the fact that the total volatility of the return, equal to the volatility of the diffusion plus the volatility of the jump times γ, also increases.

[Figure 4.8 about here]

[Table 4.6 about here]

2.5 Conclusion

In this paper we have presented an approach to the pricing of CAT event derivatives that is drastically at variance with the established methodology in earlier studies. That methodology followed the Merton (1976) assumption about a fully diversifiable CAT event risk, in which case a unique price results. Our approach relies on recent literature suggesting that economic agents trading in CAT instruments (for instance, insurance companies) specialize locally and in special types of CAT risks. Hence, the CAT event risk may not be fully diversifiable. In such a case our approach recognizes the market incompleteness introduced by the CAT event and relies on stochastic dominance arguments to develop bounds on the CAT event derivatives.

We apply our method to the pricing of a CAT event call option and a reinsurance contract, both modeled on the hurricane risk in the state of Florida; the option is similar to the ones offered by the CME, while the contract is modeled as a put option on an insurers assets. Our theoretical analysis predicts that the call option price bounds would lie above the Merton price, with the distance depending on the price of a hurricane futures contract such as the ones offered by the CME or trading over the counter. For the reinsurance contract, on the other hand, we show that the Merton price is equal to the lower bound that we develop, with the width of the bounds depending on the risk premium of the representative agent portfolio over the riskless rate.

We use realistic parameter values and show that the CAT call option produces tight bounds for all admissible values of the parameters and all strike prices of the option. We also show that the Merton price lies far below our bounds for almost all realistic values of the hurricane futures contract parameter, thus illustrating the pitfalls of neglecting the CAT event risk. On the other hand, the Merton price is part of the admissible set of prices for the reinsurance contract, even though it does form the lowest value of the set. We show the dependence of the width of the bounds on various parameters of the jump diffusion process, in particular the risk premium of the portfolio return and the dependence of this return on the CAT event.

Appendix

2.A Proof of Lemma 1

The distribution $q_0, q_1, \ldots, q_n$ of $\eta_{t+\Delta t}$ is $q_0 = 1 - \lambda\Delta t$, $q_i = \lambda\Delta t p_i$, $i = 1, \ldots, n$, where (p_i, η_i), $i = 1, \ldots, n$ is the distribution of $\eta_{t+\Delta t}$ given that the CAT event occurred. Since $\hat{Y}_t(\eta_{t+\Delta t})$ is non-decreasing, let $\eta_0 = 0$, $\hat{Y}_t(\eta_i) \equiv \hat{Y}_i$, $i = 0, \ldots, n$, with $\hat{Y}_0 \leq \hat{Y}_1 \leq \ldots \leq \hat{Y}_n$. The proof follows closely Ritchken (1985) and will only be sketched here. We set $\hat{Y}_0 = x_0$, $\hat{Y}_1 = x_0 + x_1$, $\hat{Y}_n = \sum_0^n x_i$ and replace into (3.2.3). Define also

$$
\begin{aligned}
& x_i \sum_{k=i}^{k=n} q_{kt} \equiv \overline{Y}_i \\
& \bar{c}_i = \frac{\sum_{k=i}^{n} C_{t+\Delta t}(A_t + \eta_{k,t+\Delta t}) q_{kt}}{\sum_{k=i}^{n} q_{kt}} = E[C_{t+\Delta t}(A_t + \eta_{t+\Delta t}) | \eta_{t+\Delta t} \geq \eta_{i,t+\Delta t}, A_t]
\end{aligned}
\tag{2.A.1}
$$

Replacing these values, as well as $\hat{\eta}_i = \dfrac{\sum_{k=i}^{k=n} q_k \eta_k}{\sum_{k=i}^{k=n} q_k}$, $i = 0, 1, \ldots, n$,[24] into (3.2.3) we see that it can be rewritten in the following form

[24] Note that $\hat{\eta}_0 = \hat{\eta}\lambda\Delta t$.

$$\sum_{i=0}^{n}\overline{Y}_i = 1, \quad \sum_{i=0}^{n}\overline{Y}_i\,\hat{\eta}_{i,t+\Delta t} = \phi, \quad \overline{Y}_i \geq 0,\, i = 0,,1,\ldots,n$$
$$C_t(A_t) = \frac{1}{R}\sum_{i=0}^{n}\overline{Y}_i\,\overline{c}_i \tag{2.A.2}$$

The upper and lower bounds on the call option price can then be found by solving the following linear programs (LPs).

$$\frac{1}{R}[max_{\overline{Y}_i}\sum_{i=0}^{n}\overline{Y}_i\,\overline{c}_i\,(min_{\overline{Y}_i}\sum_{i=0}^{n}\overline{Y}_i\,\overline{c}_i)] \tag{2.A.3}$$

The solution of the LPs is given in closed form by the expressions (3.2.5), which becomes (3.2.6) when the distribution of η is continuous, when the function $C(A_t)$ is convex for any t. The proof is an adaptation of Ritchken (1985), and can be seen geometrically in Figure 2.A.1. Convexity is preserved in the graph $(\hat{\eta}_i, \hat{c}_i)$, $i = 0, 1, \ldots, n$. It is also clear from the graph that both bounds are increasing functions of the parameter ϕ.

2.B Proof of Lemma 2

We prove the convergence of the discretization (2.3.1) in the *i.i.d.* case[25] $\mu_t - \lambda\mu_J = \mu$, $\sigma_t = \sigma$. Convergence in the non-*i.i.d.* case follows from the convergence criteria for stochastic integrals, presented in Duffie and Protter (1992). It is shown in an appendix, available from the authors on request.

The characteristic function of the terminal stock price at time T for a \$1 initial price under the jump-diffusion process (3.3.2) is

$$\begin{aligned}\varphi_{JD}(\omega) &= exp(i\omega\mu T - \frac{\omega^2\sigma^2 T}{2})exp(-\lambda T)\sum_{N=0}^{\infty}\frac{(\lambda T)^N}{N!}[\varphi_J(\gamma\omega)^N] \\ &= exp(i\omega\mu T - \frac{\omega^2\sigma^2 T}{2})exp[\lambda T(\varphi_J(\gamma\omega) - 1)],\end{aligned} \tag{2.B.1}$$

[25] The proof is similar to that of Theorem 2.1 in Jacod and Protter (2002). See also Oancea and Perrakis (2014)

where $\varphi_J(\omega)$ is the characteristic function of the CAT event amplitude distribution. The first exponential corresponds to the diffusion component and the second to the jump component. The characteristic function of the discretization (2.3.1) is

$$\varphi(\omega) = (\lambda \Delta t \varphi_J(\gamma\omega) + 1 - \lambda \Delta t)[exp(i\omega\mu\Delta t)\varphi_\epsilon(\omega\sigma\sqrt{\Delta t})], \tag{2.B.2}$$

where $\varphi_\epsilon(\omega)$ is the characteristic function of ϵ.[26] Since the distribution of ϵ has mean 0 and variance 1, we have

$$\begin{aligned} E[\epsilon] &= 0 = i\varphi'_\epsilon(0), \\ E[\epsilon^2] &= 1 = -\varphi''_\epsilon(0) \end{aligned}$$

By the Taylor expansion of $\varphi_\epsilon(\omega)$, we get

$$\varphi(\omega) = (\lambda \Delta t \varphi_J(\gamma\omega) + 1 - \lambda \Delta t)\Big[exp(i\omega\mu\Delta t)[1 - \frac{\omega^2\sigma^2\Delta t}{2} + \omega^2\sigma^2\Delta t h(\omega\sigma\sqrt{\Delta t})]\Big],$$

where $h(\omega) \to 0$ as $\omega \to 0$. The multiperiod convolution has the characteristic function $\varphi(\omega)^{\frac{T}{\Delta t}}$. Taking the limit, we have

$$\begin{aligned} \lim_{\Delta t \to \infty} [\varphi(\omega)]^{\frac{T}{\Delta t}} &= \lim_{\Delta t \to \infty} exp\Bigg[\frac{T}{\Delta t}\Big[ln(\lambda \Delta t \varphi_J(\gamma\omega) + 1 - \lambda\Delta t) \\ &\quad + ln\Big[exp(i\omega\mu\Delta t)[1 - \frac{\omega^2\sigma^2\Delta t}{2} + \omega^2\sigma^2\Delta t h(\omega\sigma\sqrt{\Delta t})]\Big]\Big]\Bigg] \\ &= exp\Big[\lambda T(\varphi_J(\gamma\omega) - 1) + i\omega\mu T - \frac{\omega^2\sigma^2 T}{2}\Big] \end{aligned} \tag{2.B.3}$$

after applying $l'H\hat{o}pital's$ rule. (2.B.3) is, however, the same as (2.B.1), the characteristic function of (3.3.2), and Levy's continuity theorem[27] proves the weak convergence of (2.3.1) to (3.3.2). □

[26] If, instead of (3.2.2) we have a mixture of the diffusion and jump components then the characteristic function becomes $\varphi(\omega) = \eta\Delta t\varphi_J(\omega) + (1 - \eta\Delta t)[exp(i\omega\mu_d\Delta t)\varphi_\epsilon(\omega\sigma_d\sqrt{\Delta t})]$. The multiperiod convolution, however, still converges to (2.A.3).

[27] See for instance Jacod and Protter (2002), Theorem 19.1.

2.C Proof of Proposition 2

For the proof we rely on the weak convergence of the underlying price process given by (2.3.1) to the jump diffusion given by (3.3.2). For any number m of time periods to expiration we define a sequence of stock prices $\{S_t|\Delta t, m\}$ and an associated probability measure P^m. The weak convergence property for such processes[28] stipulates that for any continuous bounded function f we must have $E^{P_m}[f(S_T^m)] \to E^P[f(S_T)]$, where the measure P corresponds to the continuous limit of the process. P_m is then said to converge weakly to P and S_T^m is said to converge in distribution to S_T. By Lemma 2 and the assumed independence of the diffusion and jump components, the distribution of S_T under the measure given n realizations of the jump component is equal to $S_{TD}(1+\gamma\eta)^n$, where S_{TD} denotes the diffusion component. By Lemma 1 and Proposition 1 the upper and lower bounds of the CAT event call option are the recursive expectations of the payoff respectively under the $U(\eta)$ and $L(\eta)$ distributions, so that the bounds are equal to the expectations with the limiting values of the distributions as $\Delta t \to 0$.

Turning first to the distribution $L(\eta)$, we note that the risk neutral transformation affects only the jump and not the diffusion component. As Δt decreases, the parameter $\phi = \alpha_t \Delta t$ lies between $\lambda\hat{\eta}\Delta t$ and $\hat{\eta}$, implying that the risk neutral transformation cannot affect the amplitude but only the intensity of the jump, which becomes equal to λ_t^L. Hence, for $\Delta t \to 0$ the expected payoff of the option tends to the discounted Poisson expectation given by $\underline{C}(A_t)$. The situation is more complex for the $U(\eta)$ distribution, where (3.2.6) implies that for $\Delta t \to 0$, $U(\eta)$ becomes a mixture of η_n with probability $\lambda_t^U \Delta t$ and $\eta_{t+\Delta t}$ with probability $1-\lambda_t^U\Delta t$.On the other hand, $\eta_{t+\Delta t}$ is equal to 0 with probability $1-\lambda\Delta t$ and to η with probability $\lambda\Delta t$. Simplifying and neglecting the terms in $o(\Delta t)$, we find that $U(\eta)$ tends to the distribution given by (3.3.4), implying that the upper bound of the option tends to $\overline{C}(A_t)$ as in (3.3.3). □

[28]For more on weak convergence for Markov processes see Ethier and Kurtz (1986), or Stroock and Varadhan (1979).

2.D Proof of Proposition 3

The proof is very similar to that of Proposition 2: the bounds are equal to the discounted expectations of the reinsurance contract payoff $(K - I_T)^+$ with the limiting distributions; we need, therefore, to find the limits of the U- and L-distributions as $\Delta t \to 0$, which in this case are going to be mixed jump-diffusion processes.[29] For the upper bound, we note that the key probability $\Theta = \frac{\hat{z} + 1 - R}{\hat{z} - z_{min}}$ tends to $\lambda_t^U \Delta t$, implying that the U-distribution becomes a mixture of $z_{t+\Delta t}$ with probability $1 - \lambda_t^U \Delta t$ and z_{min} with probability $\lambda_t^U \Delta t$. z_{min} tends to $-J_{max} = -\gamma\eta_n$, while the instantaneous mean of the U-distribution is equal to $r\Delta t$. Since $z_{t+\Delta t}$ is itself a convolution of diffusion and jump components, replacing the terms, simplifying and neglecting the terms in $o(\Delta t)$, we find that the limit distribution is the jump-diffusion process given by (3.3.5)-(2.3.7). For the lower bound we observe again that as $\Delta t \to 0$ the highest values of $z_{t+\Delta t}$ come from the diffusion component whenever the jump event does not occur. Since the instantaneous mean of the process is by assumption greater than the riskless rate, the only way to decrease it is by truncating it as in (2.6)' in order to decrease the instantaneous mean from μ_t to r. It follows that the limiting jump-diffusion process of the L-distribution is the one given by (2.3.8). □

[29] For more details see also the similar proofs of Propositions 3 and 4 in Oancea and Perrakis (2014).

Table 2.1: **Convergence of the Bounds**

n	Lower Bound	Upper Bound	Merton Price
10	1412	1468	1267
20	1527	1562	1370
30	1566	1600	1405
40	1585	1620	1421
50	1597	1632	1432
100	1621	1655	1453
150	1629	1663	1460
200	1633	1667	1463
∞	1644	1679	1473

The table shows the continuous time values of the two bounds and the Merton price for the option in case A for different number of periods. The option has 6 months left to maturity and the strike price is \$5000. The annual risk free rate is 1%, and the annual expected return on the market is equal to 5%. We set $\gamma = -0.004$, in which case the parameter α lies between 6.403 and 10; we choose $\alpha = 7$.

Table 2.2: **Effect of Strike Price on the Bounds: Case A**

Strike Price	Lower Bound	Upper Bound	Merton Price
3000	2322	2347	2103
4000	1936	1968	1743
5000	1644	1679	1473
7000	1204	1230	1070
9000	769	787	671
12000	434	440	366

The table shows the continuous time values of the two bounds and the Merton price for the option in case A for different values of the strike price. The option has 6 months left to maturity. The annual risk free rate is 1%, and the annual expected return on the market is equal to 5%. We set $\gamma = -0.004$ and we choose $\alpha = 7$.

Table 2.3: **Effect of CAT Event Risk Premium on the Bounds: Case A**

α	Lower Bound	Upper Bound	Merton Price
6.4	1474	1474	1473
7	1644	1679	1473
8	1948	2036	1473
9	2240	2390	1473
10	2574	2748	1473

The table shows the continuous time values of the two bounds and the Merton price for the option in case A for different values of alpha. The option has 6 months left to maturity and the strike price is \$5000. The annual risk free rate is 1%, and the annual expected return on the market is equal to 5%. We set $\gamma = -0.004$, in which case the parameter α lies between 6.403 and 10.

Table 2.4: **Effect of Market Risk Premium on the Bounds: Case B**

Market Return	Lower Bound	Upper Bound
0.01	2.71	2.71
0.02	2.71	2.76
0.03	2.71	2.8
0.04	2.71	2.89
0.05	2.71	2.97
0.06	2.71	3.06
0.07	2.71	3.12

The table shows the continuous time values of the bounds (in million dollars) for different values of the market return. The initial asset value is equal to \$100M. The option is at the many and has six months left to maturity. The risk free rate is 0.01, γ is set to -0.004, and the volatility of market return is 0.1.

Table 2.5: **Effect of Strike Price on the Bounds: Case B**

Strike Price	Lower Bound	Upper Bound
90	0.21	0.32
95	0.94	1.15
100	2.71	2.97
105	5.73	5.99
110	9.86	9.98

The table shows the continuous time values of the bounds (in million dollars) for different values of the strike price. The initial asset value is equal to \$100M. The option has six months left to maturity. The risk free rate is 0.01, γ is set to -0.004. Market return is 0.05 and the volatility of market return is assumed 0.1.

Table 2.6: **Effect of γ on the Bounds: Case B**

γ	Lower Bound	Upper Bound
-0.001	2.5	2.65
-0.002	2.54	2.76
-0.003	2.6	2.87
-0.004	2.71	2.97
-0.005	2.76	3.1

The table shows the continuous time values of the bounds (in million dollars) for different values of γ. The initial asset value is equal to $\$100M$. The option is at the money and has six months left to maturity. The risk free rate is 0.01. Market return is 0.05 and the volatility of market return is assumed 0.1.

Figure 2.1: **Loss Process: Case A**

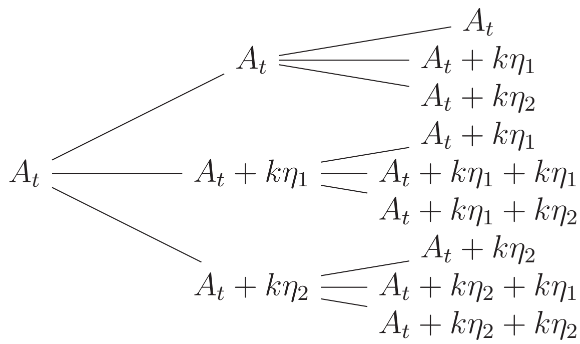

Underlying accumulated hurricane loss process for two periods

Figure 2.2: **Convergence of the Bounds**

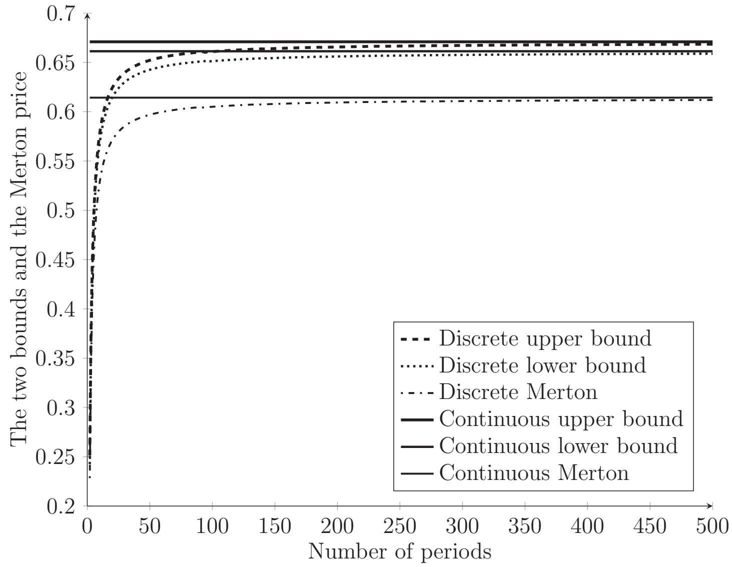

This figure shows the convergence of the discrete time bounds and the Mertons price to their continuous time counterparts for the call option in Case A. The option has 6 months left to maturity and the strike price is \$5000. The annual risk free rate is 1%, and the annual expected return on the market is equal to 5%. We set $\gamma = -0.04$, in which case the parameter α lies between 6.403 and 10; we choose $\alpha = 7$.

Figure 2.3: **Effect of Strike Price on the Bounds: Case A**

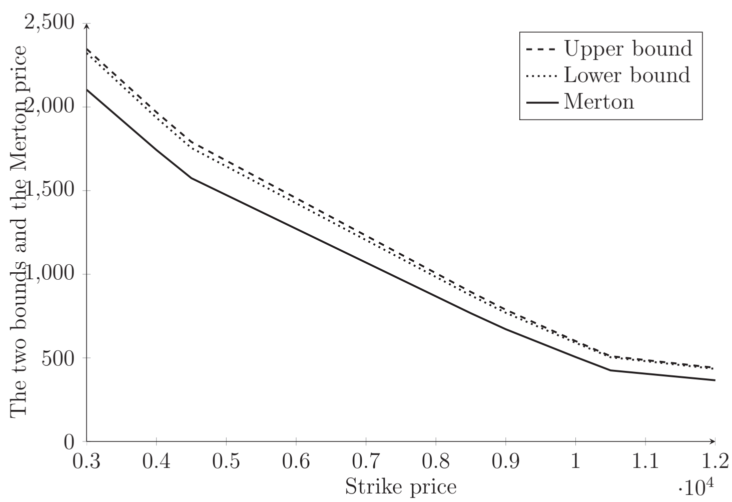

This figure shows the value of the two bounds and the Mertons price for the option considered in case A as a function of the strike price. We evaluate the option price for a strike price range from \$3000 to \$12000. The option has 6 months left to maturity. The annual risk free rate is 1%, and the annual expected return on the market is equal to 5%. We set $\gamma = -0.04$, and we choose $\alpha = 7$.

Figure 2.4: **Effect of CAT Event Risk Premium on the Bounds: Case A**

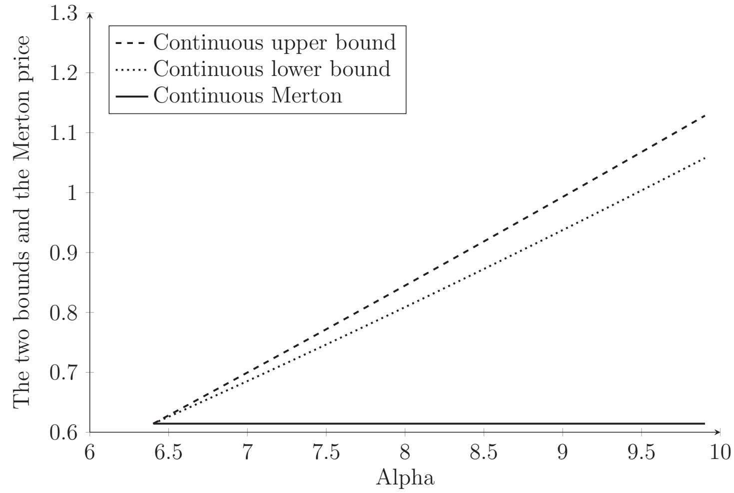

The figure shows the value of the two bounds and the Merton price for the option considered in Case A for different values of alpha. The option has 6 months left to maturity and the strike price is \$5000. The annual risk free rate is 1%, and the annual expected return on the market is equal to 5%. We set $\gamma = -0.04$, in which case the parameter α lies between 6.403 and 10.

Figure 2.5: **Loss Process: Case B**

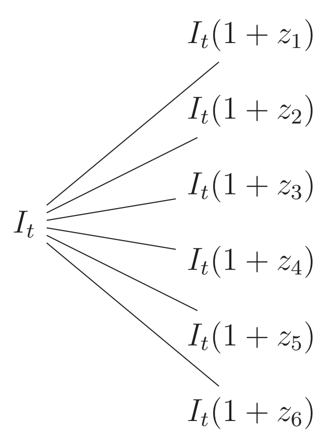

Underlying process for insurer's assets value in one period.

Figure 2.6: **Effect of Market Risk Premium on the Bounds: Case B**

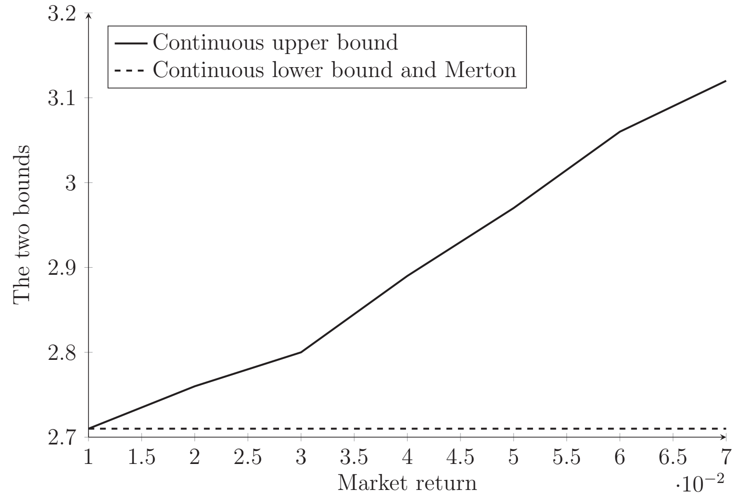

The figure depicts the continuous time values of the bounds (in million dollars) for different values of the market return. The initial asset value is equal to $\$100M$. The option is at the many and has six months left to maturity. The risk free rate is 1%, γ is set to -0.004, and the volatility of market return is 0.1.

Figure 2.7: **Effect of Strike Price on the Bounds: Case B**

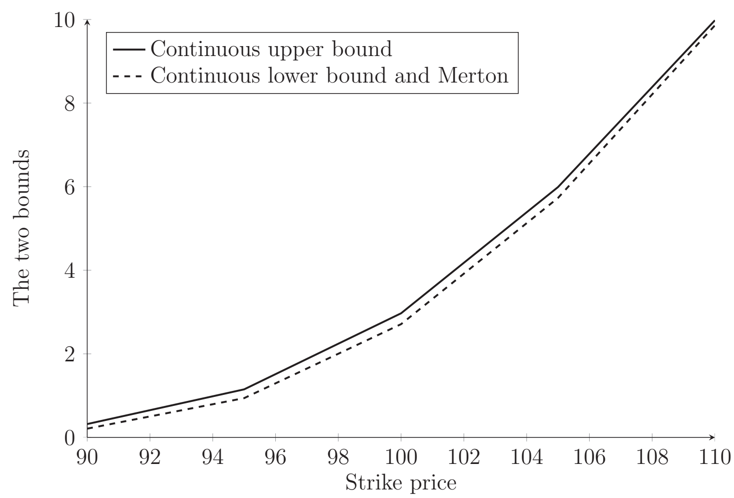

The figure shows the continuous time values of the bounds (in million dollars) for different values of the strike price. The initial asset value is equal to $\$100M$. The option has six months left to maturity. The risk free rate is 1%, γ is set to -0.004. Market return is 5% and the volatility of market return is assumed 0.1.

Figure 2.8: **Effect of γ on the Bounds: Case B**

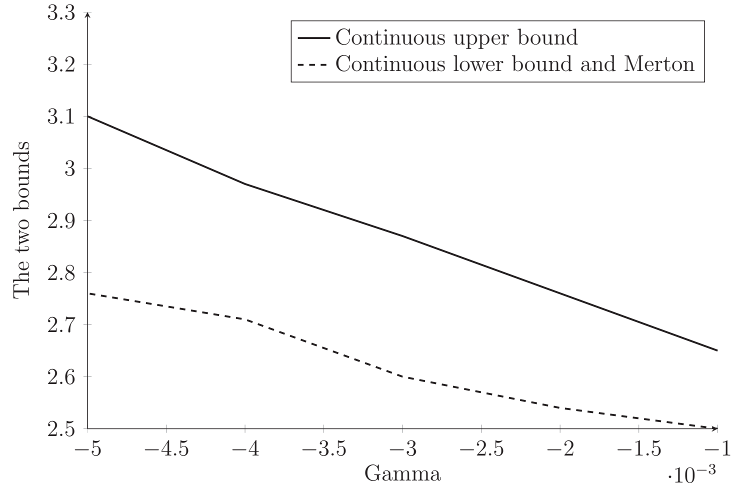

The figure shows the continuous time values of the bounds (in million dollars) for different values of γ. The initial asset value is equal to $\$100M$. The option is at the many and has six months left to maturity. The risk free rate is 1%. Market return is 5% and the volatility of market return is assumed 0.1.

Figure 2.A.1: **Geometrical representation of the bounds**

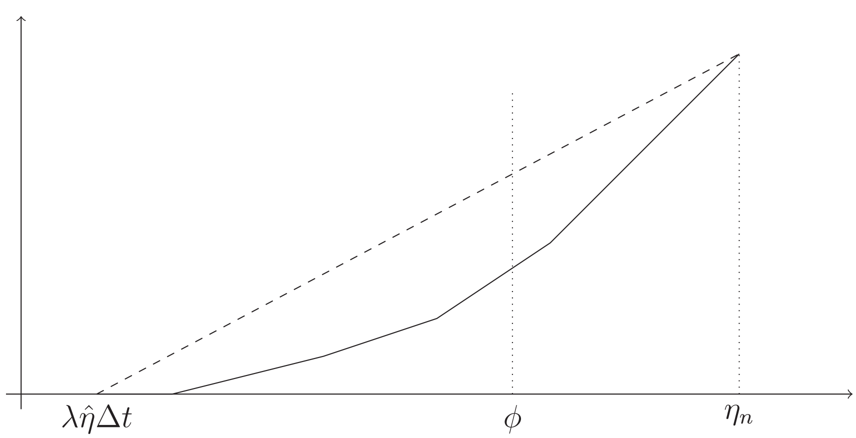

Chapter 3

Catastrophe Derivatives and Reinsurance Contracts: An Incomplete Markets Approach

3.1 Introduction

Catastrophe (CAT) derivatives are financial instruments indexed on a rare events process, a physical event whose occurrence reduces aggregate wealth and/or consumption (a catastrophe event). Such instruments have appeared often in recent years, fulfilling through securitization the financing needs of the insurance industry.[1] These CAT products typically pay a cash flow to their holders that is conditional on the catastrophe event occurring and, whose size is, in most cases, proportional to the intensity of the event or to the financial losses incurred by the holder as a result of the catastrophe event. Such financial instruments may trade over the counter or in organized exchanges. They include catastrophe bonds, whose coupon and/or principal are reduced by pre-specified amounts due to the occurrence of the CAT event, futures contracts whose payments are proportional to the difference of the events intensity from a reference value, and reinsurance contracts that most often include a deductible and a ceiling on payments as a result of the CAT event.

In this paper we model the pricing of the catastrophe derivative as a contingent claim on the underlying accumulated loss due to hurricane landings in a geographical region. We apply a recently introduced valuation methodology for such contingent claims that recognizes the fundamental incompleteness of financial markets arising from the occurrence of rare events. Unlike most studies on such derivatives this method does not assume that the CAT event risk can be diversified away, as in the classic Merton (1976) study. We argue that such an assumption is fundamentally contradicted by empirical evidence showing values of CAT-indexed financial instruments far in excess of their expected payoffs.[2] We apply the method to the valuation of a contingent claim with a non-convex payoff, for which closed form expressions are not available. We also argue that the continuous time approach that dominates most of the derivatives literature is not suitable for CAT events dependent on physical disasters, such as the hurricane contracts that we consider. For this reason we adopt a discrete time framework and derive tight bounds on the value of a reinsurance contract on CAT event losses that depend solely on the observed price of a futures contract

[1] See MMC Securities (2005).

[2] See the World Bank study by Lane and Mahul (2008).

indexed on the intensity of the event.

For our valuation we use the theoretical framework of the stochastic dominance (SD) methodology that was introduced to the study of CAT derivatives by Perrakis and Boloorforoosh (2013, PB), itself based on earlier literature.[3]. That literature was limited to the valuation of contingent claims with convex payoffs, for which closed form expressions are available. For this reason, the reinsurance contract in that earlier study could be valued only if the contract had a deductible or a cap, but not both. Further, the derived bounds on the contract only used information from the distribution of the CAT event amplitude and were, for that reason, relatively wide. Here, by contrast, we value the more realistic case of a contract with both a deductible and a cap, and we derive tight bounds by also using data from the futures market on CAT events. Last, we adopt a recursive discrete time approach, which recognizes that the CAT event of a hurricane landing develops over a number of days and, therefore, the number of such events in a given time frame is perforce limited. Our method is independent of distributional assumptions on the CAT event amplitude and can be generalized without reformulation to a Markovian process that may include dependence between the amplitude distribution and the frequency of occurrence of the event.

Although the underlying process is not a traded financial instrument, there exist futures contracts that allow the trading of the intensity of the landed hurricanes. Such contracts were introduced by the Chicago Board of Trade as early as 1992. They trade both in the organized exchanges, but also in a very active over-the-counter market in such instruments.[4] We demonstrate that we can derive tight bounds on the price of the CAT derivative from observed futures prices and the physical distribution of the CAT event amplitude. In the numerical calculation we find reasonably tight bounds for the price of the reinsurance contract conditional on the assumed price of the futures contract, which can be represented by a linear premium on the expected amplitude of the CAT event. The fact that this single parameter is

[3] See Perrakis and Ryan (1984), Levy (1985), Ritchken (1985), Perrakis (1986, 1988), and Ritchken and Kuo (1988)

[4] Hurricane futures and options contracts are trading in the Chicago Mercantile Index (CME). These contracts are indexed on the CHI, the CME hurricane index, but their prices are available only to subscribers or data purchasers. The CME indicated in a private communication that most such instruments are traded via blocks as option structures, and that brokers post markets in these niche products. For these reasons we did not use actual futures price data, as we discuss further on in Section 4 of this paper.

sufficient to accurately price the contract illustrates the major advantage of the SD method vis--vis alternative approaches that assume the knowledge of the entire *risk adjusted* distribution of the CAT event, allegedly extracted from other priced CAT instruments.[5] Aside from the fact that such an extraction assumes away the formidable data problems and accepts the efficiency of a non-transparent market of over-the-counter instruments, our numerical results show that the admissible values of the CAT reinsurance contract are strongly dependent on the characteristics of the contract and cannot be represented by an event-dependent markup over the Merton value, which often lies far outside the derived bounds.

Apart from the derivation of the reinsurance contract values on the basis of the futures contract price there are several other important contributions of this paper over and above the PB study. First, we introduce the discrete time valuation methodology for CAT instruments with non-convex payoffs and show that such instruments values cannot be replicated with options with convex payoffs, as in the arbitrage-based approaches. In turn, this methodology has applications to other commonly encountered instruments such as CAT-indexed bonds whose payoffs resemble a combination of a bond and a digital option. We also show that in many cases no closed form expression arises for continuous time valuation quite apart from the limitations of the physical CAT process. Last but not least, we show that the same bounds can be derived from arbitrage strategies using second-degree stochastic dominance considerations. These strategies can exploit the mispricing of the CAT instrument whenever it lies outside the bounds.

Catastrophe financial instruments have attracted relatively little interest in the mainstream financial literature, and most contributions have appeared in the insurance literature.[6] A catastrophe event is by definition a rare event and as such is modeled mathematically as a pure jump process, with Poisson arrivals and generally distributed amplitudes.[7] The losses inflicted by the event are generally considered proportional to the amplitude of a measurable characteristic of the event, such as the intensity of a hurricane or the extent of a flood. Thus,

[5]See, for instance, the discussion on market incompleteness in Geman and Yor (1997, p. 187) and in Bakshi and Madan (2002, pp. 107-108).

[6]See, for instance, Dassios and Jang (2002), Jaimungal and Wang (2006),Lee and Yu (2007), and Lin and Wang (2009)

[7]See Geman and Yor (1997), Froot (2001), and Muermann (2003).

the probability distribution of the losses conditional on the occurrence of the event, as well as the frequency of the event, can be extracted reasonably accurately from past data.

More contentious are the valuations of these losses, on which earlier studies have followed two divergent paths. Although it is well-known at least since Merton (1976) that rare events introduce incompleteness into the financial markets, this incompleteness has often been ignored or assumed away in mathematical insurance studies that rely on continuous time models paying scant attention to the underlying economic reasoning. In most of these models it is assumed that the CAT event risk is fully diversifiable through an efficient reinsurance market and a unique arbitrage-based price for the CAT financial instrument, equal to the expectation of the losses with the properly estimated financial process. Nonetheless, this assumption is clearly not applicable if catastrophe risk has economy-wide implications.[8] More to the point, the diversifiable CAT risk assumption is flatly contradicted by the empirical evidence presented by Lane and Mahul in a 2008 World Bank survey of the markets for CAT instruments over the period 1997-2008, in which the average instrument traded at 2.69 times the expected loss. Clearly, with such multipliers any accuracy in modelling the probabilistic structure of the event is dwarfed by the error in the economic assumption.

On the other hand, the more traditional insurance literature does recognize the existence of an insurance premium on the expected loss, but relies mostly on stylized single-period models that are not useful in pricing financial instruments on the basis of their cash flows.[9] In-between the two literature strains are some studies that pay lip service to the economic valuation of the cash flows, mostly by following the option pricing literature under jump diffusion asset dynamics as in Bates (1991).[10] Unfortunately these valuation exercises rely on market equilibrium arguments using particular utility functions to transform the probability distributions and make the values dependent on the risk aversion of a representative investor. Since there is no general agreement on the size of this parameter,[11] the transformation is left

[8]This was already known from the option pricing literature in the presence of event risk. See, for instance, Bates (1991). Further, Ibragimov et al. (2009) note that the efficient reinsurance assumption may not be satisfied in real markets even if catastrophe risk does not have economy-wide impact.

[9]See, for instance, Barrieu and Loubergé (2009) and Bernard and Tian (2009).

[10]See Christensen and Schmidli (2000), Duan and Yu (2005) and Chang et al. (2010).

[11]See Kocherlakota (1996), as well as the comments by Chang et al. (2010, p. 28, note 6).

unspecified in most empirical applications of jump-diffusion processes.

The incomplete markets bounding approach introduced by PB is an alternative to the equilibrium results that does not require fundamentally unobservable elements. It is based on the only assumption of the monotonicity of the pricing kernel used in valuing financial instruments with respect to the CAT event amplitude. It is, therefore, well-suited to the valuation of contingent claims indexed on a catastrophe event, given the market incompleteness that it gives rise to, as well as the fact that the contingent claims are negative beta assets, whose cash flows increase when aggregate wealth decreases. Here we extend the PB results in a discrete time context, by showing that their SD approach can also be applied to claims whose payoff is not necessarily convex, and by tying it explicitly to the price of a CAT event futures contract that yields tight bounds on the value of the claim solely as functions of the futures price. The discrete time representation has the added advantage of modeling reality much more accurately than its continuous time counterpart: hurricanes and floods typically takes several days to develop and land in a particular region, which puts stringent limits on the number of possible landings in any finite time interval.

In the next section we formulate the basic equations of the market model used in pricing the catastrophe instruments with a non-convex payoff in a single-period context and show that they cannot be derived from conventional call options. Section III derives an algorithm that extends the valuation results to any number of periods. Section IV discusses the estimation of the parameters and the numerical calculations of the bounds and compares the discrete to the continuous time model and to available empirical results. Section V concludes.

3.2 The Single-Period Model

It is well known since Merton (1976) that the market for rare event risk is incomplete, and the arbitrage valuation of a derivative instrument cannot yield a unique price. To value CAT derivatives we use the market equilibrium under the stochastic dominance assumptions formulated in previous studies:[12] there exists at least one utility-maximizing risk averse investor (*the trader*) in the economy who holds a portfolio containing an index and the riskless asset, this particular investor is marginal in the derivative market[13], and the riskless rate is non-random. In this particular case it is also assumed that the index return depends linearly on the intensity of the CAT event.

The market equilibrium conditions for a trader holding a portfolio of the riskless asset and the index and maximizing recursively the expected utility of final wealth over a number of periods longer than the derivatives maturity were derived in an earlier study[14] and will be only briefly summarized here. Let $x_t + y_t$ denote the current value of the traders portfolio, with x_t and y_t denoting, respectively, the amounts invested in the riskless asset and the index. Let also $R > 1$ be the return of the riskless asset per period. Time is discrete $t = 0, 1, ..., T$, with intervals of length Δt. In each interval, the return of the index has the following form:

$$\frac{y_{t+\Delta t} - y_t}{y_t} \equiv z_{t+\Delta t} = v_{t+\Delta t} + \gamma H \Delta N \tag{3.2.1}$$

We assume that $E_t[z_{t+\Delta t}] \geq R - 1$. The term H represents the level of hurricane intensity, measured in CME Hurricane Index (CHI) units, and N is a Poisson counting process with intensity with intensity λ. We consider a single-period model with horizon Δt, in which the probability of a hurricane landing in the area covered by the CAT derivative is $\lambda \Delta t$.[15] Hence, the return $z_{t+\Delta t}$ is the convolution of the process $\gamma H_{t+\Delta t}$ with probability $\lambda \Delta t$ and $H_{t+\Delta t} = 0$ with probability $1 - \lambda \Delta t$, and of $\nu_{t+\Delta t}$, which is the index return component that is independent of the CAT event and whose distribution may depend on the current stat level. The parameter γ denotes the impact of the hurricane on the traders portfolio return. If $\gamma = 0$ then the CAT event risk is diversifiable and does not affect the traders optimal invested wealth, and we have a case similar to that of Merton (1976). In such a case the CAT event risk is not priced in equilibrium and we can obtain a unique price for the contingent claim,

[12]See, in particular, Constantinides and Perrakis (2002, 2007), Perrakis and Boloorforoosh (2013), and Oancea and Perrakis (2014).

[13]This condition will be defined more precisely in Section 4.

[14]See Perrakis and Boloorforoosh (2013).

[15]Alternatively, we can set this probability equal to $1 - e^{-\lambda \Delta t} = \lambda \Delta t + o(\Delta t)$, in which case the probability of no landing is $e^{-\lambda \Delta t} = 1 - \lambda \Delta t + o(\Delta t)$.

equal to the discounted expected payoff under the physical distribution of the underlying loss process. In our analysis we assume that γ is negative, meaning that the CAT event risk is not diversifiable.[16] Let also F denote the futures price of a contract that matures at the end of the single-period horizon and whose payments are proportional to the hurricane intensity index. The contract execution is triggered by a hurricane landing, after which the contract expires and a new contract is issued till the end of the hurricane season. Without loss of generality we define H_0 as a hurricane intensity level of 0, the absence of a hurricane landing whose arrival triggers the futures contract maturity; alternatively, it is a below-hurricane-level wind intensity. With this definition we let (q_i, H_i) be the combined distribution of wind intensity level and hurricane landing, with $q_0 = 1 - \lambda\Delta t$, $q_i = p_i\lambda\Delta t, i = 1, \dots, n$.

Let $V_t \equiv \kappa \sum_{\tau=0}^{\tau=N_t} H_\tau$ denote the accumulated losses from the CAT event, assumed proportional to the hurricane intensity, where κ represents the dollar loss per CHI units. As noted in the introduction, the accumulated loss process is a discrete time process, since the formation and landing of a hurricane takes time. If the trader also takes a marginal position in the reinsurance contract, $C_t(V_t)$, that expires at time $T \leq T'$, then the following relations characterize the market equilibrium in any single trading period $(t, t + \Delta t)$, assuming no market frictions:

$$\begin{aligned} &E_t[X(z_{t+\Delta t}] = 1 \\ &E_t[(1 + z_{t+\Delta t})X(z_{t+\Delta t})] = R \\ &C_t(V_t) = \frac{1}{R} E_t[C_{t+\Delta t}(V_t + \kappa H_{t+\Delta t})X(z_{t+\Delta t})] \end{aligned} \tag{3.2.2}$$

In (3.2.2) $X(z_{t+\Delta t})$ represents the pricing kernel or normalized rate of substitution of the trader evaluated at her optimal portfolio choice. Since the trader is assumed to be risk averse and since she has a marginal position in the contingent claim, it can be easily seen that the pricing kernel would be monotone non-increasing in $z_{t+\Delta t}$. Let also $\hat{X}_{t+\Delta t} = E_t[X_{t+\Delta t}, H_{t+\Delta t}]$ denote the pricing kernel conditional on the intensity of the landed hurricane. Since the hurricane is obviously an event that negatively affects aggregate consumption, as modeled

[16] We make no assumption regarding the value of γ since its effect on the bounds is incorporated into the futures price.

by $\gamma < 0$ in (3.2.1), $\hat{X}_{t+\Delta t}$ is now non-decreasing in the intensity of the CAT event, with values $\hat{X}_i, i = 1 \ldots, n$, such that $\hat{X}_0 \leq \hat{X}_1 \leq \ldots \leq \hat{X}_n$.

First we consider a single-period case where there is one period left to the end of the hurricane season, and the reinsurance contract that expires in the next period is valid for only one hurricane landing.[17] We model the reinsurance contract C, which is a contract against wind damages, in the form of a spread. We assume that the reinsurance contract has a deductible and a ceiling, corresponding to the hurricane intensities H_l and H_h. The reinsurance contract has the following payoff, C_T, at the end of the hurricane season:

$$\begin{aligned} &C_{Ti} = 0, H_i \leq H_l \\ &C_{Ti} = \kappa(H_i - H_l), H_l < H_i \leq H_h \\ &C_{Ti} = \kappa(H_h - H_l), H_i > H_h \end{aligned} \tag{3.2.3}$$

The market equilibrium equations are, therefore, the following:

$$\begin{aligned} &\textstyle\sum_{i=0}^{n} q_i \hat{X}_i = 1, \\ &\textstyle\sum_{i=0}^{n} q_i \hat{X}_i H_i = F, \\ &\hat{X}_0 \leq \hat{X}_1 \leq \ldots \leq \hat{X}_n \end{aligned} \tag{3.2.4}$$

The last equation in (3.2.4) reflects the fact that any cash flows that accrue because of the hurricane event must be considered as an equivalent to a "negative beta" stock. Moreover, the price of this claim at a time of prior to the expiration of the contract is given by the following

$$C = R^{-1} \sum_{i=0}^{n} q_i \hat{X}_i C_{Ti} = \kappa\Big(\sum_{l+1}^{h-1} q_i \hat{X}_i (H_i - H_l) + (H_h - H_l) \sum_{h}^{n} q_i \hat{X}_i\Big) R^{-1} \tag{3.2.5}$$

Since the number of states, n, is obviously greater than 2, the market is incomplete, the

[17] Alternatively, we may consider the single-period case as referring to the aggregate landings over a given time period, the convolution of individual hurricane landed intensities covered by a single futures contract. The more common case, in which the reinsurance contract covers the total loss during a hurricane season but each futures contract covers a single landing, will be examined in the multiperiod model.

pricing kernel is not unique, and no unique price can be defined by arbitrage methods alone. Following the linear programming (LP) approach pioneered by Ritchken (1985), we develop the tightest upper and lower bounds that the market equilibrium described in (3.2.4) can support. Further, the payoff is not convex[18] with respect to the underlying random variable and the underlying asset is a negative beta security, implying that the expressions need to be modified. Nonetheless, the Ritchken approach can be easily adapted to account for the negative beta. In an appendix available from the authors on request it is shown that the bounds on the contingent claim are found by the following transformation of the market equilibrium. For a set of non-negative numbers $\epsilon_0, \ldots, \epsilon_n$, we set $\hat{X}_0 = \epsilon_0, \hat{X}_1 = \epsilon_0 + \epsilon_1, \ldots, \hat{X}_n = \sum_0^n \epsilon_i$, and we define $\tilde{X}_i = \epsilon_i \sum_{k=i}^{k=n} q_k$. We also define the following conditional moments:

$$\overline{H}_i = \frac{\sum_{j=i}^n q_j H_j}{\sum_j^n q_j}, \overline{C}_{Ti} = \frac{\sum_{j=i}^n q_j C_{Tj}}{\sum_j^n q_j}, i = 0, \ldots, n. \tag{3.2.6}$$

Replacing these relations into (3.2.4), we have the following transformed market equilibrium conditions:

$$\begin{aligned} &\textstyle\sum_0^n \tilde{X}_i = 1, \\ &\textstyle\sum_0^n \tilde{X}_i \overline{H}_i = F, \\ &\tilde{X}_i \geq 0, \; i = 0, \ldots, n. \end{aligned} \tag{3.2.7}$$

Clearly $\overline{H}_0 = E[H_i]$, $\overline{H}_n = H_n$, and similarly $\overline{C}_{T0} = E[C_T]$, $\overline{C}_{Ti} = \kappa(H_h - H_l), i \geq h$. The price of the contingent claim represented by the reinsurance contract in (3.2.5) becomes now equal to

$$C = R^{-1} \sum_0^n \tilde{X}_i \overline{C}_{Ti}, \tag{3.2.8}$$

for a set $\{\tilde{X}_i\}$ satisfying relations (3.2.7), which represents an admissible martingale prob-

[18]The bounds are given by closed form expressions when the claims payoff is convex. See Perrakis and Boloorforoosh (2013, p. 3160).

ability. We want to determine the option bounds that this equilibrium supports, which are given as the solution of the following programs:

$$C_{max} = R^{-1} \max_{\tilde{X}_i} \sum_0^n \tilde{X}_i \overline{C}_{Ti} \ , \ C_{min} = R^{-1} \min_{\tilde{X}_i} \sum_0^n \tilde{X}_i \overline{C}_{Ti} \tag{3.2.9}$$

subject to (3.2.7).

If the contract has no deductible or no ceiling then the payoff is respectively concave or convex with respect to the hurricane intensity. In such cases the solution to (3.2.9) can be found by an application of the dual of the LP (3.2.9), as in Ritchken (1985), which relies on such payoff shapes. Here it is possible to extend the the LP approach by means of the following result.

Lemma 3.1. *The graph of the expected conditional payoff, $\overline{C}_{Ti}$, as a function of the expected conditional intensity, $\overline{H}_i$, is concave over the region $\overline{H}_i > \overline{H}_l$, while it is convex for $\overline{H}_i \leq \overline{H}_l$.*

Proof. See Appendix A. □

Given this result we may now find the bounds on the admissible values of the reinsurance contract as functions of the futures price F. This price plays the role of the stock price in conventional financial derivatives, and it turns out that the effects of market incompleteness can be represented by the excess of F over $\overline{H}_0$, the expected hurricane intensity. The bounds of the reinsurance contract are given by the following result.

Proposition 3.1. *The upper and lower bounds $C_{max}(F)$ and $C_{min}(F)$ of the reinsurance contract, the solutions of the LP (3.2.6)-(3.2.9), depend on the size of the futures price F relative to the deductible and the ceiling of the contract. The bounds are found as the intersection of the vertical line stemming from F and the boundaries of the convex hull of the conditional payoff, as illustrated in Figure 3.2.1 and are described in equations (3.B.1)-(3.B.3) in the appendix..*

Proof. See Appendix B. □

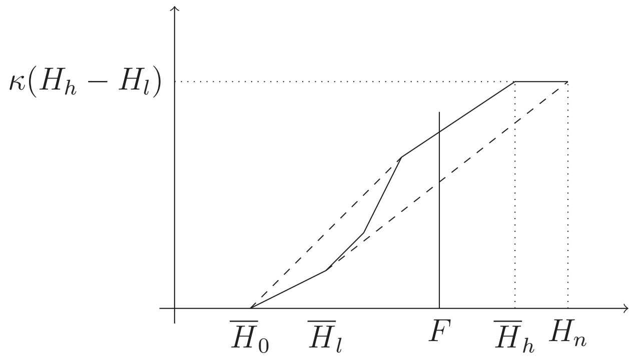

Figure 3.2.1: **Convex Hull of the Conditional Payoff**

The results of Proposition 1 apply to the valuation of a reinsurance contract on the accumulated hurricane losses, whose payoff has the shape of a vertical spread. This payoff, however, can be replicated with a long position in a call option with strike price equal to the deductible of the reinsurance contract, and a short position in a similar call option with strike price equal to the ceiling of the reinsurance contract. In complete markets, and in valuation methods that rely on a representative investor, this replication is sufficient to price the vertical spread. In the stochastic dominance approach the replicating call options upper and lower bounds can be calculated using the method presented by Perrakis and Boloorforoosh (2013), as the discounted expected option payoff under the upper and lower bound distributions which are available in closed form. The bounds on the reinsurance contract can thus be obtained from the bounds on the two calls. In Appendix C we show, both theoretically and numerically that except for the trivial cases where the deductible is very low and/or the ceiling is very high, the method presented in this paper yields considerably tighter bounds.

Although the bounds were derived under the assumption that the amplitude distribution is discrete, the derived expressions can be obviously adapted without reformulation to a distribution with compact support. Further, the valuation method presented here can be easily extended with very little reformulation to all other rare event instruments with non-convex payoff. These include the important case of catastrophe bonds issued by an insurer. The latter are straight bonds for which the issuer retains a digital call option in the form of a fixed reduction from the principal in case of a hurricane landing.

3.3 Multiperiod Analysis

In a multiperiod analysis the market equilibrium equations (3.2.4) or (3.2.7) remain the same.[19] What changes is the nature of the payoffs to the reinsurance contract, which are similar to those of an *Asian* option. The reinsurance contract typically applies to the cumulative losses arising during the entire hurricane season. Hence, at any time during the season the state of the contract is described by the accumulated losses up to that point, whose amount may change during any time period due to a new hurricane landing. As for the CME hurricane futures, we assume that after each hurricane landing the existing futures contract matures, and a new futures contract is initiated. We assume that the reinsurance contract is of the European type and is exercisable only at T.

Let $t = 0, 1, \ldots, T$ denote the sequence of time points from the beginning to the end of the hurricane season. As defined in the previous section, V_t denotes the accumulate losses till time t and $C_t(V_t|F_t)$ is the corresponding value of the contract at t if the observed futures price is F_t. We assume that the hurricane landings arrive independently with probabilities $\lambda\Delta t$ per period and their intensities are independent and identically distributed (iid) random variables with distributions (p_i, H_i), $i = 0, 1, \ldots, n$. Further, we assume that in any period there cannot be more than one hurricane landing. Letting again $q_0 = 1-\lambda\Delta t$, $q_i = p_i\lambda\Delta t$, $i = 1, \ldots, n$ denote the combined probabilities of a hurricane landing and of the hurricane intensity we have

$$V_0 = 0,\ V_t = V_{t-1} + \kappa H_t,\ H_t \sim\ (q_i, H_i),\ i = 1, \ldots, n,\ t = 1, \ldots, T \tag{3.3.1}$$

Observe that barring a shift in the pricing kernel due to a shift in preferences the iid assumption on hurricane occurrence and intensity implies from (3.2.4) that F_t will remain the same for all t. Nonetheless, our formulation is sufficiently flexible to accommodate predictably seasonal shifts in the hurricane intensity distribution within the reinsurance contracts maturity,

[19]This follows from the fact that the value of the futures contract is zero at any time point and the contract can be closed at any time point. Hence, the futures price is equal to the value of all possible cash flows that accrue to the contracts position.

which will also be reflected by predictable changes in the futures price. It is also possible to extend the model by relaxing the iid assumption of hurricane intensities and replace it by a Markovian one, in which the intensity of a hurricane depends on that of the previous one.

At time $t = T-1$ the bounds on the value of the reinsurance contract can be found by applying the analysis of the previous section, with the important *caveat* that the terminal payoffs depend on the accumulated losses V_{T-1} which is the state variable. Both the deductible and the ceiling on cash flows are imposed on the terminal accumulated losses V_T, at two levels $\overline{V}_l$ and $\overline{V}_h$. This simplifies the single-period analysis. Thus, for $V_{T-1} \geq \overline{V}_h$ we clearly have $C_T = \overline{V}_h - \overline{V}_l$ for all values of V_T, while for $V_{T-1} \geq \overline{V}_l$ the terminal payoff is concave (linear and then constant) in H_i for all values of V_T. For $V_{T-1} < \overline{V}_l$ we redefine the values $H_l(V_{T-1}) \geq H_0$ and $H_h(V_{T-1}) < H_n$, with $l(V_{T-1})$ and $h(V_{T-1})$ being the smallest integers such that $\kappa H_l + V_{T-1} \geq \overline{V}_l$, $\kappa H_h + V_{T-1} \geq \overline{V}_h$, with the payoff having the same shape as in the single-period problem; note that both $l(V_{T-1})$ and $h(V_{T-1})$ are decreasing functions. With these redefinitions the payoff of the contingent claim $C_{T-1}(V_{T-1}|F_{T-1})$ becomes at time T a function that has the same shape as in Figure 3.2.1, with the starting point displaced by the amount V_{t-1}:

$$\begin{aligned}
&V_{T-1} < \overline{V}_l : \\
&\quad C_T = 0,\ H_T \leq H_l;\ C_T = V_{T-1} + \kappa H_T - \overline{V}_l,\ H_T \in (H_l, H_h]; \\
&\quad C_T = \overline{V}_h - \overline{V}_l,\ H_T > H_h \\
&\overline{V}_l \leq V_{T-1} < \overline{V}_h : \\
&\quad C_T = V_{T-1} + \kappa H_T - \overline{V}_l,\ H_T \leq H_h;\ C_T = \overline{V}_h - \overline{V}_l,\ H_T > H_h \\
&V_{T-1} \geq \overline{V}_h : \\
&\quad C_T = \overline{V}_h - \overline{V}_l, \forall H_T
\end{aligned} \tag{3.3.2}$$

Given this payoff, the following results allow us to estimate the upper and lower bounds $C_{T-1,max}(V_{T-1}|F_{T-1})$ and $C_{T-1,min}(V_{T-1}|F_{T-1})$ on the value $C_{T-1}(V_{T-1}|F_{T-1})$.

Lemma 3.2. *The bounds $C_{T-1,max}(V_{T-1}|F_{T-1})$ and $C_{T-1,min}(V_{T-1}|F_{T-1})$ on the reinsurance contract lie on the convex hull of the conditional payoffs, and their values depend on V_{T-1}.*

Proof. Proof. The equations representing the boundary of the convex hull and the proof are presented in Appendix D. □

Lemma 3.3. *For any given futures price F_{T-1} the bounds on the contingent claim $C_{T-1}(V_{T-1}|F_{T-1})$ are increasing functions of V_{T-1}, concave for the upper and convex for the lower, and are identified by the intersection of the vertical line stemming from the futures price with the convex hull of the conditional payoff.*

Proof. The equations representing the bounds and the proof are presented in Appendix E. □

We may now formulate the recursive problem that derives the multiperiod bounds for the value of the contingent claim $C_t(V_t|F_t)$ corresponding to the reinsurance contract. Assuming no shift in preferences, so that the futures price remains the same, $F_t = F$ for all t, at any time $t < T-1$ the capital market equilibrium must satisfy the following system:

$$\begin{aligned}
&\sum_0^n q_i \hat{X}_i = R^{-1}, \ \sum_0^n q_i \hat{X}_i H_i = FR^{-1}, \\
&C_t(V_t|F) = R^{-1} \sum_0^n q_i X_i C_{t+1}(V_{t+1}|F) = R^{-1} \sum_0^n q_i \hat{X}_i C_{t+1}(V_t + \kappa H_i|F), \\
&\hat{X}_0 \leq \hat{X}_1 \leq \ldots \leq \hat{X}_n
\end{aligned} \tag{3.3.3}$$

Define also the following values, the counterpart of (3.2.6):

$$\overline{C}_{t+1,\alpha,i} = \frac{\sum_{j=i}^n q_i C_{t+1,\alpha,j}(V_t + \kappa H_j|F)}{\sum_{j=i}^n q_j}, \ \alpha = \max, \min, \ i = 0, \ldots, n. \tag{3.3.4}$$

We may then prove the following result.

Proposition 3.2. *The value $C_t(V_t|F)$ of the reinsurance contract lies between the following recursive bounds $C_{t,min}(V_t|F)$ and $C_{t,max}(V_t|F)$ for all $t < T$:*

$$
\begin{aligned}
C_{t,max}(V_t|F) &= \frac{\overline{H}_{i^*+1} - F}{\overline{H}_{i^*+1} - \overline{H}_{i^*}} \overline{C}_{t+1,max,i^*} + \frac{F - \overline{H}_{i^*}}{\overline{H}_{i^*+1} - \overline{H}_{i^*}} \overline{C}_{t+1,max,i^*+1}, \\
C_{t,min}(V_t|F) &= \frac{\overline{H}_n - F}{\overline{H}_n - \overline{H}_0} \overline{C}_{t+1,min,\underline{i}} + \frac{F - \overline{H}_0}{\overline{H}_n - \overline{H}_0} \overline{C}_{t+1,min,n}, \;\; V_{T-1} < \overline{V}_l \\
C_{t,min}(V_t|F) &= \frac{\overline{H}_n - F}{\overline{H}_n - \overline{H}_0} \overline{C}_{t+1,min,0} + \frac{F - \overline{H}_0}{\overline{H}_n - \overline{H}_0} \overline{C}_{t+1,min,n}, \;\; V_{T-1} \geq \overline{V}_l
\end{aligned}
\tag{3.3.5}
$$

where $C_{T-1,max}(V_{T-1}|F)$ *and* $C_{T-1,min}(V_{T-1}|F)$ *are as defined in Lemma 3, and* i^* *is a state such that* $\overline{H}_{i^*} \leq F \leq \overline{H}_{i^*+1}$.

Proof. See Appendix F. □

The recursive evaluation of the reinsurance contracts bounds is computationally very simple, in spite of the complexity of the notation in representing the convex hull of the one-period payoff. This hull remains the same at every recursion and every state space node, and what changes is the starting point that is a function of the cumulative losses to that node, which determines the location of the futures price F. The convexification of the bounds is also maintained in every recursion, thus allowing the closed form expression of the LP solution as derived by Ritchken (1985). Unlike the continuous time derivatives with convex payoffs examined in Perrakis and Boloorforoosh (2013), there is no general closed form solution for the bounds on the value of the contract, because the upper and lower bound distributions are now state dependent. Nevertheless, the evolution of the value is Markovian and can be easily estimated for realistic numbers of partitions in the time subdivision.[20]

3.4 The Arbitrage Derivation of the CAT Bounds

In this section we prove that the derived bounds on the CAT reinsurance contract are also arbitrage bounds, in the sense that their violation indicates a profit opportunity for the trader. Such an opportunity differs from conventional arbitrage, insofar as the SD bounds violations are exploited by adding to the traders wealth zero net cost portfolios that contain

[20] For the same reason it is not easy to visualize the continuous time limit of the bounds for such derivatives with non-convex payoffs. Such continuous time limits, however, are not realistic for physical CAT events such as hurricanes, earthquakes or floods.

the mispriced contract and that offer superior *risk-adjusted* returns to the trader independent from his/her wealth or attitude towards risk. The derivation of the bounds by this type of arbitrage is an alternative to the LP approach used in this paper and was presented in detail in Oancea and Perrakis (2014). Nonetheless, that derivation was applicable to positive beta securities and relied heavily on the convexity of the payoffs, resulting in closed form expressions for the arbitrage portfolios. For this reason we present here the single-period derivation of the reinsurance contract bounds by arbitrage, since the zero net cost portfolios necessary for the exploitation of contract mispricing need to be adapted to the specific problem. We shall derive the arbitrage strategies for the case shown in Figure 3.2.1, with other cases left as an exercise.

For expository purposes and without loss of generality we assume, unlike the LP approach, that the landed hurricane intensity has a distribution $P(H)$ with compact support $H \in [H_1, H_n]$, $H_1 \geq 0, H_n < \infty$. The CAT event has amplitude 0 with probability $1 - \lambda\Delta t$ and H with distribution $\lambda\Delta t P(H)$. We also assume, without loss of generality, that the universe of traders consists of Π identical agents, in which case the marketed contract is $C_\Pi(V_{T-1}) = \frac{1}{\Pi}C(V_{T-1})$. Setting for simplicity $V_{T-1} = 0$ and defining the multipliers $\kappa_\Pi = \frac{\kappa}{\Pi}$, our valuation reduces to the estimation of the bounds of the contingent claim C_Π. The payoff, denoted by $C_{\Pi T}(H)$, has the same functional form as in (2.3) with κ_Π replacing κ conditional on the occurrence of a hurricane landing, and 0 otherwise. Let also $\Omega(x_t + y_t)$ denote the traders value function, the maximized expected utility of her portfolio at time $t = T-1$, which is increasing and concave. By definition, $\Omega(x_{T-1} + y_{T-1}) = E_{T-1}[\Omega(x'_{T-1}R + y'_{T-1}(1 + z_T))]$, where the portfolio $(x'_{T-1}, y'_{T-1})$is the optimally selected asset allocation at $T - 1$. For the index return given by (3.2.1) we also define the function $W_{T-1}(H)$ as follows:

$$\Omega(x_{T-1} + y_{T-1}) = E_{T-1}[E[\Omega(x'_{T-1}R + y'_{T-1}(1 + z_T)) \,|v_T]] = E_{T-1}[W_{T-1}(H)] \qquad (3.4.1)$$

From the properties of the value function and the index return (3.2.1) it follows that the derivative of the function $W_{T-1}(H)$ is increasing in the hurricane intensity, a property that will be important in deriving the bounds.

To derive the upper bound we assume that the trader can short the contract for a price of C, which is invested in the cash account. Suppose the trader shorts the contract and also adopts a long position in $\frac{CR}{MF}$ futures contracts, whose price is F and whose multiplier is denoted by M. This is a zero net cost position, corresponding to the short contract together with a long position in an instrument with payoff proportional to landed hurricane intensity, with coefficient $\frac{CR}{F}$, which is strictly less than κ_N to avoid arbitrage between the futures market and the reinsurance contract; this trader is termed the C-trader and the zero net cost portfolio is termed $h(H)$. For a properly priced contract such a C-trader should not be able to increase her value function over that of an unspecified risk averse trader with identical characteristics and wealth positions. Let $\Omega^C(x_t + y_t)$ denote the value function of the C-trader and set

$$\Delta_{T-1} = \Omega^C(x_{T-1} + y_{T-1}) - \Omega(x_{T-1} + y_{T-1}) \tag{3.4.2}$$

This difference will certainly not increase if the C-trader adopts the same portfolio revision policy as the regular trader. We then have:

$$\begin{aligned}\Delta_{T-1} =& \Omega^C(x_{T-1} + y_{T-1}) - \Omega(x_{T-1} + y_{T-1}) = E_{T-1}[W^C_{T-1}(H)] - E_{T-1}[W_{T-1}(H)] \geq \\ & E_{T-1}[W_{T-1}(H + h(H)) - W_{T-1}(H)] \geq E_{T-1}[W^1_{T-1} h(H)]\end{aligned} \tag{3.4.3}$$

In (3.4.3) the term W^1_{T-1}denotes the derivative of $W_{T-1}(H + h(H))$. Since the trader is by assumption marginal in the derivative market, W^1_{T-1} is also increasing in the hurricane intensity$H$. Since $h(H) \geq 0$ by construction when there is no hurricane landing, we concentrate on the shape of $h(H)$ conditional on a landing. Note that $h(H)$ is always initially increasing for $H \in [H_1, H_l]$, decreasing for $H_l < H_i \leq H_h$, and increasing again for $H \in [H_h, H_n]$. Its zeroes are at most two beyond the origin and are parameter-dependent, but for the case shown in Figure 3.2.1 there are exactly two zeroes, denoted by$H^*$and $\tilde{H}$.

We then have

$$\Delta_{T-1} \geq E_{T-1}[W^1_{T-1} h(H)] \geq \lambda \Delta t \int_{H^*}^{H_n} W^1_{T-1} h(H) dP(H) \tag{3.4.4}$$

So, to assess whether $\Delta_{T-1} \geq 0$ it suffices to consider only the last term of (3.4.4). We have:

$$\int_{H^*}^{H_n} W^1_{T-1} h(H) dP(H) \geq W^1_{T-1}(\tilde{H}) \int_{H^*}^{H_n} h(H) dP(H) \tag{3.4.5}$$

Substituting $h(H) = \frac{CRH}{F} - C_{NT}(H)$ into the last integral of (4.5) and dividing by$P(H^*)$ we conclude that $\Delta_{T-1} \geq 0$ unless $C \leq \frac{1}{R} \frac{FE[C_{NT}(H)|H \geq H^*]}{E[H|H \geq H^*]} = \frac{E[C_{NT}(H)|H \geq H^*]}{R}$. This last expression is, however, identical to the upper bound derived in Proposition 1 and shown in the second expression of (3.B.1) and (3.B.3).

The derivation of the lower bound is also straightforward, albeit a little more complex. The arbitrage strategy here is to purchase the reinsurance contract from the cash account and adopt a short position in $\frac{CR}{MF}(1+\alpha)$, $\alpha > 0$ futures contracts, with$\frac{CR}{F}(1+\alpha) < \kappa_N$. For the no landing case this strategy clearly dominates the no action case, so we concentrate on a landed hurricane with intensity H. The payoff function $h(H)$ is initially decreasing and equal to $\frac{CR}{F}\alpha - \frac{CR}{F}(1+\alpha)H$, becomes increasing at $H_l$and decreasing again at H_h. Hence, it can have at most three zeroes, and we choose α by imposing the condition that the third zero should be at H_n, so that

$$\frac{CR}{F}\alpha - \frac{CR}{F}(1+\alpha)H_n = C_{NT}(H_n) \tag{3.4.6}$$

Letting H^* and $\tilde{H}$ denote again the first and second zeros, we repeat the reasoning leading to the derivation of the upper bound, which was given by the no stochastic dominance condition that implies $E[h(H)\,|H \geq H^*] = 0$. This same condition, together with (3.4.6), allows the derivation of the lower bound, which is left as an exercise, together with the demonstration that it is the same expression given in (3.B.1) and (3.B.3).

The derivation procedure presented in this section also illustrates the trading strategies that can be used to exploit the mispricing in the reinsurance contract. When the price of the reinsurance contract violates the bounds derived in section III, the trader can increase her expected utility by adding a zero net cost position in the reinsurance contract and the futures contract to her portfolio, using the strategies described above.

In the next section we apply the bounds estimation procedure to a notional reinsurance contract on hurricane wind damage in the state of Florida. Without loss of generality it is assumed that the contract covers *statewide* assessed damage, implying that the underlying random process is the landed intensity in the entire state.

3.5 Estimation and Numerical Results

The basic traded futures contract that will be used to value the hurricane reinsurance contract is the hurricane futures contract traded by the CME. It is indexed on the CHI, initially denoting the Carvill Hurricane Index, but the index was subsequently purchased by the CME and renamed CME Hurricane Index. Several types of contracts indexed on the CHI trade in the CME. The contract that is the most relevant for our purposes is the hurricane seasonal contract, quoted in CHI index points and representing the accumulated CHI total for all hurricanes that occur in a specified location within a given calendar year. The contract multiplier is $1000 per CHI point.

We start by the following discretization of the time and state space. We divide the time $T-t$ until the maturity of the reinsurance into N subdivisions of equal length Δt. We then construct a multinomial lattice with $n+1$ branches emanating from each node to represent the cumulative catastrophe loss associated with the hurricane landings. Figure 3.1 depicts the cumulative loss process for two periods.

[Figure 5.1 about here]

Hence, at any time t where the cumulated loss is V_t we have $n+1$ possible outcomes for the next period cumulative loss at time $t+\Delta t$. The cumulative loss could stay the same, corresponding to the event of no hurricane landing, ($H_0 = 0$), or it could go up to $V_t + \kappa H_i$, $i = 1, \ldots, n$. After N periods we will have the probability distribution of the cumulative losses at time T.

We calibrate the underlying loss process based on the data available from CME for hurricane landings in the state of Florida for the period 1998-2007. We extract the conditional distribution of intensities, (p_i, Hi), $i = 0, 1, \ldots, n$, from the histogram of CHI values for Florida landings for $n = 3$.[21] Hence, the combined probability distribution of landing and intensities, (q_i, H_i), results in a quadrinomial lattice for the underlying cumulative loss process. Moreover, we assume that in every period there can be only one hurricane landing.[22] So, when the number of time subdivisions is N, the minimum and maximum number of hurricane landings are zero and N, respectively.

Conditional on landing, the hurricane intensity could be 3.42, 7.45, or 11.48 with equal probabilities, yielding an expected hurricane intensity conditional on landing equal to 7.45. The intensity of the Poisson process, λ, representing the rate of hurricane landing arrival is estimated as the average number of landings per hurricane season in Florida and is set equal to 0.9. It is assumed, without loss of generality, that λ is constant (no seasonality within the contract period). Moreover, we estimate the loss coefficient $\kappa = 0.41$ (in billion dollars per CHI) from the reported losses associated with nationwide landings during the period of $1975 - 2005$.[23] The average annual losses associated with hurricane landings in the state of Florida are equal to \$6.7 billion. These numbers imply that if there were 100 insurers in the region with similar policies, each would be exposed to \$67 million of losses per annum.

In order to calculate the bounds as described in the previous sections we need the price of the seasonal hurricane futures. Unfortunately, these contracts trade in a non-transparent

[21] Since the lattice is non-recombining, after N periods we have $(n+1)^N$ states. This makes the computations very intensive and even infeasible for large values of N. Nevertheless, since it is not reasonable to expect large number of hurricane landings in a geographical region in every hurricane season, there is no need to assign large values to N.

[22] This assumption is justified because the formation and landing of hurricanes takes time, and it is a reasonable to assume that, in general, there cannot be more than one hurricane landing over a period of 15-20 days.

[23] The coefficient was estimated from a linear regression of the hurricane losses on the intensity of the landed hurricanes in US, based on the data provided by CME for the period of $1975 - 2005$.

market and mostly in block trades, implying that the market prices are illiquid and only approximately reliable as reflecting the "true" value of the hurricane event. Therefore, in our calculations we assume that the price of the seasonal futures contract is linked to the expected intensity of the hurricane according to the following linear relation.

$$F = g \cdot E[H],\ g > 1 \tag{3.5.1}$$

The expectation in (3.5.1) is taken with respect to the combined distribution of landing and intensities. Since the distribution of intensities conditional on landing is assumed to be iid, F is independent of t, and only depends on the time to maturity of the contract, and thus on the combined distribution of H at the end of the period. We report the bounds for a range of value of g, which cover the values observed in past financial instruments indexed on wind damages.[24]

After the lattice is constructed and calibrated, as described in the previous sections, we calculate the bounds recursively starting at time $T-1$. Figure 3.2 shows the recursive process for the underlying lattice depicted in Figure 3.1. As shown in the lattice we start from the maturity of the contract, for which the conditional payoff is known to be c. Then by constructing the convex hull of the conditional payoffs as discussed in the previous sections, we calculate the bounds at time $T-1$. These bounds are then used to construct the convex hull and obtain the bounds at time $T-2$. This recursive estimation of the bounds is repeated until we get the bounds at $T-t$.

[Figure 5.2 about here]

Figure 3.3 and Table 3.1 show the multiperiod bounds, along with the Merton price, for a reinsurance contract with six months[25] left to maturity and with a deductible and ceiling of 1 and 10 billion dollars, respectively. The riskless return is set to $R = 1.01$ and we choose the parameters values for our base case as $g = 2$, $\lambda = 0.9$, and $\kappa = 0.41$. The two bounds and the Merton price increase as we increase the number of periods. This is because as we

[24]See Lane and Mahul (2008, Table 5). In fact the average coefficients for wind risk insurance range from 1.79 to 3.97, depending on time and place.

[25]We assume that the length of the hurricane season in Florida is six months.

increase N we also increase the number of potential hurricane landings during the term of the reinsurance contract. Moreover, as the number of subdivisions increases, the bounds become tighter. Furthermore, we can observe that the Merton price always lies below the lower bound, and thus the assumption of a diversifiable CAT risk, results in serious underpricing of the reinsurance contact whenever $g > 1$. Last but not least, we note that the value of the parameter g, the excess premium on the average intensity of the CAT event reflected in the CAT futures price, is not sufficient to determine on its own the value of the reinsurance contract as a function of its Merton bound. For instance, for $N = 2$ the midpoint of the bounds is 3.53, about 9% higher than the value of 3.24, the Merton value multiplied by g, which understates even the lower bound of 3.47.

[Table 5.1 about here]

[Figure 5.3 about here]

The extent of underpricing of the contract as a function of for our base case parameters is shown in Figure 3.4 and Table 3.2. When $g = 1$, the futures price is proportional to expected hurricane intensity, corresponding to full diversifiability of the CAT event risk. In this case, the two bounds coincide with the Merton price, which is the discounted expected payoff under the physical distribution of the underlying process. As noted in the introduction, such an assumption is at variance with the observed facts, as well as with recent theoretical models that show the regional specialization of insurance firms that are the purchasers of the reinsurance contract. As we increase g, thus increasing the premium of the future price over the expected hurricane intensity, the two bounds increase and become wider, indicating the market incompleteness with respect to the CAT event risk, while the Merton price remains the same. Further, the relation of the reinsurance value as a function of g is not linear: as Figure 3.4 and Table 3.2 show, for the base case $N = 8$ the bounds midpoint is initially approximately proportional to g, but for higher values of g, multiplying g by the Merton value overstates significantly the value of the reinsurance contract. This result, as well as the previously noted understating of the reinsurance contract value when $N = 2$ and $g = 2$, validate our approach in valuing such contracts by recognizing market incompleteness and the physical attributes of the CAT event.

[Table 5.2 about here]

[Figure 5.4 about here]

Figure 3.5 and Table 3.3 evaluate the bounds for different numbers of recursions. All the bounds are calculated for the same set of parameters as above, and for $N = 8$, implying that there can be up to 8 hurricane landings in the time left to the maturity of the reinsurance contract in all cases. We calculate the bounds when there are 1, 2, 4, and 8 recursions (trading opportunities) available during the life of the reinsurance contract, where these trading opportunities are spaced with equal length from each other. One recursion corresponds to the single-period bounds, with the remaining entries corresponding to the indicated number of recursions. It is clear that increasing the number of recursions results in progressively tighter bounds for a given number of potential hurricane landings.

[Table 5.3 about here]

[Figure 5.5 about here]

The comparison of single-period and multiperiod bounds for a given number of potential landings is explored further in Table 3.4 and Figure 3.6. We calculate the single-period bounds with the same underlying loss process as for the multi-period model and for the indicated number of potential landings. We calculate the terminal accumulated intensities and the associated probabilities and aggregate the states by summing over the probabilities of the states with the same intensity. Hence, we end up with the terminal distribution of landing and intensities that is the equivalent to the multiperiod convolution of (q_i, H_i). We denote this distribution $({q'}_i, {H'}_i)$ and set the price of the seasonal futures contract to $F' = g \cdot E[H']$. It is clear that the multi-period bounds are much tighter than their single-period counterparts. This indicates as expected, that recursive estimation reduces the gap between the two bounds. This gap would converge to a minimum at the continuous time limit if such a convergence were not precluded by the physical restrictions of the hurricane development process. Even if we ignore these restrictions, the lack of closed form expressions for the bounds precludes the estimation of their limiting values.

[Table 5.4 about here]

[Figure 5.6 about here]

3.6 Summary and Conclusions

In this paper we have presented an approach to the pricing of CAT derivatives with non-convex payoffs that is significantly different from the established methodology in earlier studies. Almost all these studies follow the Merton (1976) assumption that the rare event risk is fully diversifiable and should not be priced in equilibrium. Alternatively, they assume that a unique risk-adjusted distribution of the CAT event can be extracted from other traded financial instruments indexed on the event distribution. In such a case a unique CAT derivative price can be obtained in all cases and non-convexity is not an issue, since the unique price can be replicated by the prices of derivatives whose payoff replicates that of the valued derivative. Our approach relies on recent literature suggesting that economic agents trading in CAT instruments (for instance, insurance companies) specialize locally and in special types of CAT risks, and is consistent with empirical evidence that rejects decisively the Merton assumption. Hence, the CAT event risk is in general not fully diversifiable. In such a case our approach recognizes the market incompleteness introduced by the non-diversifiable CAT event and relies on stochastic dominance arguments to develop bounds on the CAT event derivatives relying solely on traded futures contracts on the CAT event. We show that our bounds cannot be derived by replicating the derivative with plain vanilla call and put options, and we present an efficient discrete time algorithm that can handle several empirically meaningful CAT event derivatives. We also argue that the continuous time approach is not relevant to financial instruments indexed on CAT events arising from physical phenomena. We apply our method to the pricing of a catastrophe reinsurance contract on the underlying cumulative hurricane losses in the state of Florida. We assume that the reinsurance contract has a deductible and a ceiling, and so is in the form of a vertical spread. Our theoretical analysis predicts that the call option price bounds would lie above the Merton price, with the distance depending on the price of a hurricane futures contract such as the ones offered by the CME or trading over the counter. We use realistic parameter values and show that the reinsurance contract produces tight bounds for all admissible values of the parameters. We show that the Merton price lies far below our bounds for almost all realistic values of the hurricane futures contract parameter, expressed as a multiple of the expected intensity of the CAT event. We also show that the dependence of our bounds on this parameter is non-linear and varies with contract characteristics, thus illustrating the pitfalls of neglecting or minimizing the importance of the CAT event systematic risk. Last but not least, our method is applicable with minimal adaptation to other important CAT financial instruments such as bonds indexed to CAT events.

Appendix

3.A Proof of Lemma 1

For the first part, we note that for $H_i > H_h$the unconditional payoff C_{Ti} is concave over the set of points H_i, while the points (H_i,C_{Ti}) for $H_i \leq H_h$ play no role in the computation of the points $(\overline{H}_i,\overline{C}_{Ti})$. Hence, concavity is preserved, since the points of the latter points are linear combinations of the former ones. For the points $(\overline{H}_i,\overline{C}_{Ti})$ in the domain $\overline{H}_i \leq \overline{H}_l$ we consider the slope $\frac{\Delta\overline{C}_{Ti}}{\Delta\overline{H}_i} = \frac{\overline{C}_{Ti}-\overline{C}_{Ti-1}}{\overline{H}_i-\overline{H}_{i-1}}$ where the conditional payoffs and intensities are replaced from (3.2.6). Simplifying, we find that this quantity is a fraction with the same numerator but with a denominator that decreases as a function of i in the region $\overline{H}_i \leq \overline{H}_l$, implying the convexity of the function. □

3.B Proof of Proposition 1

The following relations characterize the lower and upper boundaries of the convex hull of the conditional contract payoff as presented in Figure 3.2.1.

$$
\begin{aligned}
C_{i,upper} &= \frac{\overline{C}_{T\bar{i}} - \overline{C}_{T0}}{\overline{H}_{\bar{i}} - \overline{H}_0} \cdot (\overline{H}_i - \overline{H}_0) + \overline{C}_{T0},\ i \leq \bar{i} \\
C_{i,upper} &= \overline{C}_{Ti},\ i \in (\bar{i}, h) \\
C_{i,upper} &= \kappa(H_h - H_l),\ i \geq h \\
C_{i,lower} &= \overline{C}_{Ti},\ i \leq \underline{i} \\
C_{i,lower} &= \overline{C}_{T\underline{i}} + \frac{\kappa(H_h - H_l) - \overline{C}_{T\underline{i}}}{H_n - \overline{H}_{\underline{i}}} \cdot (\overline{H}_i - \overline{H}_{\underline{i}}),\ i > \underline{i}
\end{aligned}
\tag{3.B.1}
$$

where, l and h represent states corresponding to the hurricane intensities H_l and H_h, respectively. Moreover, $\underline{i}$ and $\bar{i}$ are the tangency points of the non-convex payoff on the lower and upper boundaries of the convex hull, respectively, and can be found as follows:

$$
\bar{i} = \arg\max_i \left(\frac{\overline{C}_{Ti} - \overline{C}_{T0}}{\overline{H}_i - \overline{H}_0}\right)\ ,\ \underline{i} = \arg\max_i \left(\frac{\kappa(H_h - H_l) - \overline{C}_{Ti}}{H_n - \overline{H}_i}\right)
\tag{3.B.2}
$$

The bounds on the contingent claim can be found based on the value of F relative to the deductible and ceiling of the payoff, as follows:

For $F \leq \overline{H}_h$ we have:

$$
\begin{aligned}
&C_{max} = R^{-1} C_{iF,upper},\ C_{min} = R^{-1}[\overline{C}_{T\underline{i}} + \frac{\kappa(H_h - H_l) - \overline{C}_{T\underline{i}}}{H_n - \overline{H}_{\underline{i}}}(F - \overline{H}_{\underline{i}})] \\
&C_{iF,upper} \equiv \frac{\overline{H}_{i^*+1} - F}{\overline{H}_{i^*+1} - \overline{H}_{i^*}} C_{i^*,upper} + \frac{F - \overline{H}_{i^*}}{\overline{H}_{i^*+1} - \overline{H}_{i^*}} C_{i^*+1,upper}, \\
&i^* : \overline{H}_{i^*} \leq F \leq \overline{H}_{i^*+1}
\end{aligned}
\tag{3.B.3}
$$

For $F \in (\overline{H}_h, H_n)$the lower bound remains unchanged while the upper bound becomes $C_{max} = R^{-1}\kappa(H_h - H_l)$. Similarly, for $F \in (\overline{H}_{\underline{i}}, \overline{H}_{\bar{i}})$ the lower bound remains unchanged while the upper bound is still given by (3.B.3), but now $C_{i,upper} = \frac{\overline{C}_{T\bar{i}} - \overline{C}_{T0}}{\overline{H}_{\bar{i}} - \overline{H}_0} \cdot (\overline{H}_i - \overline{H}_0) + \overline{C}_{T0}$. Last, for $F \in (\overline{H}_0, \overline{H}_{\underline{i}})$ the upper bound is the same as for $F \in (\overline{H}_{\underline{i}}, \overline{H}_{\bar{i}})$, but for the lower bound we now have $C_{min} = R^{-1}\overline{C}_{iF,lower}$ and $C_{iF,lower} \equiv \frac{\overline{H}_{i^*+1} - F}{\overline{H}_{i^*+1} - \overline{H}_{i^*}} \cdot C_{Ti^*} + \frac{F - \overline{H}_{i^*}}{\overline{H}_{i^*+1} - \overline{H}_{i^*}} \cdot C_{Ti^*+1}$, where i^* is chosen such that $\overline{H}_{i^*} \leq F \leq \overline{H}_{i^*+1}$.

Proof: By conditions (3.2.7)-(3.2.8) all admissible points C belong to the convex hull of the points $\overline{C}_{Ti}$, the smallest convex space containing all such points. The boundaries of the convex hull is defined by equations (3.B.1) and (3.B.2), which are illustrated in Figure refconvexhull.

Similarly, all admissible sets $\{\tilde{X}_i\}$ satisfying (3.2.7) lie within the convex hull and on the vertical line shown, which satisfies the constraint $\sum_0^n \tilde{X}_i \overline{H}_i = F$. Hence, the upper and lower bounds of C are found from the intersection of the line with the upper and lower boundaries of the hull, $C_{i,upper}(\overline{H}_i)$ and $C_{i,lower}(\overline{H}_i)$ respectively. These intersections, however, yield (3.B.3), as well as the alternative cases that arise from the size of the futures price F as described in the Proposition. □

3.C Tightness of the Bounds

Let C denote the reinsurance contract on the accumulated hurricane losses with a deductible and a ceiling equal to H_l and H_h, respectively. The payoff to the reinsurance contract has the shape of a spread and can be replicated by two call options on the same underlying process. Consider a long position in a call option, C_1, with strike price equal to H_l, and a short position in a call option, C_2, with strike price H_h. The payoff to the combined position, $c_1 - c_2$, replicates the payoff to the reinsurance contract, c.

The upper bound of C_1 can be found as the solution to the following LP, subject to the market equilibrium conditions, as described in the single-period model.

$$C_{1,\max} = Max_{\{\widehat{X}_i\}} \sum q_i \widehat{X}_i c_{1,i} = \sum q_i X^*_{1,i} c_{1,i} \tag{3.C.1}$$

Similarly, the lower bound of C_2 is the following:

$$C_{1,\min} = Min_{\{\widehat{X}_i\}} \sum q_i \widehat{X}_i c_{2,i} = \sum q_i X^*_{2,i} c_{2,i} \tag{3.C.2}$$

where, X_1^* and X_2^* are the values of the pricing kernel corresponding to the solution of the maximization and minimization, respectively. Thus, the upper bound of the reinsurance contract, obtained from the replicating portfolio, would be equal to $C_{2,\max} - C_{1,\min}$.

On the other hand, the upper bound of the reinsurance contract can be obtained as the solution to the following LP, subject to the market equilibrium conditions:

$$\begin{aligned} C_{\max} &= Max_{\{\hat{X}_i\}} \sum q_i \hat{X}_i c_i = \sum q_i X_i^* c_i = \sum q_i X_i^* (c_{1,i} - c_{2,i}) \\ &= \sum q_i X_i^* c_{1,i} - \sum q_i X_i^* c_{2,i} < C_{1,\max} - C_{2,\min} \end{aligned} \tag{3.C.3}$$

Where, the last inequality follows from the fact that X^* is the solution to LP in (3.C.3) and produces suboptimal results compated to X_1^* and X_2^* as in (3.C.1) and (3.C.2). Similarly, it can be shown that the lower bound derived from the convex hull, $C_{\min}$, is greater than the lower bound obtained from the replicating portfolio, $C_{1,\min} - C_{2,\max}$. □

We compare the tightness of the bounds obtained from the convex hull to those obtained from the replicating portfolio in the following numerical analysis. As in Section 4, we consider a reinsurance contract on the accumulated hurricane losses in the state of Florida, with a six month time to maturity and with a deductible and a ceiling equal to 2 and 5 billion dollars, respectively. The combined distribution of the hurricane landing and intensities are derived from the actual hurricane landing data in the state of Florida, as discussed in Section 5. We also assume that there is a futures contract on the same underlying loss process. The futures price is set equal to a multiple of the expected hurricane loss: $F = g.E[H]$. Table 3.C.1 shows the percentage difference between the upper and lower bounds obtained from the convex hull and those found from the replicating portfolio.

As evident from the table, the bounds obtained from the convex hull method are considerably tighter than those obtained from the replicating portfolio, especially for the more realistic midrange values of g.

Table 3.C.1: **Tightness of the Bounds**

g	$((C_{1,\max} - C_{2,\min}) - C_{\max})/C_{\max}$	$(C_{\min} - (C_{1,\min} - C_{2,\max}))/C_{\min}$
1.1	3.1	5.5
1.3	7.6	13.9
1.5	10.6	20.1
1.7	12.8	24.6
1.9	14.1	25.4
2.1	14.7	23.9
2.3	13.4	17.7
2.5	10.9	12.6
2.7	7.3	8.2
2.9	4.1	4.4
3.0	2.7	2.8

3.D Proof of Lemma 2

The boundary of the convex hull of the conditional payoff one period prior to the contract expiration can be characterized by the following relations.[26]

a) For $V_{T-1} < \overline{V}_l$:

$$\begin{aligned}
&C_{T-1,upper}(V_{T-1}, \overline{H}_i) = R^{-1}[\frac{\overline{C}_{T\bar{i}} - \overline{C}_{T0} - V_{T-1}}{\overline{H}_{\bar{i}} - \overline{H}_0}(\overline{H}_i - \overline{H}_0) + V_{T-1} + \overline{C}_{T0}],\ i \leq \bar{i};\\
&C_{T-1,upper}(V_{T-1}, \overline{H}_i) = R^{-1}[V_{T-1} + \overline{C}_{Ti}],\ i \in (\bar{i}, h);\\
&C_{T-1,upper}(V_{T-1}, \overline{H}_i) = R^{-1}[\overline{V}_h - \overline{V}_l],\ i \geq h\\
&C_{T-1,lower}(V_{T-1}, \overline{H}_i) = R^{-1}(V_{T-1} + \overline{C}_{Ti}),\ i \leq \underline{i};\\
&C_{T-1,lower}(V_{T-1}, \overline{H}_i) = R^{-1}[V_{T-1} + \overline{C}_{T\underline{i}} + \frac{\overline{V}_h - \overline{V}_l - \overline{C}_{T\underline{i}} - V_{T-1}}{H_n - \overline{H}_{\underline{i}}}(\overline{H}_i - \overline{H}_{\underline{i}})],\ i > \underline{i}.
\end{aligned} \tag{3.D.1a}$$

b) For $\overline{V}_l \leq V_{T-1} < \overline{V}_h$ relations (3.3.2) yield now a graph $\overline{C}_{Ti}(\overline{H}_i)$ that is concave. The convex hull now becomes

$$\begin{aligned}
&C_{T-1,i,upper}(V_{T-1}, \overline{H}_i) = R^{-1}(V_{T-1} + \overline{C}_{Ti}) = R^{-1}[V_{T-1} + \kappa\overline{H}_i - \overline{V}_l],\ i \leq h;\\
&C_{T-1,i,upper}(V_{T-1}, \overline{H}_i) = R^{-1}[\overline{V}_h - \overline{V}_l],\ i \geq h;\\
&C_{T-1,i,lower}(V_{T-1}, \overline{H}_i) = R^{-1}[(\overline{H}_i - \overline{H}_0)\frac{\overline{V}_h - \overline{V}_l - V_{T-1} - \overline{C}_{T0}}{H_n - \overline{H}_0} + V_{T-1} + \overline{C}_{T0}].
\end{aligned} \tag{3.D.1b}$$

c) For $V_{T-1} \geq \overline{V}_h$ the convex hull is clearly the line

$$C_{T-1,i,upper}(V_{T-1}, \overline{H}_i) = C_{T-1,i,lower}(V_{T-1}, \overline{H}_i) = R^{-1}[\overline{V}_h - \overline{V}_l],\ \forall i. \tag{3.D.1c}$$

[26] Note that in (3.D.1) the values $\bar{i}$ and $\underline{i}$ are non-increasing functions $\bar{i}(V_{T-1})$ and $\underline{i}(V_{T-1})$.

Proof: Conditional on V_{T-1}, $C_{T-1}(V_{T-1}|F_{T-1})$ satisfies the market equilibrium conditions (3.2.4)-(3.2.5). Applying the transformations (3.2.6) to its payoff given by (3.3.2), we note that the admissible values of $C_{T-1}(V_{T-1}, H_T)$ satisfy (3.2.7)-(3.2.8), implying that they must lie within the convex hull of the points $\overline{C}_{Ti}$, whose form is now conditional on V_{T-1} as in (3.3.2). For case (a) the function $\overline{C}_{Ti}(\overline{H}_i)$ has the same shape as in Lemma 1, with the convex hull being, therefore, given by(3.D.1a). For the case (b), $\overline{C}_{Ti}(\overline{H}_i)$ is concave, while for (c) it is constant, thus proving the remaining part of the Lemma. □

3.E Proof of Lemma 3

Given the value of the state variable V_{T-1}, and value of the futures price F_{T-1}, the two bounds can be found as the intersection of the vertical line stemming from F_{T-1} with the convex hull, whose boundaries are presented in Lemma 2:

$$
\begin{aligned}
&C_{T-1,max}(V_{T-1}|F_{T-1}) = R^{-1}C_{iF,upper},\ V_{T-1} \le \overline{V}_{T-1},\\
&C_{T-1,max}(V_{T-1}|F_{T-1}) = R^{-1}(\overline{V}_h - \overline{V}_l),\ V_{T-1} > \overline{V}_{T-1}\\
&C_{T-1,min}(V_{T-1}|F_{T-1}) = R^{-1}C_{iF,lower},\ V_{T-1} \le \underline{V}^1_{T-1},\\
&C_{T-1,min}(V_{T-1}|F_{T-1}) = R^{-1}C_{iF,lower},\ \underline{V}^1_{T-1} < V_{T-1} < \underline{V}^2_{T-1},\\
&C_{T-1,min}(V_{T-1}|F_{T-1}) = R^{-1}(\overline{V}_h - \overline{V}_l),\ V_{T-1} > \underline{V}^2_{T-1}\\
&C_{iF,upper} \equiv \frac{\overline{H}_{i^*+1} - F_{T-1}}{\overline{H}_{i^*+1} - \overline{H}_{i^*}} C_{T-1,i^*,upper} + \frac{F_{T-1} - \overline{H}_{i^*}}{\overline{H}_{i^*+1} - \overline{H}_{i^*}} C_{T-1,i^*+1,upper},\\
&C_{iF,lower} \equiv \frac{\overline{H}_{i^*+1} - F_{T-1}}{\overline{H}_{i^*+1} - \overline{H}_{i^*}} C_{T-1,i^*,lower} + \frac{F_{T-1} - \overline{H}_{i^*}}{\overline{H}_{i^*+1} - \overline{H}_{i^*}} C_{T-1,i^*+1,lower},\\
&i^* : \overline{H}_{i^*} \le F \le \overline{H}_{i^*+1}
\end{aligned}
\tag{3.E.1}
$$

In the above expressions $C_{T-1,i,upper}$ and $C_{T-1,i,lower}$ are as defined in (3.D.1). Moreover, for a given F_{T-1} the boundary values $\overline{V}_{T-1}$, $\underline{V}^1_{T-1}$ and $\underline{V}^2_{T-1}$, that determine the break points of the concave functions $C_{T-1,max}(V_{T-1}|F_{T-1})$ and $C_{T-1,min}(V_{T-1}|F_{T-1})$, are defined by the values of V_{T-1} at which the intersection of the futures price with the convex hull given by (3.D.1) switches between two successive portions of the piecewise linear (convex) or concave (linear) functions of the upper (lower) boundaries of the convex hull.

Proof: The proof is similar to Proposition 1. As V_{T-1} increases the upper bound is initially at some $i \leq \bar{i}$, then at $i \in (\bar{i}, h)$ and eventually at an $i > \bar{i}$; this proves the two first equations of (3.E.1). The concavity of $C_{T-1,max}(V_{T-1}|F_{T-1})$ follows directly from the concavity of the upper boundary of the convex hull for $i > \bar{i}$. Similarly, for $C_{T-1,min}(V_{T-1}|F_{T-1})$ the properties of the bound with respect to V_{T-1} stem from the definition of $\underline{i}(V_{T-1})$, with $\underline{i}(V^1_{T-1}) = 0$ and $\underline{i}(0) > \bar{i}$, which imply that the lower bound lies also at some $i \in (\underline{i}, h)$, thus yielding the piecewise linearity and concavity of $C_{T-1,i,lower}(V_{T-1}, \overline{H}_i)$ with respect to V_{T-1}. □

3.F Proof of Proposition 2

We use induction and we rely on the concavity of the bounds.[27] At $T-1$ the proposition holds by Lemma 3. Assume that it holds at $t+1$. Then the bounds on $C_t(V_t|F_t)$ are found by the following program:

$$\begin{aligned} &\max(\min)_{\hat{X}_i} \ \{R^{-1} \textstyle\sum_0^n q_i \hat{X}_i C_{t+1}(V_{t+1}|F)\} = \\ &\max(\min)_{\hat{X}_i} \ \{R^{-1} \textstyle\sum_0^n q_i \hat{X}_i C_{t+1}(V_t + \kappa H_i|F)\} \end{aligned} \tag{3.F.1}$$

subject to the constraints in (3.3.3). Applying the transformation (3.2.6) and (3.3.4) to this program, we find that it is reduced to $\max(\min)_{\tilde{X}_i}\{R^{-1}\sum_0^n \tilde{X}_i \overline{C}_{t+1,i}\}$ subject to (3.2.7). Since by the induction hypothesis $C_{t+1,min}(V_{t+1}|F) \leq C_{t+1}(V_{t+1}|F) \leq C_{t+1,max}(V_{t+1}|F)$, the following programs yield upper and lower bounds on $C_t(V_t|F)$, subject to (3.2.7):

$$\max_{\tilde{X}_i}\{R^{-1}\sum_0^n \tilde{X}_i \overline{C}_{t+1,max,i}\}\ ,\ \min_{\tilde{X}_i}\{R^{-1}\sum_0^n \tilde{X}_i \overline{C}_{t+1,min,i}\} \tag{3.F.2}$$

Since both bounds at $t+1$ are increasing and concave functions and the transformation (3.2.6)-(3.3.4) preserves the concavity of the functions $\overline{C}_{t+1,\alpha,i}$, $\alpha = \max, \min$, $i = 0, \ldots, n$ with respect to $\tilde{X}_i$, $i = 0, \ldots, n$, the results of the programs (3.F.2) yield a value $C_{t,max}(V_t|F)$

[27] If the bounds are neither convex nor concave we use the convexification described in Section 2 to derive the convex hull and then proceed as in proposition 1

at the intersection of the futures price with the upper boundary of the convex hull of the points $(\overline{C}_{t+1,max,i}, \overline{H}_i)$ as in (3.E.1). The value $C_{t,min}(V_t|F)$, on the other hand, would lie on the lower boundary of the convex hull of the points $(\overline{C}_{t+1,min,i}, \overline{H}_i)$, which is the straight line connecting the initial and final points, as in (3.E.1). Last, the concavity of both bounds follows from the fact that as V_t increases both bounds trace the shape of the convex hull in its concave and linear parts respectively, which eventually degenerates into a straight line. $\square$

Table 3.1: **The Bounds and the Number of Periods**

N	Lower Bound	Upper Bound	Merton Price
1	2.60	2.75	1.22
2	3.47	3.59	1.62
3	3.70	3.81	1.77
4	3.80	3.91	1.85
5	3.86	3.97	1.89
6	3.89	4.00	1.93
7	3.92	4.02	1.95
8	3.93	4.04	1.97

This table shows the multi-period bounds and the Merton price for different numbers of periods. The bounds are calculated for a reinsurance contract with six months left to maturity, and with a deductible and ceiling equal to 1 and 10 billion dollars, respectively. The parameters are $R = 1.01$, $\lambda = 0.9$, $g = 2$, and $\kappa = 0.41$.

Table 3.2: **Effect of CAT Event Risk Premium on the Bounds**

g	Lower Bound	Upper Bound	Merton Price	$g \cdot$ Merton Price
1	1.97	1.97	1.97	1.97
1.2	2.36	2.38	1.97	2.36
1.5	2.96	3.02	1.97	2.96
1.7	3.36	3.44	1.97	3.35
2	3.93	4.04	1.97	3.94
2.3	4.48	4.61	1.97	4.53
2.5	4.82	4.97	1.97	4.93
2.7	5.15	5.32	1.97	5.32
3	5.61	5.81	1.97	5.91
3.5	6.29	6.54	1.97	6.90
4	6.86	7.16	1.97	7.89

This table shows the multi-period bounds and the Merton price for different values of g. The bounds are calculated for a reinsurance contract with six months left to maturity, and with a deductible and ceiling of 1 and 10 billion dollars respectively. The parameters are $N = 8$, $R = 1.01$, $\lambda = 0.9$, and $\kappa = 0.41$.

Table 3.3: **Effect of the Number of Recursions on the Tightness of the Bounds**

Number of Recursions	Lower Bound	Upper Bound
1	2.46	4.21
2	2.94	4.19
4	3.64	4.10
8	3.93	4.04

This table shows the values of the two bounds for different number of recursions. The bounds are calculated for reinsurance contract with six monthsleft to maturity, and with a deductible and ceiling of 1 and 10 billion dollars respectively. The parameters are $N = 8$, $R = 1.01$, $\lambda = 0.9$, $g = 2$, and $\kappa = 0.41$.

Table 3.4: **Multi-Period and Single-Period Bounds**

	Single-Period		Multi-Period	
N	Lower Bound	Upper Bound	Lower Bound	Upper Bound
1	2.60	2.75	2.60	2.75
2	3.47	3.59	3.39	3.67
3	3.70	3.81	3.20	3.90
4	3.80	3.91	2.89	4.02
5	3.86	3.97	2.72	4.09
6	3.89	4.00	2.60	4.15
7	3.92	4.02	2.52	4.19
8	3.93	4.04	2.46	4.21

This table compares the multi-period and single-period bounds for different numbers of periods. The bounds are calculated for a reinsurance contract with six months left to maturity, and with a deductible and ceiling equal to 1 and 10 billion dollars, respectively. The parameters are $R = 1.01$, $\lambda = 0.9$, $g = 2$, and $\kappa = 0.41$.

Figure 3.1: **Underlying Loss Process**

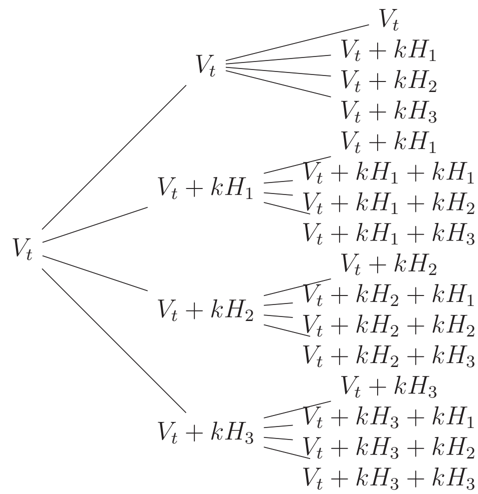

This figure shows the discrete time underlying loss process for two periods. V is the accumulated loss, H is the intensity of the landed hurricane, and κ is the multiplier translating the hurricane intensity into dollar losses.

Figure 3.2: **Recursive Evolution of the Bounds**

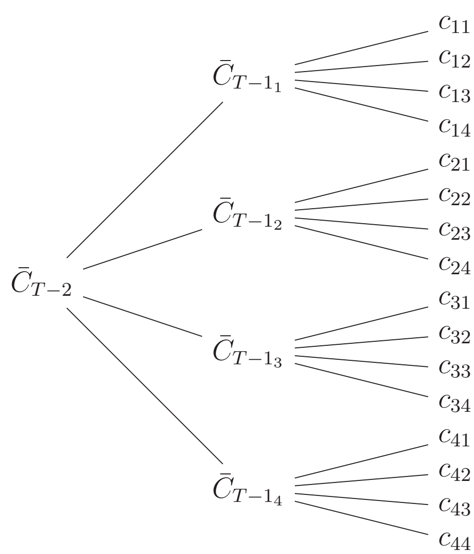

The multi-period lattice showing the recursive process for the upper bound starting from the maturity of the contract.

Figure 3.3: **The Bounds and the Number of Periods**

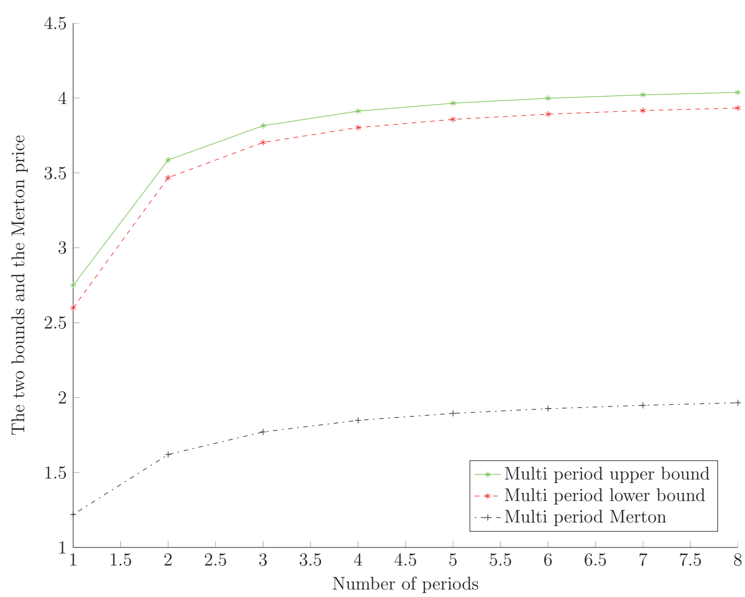

Multi-period bounds and the Merton's price for different number of periods. The bounds are calculated for a reinsurance contract with six months left to maturity, and with a deductible and ceiling equal to 1 and 10 billion dollars, respectively. The parameters are $R = 1.01$, $\lambda = 0.9$, $g = 2$, and $\kappa = 0.41$.

Figure 3.4: **Effect of CAT Event Risk Premium on the Bounds**

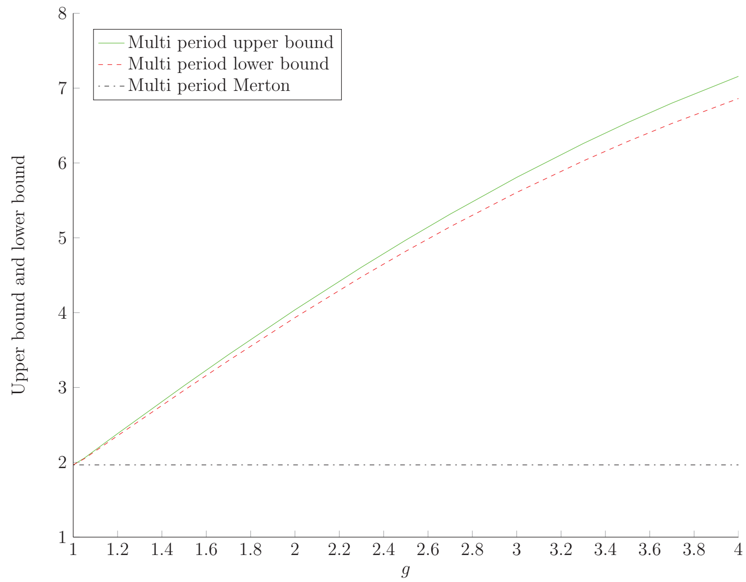

Multi-period bounds and Mertons price for different values of g in the relation $F = g.E[H]$. The bounds are calculated for a reinsurance contract with six months left to maturity, and with a deductible and ceiling of 1 and 10 billion dollars respectively. The parameters are $N = 8$, $R = 1.01$, $\lambda = 0.9$, and $\kappa = 0.41$.

Figure 3.5: **Effect of the Number of Recursions on the Tightness of the Bounds**

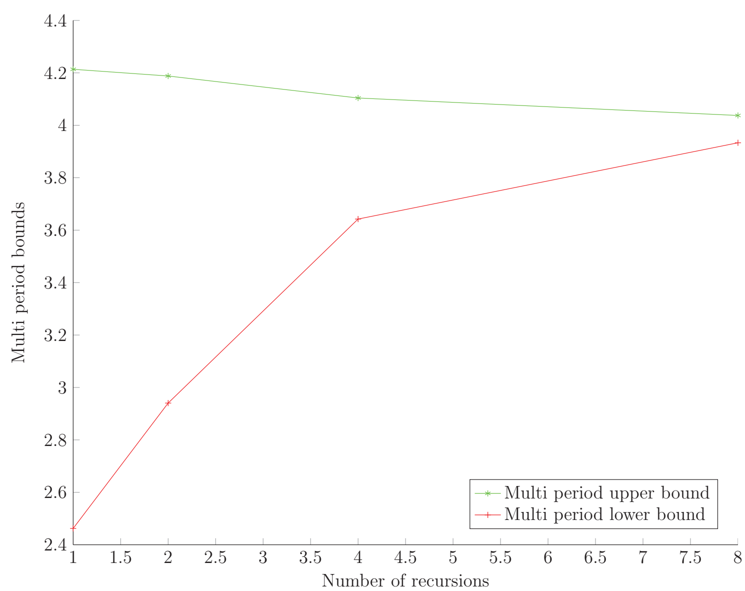

The two bounds for different numbers of recursions. The bounds are calculated for a reinsurance contract with six months left to maturity, and with a deductible and ceiling of 1 and 10 billion dollars respectively. The parameters are $N = 8$, $R = 1.01$, $\lambda = 0.9$, $g = 2$, and $\kappa = 0.41$.

Figure 3.6: **Multi-Period and Single-Period Bounds**

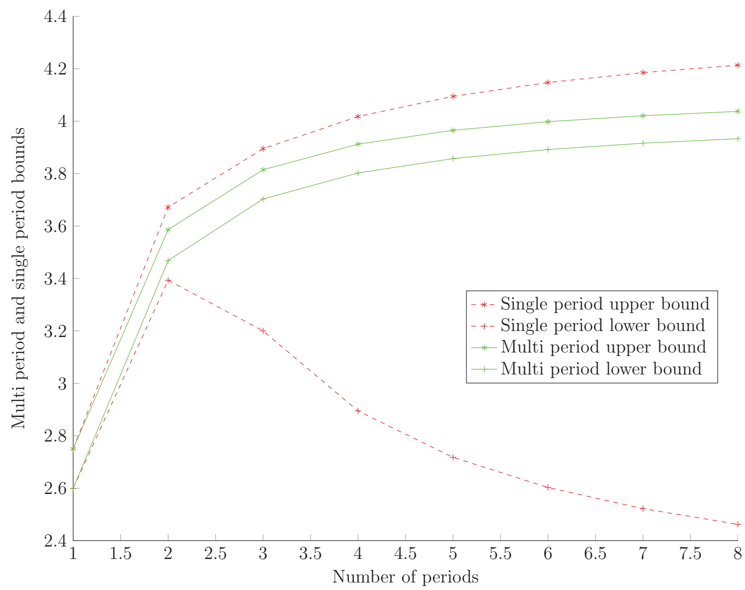

The figure compares the multi-period and single-period bounds for different numbers of periods. The bounds are calculated for a reinsurance contract with six months left to maturity, and with a deductible and ceiling equal to 1 and 5 billion dollars, respectively. The parameters are $R = 1.01$, $\lambda = 0.9$, $g = 2$, and $\kappa = 0.41$.

Bibliography

Ait-Sahalia, Y. and R. Kimmel (2007). Maximum likelihood estimation for stochastic volatility models. *Journal of Financial Economics 83*(413).

Amin, K. (1993). Jump diffusion option valuation in discrete time. *Journal of Finance 48*(5), 1833–63.

Amin, K. and V. K. Ng (1993). Option valuation with systematic stochastic volatility. *Journal of Finance 48*(3), 881–910.

Ang, A., R. J. Hodrick, Y. Xing, and X. Zhang (2006). The cross-section of volatility and expected returns. *Journal of Finance*, 259–299.

Bakshi, G., C. Cao, and Z. Chen (1997). Empirical performance of alternative option pricing models. *Journal of Finance 52*(5), 2003–49.

Bakshi, G., N. Kapadia, and D. Madan (2003). Stock return characteristics, skew laws, and the differential pricing of individual equity options. *Review of Financial Studies 16*(1), 101–143.

Bakshi, G. and D. Madan (2002, March). Average rate claims with emphasis on catastrophe loss options. *Journal of Financial and Quantitative Analysis 37*(01), 93–115.

Bali, T. G. and N. Cakici (2008, 3). Idiosyncratic volatility and the cross section of expected returns. *Journal of Financial and Quantitative Analysis 43*, 29–58.

Barrieu, P. and H. Loubergé (2009). Hybrid cat bonds. *Journal of Risk & Insurance 76*(3), 547–578.

Bates, D. S. (1991, July). The crash of '87: Was it expected? the evidence from options markets. *Journal of Finance 46*(3), 1009–44.

Bates, D. S. (2000). Post-'87 crash fears in the s&p 500 futures option market. *Journal of Econometrics 94*(1-2), 181–238.

Bates, D. S. (2006). Maximum likelihood estimation of latent affine processes. *Review of Financial Studies 19*(3), 909–965.

Bernard, C. and W. Tian (2009). Optimal reinsurance arrangements under tail risk measures. *Journal of Risk & Insurance 76*(3), 709–725.

Buss, A. and G. Vilkov (2012). Measuring equity risk with option-implied correlations. *Review of Financial Studies 25*(10), 3113–3140.

Carr, P. and L. Wu (2009). Variance risk premiums. *Review of Financial Studies 22*(3), 1311–1341.

Chang, B.-Y., P. F. Christoffersen, K. Jacobs, and G. Vainberg (2011). Option-implied measures of equity risk. *Review of Finance 16*(2), 385–428.

Chang, C. W., J. S. Chang, and W. Lu (2010). Pricing catastrophe options with stochastic arrival intensity in claim time. *Journal of Banking & Finance 34*(3), 695–695.

Chen, Z. and R. Petkova (2012). Does idiosyncratic volatility proxy for risk exposure? *Review of Financial Studies 25*(9), 2745–2787.

Chernov, M. and E. Ghysels (2000, June). A study towards a unified approach to the joint estimation of objective and risk neutral measures for the purpose of options valuation. *Journal of Financial Economics 56*(3), 407–458.

Christensen, C. V. and H. Schmidli (2000). Pricing catastrophe insurance products based on actually reported claims. *Insurance: Mathematics and Economics 27*(2), 189–200.

Christoffersen, P., M. Fournier, and K. Jacobs (2013). The factor structure in equity options. *Working Paper*.

Christoffersen, P. F., S. Heston, and K. Jacobs (2013). Capturing option anomalies with a variance-dependent pricing kernel. *Review of Financial Studies 26*(8), 1963–2006.

Christoffersen, P. F., K. Jacobs, and K. Mimouni (2010). Volatility dynamics for the s&p500: Evidence from realized volatility, daily returns, and option prices. *Review of Financial Studies 23*(8), 3141–3189.

Constantinides, G. M., M. Czerwonko, J. C. Jackwerth, and S. Perrakis (2011). Are options on index futures profitable for risk averse investors? empirical evidence. *Journal of Finance forthcoming*.

Constantinides, G. M., J. C. Jackwerth, and S. Perrakis (2009, March). Mispricing of s&p 500 index options. *Review of Financial Studies 22*(3), 1247–1277.

Constantinides, G. M. and S. Perrakis (2002, July). Stochastic dominance bounds on derivatives prices in a multiperiod economy with proportional transaction costs. *Journal of Economic Dynamics and Control 26*(7-8), 1323–1352.

Constantinides, G. M. and S. Perrakis (2007). Stochastic dominance bounds on american option prices in markets with frictions. *Review of Finance 11*, 71–115.

Dassios, A. and J.-W. Jang (2002). Pricing of catastrophe reinsurance and derivatives using the Cox process with shot noise intensity. *Finance and Stochastics 7*(1), 73–95.

Dennis, P. and S. Mayhew (2002). Risk-neutral skewness: Evidence from stock options. *Journal of Financial and Quantitative Analysis 37*(03), 471–493.

Driessen, J., P. J. Maenhout, and G. Vilkov (2009). The price of correlation risk: Evidence from equity options. *Journal of Finance 64*(3), 1377–1406.

Duan, J.-C. and J. Wei (2009). Systematic risk and the price structure of individual equity options. *Review of Financial Studies 22*(5), 1981–2006.

Duan, J.-C. and M.-T. Yu (2005). Fair insurance guaranty premia in the presence of risk-based capital regulations, stochastic interest rate and catastrophe risk. *Journal of Banking & Finance 29*(10), 2435–2454.

Duarte, J., A. Kamara, S. Siegel, and C. Sun (2012). Does idiosyncratic volatility proxy for risk exposure? *Working Paper*.

Duffie, D. and P. Protter (1992). From discrete- to continuous-time finance: Weak convergence of the financial gain process. *Mathematical Finance 2*(1), 1–15.

Eraker, B. (2004). Do stock prices and volatility jump? reconciling evidence from spot and option prices. *Journal of Finance 59*, 1367403.

Ethier, S. N. and T. G. Kurtz (1986). *Markov Processes*. Wiley, New York.

French, D. W., J. C. Groth, and J. W. Kolari (1983). Current investor expectations and better betas. *Journal of Portfolio Management 10*, 12–18.

Froot, K. A. (2001). The market for catastrophe risk: A clinical examination. Working Paper 8110, National Bureau of Economic Research.

Fu, F. (2009). Idiosyncratic risk and the cross-section of expected stock returns. *Journal of Financial Economics 91*(1), 24–37.

Geman, H. and M. Yor (1997). Stochastic time changes in catastrophe option pricing. *Insurance: Mathematics and Economics 21*(3), 185 – 193.

Gordon, N. J., D. J. Salmond, and A. F. M. Smith (1993, April). Novel approach to nonlinear/non-Gaussian Bayesian state estimation. *Radar and Signal Processing, IEE Proceedings F 140*(2), 107–113.

Han, Y. and D. Lesmond (2011). Liquidity biases and the pricing of cross-sectional idiosyncratic volatility. *Review of Financial Studies 24*(5), 1590–1629.

Herskovic, B., B. T. Kelly, H. N. Lustig, and S. Van Nieuwerburgh (2013). The common factor in idiosyncratic volatility. *Working Paper*.

Heston, S. L. (1993). A closed-form solution for options with stochastic volatility with applications to bond and currency options. *Review of Financial Studies 6*(2), 327–43.

Huang, J.-z. and L. Wu (2004). Specification analysis of option pricing models based on time-changed lvy processes. *Journal of Finance 59*(3), 1405–1440.

Huang, W., Q. Liu, S. G. Rhee, and L. Zhang (2010). Return reversals, idiosyncratic risk, and expected returns. *Review of Financial Studies 23*(1), 147–168.

Ibragimov, R., D. Jaffee, and J. Walden (2009). Nondiversification traps in catastrophe insurance markets. *Review of Financial Studies 22*(3), 959–993.

Jacod, J. and P. Protter (2002). *Probability Essentials*. Springer.

Jaimungal, S. and T. Wang (2006). Catastrophe options with stochastic interest rates and compound poisson losses. *Insurance: Mathematics and Economics 38*(3), 469–483.

Johannes, M. S., N. G. Polson, and J. R. Stroud (2009). Optimal filtering of jump diffusions: Extracting latent states from asset prices. *Review of Financial Studies 22*(7), 2559–2599.

Jones, C. S. (2003). The dynamics of stochastic volatility: evidence from underlying and options markets. *Journal of Econometrics 116*(1-2), 181–224.

Kocherlakota, N. (1996). The equity premium: It's still a puzzle. *Journal of Economic Literature 34*(1), 42–71.

Lane, M. and O. Mahul (2008). Catastrophe risk pricing: An empirical analysis. *World Bank, Washington, DC*.

Lee, J.-P. and M.-T. Yu (2007). Valuation of catastrophe reinsurance with catastrophe bonds. *Insurance: Mathematics and Economics 41*(2), 264–278.

Levy, H. (1978, September). Equilibrium in an imperfect market: A constraint on the number of securities in the portfolio. *American Economic Review 68*(4), 643–58.

Levy, H. (1985, September). Upper and lower bounds of put and call option value: Stochastic dominance approach. *Journal of Finance 40*(4), 1197–1217.

Lin, X. S. and T. Wang (2009). Pricing perpetual american catastrophe put options: A penalty function approach. *Insurance: Mathematics and Economics 44*(2), 287–295.

Merton, R. C. (1976). Option pricing when underlying stock returns are discontinuous. *Journal of Financial Economics 3*(1-2), 125–144.

Merton, R. C. (1987). A simple model of capital market equilibrium with incomplete information. *Journal of Finance 42*(3), 483–510.

MMC Securities (2005). The growing appetite for catastrophe risk. *Guy Carpenter and Company Inc.*.

Muermann, A. (2003). Actuarially consistent valuation of catastrophe derivatives. Working paper, The Wharton Financial Institutions Center.

Oancea, I. M. and S. Perrakis (2014). Jump-diffusion option valuation without a representative investor: a stochastic dominance approach. Working paper, Concordia University.

Pan, J. (2002). The jump-risk premia implicit in options: evidence from an integrated time-series study. *Journal of Financial Economics 63*(1), 3–50.

Perrakis, S. (1986, January). Option bounds in discrete time: Extensions and the pricing of the american put. *Journal of Business 59*(1), 119–41.

Perrakis, S. (1988). Preference-free option prices when the stock return can go up, go down, or stay the same. *Advances in Futures and Options Research 3*.

Perrakis, S. and A. Boloorforoosh (2013). Valuing catastrophe derivatives under limited diversification: A stochastic dominance approach. *Journal of Banking & Finance 37*(8), 3157–3168.

Perrakis, S. and P. J. Ryan (1984). Option pricing bounds in discrete time. *Journal of Finance 39*(2), 519–25.

Pitt, M. K. (2002). Smooth particle filters for likelihood evaluation and maximisation. The warwick economics research paper series (twerps), University of Warwick, Department of Economics.

Ritchken, P. H. (1985). On option pricing bounds. *Journal of Finance 40*(4), 1219–33.

Ritchken, P. H. and S. Kuo (1988). Option bounds with finite revision opportunities. *Journal of Finance 43*(2), 301–08.

Santa-Clara, P. and S. Yan (2010). Crashes, volatility, and the equity premium: Lessons from s&p 500 options. *The Review of Economics and Statistics 92*(2), 435–451.

Siegel, A. F. (1995). Measuring systematic risk using implicit beta. *Management Science 41*(1), 124–128.

Stroock, D. W. and S. S. Varadhan (1979). *Multidimensional Diffusion Processes*. Springer.

Xu, Y. and B. G. Malkiel (2003). Investigating the behavior of idiosyncratic volatility. *The Journal of Business 76*(4), 613–644.