“Uncommon” Risks: Using the Probabilities of Extreme Events to Enhance Returns

Dr. Guillermo Bublik, CRO, LJM Partners, Ltd.

Introduction

The market crash of 1987 (“Black Monday”) confirmed that statistical models of the time underestimated the likelihood of extreme market movements. This prompted market participants to develop new paradigms designed to achieve a more accurate estimation of the probabilities of large market sell-offs. After Black Monday, investors increased their requirements for down-side insurance so that out-of-the-money (OTM) put options began trading at much higher levels of premium.

The crash of October 2008 reinforced the message that market practitioners and risk managers must base their strategies on probability distributions which accurately measure the likelihood of extreme events. In addition, this recent crisis highlighted the need for simulation techniques that realistically model market conditions during extreme events such as liquidity squeezes, correlations of large swaths of sectors quickly becoming 100%, and contagion effects when a high percentage of large firms hold most of the existing contracts of highly leveraged instruments.

Furthermore, widely used risk parameters such as Value-at-Risk (VaR) proved inadequate to gauge the monetary risks of portfolios of derivative instruments under the conditions unleashed by the financial crisis of 2008.

A thorough understanding of the risks from extreme market moves is essential to a realistic estimation of a strategy’s risk/reward profile. Every market professional is aware that with risk come opportunities. In what follows we show that a process can be built to use the heightened premiums of OTM S&P options to enhance returns.

Elevated Option Premium

Before the October 1987 crash, it was believed that the probability of very large market moves was negligibly small, and that equity markets behaved symmetrically such that rallies and sell-offs had reasonably similar chances of occurring.

With the realization that large negative gap moves were much more likely to occur than previously thought, demand for OTM puts increased dramatically. These contracts have remained "bid up" ever since.

The usual way to describe this effect is through the implied volatilities as a function of strike of all liquid contracts for a given expiration, as shown below for 3 week options on the front S&P 500 future:

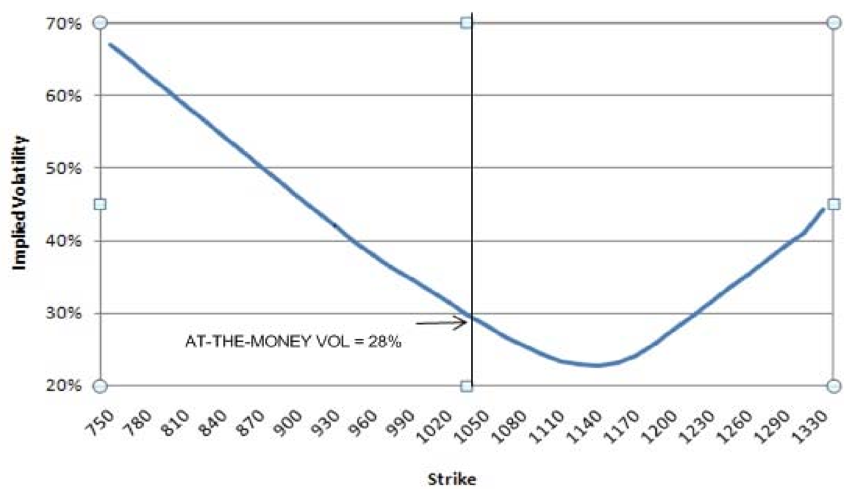

Notice that a put struck at $810 trades at an implied volatility of 60%, more than double the at-the-money (ATM) value. Hence, the manager of an options writing strategy has the opportunity to sell options quite far from the existing market level at a high enough premium to achieve very good returns as long as he hedges his exposure to an acceptable level.

A put option writer is selling insurance to someone holding a basket of long equities. Since markets can become very volatile, the seller needs to model risk accurately and either minimize risk relative to a target return or maximize return relative to a target level of risk.

Behavioral finance suggests that investors will pay a premium for insurance to mitigate losses; losing money is more painful than making money is pleasurable. This "skewness" suggests that, over time, disciplined option writer can be expected to generate alpha with risk limited to acceptable levels. To flesh this out, let us imagine that the implied volatilities for all strikes were 28%. A quick calculation would result in a premium of $1,000 for a put struck at $985 (~6% below the market). On the other hand, from the plot above, a put struck at $940 (~11% below the market) trades at an implied volatility of 41%, which results in a premium of $1,000.

Therefore, the seller of S&P put options can initiate a short position almost twice as far from the existing market level and still obtain the same option premium otherwise available were option volatilities to trade at a constant level. .

Before going on let's turn our attention to some of the tools used by LJM Partners, Ltd., to estimate the risks stemming from large market moves.

Value-at-Risk (VaR)

Value at risk is defined with respect to a confidence level and a time horizon. If the X% confidence level, Y time horizon VaR of a portfolio is given by V, the probability of suffering losses exceeding V over time Y is less than or equal to (100 – X)%. For instance, if the 1 month, 99% confidence level VaR is $10 million, the probability of incurring losses exceeding $10 million is less than or equal to 1%.

VaR is the most widely accepted risk metric implemented throughout the financial industry, from firm-wide risk calculations to regulatory reports on which capital requirements are based.

However, the VaR metric is not fail-safe. Because VaR only uses 99% of the probability distribution (it ignores the 1% scenarios where market crashes live!), creative traders can actually move risk to the tails, dramatically increasing risk, without having an effect on the VaR level.

Furthermore, there exist trades than can be used to lower the VaR of a portfolio without actually lowering the risk.

Additionally, this risk parameter is not "sub-additive", i.e., VaR of the sum of 2 portfolios may be greater than the sum of their individual VaR's. This issue is not just a mathematical curiosity but can have a real life effect, as diversification, believed to lower the risk, may in some cases result in higher VaR levels.

Expected Shortfall

A concept that is gaining traction among risk managers is that of Expected Shortfall ("ES"), which considers the entire probability curve. Coupled with a distribution that provides "fat tails", ES provides a risk manager capable tool capable of more accurately modeling the effects of extreme events such as market crashes.

If for a given portfolio, the X% confidence interval, Y time horizon ES is given by E, the "expected (average)" loss if the worst (100-X)% outcomes are realized equals E.

For instance, if the 99% confidence level, 1 week ES of a portfolio is $10 million, the average loss of the worst 1% outcomes equals $10 million.

If an investment portfolio is such that the worst losses would be incurred during sharp market sell-offs, the expected shortfall, ES, will be the average of the losses in the red portion of the distribution as shown below.

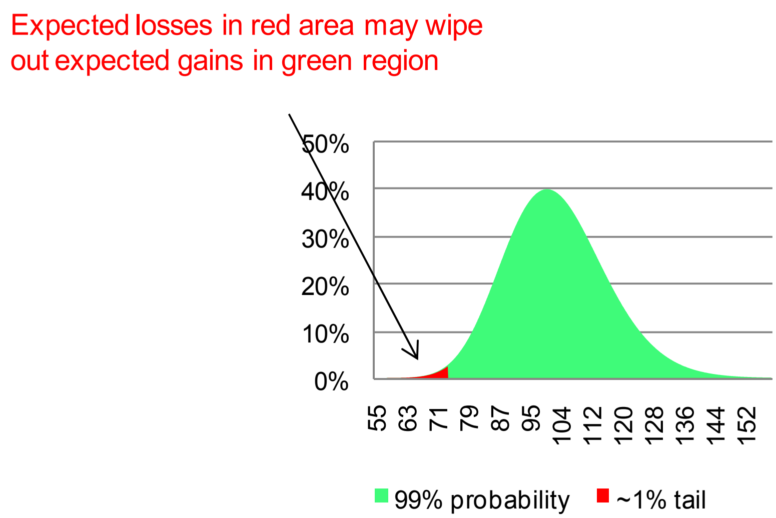

It is very important to have a good estimate of ES to make sure that our positions are not such that the value of the expected losses in dramatic market moves (red tail) exceeds the expected profits under "normal" market conditions.

The estimation of ES may involve time intensive calculations, as it may not be clear prior to performing the actual computations what the worst 1% outcomes are.

Expected shortfall is sometimes called Conditional VaR, since it is a measure of the average losses given that a given percentage of worst case scenarios are realized. Unlike VaR, it can be proven that ES is sub-additive.

Those interested in a numerical example of ES are referred to the Appendix. We will return to ES later on when we discuss the details of an actual strategy.

Market fitting

Another component of a sound risk management system is a probability distribution that reliably captures the likelihood of extreme events.

LJM Partners, Ltd. (“LJM”) has invested heavily to develop the **LJM STORM^{SM} System** proprietary technology that fits the historical prices of the S&P 500 index with an asymmetrical, “fat-tailed” distribution. Besides providing a very good fit of the market as shown below, the model incorporates the effect of volatility clustering: highly volatile days are more likely to be followed by additional days of higher volatility rather than by quiet days. This is a key feature necessary to describe the transitions from quiet(volatile) to volatile(quiet) days and vice-versa.

In the plot below we compare a market fit obtained by **LJM STORM^{SM} System** (red line) to that using a Gaussian (“Bell”) curve shaped distribution (green line). The actual market moves are depicted by the blue dots. The y axis represents the natural logarithm of the probability distribution while the x axis is the natural logarithm of the market moves.

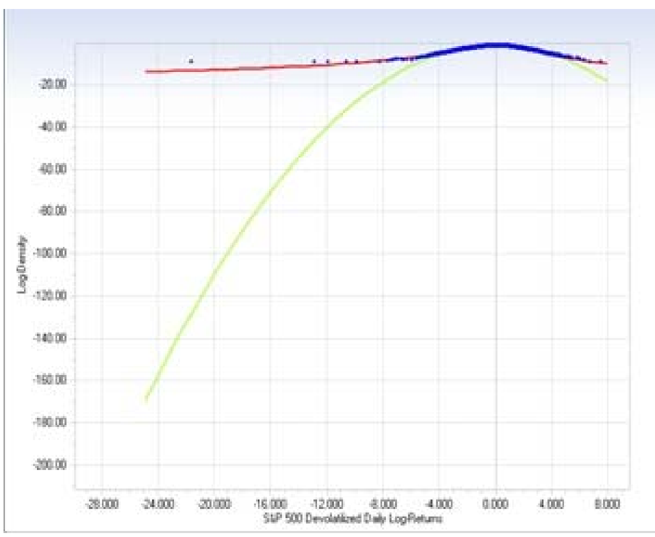

The green line does a good job for small moves but it assigns a negligibly small probability to extreme market sell-offs. In contrast, the red line generated by the **LJM STORM^{SM} System** estimates the likelihood of such events very accurately.

In order to highlight the magnitude of the differences between these distributions we calculated the probability of having at least one intra-day move of -10% since the beginning of 1950. The LJM STORM^{SM} distribution assigns a probability equal or greater than 90% to such an event. In contrast, the green distribution resulting from a Bell shaped distribution estimates the probability to be approximately 0.0003%. Anyone using software based on Gaussian probabilities will invariably feel inordinately unlucky as extreme market moves seem to occur with much greater regularity than their models would otherwise suggest!

How sensitive is ES to the accuracy of the underlying probability distribution? We calculated its value for a very simple portfolio of 100 put contracts with 1 month to expiration and a strike 22% below the market.

Our computations show that the Gaussian (green) distribution results in ES (green line) = 75% of the initial value of the portfolio, whereas the red distribution obtains an ES = 125% of the initial value.

Using Gaussian distributions can lull an investor into a false sense of security: The ES resulting from the fat-tailed distribution is 66% higher than the value obtained using a Gaussian model. Furthermore, the ES calculated with the red distribution tells us that we would go debit were the worst 1% of outcomes (i.e. market crashes) realized. The Gaussian distribution reflected by the green line,on the other hand, gives us a false sense that we would probably be able to survive a crash event.

Monte Carlo Simulations

Once a risk manager is armed with probability densities that correctly capture the likelihood of unusual, large market moves, an estimation of the risk profile is done by using these probabilities to generate a multitude of market outcomes and to compute the profit and loss under each of these scenarios.

For these simulations to be relevant not only do the probability distributions need to be fairly accurate, but the parameters must realistically describe all market conditions.

The last point cannot be overemphasized. When analyzing the consequences of extreme events a risk process must stretch the market behavior beyond the relationships that usually hold. Examples of this are the correlations of large swaths of sectors tending to 100% during such events and the inability to transact out of certain instruments due to liquidity squeezes.

As shown during the crisis of 2008, a risk manager should also pay close attention to the risks of having a high concentration of large institutions holding similar, highly leveraged positions.

Harnessing risk to enhance returns

As stated at the beginning of this note, out-of-the-money puts trade at implied volatilities much higher than those prevalent for at-the-money contracts.

Given the behavioral aspects that investors tend to be risk averse, is it possible for a premium writer such as LJM to capture the heightened premium to achieve higher returns by selling these options while still managing risk to finite and acceptable levels?

A common error that portfolio managers (PM) may make at times is to use different assumptions for estimating expected returns from those used to calculate risks. Practitioners who accurately identify the potential for enhanced returns stemming from exposures to extreme events may find themselves using risk management processes that under-estimate risks by sacrificing accuracy for speed.

Keeping this in mind, one can build a system that accurately estimates the risks from tail events and design a program to sell the heightened premium of OTM options while containing the overall risk levels.

Since the strategy accrues value while the short options decay as they approach expiration, it is crucial for the risk management system to estimate the expected profit and loss profile to multiple time horizons.

Additionally, in today's fast moving markets risk reports can become obsolete very quickly, so the system must be capable of computing the risk profile in real time to provide the PM with the necessary information to hedge.

The steps and calculations that are necessary for a successful option writing program are:

1. Fit the market with an asymmetric, fat-tailed distribution that incorporates volatility clustering and accurately estimates the likelihood of large market sell-offs
2. Run Monte Carlo simulations to calculate the probability of a range of market moves over time intervals of: intra-day, 1, 2, and 3 weeks
3. In real-time, as the prices move, using the results from the Monte Carlo simulations, calculate the profit and loss profile for a wide range of market levels and the 4 time intervals mentioned above.
4. Calculate the Expected Shortfall to ensure extreme events risk is within strategy's limits
5. Calculate the expected value and standard deviation of the profit and loss to check that winning scenarios constitute a majority of possible outcomes.

Below we show an example of a risk report for an option writing fund targeting returns of 8-12% per year. The strategy includes a high level of downside risk hedging, accomplished by buying OTM put options.

The screen shot displays the expected profit and loss of the positions for a number of market moves over a period of 1 week. The report was obtained in mid May of 2009, when the short option positions included May, June and July option contracts.

Market Move	-22.73%	-20.63%	-18.57%	-16.55%	-14.57%	-12.63%	-10.72%	-8.85%	-7.02%	-5.22%	-3.45%	-1.71%	0.00%	1.68%	3.33%
TOTAL	-6.189%	-4.532%	-3.124%	-1.944%	-0.968%	-0.172%	0.464%	0.959%	1.331%	1.589%	1.738%	1.771%	1.664%	0.500%	-0.500%
MAY-09	0.063%	0.063%	0.063%	0.063%	0.063%	0.063%	0.063%	0.063%	0.063%	0.063%	0.063%	0.063%	0.063%	0.063%	0.063%
JUN-09	-2.634%	-1.773%	-1.064%	-0.487%	-0.027%	0.333%	0.607%	0.809%	0.952%	1.042%	1.084%	1.072%	0.991%	0.809%	0.469%
JUL-09	-3.618%	-2.821%	-2.123%	-1.519%	-1.003%	-0.568%	-0.206%	0.087%	0.316%	0.484%	0.591%	0.636%	0.610%	0.501%	0.286%
Cum Prob	0.022%	0.024%	0.030%	0.039%	0.051%	0.084%	0.151%	0.380%	1.098%	3.845%	11.003%	27.234%	50.373%	73.135%	88.839%
P&L x Prob	-0.0014%	-0.0001%	-0.0002%	-0.0002%	-0.0001%	-0.0001%	0.0003%	0.0022%	0.0096%	0.0437%	0.1244%	0.2874%	0.3851%	0.1138%	-0.0785%
Expected P&	0.7064%														
Stdev EP&L	1.1668%														
ES	-2.7791%														

The green(pink) background represent gains(losses), while the yellow cells display the probability and probability times profit/loss as calculated using an asymmetric, fat-tailed distribution.

First off note that ES = -2.78%. Recall that this is the average loss that would be sustained if the worst 1% of scenarios were realized. At a level of a quarter of the target returns, this is considered to be a reasonable balance of risk and rewards.

In addition, the expected value and standard deviation show that over 70% of outcomes (1-week market moves) will result in gains.

For the sake of comparison, VaR for this portfolio is 1.9%, quite lower than the more relevant ES.

Conclusions

We have shown how one would design a strategy that harnesses the risk of extreme events to enhance returns. We have detailed how the success of such a program must use probabilistic techniques that accurately describe tail events.

LJM Partners harnesses this expertise to manage risk to capture the excess returns priced into OTM S&P put and call options.

Appendix – An Illustration of Expected Shortfall

Let us analyze a game where we pay $1,000 for the chance to draw a ball from a fishbowl containing 1,000 balls, numbered 1 to 1,000. Each time the ball is returned to the fishbowl before drawing again. The payout depends on the number drawn as follows:

Number drawn	Receive	Profit(Loss)
1 to 500	$1,005	$5
501 to 800	$1,010	$10
801 to 900	$1,015	$15
901 to 980	$1,000	$0
981 to 990	$995	($5)
991 to 995	$750	($250)
996 to 999	$500	($500)
1,000	-$3,000	($4,000)

A calculation of 99% confidence level VaR would ignore the outcomes displayed in red, as they comprise the worst 1% outcomes and would result in a value of $5. Losses are displayed in parentheses.

We can quickly compute ES to obtain

$$\text{ES} = \frac{0.5\% * (-\$250) + 0.4\% * (-\$500) + 0.1\% * (-4000)}{1\%} = -\$725$$

So in this made up example the expected shortfall is 145 times the value at risk.

We can calculate the expected profit and loss, EP&L, by performing the same probability weighed outcomes to obtain EP&L = -$0.30. This underscores the fact that, were to draw 1,000 times, we would expect to lose $300. The risks outweigh the rewards.