Is the gamma risk of options insurable?

Egli, Daniel and Blum, Peter and Dacorogna, Michel M and Müller, Ulrich A

SCOR Switzerland

24 August 2005

Is the Gamma Risk of Options Insurable?

Daniel Egli, Peter Blum, Michel M. Dacorogna,

and Ulrich A. Müller

Converium Ltd, Zurich, Switzerland

August 24, 2005

Abstract

In this article we analyze the risk associated with hedging written call options. We introduce a way to isolate the gamma risk from other risk types and present its loss distribution, which has heavy tails. Moving to an insurance point of view, we define a loss ratio that we find to be well behaved with a slightly negative correlation to traditional lines of insurance business, offering diversification opportunities. The tails of the loss distribution are shown to be much fatter than those of the underlying stock returns. We also show that badly estimated volatility, in the Black-Scholes model, leads to considerably biased values for the replicating portfolio. Operational risk is defined as caused by imperfect delta hedging and is found to be limited in today's markets where the autocorrelation of stock returns is small.

Contents

1 Introduction

There are two entirely different reasons for writing call options. On one hand, writing covered calls can be a profitable business in an upward market for a (private) investor. Banks, on the other hand, want to make their profit when writing calls only on the margin they charge for this service. Thus they are looking for a risk neutral situation where they would not be exposed to market risks. In theory, they can achieve this goal by means of delta hedging. However, perfect risk neutrality is attainable only with continuous time hedging, which cannot be done in practice. Using a discrete hedging strategy, banks are actually exposed to losses stemming from non-linear effects in price movements. In this article we examine the risk of banks writing options with only daily updates of the delta hedge. In this article, we examine the statistical properties of this risk and show that its nature makes it transferable to insurance companies and would give the latter good diversification possibilities over traditional lines of business.

The article is organized as follows: In Section 2 we explain briefly the main concepts of option valuation, delta hedging and gamma risk. Section 3 describes the hedged portfolio of call options that we choose for our model and Section 4 the data and the variable definitions used in the study. Results are presented in Section 5, Section 6 covers the insurance related topics and we conclude in Section 7.

2 Option Theory

A call (put) option is a contract giving its owner the right but not the obligation to buy (sell) a fixed number of shares of a specified common stock at a fixed price at any time on or before a given date [Cox and Rubinstein, 1985]. The act of making this transaction is referred to as *exercising* the option, the specified stock is known as the *underlying security*, the fixed price is termed the *strike price*, and the given date is called the *maturity date*. If exercising is only allowed on the maturity date, the option is called *European*, otherwise *American*. The great flexibility of options only becomes evident when combined positions are considered. The most popular of these is a *hedge*, which combines a written position in calls (or a purchased position in puts) with a long position in the underlying stock.

The fundamental problem of option theory is to determine the option value as time passes by. There is one date when the value of the call can be easily stated, namely on the date of expiration:

$$C^* = \max(S^* - K, 0) \ ,$$

where S^* is the stock price on the date of expiration, K the strike price and C^*

the price of the call on this date. But what is the value in the meantime? The basic idea is to replicate the call by a self financing portfolio consisting entirely of stocks and borrowing. This is of course not feasible with a static portfolio, a consequence of the great flexibility of an option. In fact, a dynamic strategy has to be adopted. Very roughly, the procedure is as follows. With the proceeds of writing a call one buys a certain amount of stock needed to optimally hedge the option. (The exact amount needed is the actual problem and will be discussed below.) If the proceeds are not sufficient, the rest is borrowed at the risk free rate. As time goes by, one has to adjust the value of the portfolio to mimic the call by buying or selling more stocks. This is financed through new borrowing, which is continually repaid in case of selling. If one achieves to perfectly imitate a call, then its value would be given by the value of the corresponding portfolio, because any deviation therefrom would imply immediate arbitrage opportunities. Option pricing theory underwent a revolutionary change in 1973 when Fischer Black and Myron Scholes presented their option pricing model [Black and Scholes, 1973]. In the same year, Robert Merton extended their model in several important ways. In particular is it his credit to clearly introduce the concept of the replicating portfolio and no-arbitrage condition [Merton, 1973]. The assumptions underlying their differential equation are:

- The underlying stock pays no dividends during the life of the option.
- The price of the stock one period ahead (given its current price) has a log-normal distribution with mean μ and volatility σ, which are both constant over the life of the option.
- There exists a risk-free interest rate which is constant over the life of the option.
- Individuals can borrow as well as lend at the risk-free interest rate.

In a continuous-time setting, we assume that the behavior of the stock price is determined by the following stochastic differential equation:

$$\mathrm{d}S = S(\mu \mathrm{d}t + \sigma \mathrm{d}W) ,$$

where W is a standard Brownian motion (and $\mathrm{d}W$ a Wiener process). Let us now assume that we have a portfolio consisting of one option of value $C(S,t)$ and Δ amount of stock. The value of the portfolio at any time t is

$$\Pi = C(S,t) + \Delta\ S .$$

The gain of the portfolio is thus given by

$$\mathrm{d}\Pi = \mathrm{d}C(S,t) + \Delta\ \mathrm{d}S . \tag{1}$$

It can now be shown that in order for the gain to be deterministic (i.e. no dependence on the stochastic variable W) the following relation must be fulfilled:

$$\Delta = -\frac{\partial C}{\partial S} \ .$$

If the gain in value of the portfolio is deterministic, then it can be no more or less than the gain in value were it invested at the risk free interest rate r:

$$\mathrm{d}\Pi = r\Pi \mathrm{d}t = r\left(C - \frac{\partial C}{\partial S}S\right)\mathrm{d}t \ . \tag{2}$$

By comparison the differential equation of Black and Scholes (BS) follows directly and can be solved to yield the value of a call option (since we assume that there are no dividends, the values of American and European calls are identical)

$$C = SN(x) - Ke^{-rt}N(x - \sigma\sqrt{t}) \tag{3}$$

with

$$x \equiv \frac{\log(S/K) + r}{\sigma\sqrt{t}} + \frac{1}{2}\sigma\sqrt{t} \ ,$$

where N is the normal cumulative distribution function and t the current time to expiration.

Here is one interpretation of the formula: If we exercise the call on the expiration date we will receive the stock, but in return we will have to pay the strike price. The first term in the formula, $SN(x)$ is the present value of receiving the stock if and only if $S^* > K$, and the second term, $-Ke^{-rt}N(x - \sigma\sqrt{t})$, is the present value of paying the strike price if and only if $S^* > K$. Returning to the initial idea of replicating the call by a portfolio of stocks and borrowing, we see that the formula has another interpretation. It can be shown that $N(x) = \Delta$ is the number of shares in the equivalent portfolio and $Ke^{-rt}N(x - \sigma\sqrt{t})$ is the amount borrowed. So in principle, if the writer of a call wants to protect himself from loss he could arrive at a risk free situation with a perfect and continuously adjusted delta hedge (i.e. perfectly imitate a written call with a portfolio consisting of Δ amount of shares partly financed by borrowing), given by

$$\Delta = \frac{\partial C}{\partial S} = N(x) \ . \tag{4}$$

But in practice it will not be possible and not even desirable due to transaction costs to adjust the delta hedge continuously. The writer is therefore exposed to quick changes of delta with varying stock price. The risk that delta changes significantly between the discrete hedging times is termed gamma risk, where gamma is the sensitivity of delta to the underlying asset price:

$$\Gamma = \frac{\partial \Delta}{\partial S} = \frac{\partial^2 C}{\partial S^2} = \frac{1}{S\sigma\sqrt{t}}N'(x) \ , \tag{5}$$

where

$$N'(x) = \frac{1}{\sqrt{2\pi}}\, e^{-x^2/2} \ .$$

If gamma is high, i.e. delta changes quickly, it may be close to impossible to keep the hedge dynamically adjusted.

Apart from delta and gamma, there are other interesting quantities that can easily be computed from the BS formula:

$$\begin{aligned}
\Theta &= -\frac{\partial C}{\partial t} \qquad \text{(theta: sensitivity of call price to remaining time)} \\
\nu &= \frac{\partial C}{\partial \sigma} \qquad \text{(vega: sensitivity of call price to volatility)} \\
\rho &= \frac{\partial C}{\partial r} \qquad \text{(rho: sensitivity of call price to interest rate)}
\end{aligned}$$

3 Constructing a portfolio of hedged call options

To model the risk of writing call options, we proceed as follows. A certain number (inversely proportional to the call price) of call options is sold every day. All calls issued have the same time to maturity and same kind of strike price (either all at-the-money, all in-the-money or all out-of-the-money). We compute the option price with historic or artificial stock data using BS. Our study is based on daily data, so all computations refer to one daytime, typically at "market close" of a given day. At that time of each trading day, we hedge the issued calls according to the delta gained from BS. We keep track of all options individually, and can thus look at different quantities, such as the performance of a call during its lifetime, or the daily total position of the portfolio comprising all calls. An especially important quantity is the loss/profit distribution. As we explained above, if we hedged continuously, profit would be identically zero. Any deviation from zero bears therefore information about the risk associated with delta hedging. The loss distribution is computed in the following way. When we look at a single call, we denote by Π_i the difference between the value of the hedge portfolio on day i and the value of the call on this day, before updating the hedging:

$$\Pi_i \;=\; \Delta_i\, S_i - B_i - C_i \qquad\qquad i = 1, ..., n \tag{6}$$

where the subscript i denotes the respective value on day i in the life of the option (n is the expiration day) and B_i the current amount borrowed:

$$\begin{aligned}
B_1 &= \Delta_1\, S_1 - C_1 \\
B_{i+1} &= B_i e^{r/252} + (\Delta_{i+1} - \Delta_i) S_{i+1}
\end{aligned}$$

The total value of our position on day i is then simply the sum over all calls issued:

$$P_i = \sum_{calls} \Pi_i \tag{7}$$

The quantity we are looking for, the loss distribution, finally is the statistical distribution of the returns $\ln(P_i/P_{i-1})$. In this paper, we are also looking at first differences $P_i - P_{i-1}$.

If we compute the loss distribution in the above way, the daily losses (gains) are a mixture of different quantities. They are obtained from the Taylor expansion of the option pricing formula:

$$\Pi_i - \Pi_{i-1} \simeq \frac{1}{2}\Gamma(\delta S)^2 - \Theta\delta t + \nu\delta\sigma + \rho\delta r \ ,$$

but what we really want to analyze is the contribution of gamma alone: $\frac{1}{2}\Gamma(\delta S)^2$. In our world of the replicating portfolio, we define the following quantity which corresponds to the the gamma contribution:

$$\lambda^{\Gamma}_{i+1} = \Delta_i \ S_{i+1} - B_i - C(S_{i+1}, K, r, t_i, \sigma_i) - (\Delta_i \ S_i - B_i - C(S_i, K, r, t_i, \sigma_i))$$

which simplifies into

$$\lambda^{\Gamma}_{i+1} = \Delta_i \ \delta S - C(S_{i+1}, K, r, t_i, \sigma_i) + C(S_i, K, r, t_i, \sigma_i)$$

Note that the correspondence to the gamma contribution above follows easily if we expand C_{i+1}:

$$C_{i+1} \simeq C_i + \Delta_i \ \delta S + \frac{1}{2}\Gamma(\delta S)^2 \ .$$

Notice, however, that λ^{Γ} actually comprises not only the second order Taylor term, but all higher order terms proportional to a power of δS as well. If those higher order terms are not negligible, which can be the case for highly non linear options (near the money, short time to maturity), the quantity $\frac{1}{2}\Gamma(\delta S)^2$ is meaningless and λ^{Γ} must be used [Estrella, 1995]. Instead, we want to study the gamma risk in terms of the quantity

$$\Lambda^{\Gamma}_i = \sum_{calls} \lambda^{\Gamma}_i \tag{8}$$

4 Data

There are five quantities in the Black-Scholes formula (3) for which we need inputs to conduct an empirical study: stock and strike prices, time to maturity, volatility values and risk-free interest rates. In this section, we describe the data we use for those quantities.

4.1 Stock market data

We use daily stock-price data from Bloomberg, where "daily" means one observation per business day, excluding weekends and holidays. We take stock indices rather than individual stocks because indices are more representative, and corresponding volatility data are better available. For different stock indices, we have different data samples as listed here:

- S&P 500 (1941-2003)
- SPI (1988-2000)
- DAX (1988-2000)
- Nikkei (1987-2000)
- FTSE-100 (1984-2000)

4.2 Strike prices

The strike prices K used in the study are not taken from an external data source. We assume strike prices in relation to the stock price S on the issue date. We normally use five different strike prices, one at the money ($K = S$) and the others symmetrically distributed around it (nearest integer values):

$$K = S(1 \pm k\sigma\sqrt{\tau}), \quad k = 0, \frac{1}{2}, 1 \tag{9}$$

In the text, we call $k = \pm 1/2$ out of/in the money, and $k = \pm 1$ deep out of/in the money.

4.3 Times to maturity

As in the case of strike prices, we do not rely on an external data source. We rather assume that options can be written for a predefined set of times to maturity τ, namely $\tau = 1/4$, $1/2$, 1 and 1 1/2 years.

4.4 Volatility data

In a study on options, it is natural to take implied volatilities from option markets. Implied volatilities result from the Black-Scholes formula (3), which relates the value C of a call option to given values of S, K, t, r, and σ. We can (numerically) invert the formula to get σ from S, K, t, r, and C. This value for σ is therefore termed 'implied' volatility. In principle, different options on the same underlying stock may have different implied volatilities. If the volatility is changing over time, options with different expiration dates would not be expected to have the same implied volatility. Options with different strike prices may also have different implied volatilities, a phenomenon that is called the volatility "smile".

The data provider Bloomberg therefore quotes weighted averages of implied volatility:

$$\sigma = \frac{\sum_{j=1}^{J} w_j \sigma^*(t_j, K_j)}{\sum_j w_j} \ ,$$

where w_j is the weight given to option j, J is the total number of options used for the average and $\sigma^*(t_j, K_j)$ represents the implied volatility for option j with time to maturity t_j and strike price K_j. There are different weighting schemes in use; our volatility data result from Bloomberg's scheme.

We are using implied volatilities for times to maturity of 3,6 and 12 months as obtained from Bloomberg, in (business-)daily resolution.

For dates earlier than 02.01.1990, implied volatility is not available. We replace it by statistically computed historical volatility of logarithmic stock returns. The annualized historical (or realized) volatility σ is defined as follows:

$$\begin{aligned}
\sigma^2 &= \frac{1}{T-1} \sum_{t=1}^{T} (R_t - m)^2 \\
m &= \frac{1}{T} \sum_{t=1}^{T} R_t \\
R_t &= \ln(S_t / S_{t-1}) \ ,
\end{aligned}$$

where T is the length of the historical sample in days. Too large a T may lead to irrelevant results, since volatility can change considerably over large time scales, while too small a T could yield insufficient accuracy. A reasonable value for T lies between 100 and 200 days, leading to volatility estimations that are best attainable proxies for the missing implied volatility.

4.5 Interest rate data

When analyzing options, we use interest rates with the same times to maturity as those of the options and the implied volatilities. Implied US LIBOR rates for 3,6 and 12 months are obtained from Bloomberg. (We also use the LIBOR rates for 12 months when options with maturity time of 18 months are analyzed).

Interest rates have a minor influence on the investigated behavior of options and gamma risk when compared to other factors.

Figure 1: Cumulative Distributions of Daily Logarithmic Stock Returns

Left: stock index S&P 500, right: artificially constructed, log-normal stock index. The left graph shows a clear deviation from a straight line, indicating a fat-tailed distribution.

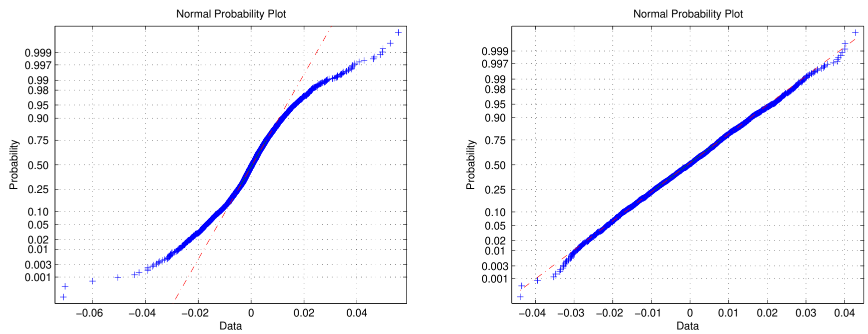

4.6 Artificial stock index data

In order to eliminate effects of varying volatility in some special types of analysis, we also construct and investigate artificial stock data, the returns of which are made to be log-normally distributed with constant volatility ($\sigma = 0.2$, annualized):

$$S_1 = 1000 \ , \quad S_i = S_{i-1}\, e^{X_i} \text{ with } X_i \in \mathcal{N}(\prime, \sigma^{\in})$$

This means that the BS model is fully applicable and call prices are in principle (continuous hedging) perfectly matchable by the hedging portfolio. By using artificially created stock data as in a Monte-Carlo simulation, we are able to generate arbitrarily large samples. This allows us to compute results for several different stock-price series and then to average over the whole set, in order to minimize stochastic noise.

The assumption that the evolution of stock price is governed by a geometric Brownian motion with *constant* volatility is not fulfilled in practice, as is clearly shown in Figure 1. Real log-normal data should yield a straight line.

5 Results

5.1 Properties of the hedged portfolio

As discussed in Section 3, we are interested in a full portfolio of call options with a hedging portfolio consisting of stock positions for delta hedging and loans or deposits. At each business day, a new call option is written (at-the-money),

Figure 2: Distribution of Daily Option Portfolio Returns, Based on the S&P 500

The distribution of $P_i - P_{i-1}$ is shown, where P_i is the value of the complete, hedged portfolio. The first call option of the portfolio is assumed to be issued on the 1.1.1990, leading to a total of ≈ 3500 call options (not at the same time in the portfolio, as old options continuously expire). Time to maturity is $\tau = 1/4$ (in years), K is at the money.

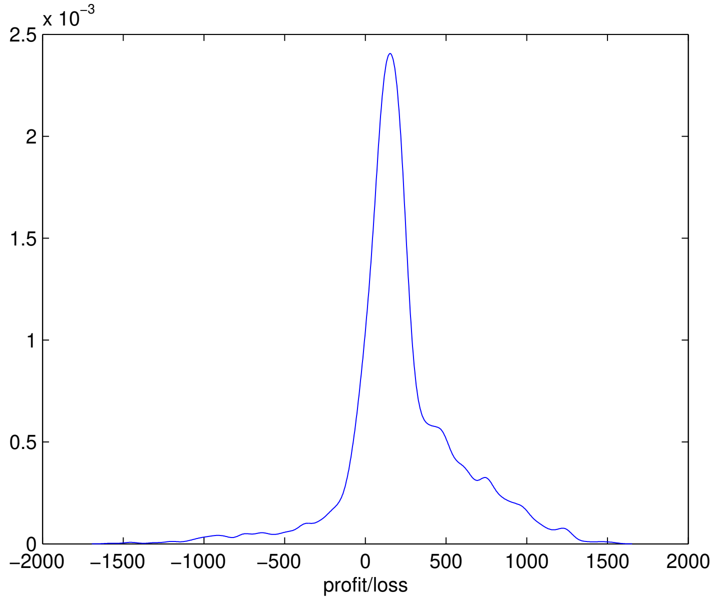

hedged and added to the portfolio, and an old one from the portfolio expires. The value of the whole portfolio as defined by Equation 7 is a new, basic object of study.

Figure 2 shows the distribution of daily changes (profits and losses) of the portfolio value. The distribution of $P_i - P_{i-1}$ is rather peculiar. The tail on the positive side is attributable to the lack of exact knowledge concerning the volatility. Fluctuations in the volatility σ widen the distribution. Different factors determining the shape of the distribution will be identified in the further course of the analysis.

Successful hedging crucially depends on the fact that call price can be perfectly replicated by the hedging portfolio. As we can see in Figure 3, we miss the goal when using a bad guess of implied volatility. If implied volatility is too high, the hedging portfolio does not exactly duplicate the call price, but rather follows it on a systematically higher level (and analogously on a lower level for too low a volatility). Our model interprets that as a profit according to Equation 6. This is reflected in the fat tail of the distribution in Figure 2.

Figure 3: Values of the Call Option Portfolio and the Replicating Portfolio, Based on the S&P 500

Top: $\sigma = 0.1949$ (implicit). Middle: $\sigma = 0.122$. Bottom: $\sigma = 0.05$. The call option value (blue) deviates from that of the replicating portfolio (green) as soon as an inaccurate volatility figure is used for hedging. If constant volatility is assumed, we see that too high or too low a volatility leads to systematic bias for the replicating portfolio.

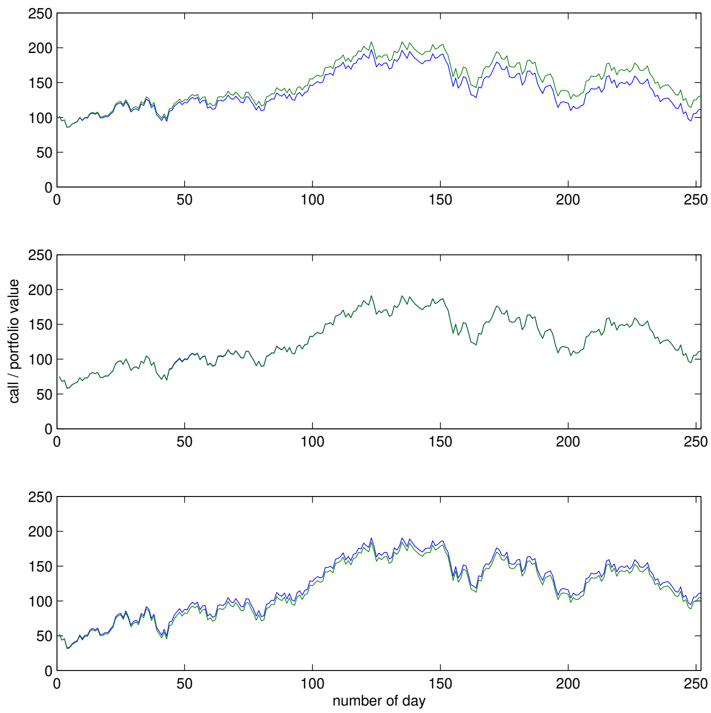

5.2 Loss distributions of portfolio returns

We now present the loss distributions of hedged call option portfolios based on the artificial stock data as described in Section 4.6. Figure 4 shows the loss distributions for different strike prices and times to maturity. The loss distributions were computed for 200 artificial stock-price series, each covering 3000 business days. The total corresponds to about 2400 business years. A portfolio based on real data, for comparison, just covers 24 years.

Considering their overall shape, all distributions are quite similar. They are centered slightly in the positive range but asymmetric: Their negative tails are much fatter and more extended than the positive tails. We can observe that the

Figure 4: Distributions of Daily Option Portfolio Returns, Using Artificial Data

The distributions are skewed with fat tails on the negative sides, (representing the losses). They are shown for different times to maturity: Top left: $\tau = 1/4$, top right: $\tau = 1/2$, bottom left: $\tau = 1$, bottom right: $\tau = 3/2$. The initial strike price K is always at the money ($= S$). Monetary units are for these graphs arbitrarily chosen, as we are here just considering the shape in general.

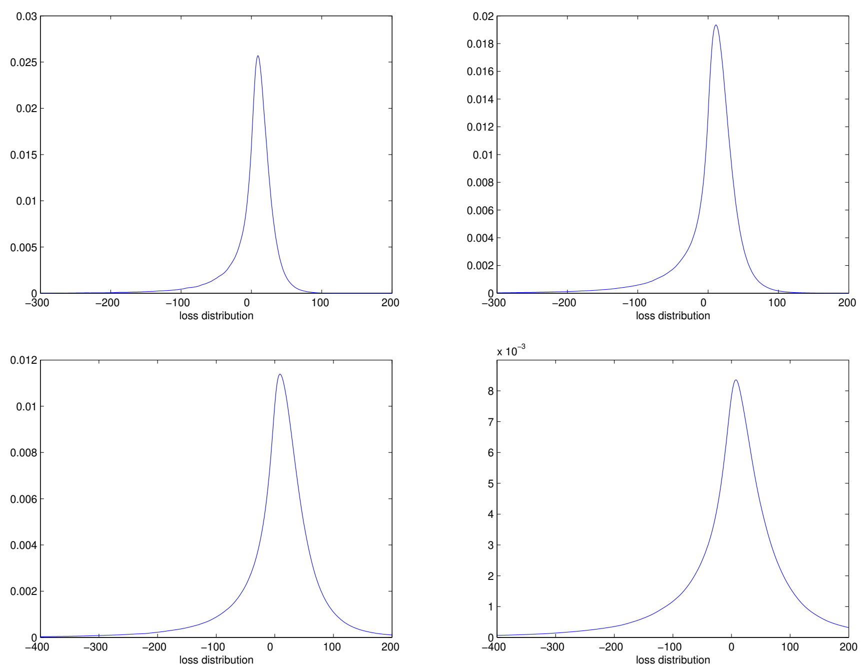

distributions get flatter and less peaked with longer time to maturity and strike prices nearer to the money.

Figure 5 shows the corresponding return distributions for the hedged call option portfolio based on the S&P 500. Unlike Figure 4, which is based on constant volatility, this distribution is dominated by effects of fluctuating volatility. Notice some rare extreme losses in the far negative region of the distribution. Only in this region of high losses, we observe some rare events representing the skewness of the gamma risk.

In order to fully understand these loss distributions, we need to have a closer look at the different terms that drive the behavior of the portfolio returns.

Figure 5: Distributions of daily option portfolio returns, using S&P 500 data

This distribution of daily returns of the option portfolio is based on S&P 500 data. The initial strike price K is always at the money ($= S$). The portfolio is analyzed from 2.1.1980 to 28.6.2004. This sample is smaller than those of Figure 4 and leads to a more jagged graph with more stochastic noise.

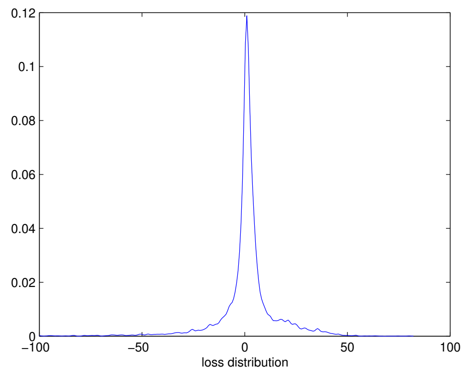

5.3 Components of portfolio returns and risks

The features of those distributions are readily explained if we look at the Taylor-expansion of the call value:

$$C - C_0 \cong \Delta \, \mathrm{d}S + \underbrace{\frac{1}{2}\Gamma \left(\mathrm{d}S\right)^2}_{>0} \underbrace{- \, \Theta \mathrm{d}t}_{<0} + \nu \mathrm{d}\sigma$$

The second derivative Γ as defined by Equation 5 is always positive. We see that, on one hand, Γ always increases the call value beyond the hedge portfolio value, but Θ, on the other hand, always decreases it ($-\Theta > 0$, but $\mathrm{d}t < 0$). It is of course bad for the writers of the call to underestimate a call price (as we see from equation (6)). Thus Γ always works against the writers, while Θ always benefits them. The crucial question now is the relative influence of these two antagonistic quantities. As a matter of fact, on usual days (meaning no large jumps in stock prices) the following relation holds:

$$\frac{1}{2}\Gamma \left(\mathrm{d}S\right)^2 < \Theta \mathrm{d}t \; ,$$

which means that usual days should tend to be profitable. That the peak of the loss distributions is in the positive clearly reflects this fact. If, however, the stock price changes substantially, then $\frac{1}{2}\Gamma \left(\mathrm{d}S\right)^2$ outweighs $\Theta \mathrm{d}t$ because of its quadratic dependence on the change in stock price. The fat negative tail shows these rarer but more sizeable events.

Figure 6: Comparing return components for a single hedged call option

Left: Two return components, $\frac{1}{2}\Gamma\,(\mathrm{d}S)^2$ (solid) and $\Theta\mathrm{d}t$ (dashed) for a single, hedged at-the-money call of S&P 500 with $\tau = 1$, using implied volatility. Right: Same as left, but the figure superimposes on top $\nu\mathrm{d}\sigma$.

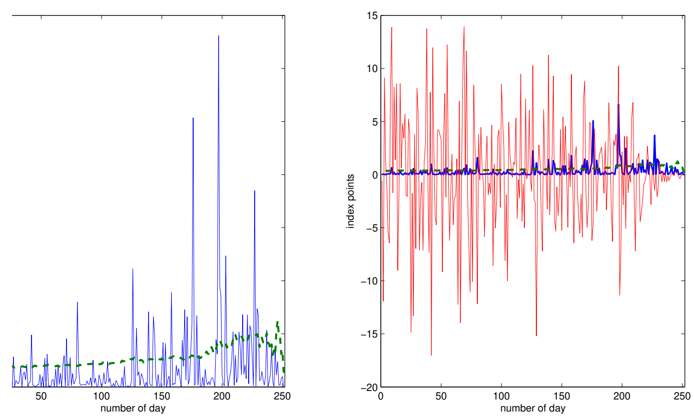

For real stock data as analyzed in Figure 5, we have to include another term into the Taylor approximation, $\nu\mathrm{d}\sigma$, the vega component stemming from variable volatility (we actually should also account for the changing interest rate term, $\rho\mathrm{d}r$, but $\mathrm{d}r$ is so small as to make this term negligible). Unlike Γ and Θ, we cannot ascribe a definitely detrimental or beneficial effect for the call writer. Nevertheless, since ν is always positive, we can state that increases in volatility decrease profit of writer and vice versa. This so called vega-risk is the dominating effect. Figure 6 compares the three quantities and corroborates the statements made above. Note on the right figure how the other effects are dwarfed by the superimposed term: $\nu\mathrm{d}\sigma$. Be also aware that these huge volatility-based amplitudes are mainly due to $\mathrm{d}\sigma$, which can be very high because of implied volatility fluctuations. The picture would look different if using historical volatility, which has a strong smoothing effect, instead of implied volatility. In this case, the gamma and vega risks would be of the same order of magnitude.

The situation becomes clearer when we look at the dependence of Γ and Θ on the volatility (Figure 7). Θ grows (in absolute terms) strongly with volatility, while in marked contrast Γ drops quickly. From this figure it is clear that too high a volatility will lead to consistently higher hedge portfolio values (theta dominates). Exactly the opposite is true if the volatility is too small (gamma dominates). This explains the beforehand peculiar systematic discrepancy between the replicating portfolio and the call value if volatility is wrongly estimated.

Figure 7: Dependence of gamma and theta on volatility

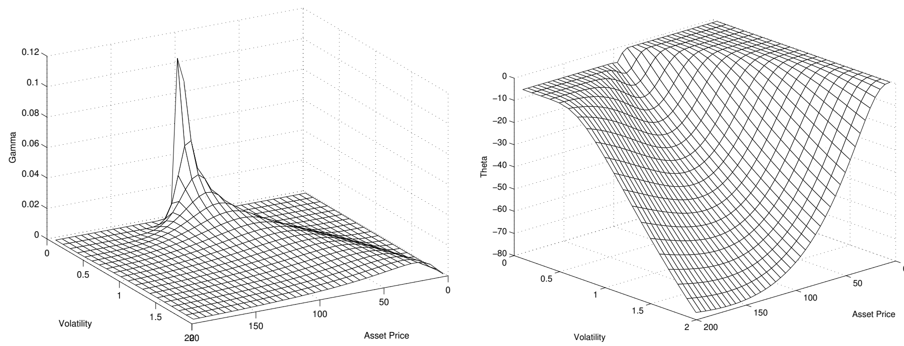

6 From Risk to Insurance

6.1 Focusing on the gamma risk

So far we have not specified the nature of the call writer. Normally, this will be a bank or another big financial institution, as writing call options is very risky for a private investor who does not delta hedge. Moreover, frequent hedging will be too costly and time intensive, as trading capabilities of an individual are very limited. But for a bank it is no problem to follow exactly the delta hedging strategy we used in our model. In addition, since the influence of time (Θ) is predictable and always beneficial, the bank wants to exclude this effect from risk considerations. Also vega risk should be manageable, because it is part of the expertise of banks to judge and predict market volatility. Yet, sudden changes in stock price are unpredictable, and thus gamma risk shares many characteristics with a traditional insurable risk: It normally entails small losses on most days and huge losses on some rare and random occasions. Thus from an insurance point of view, it is interesting to look at the isolated gamma risk and see if it would be possible to take it on an insurance book. We measure the gamma risk in terms of the gamma losses Λ_i^Γ as defined in Equation 8 and explained at the end of Section 3.

In Figure 8 we present the loss distributions due to the gamma risk for options on the S&P 500 stock index. We use real data because volatility risk is excluded in this setting. We plot histograms which stress more the discreteness of worst loss events.

Figure 8: Distributions of gamma losses

The histograms show logarithmic gamma losses in USD for options on S&P 500 (see Section 6.2 for discussion of reference capital). Times to maturity is $\tau = 1/4$, $K = 0, -0.5, 0.5, -1, 1$ from left to right and top to bottom.

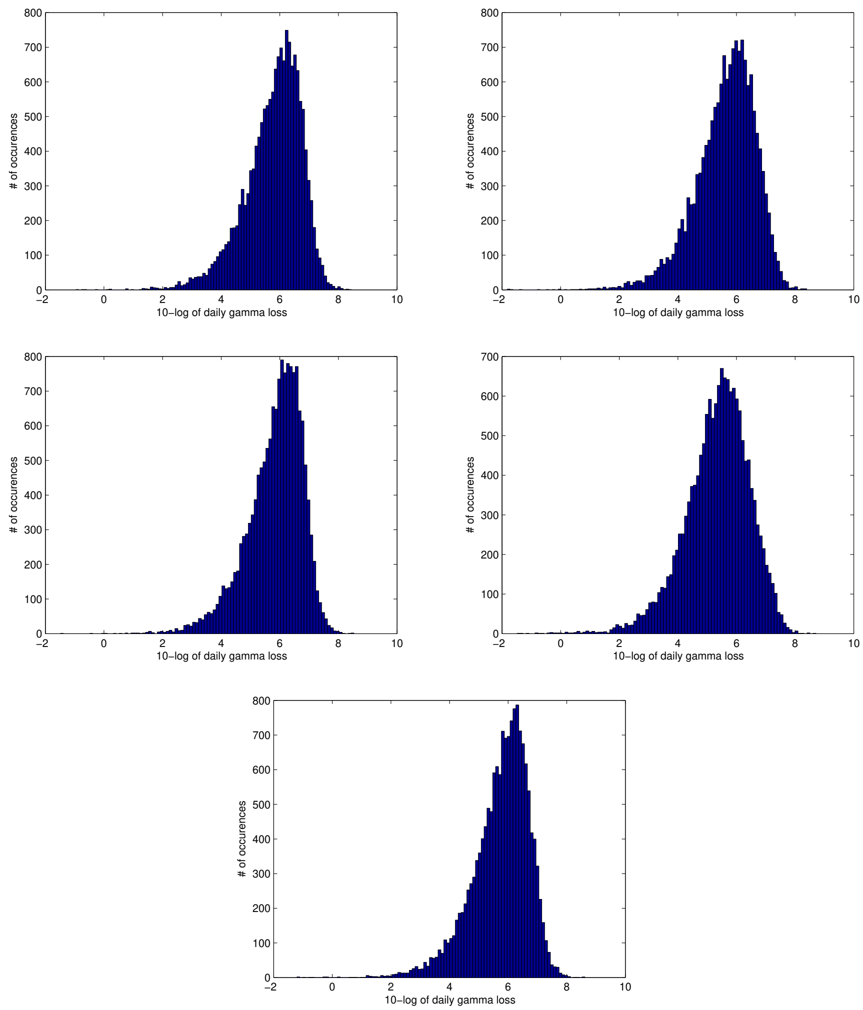

6.2 Loss Ratio

In order to gain an intuition about the orders of magnitude involved, we have to specify the capital that a company needs to have to underwrite gamma losses. However, our model assumed a self-financing strategy so far, with zero capital. In reality, some capital is needed to cover the gamma risk. An often used

Table 1: Comparison of daily premiums (in thousand USD).
All options considered are issued at the money. Recall that the reference capital is a daily tail value at risk of 50 million US Dollars.

Time to Maturity	Our Premium	PRIME® Premium
$\tau = 1/4$	3'775	3'414
$\tau = 1/2$	3'683	3'168
$\tau = 1$	3'593	2'626
$\tau = 3/2$	3'468	2'557

type of risk measure to compute the risk-based-capital is the so-called Value-at-Risk (VaR), i.e. the loss that will not be exceeded with a certain probability (99% in our case) during a given holding period (the 99% quantile of the loss distribution). The portfolio is then composed in such a way that the total exposure (i.e. the value of all calls combined on a day) gives rise to a VaR of a certain specified amount. Alternatively, we can use a stronger risk measure for the capital, namely tail value at risk (TVaR) (also called expected shortfall). This is the average of all values beyond a certain threshold (the 99%-quantile in our case). In our model, we cap the number of written calls such that daily TVaR equals 50 million.

A quantity often used in insurance to measure profitability is the so-called loss ratio that compares the losses to the premiums and is of primary interest for insurers. It is defined as follows:

$$\text{Loss Ratio} = \frac{\text{losses}}{\text{premiums}} \, ,$$

where in our model we compute premiums in the following way. The premium P is assumed to be due for a certain time period which we choose to be one year, following a common practice of the insurance industry.

$$\frac{\mathrm{E}(P - C - \Lambda^{\Gamma})}{\mathrm{TVaR} - P + C} = \mathrm{RoRBC} = 0.2 \, , \tag{10}$$

E denotes expectation value, Λ^{Γ} is the yearly gamma loss and C denotes the costs of the insurer. The numerator is the insurer's expected return, the denominator the insurer's risk-based capital (RBC) including premium and costs. The quotient is the return on RBC, called RoRBC, which we assume to be 20% in our model computations. For lack of an exact knowledge we take a conservative estimate of the cost ratio CR:

$$\mathrm{CR} = \frac{C}{P} = 0.2\ , \quad C = \mathrm{CR} \cdot P = 0.2P \tag{11}$$

Tail value at risk (TVaR) is of course also computed on a yearly basis. While yearly gamma loss is simply the sum of all losses during a year, the aggregate TVaR is not so straightforward to calculate, since our basic values are daily. We choose a Monte-Carlo approach by taking 10,000 re-samples of accumulated losses of randomly selected 252 days (mimicking business days in a year) and compute the mean of the 100 worst aggregate losses.

One can easily solve Equations 10 and 11 for the premium P:

$$P = \frac{0.2 \cdot \mathrm{TVaR} + \mathrm{E}(\Lambda^{\Gamma})}{1.2} + C = \frac{0.2 \cdot \mathrm{TVaR} + \mathrm{E}(\Lambda^{\Gamma})}{0.96}$$

These premiums are computed for the historical data. Thus, we can compute the loss ratios for different stock indices, as shown in the graphs of Figure 10). Premiums computed with Converium's pricing tool PRIME® are always lower, because they account for diversification effects in the global Converium portfolio, see Table 1. Premium calculated by PRIME® decrease more rapidly with increasing τ, reflecting increasing diversification with other insurance business when the TVaR of the gamma risk alone decreases.

The loss ratios of Figure 9 behave differently for the investigated indices over the analyzed years. Insuring the gamma risks of different countries thus offers a potential for diversification as shown by the correlation coefficients of Table 3. SPI and DAX are highly correlated, whereas the other three form a group of their own. However, all the indices are expressed in their own currencies. A deeper study of international indices should account for currency effects (which we expect to be rather small).

Loss ratios for the same stock index depend on time to maturity and moneyness of the options issued. Longer time to maturity means that there is a greater chance for the price to evolve away from and back to the issuing call price, offering both the possibility of small and large gamma. Again, we can have a look at the correlation of loss ratios of differently defined option portfolios series in order to check for diversification possibilities. In Table 2 we display the linear correlation coefficients between ten different loss ratios of option portfolios on S&P 500. If we have a closer look at this table, a clear pattern is apparent: Option series with positive k (out of the money, see Section 4.2) are only weakly or even negatively correlated to those with negative k (in the money). A more careful analysis shows that in addition that short-term options exhibit a reduced correlation with long-term options.

Figure 9: Loss Ratio for Different Stock Indices

All loss ratios are obtained for options of time to maturity $\tau = 1/4$ and strike price K originally at the money.

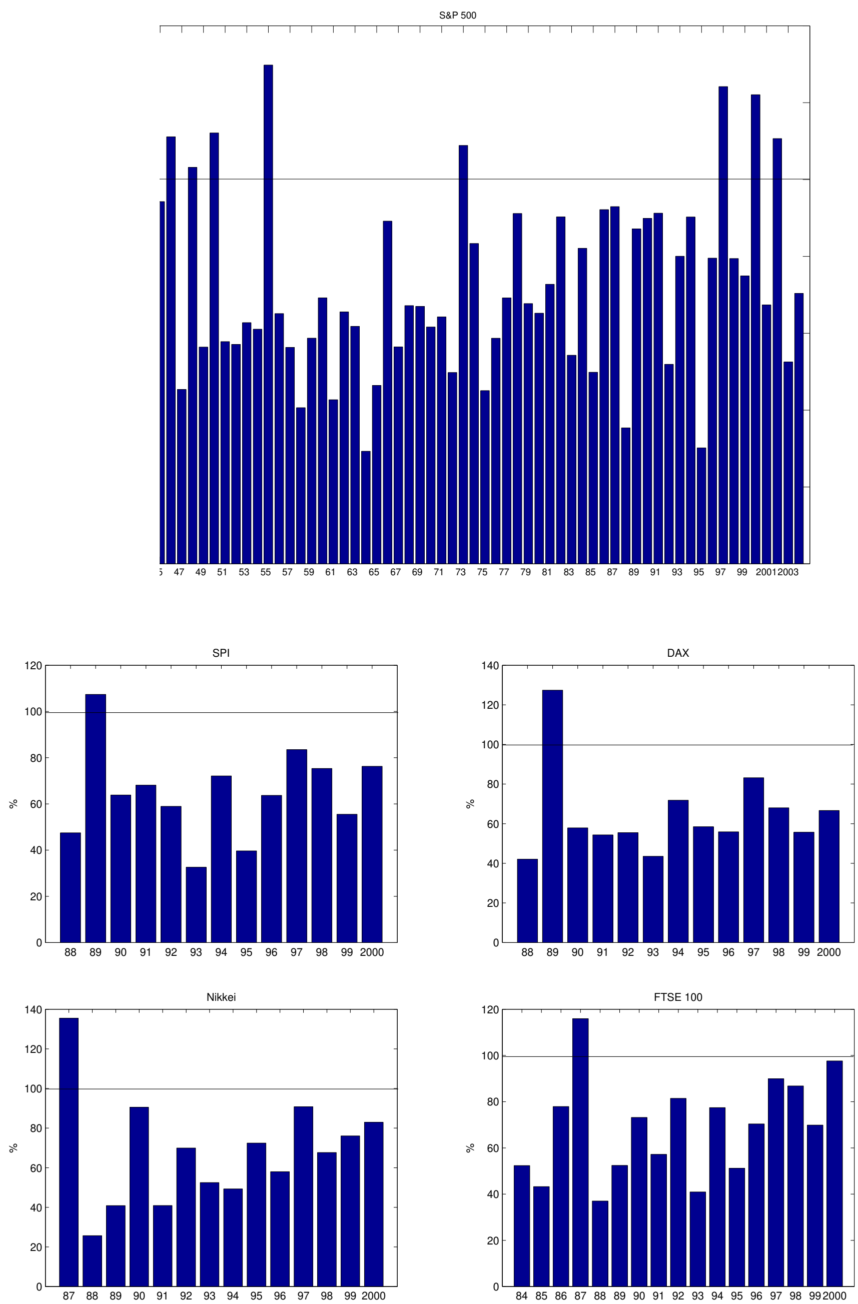

Table 2: Loss ratio correlation for different maturities and strike prices

k refers to the strike price: 0 means at the money ± 0.5 out of / in the money and ± 1 deep out of / in the money (see method section for details).

τ, k	$1/4, 0$	$1/4, -0.5$	$1/4, 0.5$	$1/4, -1$	$1/4, 1$	$1/2, 0$	$1/2, -0.5$	$1/2, 0.5$	$1/2, -1$	$1/2, 1$
$1/4, 0$	1.00	0.79	0.79	0.32	0.55	0.74	0.39	0.77	0.20	0.65
$1/4, -0.5$	0.79	1.00	0.40	0.79	0.16	0.76	0.70	0.48	0.58	0.26
$1/4, 0.5$	0.79	0.40	1.00	0.00	0.89	0.53	0.12	0.86	-0.04	0.92
$1/4, -1$	0.32	0.79	0.00	1.00	-0.17	0.61	0.84	0.18	0.83	-0.09
$1/4, 1$	0.55	0.16	0.89	-0.17	1.00	0.27	-0.10	0.68	-0.19	0.94
$1/2, 0$	0.74	0.76	0.53	0.61	0.27	1.00	0.81	0.78	0.60	0.41
$1/2, -0.5$	0.39	0.70	0.12	0.84	-0.10	0.81	1.00	0.33	0.92	0.01
$1/2, 0.5$	0.77	0.48	0.86	0.18	0.68	0.78	0.33	1.00	0.13	0.82
$1/2, -1$	0.20	0.58	-0.04	0.83	-0.19	0.60	0.92	0.13	1.00	-0.10
$1/2, 1$	0.65	0.26	0.92	-0.09	0.94	0.41	0.01	0.82	-0.10	1.00

Table 3: Correlation of loss ratios between indices

	SPI	DAX	Nikkei	FTSE 100	S&P 500
SPI	1.00	0.87	0.05	0.46	0.01
DAX	0.87	1.00	-0.00	0.18	-0.34
Nikkei	0.05	-0.00	1.00	0.70	0.52
FTSE 100	0.46	0.18	0.70	1.00	0.68
S&P 500	0.01	-0.34	0.52	0.68	1.00

We have mentioned earlier that gamma losses of options crucially depend on the time to maturity τ, and it is most natural that they also depend on the issuing strike price K. A further analysis shows that loss ratios tend to decline with increasing time to maturity. But this decrease is overshadowed by a large increase in uncertainty. The influence of K is less apparent, but there are some regularities: Both loss ratio and standard deviation increase mostly with options being farther away from the money, but the dependence is of course more intricate. During periods of downward trends in the stock price periods, in-the-money calls are more dangerous, since their gamma is larger, and the reverse is true for upward periods. In addition, the length of these trending periods is a factor. If the length of a trending period approximately coincides with the time to maturity τ, dangerously large gamma values are most likely to occur. (A Fourier analysis, however, yielded no insights.)

6.3 Tail Indices

To quantify the riskiness of insuring gamma losses, it is worthwhile to look at tail indices because the risk is well expressed by the heaviness of the tail. We say that a variable X follows a heavy-tailed distribution if

$$P[X > x] \sim x^{-\alpha} , \qquad \text{as } x \to \infty ,$$

where α is called the tail index. Estimating the tail index of an unknown distribution is not straightforward and is a much discussed issue for deciding what is the best way (see e.g. [Dacorogna et al., 2001]). For sake of simplicity we use Hill's estimator, defined by

$$\hat{\gamma}^H_{n,m} = \frac{1}{m-1} \sum_{i=1}^{m-1} \log X_{(i)} - \log X_{(m)} ,$$

Figure 10: Loss ratios plus one standard deviation

Top left: $\tau = 1/4$. Top right: $\tau = 1/2$. Bottom left: $\tau = 1$. Bottom right: $\tau = 3/2$. Loss ratios with one standard deviation (in lighter color) are shown.

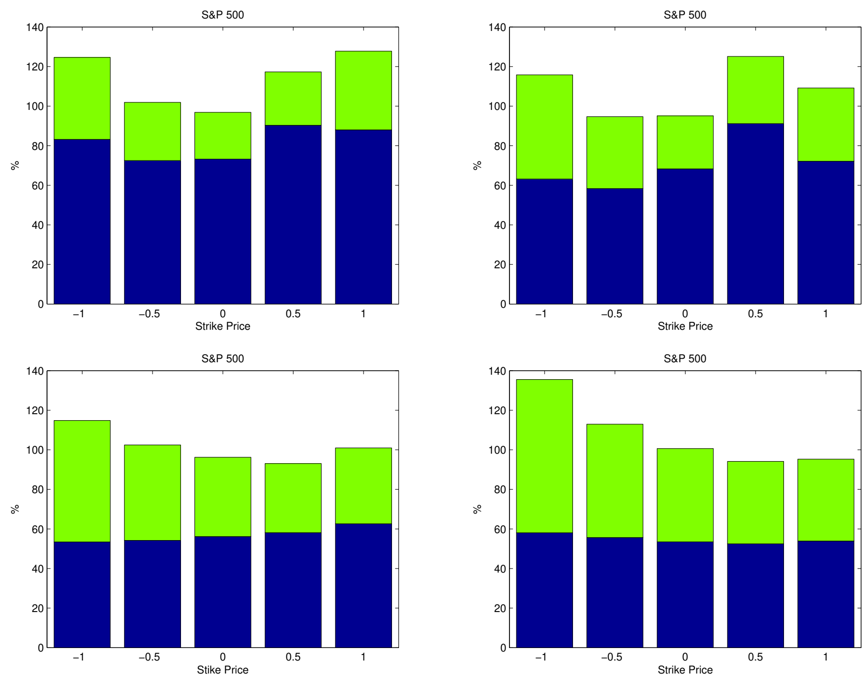

which is a consistent estimator of $\gamma = 1/\alpha$. In this formula, $X_{(i)}$ are the decreasingly ordered observations and n is total sample size. It is found that $m = \sqrt{n}$ yields a good compromise between systematic bias when m is too large, and stochastic noise when it is too small [Blum and Dacorogna, 2003].

Since gamma losses are proportional to squared returns, we would expect fatter tails than assumed by market returns. Table 4 confirms this view. While the α's of the Γ losses are around 2, those of stock returns are consistently higher between 3 and 4.

To obtain a clearer picture of the relation between the tail indices of the gamma losses and the market, we can again resort to artificial data. This time we create stock price series whose logarithmic returns are Student-t distributed with varying parameter to[1]. We generate 10 stock price series, each with 4000 business days, compute tail indices of these synthetic returns and their associated gamma risks, and then average. In this way we can plot the two α's against each other. If we know the asymptotic behavior of market returns to be

[1]Note that the tail index of a student-t distribution t_α-variable is α.

Table 4: Tail indices of Stock Indices and Γ risk

The first column K for the Γ risk refers to the strike price: 0 means at the money ± 0.5 out of / in the money and ± 1 deep out of / in the money (see Section 4.2 for details).

Index	α	K	τ	α	τ	α
S&P 500	**3.68**	1	1/4	**1.95**	1/2	**1.94**
SPI	**2.77**	0.5	1/4	**2.20**	1/2	**2.07**
DAX	**3.24**	0	1/4	**2.06**	1/2	**1.94**
Nikkei	**3.23**	-0.5	1/4	**2.05**	1/2	**1.83**
FTSE-100	**3.40**	-1	1/4	**2.05**	1/2	**1.96**
Nasdaq	**3.81**					

$\sim (\delta S)^{-\alpha-1}$, we may expect that gamma losses $1/2\Gamma(\delta S)^2$ asymptotically have an exponent of $\alpha^\Gamma = -(\alpha+1)/2$, as long as the distribution of Γ is independent of that of δS. Figure 11 confirms this expectation. (Since we look in the end at cumulative distributions, the values for market and gamma risk should be α and $\alpha/2$, respectively.)

6.4 Serial Correlation

To know whether high-loss days are randomly distributed or clustered is an important characteristic of an insured risk. Gamma risk is essentially composed of two different entities: δS and Γ. Absolute values of returns are known to be autocorrelated and total gamma on each day should as well show autocorrelation, because the total portfolio composition changes little from day to day (one pack expires and a new one is sold). Therefore we would expect some autocorrelation effects, which is confirmed by left part of Figure 12: Gamma losses show considerable autocorrelation for lags up to 5 days. Note also on the right graph that autocorrelation fades for large powers. This indicates that the large gamma losses are less autocorrelated than the small ones.

We are mainly interested in days where the gamma loss is above average. We can analyze this by computing the autocorrelation of $\left(\Gamma \cdot \delta S^2\right)^p$, thus putting more emphasis on large losses with increasing p. The result is displayed on the right side and shows that the autocorrelation coefficients peak at exponents less than 1. This fact indicates that large losses are less correlated than average ones. However, the range of powers > 1 should be looked at with caution because the variance of gamma risks, which is in the denominator of the correlation

Figure 11: Relating the tail indices of stock returns and gamma risk.

The tail index α^{Γ} of the gamma risk is plotted against the tail index α of the underlying artificial stock index (with Student-t returns). The dotted line indicates the transition where gamma risk becomes uninsurable, since the expected loss goes to infinity.

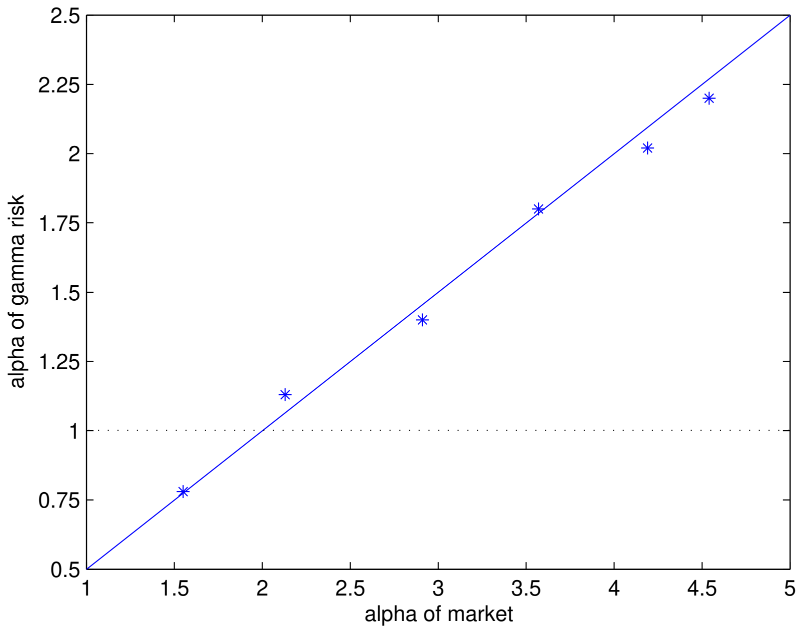

coefficients, diverges to infinity because of the low tail index values of Table 4. An infinite denominator means zero correlation in the asymptotic limit.

Since the dependence of gamma losses on changes of stock price is quadratic, one should think that catastrophic events on the stock market are also the worst days for gamma risk. This is not necessarily the case, as a low value of gamma can substantially dampen effects of sudden stock price jumps. This can be demonstrated by the example of the DJI stock-index on 'Black Monday' and on a much more benign day two years later. The gamma losses incurred on the latter day are roughly 40% higher (for options with $\tau = 1/4$) than those of the worst stock-market day in history. This astounding result is explained by the following graphs, which show that the previous history is of equal importance as the magnitude of the event itself. In the week before 'Black Monday', stock prices went down in little steps to a level where all issued calls were deep out of the money, implying a small gamma. Figure 13 illustrates the effect. As one clearly sees in the lower graph, gamma peaks when stock price is slightly below strike price and decreases rapidly when the call is out of the money. This explains why gamma losses in the right picture are much higher than those on

Figure 12: Autocorrelation of Gamma Losses

Left: Autocorrelation at various lags of Gamma losses. Right: Autocorrelation of gamma losses raised to a certain power for lags 1 to 10 days. (The dashed line is the upper limit of the 95% confidence band under the hypothesis of normal distribution.)

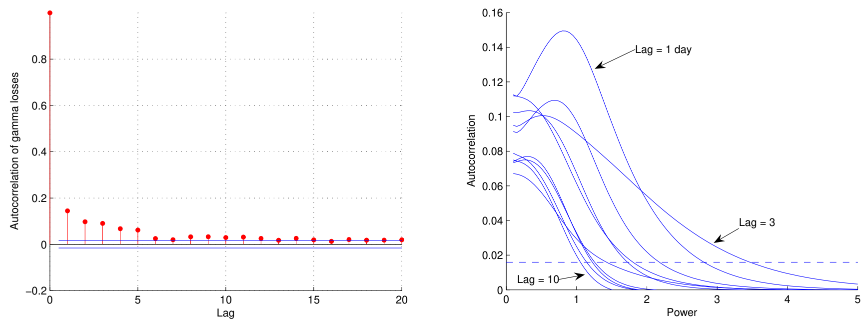

Figure 13: Influence of Stock-Price Evolution Prior to a Jump

Top: The dotted line shows the average strike price of all options issued that were in the portfolio on the day before the crash, while the dashed line is the stock price on this day.

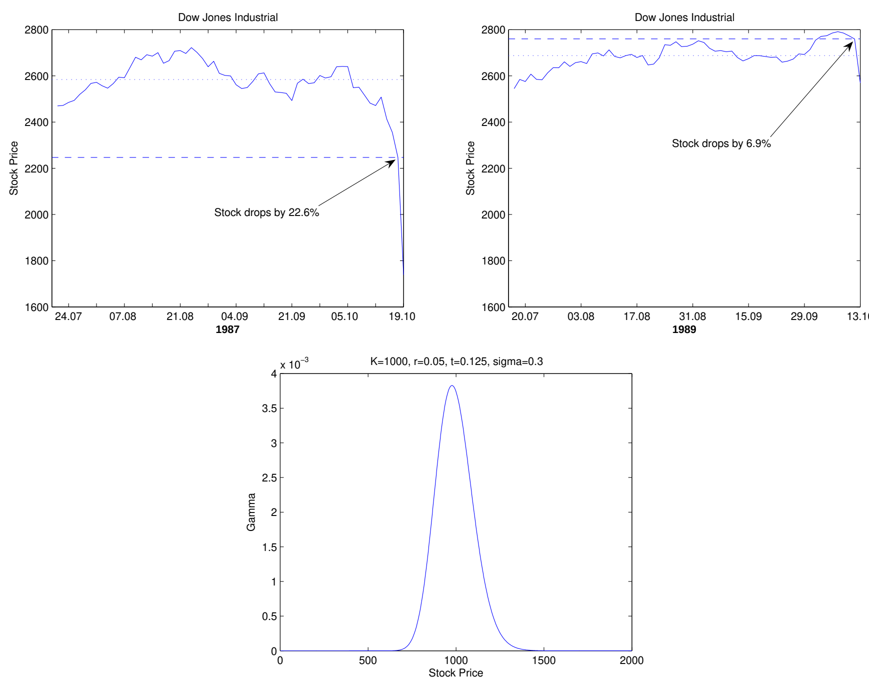

Figure 14: Imperfect Delta Hedge and Autocorrelation of S&P 500 Returns

Left: Values of a call option (solid line) and an imperfect hedge (dashed line) as functions of the stock price. A perfect hedge would appear as a tangent to the solid curve.
Right: Autocorrelation function of logarithmic returns of the S&P 500.

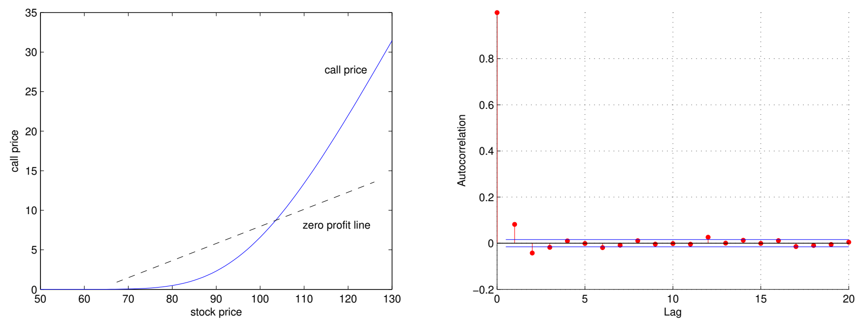

the left, although the crash was significantly smaller. We can conclude that the greatest danger for gamma losses lies in extended static periods, followed by a sudden jump of the stock price. (But be aware that this influence of prehistory decreases as the time to maturity increases, since the probability of a static period over, say, 18 months - which is the longest τ considered here - is very small. This is reflected by the fact that for these longer term options, 'Black Monday' was black indeed.)

6.5 Operational Risk

Another thing we should quantify is the operational risk due to imperfect hedging, as opposed to the inherent risk due to gamma in case of accurate hedging. A basic calculation completely analogous to the one in section 3 yields the loss[2] λ_{i+2} if hedging is done on day i, but not done on day $i+1$:

$$\lambda_{i+2} = \delta S_{i+1} \underbrace{(\Delta_i - \Delta_{i+1})}_{-\Gamma_i \delta S_i} - \frac{1}{2}\Gamma_{i+1}(\delta S_{i+1})^2 \, , \qquad \delta S_{i+1} = S_{i+2} - S_{i+1}$$

The first term on the right side, $-\Gamma_i \, \delta S_i \, \delta S_{i+1}$, is in contrast to the second term not always negative, thus yielding losses and gains, as is explained in Figure 14. A big gamma loss due to a sudden jump of stock price can be aggravated or partly offset, if on the next day the movement of stock price goes

[2] In case of continuous Δ-hedging, there is no Γ-risk at all. The case presented here is a variation of the base case presented before (hedging at discrete time points only, just with longer time intervals in between.

Figure 15: Operational Risk

Top left: Loss distribution with daily hedging. Top right: Loss distribution, where hedging is not done every second day. Bottom: Moving-window (of 5000 days) autocorrelation of S&P 500.

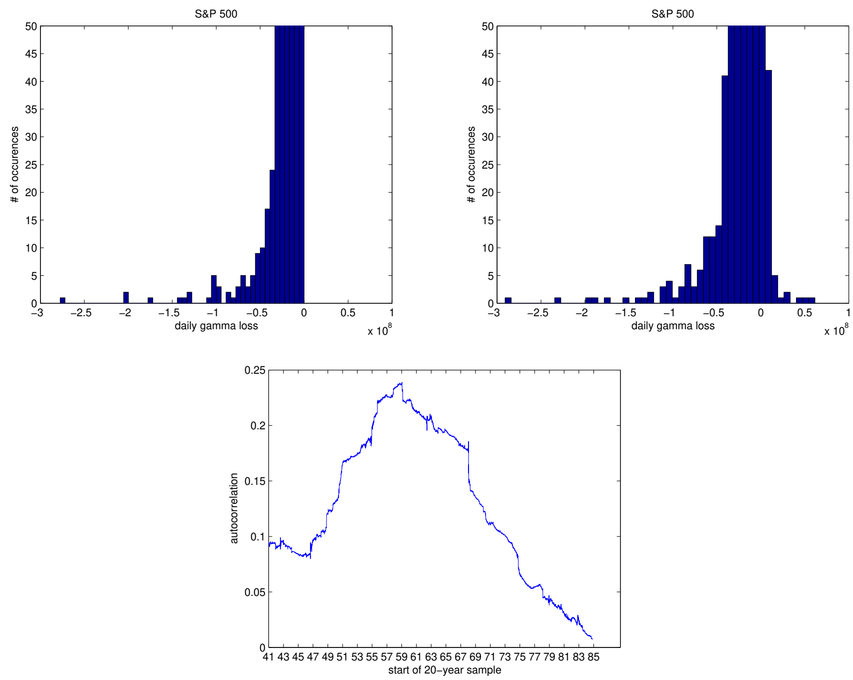

in the same or other direction, respectively. Thus in a market that exhibits reactions, operational imperfections turn out to be even beneficial. But in a market where trends tend to be continued over days, operational risk is detrimental. The right graph of Figure 14 indicates that the latter is often the case, as the autocorrelation of S&P 500 returns at a lag 1 is slightly positive. Thus the operational risk may often be accompanied by an unfavorable bias. The positive autocorrelation is further discussed below.

An imperfect hedge, as the one shown in the left graph of Figure 14, bears the risk of large losses if the stock price moves (in this case) up. On the other hand if it moves down, profit is realized. The perfect delta hedge in contrast (tangent to the call price line), is the only position, where either upwards or downwards movement of stock rice result in a loss, albeit a minimized loss. Less frequent hedging leads to less neutral positions in between hedging times and therefore to larger losses *and* gains.

In Figure 15 we show for comparison the loss distributions with and without operational risk. We see that peak losses are aggravated, but there are some

Figure 16: Loss ratios: Gamma Risk Insurance vs P&C Insurance

Loss ratios *plus one standard deviation of the loss distribution* for twenty years of gamma risk of delta-hedged option portfolios based on S&P 500. This graph is an average over option series of different times to maturity ($\tau = 1/4, 1/2, 1, 3/2$) and initial strike prices ($k = 0, \pm 0.5, \pm 1$). It should therefore represent best an actual picture of loss ratios that insurance companies would face. For comparison, a loss ratio curve representing the P&C insurance industry is added.

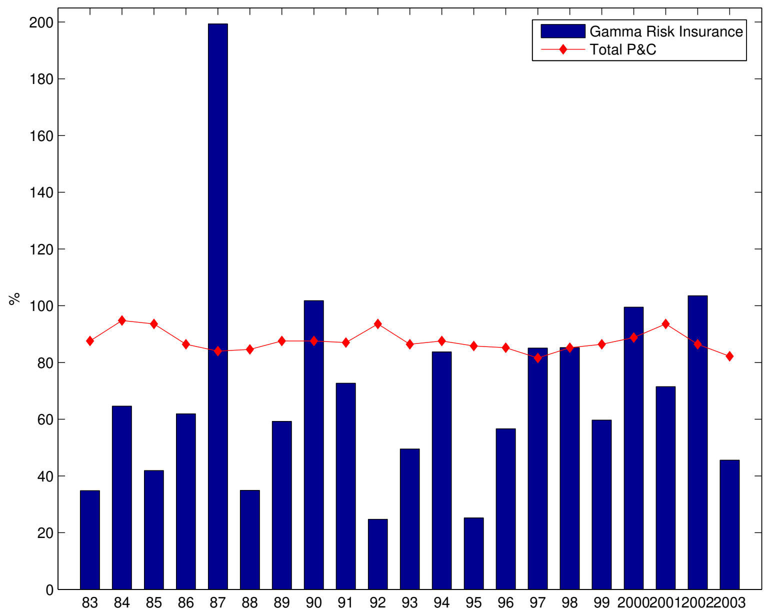

gains, too. On average, expected loss is increased by about 12%, expected shortfall by 18% and peak losses between 5 and 82%.

The positive autocorrelation of logarithmic returns as found in Figure 14 should not be expected and one could suspect data manipulation. To get a clearer picture we plot in Figure 15 the autocorrelation for a moving window of 5000 days (approximately 20 years). The surprising result is that there was indeed quite high autocorrelation roughly during the 1960s and 1970s, gradually going down to zero. This autocorrelation was exploitable in terms of profitable trading, so we interpret its fading as an increase in market efficiency. Periods of high autocorrelation directly correspond to increased operational risk.

7 Summary

In this paper we have analyzed in detail the properties of the so-called gamma risk arising from delta hedging written call options. Figure 16 summarizes our efforts. So far, we have treated different option series separately, but a more realistic approach is to allow for a portfolio with options of varying times to maturity and strike prices. We present results obtained from averaging over different option series on S&P 500. The average loss ratio plus one standard deviation is 107 (for the entire time range from 1941 until 2003 this figure is 94). An industry curve for total P&C is plotted for comparison, and it is apparent that this new kind of insurance is not correlated to the traditional ones. As a matter of fact, linear correlation is -0.24. All things considered, we think that although the tail index is with around 2 quite low, this financial risk insurance is a good diversification opportunity, with operational risk being low in recent years.

References

[Black and Scholes, 1973] **Black F. and Scholes M.**, 1973, *The pricing of option and corporate liabilities*, Journal of Political Economy, **81**, 637–659.

[Blum and Dacorogna, 2003] **Blum P. and Dacorogna M. M.**, 2003, *Extreme forex moves*, Risk Magazine, **16**(2), 63–66.

[Cox and Rubinstein, 1985] **Cox J. C. and Rubinstein M.**, 1985, *Options Markets*, Prentice-Hall, New Jersey.

[Dacorogna et al., 2001] **Dacorogna M. M., Gençay R., Müller U. A., Olsen R. B., and Pictet O. V.**, 2001, *An Introduction to High Frequency Finance*, Academic Press, San Diego,CA.

[Estrella, 1995] **Estrella A.**, 1995, *Taylor, black and scholes: Series approximation and risk management pitfalls*, Presented at Research Conference on Risk Management and Systemic Risk, Federal Reserve Board, Washington.

[Merton, 1973] **Merton R. C.**, 1973, *Theory of rational option pricing*, Bell Journal of Economics and Management Science, **4**, 141–183.