



Research

COMMENT

Portfolio Strategy & Quantitative Research

Black swan dive: fat tails, stochastic volatility and 10-sigma corrections

STRATEGYPLUS

RESEARCH COMMENTS

PORTFOLIO STRATEGY &
QUANTITATIVE RESEARCH


BLACK SWAN DIVE



Peter Gibson,
Vice-Chairman
(416) 867-3774



Ed Sollbach, CFA, PEng
(416) 867-6135



Jeff Evans, CFA
(416) 867-3776

- *In the past month, investors have observed two events so rare that, when last witnessed, Franklin Roosevelt was midway through his second term as President of the United States—only eight single-day rallies in the S&P 500 have exceeded 10% since 1928, with the 11.58% gain on October 13 representing a ‘10.06-sigma’ event, and the 10.79% rally on October 28 being a similarly unthinkable 9.34 sigmas*
- *As rare as these events may seem, they are not nearly rare enough (at least according to the normal distribution). To put the recent experience in context, daily returns exceeding 7.5 standard deviations should only occur roughly once every 33 trillion trading days; to have observed even a single such event, the universe would have to have been approximately 10x older than it actually is!*
- *However, if returns are distributed according to a fat-tailed distribution (instead of the normal distribution), extreme events become far more likely. Moreover, if volatility is random (but mean-reverting in the long run), investors should expect extreme returns to ‘cluster’, as seen in the past month. In a stochastic volatility world, with instantaneous volatility of 60%, the 11.58% return on October 13 was only a 3-sigma event—rare, but far more likely than traditional theories would suggest*
- *For those investors less familiar with fat tails and/or stochastic volatility concepts, this article provides a brief introduction and discusses scenarios where ‘conventional’ investment theories fail. Most importantly, there has been significant debate regarding risk models such as value at risk (VaR) in recent months, with some commentators arguing for the complete elimination of VaR from Basel II. A similar scenario arose in the 1950s and 1960s when civil engineers believed it was possible to predict bridge loadings with high precision: safety margins were reduced substantially (to save on costs and materials), yet when actual loads exceeded expectations, structural failures occurred and several bridges collapsed. The error in finance has not been excessive reliance on VaR models, but in assuming that VaR models had far greater precision than that of which they were actually capable. VaR and similar risk measures will continue to have an important role in successful investment management, but risk managers should not rely on estimates of the 99th percentile: in the presence of fat tails, estimating the 90th percentile and applying a safety-margin multiplier should prove to be more robust*

In the past month, investors have observed two events so rare that, when last witnessed, Franklin Roosevelt was President of the United States. In total, only eight single-day rallies in the S&P 500 index have exceeded 10% since 1928; the 11.58% rally on October 13, 2008 represented an extremely rare '10.06-sigma' event, while the October 28, 2008 10.79% rally was a similarly unthinkable 9.34 sigmas.

As rare as these events may seem, they are not nearly rare enough (at least according to the normal distribution). To put the recent experience in context, the universe is generally assumed to be somewhere between 13.5–14b years old, equivalent to approximately 3.5775 trillion trading days. Had stocks traded since the beginning of time, the likelihood of observing a single 7.5-standard-deviation (or greater) one-day return would be roughly 1 in 33 trillion: the universe would need to be almost 10x older than it likely is for there to be any reasonable chance of observing such an extreme event!

The fact that two such events occurred in the past month alone suggests that something is wrong...and indeed, something certainly is: the normal distribution and the assumption of constant volatility.

Historically, it was common practice among academics to assume that stock returns were inherently unpredictable, governed by a mathematical process referred to as 'geometric Brownian motion' (more commonly known as a random walk). Under this assumption, changes in stock prices are fully described using the normal distribution; each day's return is random and independent of returns on all other days. This assumption was employed for its simplicity: the mathematics behind portfolio and risk management theories becomes greatly simplified if the normal distribution can be assumed.

The collapse of Long-Term Capital Management (LTCM) made clear, however, that the normal distribution did not represent an adequate model of stock prices. Subsequently, models permitting 'fat tails' and 'stochastic volatility' have become more popular, though these concepts often remain inaccessible due to the lack of suitable tools (software, analytical tools and so forth). Given recent events, and the fact that the VIX continues to trade in the 50%+ range, it is worth discussing why 10-sigma events can occur with far greater frequency than one might expect (or prefer).

In a normal world...

It remains quite popular to model stock returns using the normal distribution—for example, anyone who calculates the average daily (or monthly or yearly) return on a stock, or the standard deviation of stock returns, is implicitly assuming that stock returns can be described using a normal distribution; this occurs in a surprisingly large number of contexts: mean/variance optimization, linear regression (CAPM and beta, or other linear multi-factor models), first-generation value-at-risk models, and so forth.


When stock returns are said to 'follow a normal distribution', it means that if one calculated the daily percentage returns and analyzed the frequency at which returns of different magnitudes occurred, the frequencies observed in actual returns would be identical¹ to those predicted by a normal distribution adjusted to the same average return and volatility.

The following chart plots the normal distribution calibrated to the daily returns of the S&P 500 index (eg the normal distribution is adjusted to have the same mean and standard deviation as the observed stock returns). The average daily return since 1928 was approximately 0.02539%, with a standard deviation of 1.1548%.

According to statistical theory, approximately 68% of all observations should fall between -1.12941% (average minus standard deviation) and +1.1802% (average plus standard deviation), assuming that the observations are drawn from a normal distribution.

However, approximately 81.8% of all S&P 500 daily returns lie **within** one standard deviation of the average, far more than one would expect if stock returns were distributed according to a normal distribution. Similarly, there are also more extreme/tail events than normal distribution predicts, but fewer 'intermediate' returns.

¹ With allowances for sampling error



For S&P 500 returns, the normal distribution implies that daily returns larger than $\pm 4\%$ should almost never occur, and extreme moves (such as single-day returns exceeding $\pm 10\%$) **should not have occurred at all** given the current age of the universe. Given that they do, the normal distribution is clearly incapable of adequately explaining the frequency of stock returns, and one must search for a distribution offering a better 'fit' to the observed data.

Alternative distributions and 'fitting/calibrating'


The previous chart plots the normal distribution, 'calibrated' or 'fit' to S&P 500 returns. The normal distribution is a mathematical function with two parameters, the mean and the standard deviation. To 'fit' the normal distribution to a particular dataset, one calculates the mean daily return and the standard deviation of daily returns, and substitutes these into the normal distribution equation. **The goal of calibration/fitting is to find the mathematical function that most accurately predicts the frequency of observed stock returns.**

In theory, any mathematical function can be used as a probability distribution, provided that all values are positive and that the integral (from negative infinity to positive infinity) equals 1.0. The 'normal distribution' is one function satisfying these criteria, and moreover, it is easily described using only primitive mathematical functions (explaining its popularity amongst mathematicians and model-building economists).

As noted, however, the normal distribution mispredicts the likelihood of stock returns, under-predicting the frequency of small moves and extreme moves, while over-predicting the frequency of 'medium' returns—therefore, using a different function for the probability density will likely improve the fit.

One alternative is the 'alpha stable' distribution. This distribution is related to the normal distribution (and, in certain cases, is identical to the normal distribution), but adds two parameters beyond the mean and standard deviation. The 'alpha' parameter determines the amount of probability associated with extreme tail events, while the 'beta' parameter allows for positively or negatively skewed distributions. If data is normally distributed, the alpha parameter will equal 2.0—smaller alpha values imply increasingly large-tail events, while alphas below 1.0 shift so much weight to the tails that variance (volatility) becomes infinite! (a typical value for stock returns is roughly 1.5).

The following chart plots the normal distribution and the alpha-stable distribution. In both cases, the distributions are calibrated to S&P 500 daily return data. Due to the alpha and beta parameters, calibrating the alpha-stable distribution is more complex than simply calculating the mean and standard deviation—a technique known as maximum likelihood must be employed to estimate parameters (using an optimization routine to identify the alpha, beta, mean and standard deviations that, when substituted into the alpha-stable distribution function, produce the closest match to observed data).

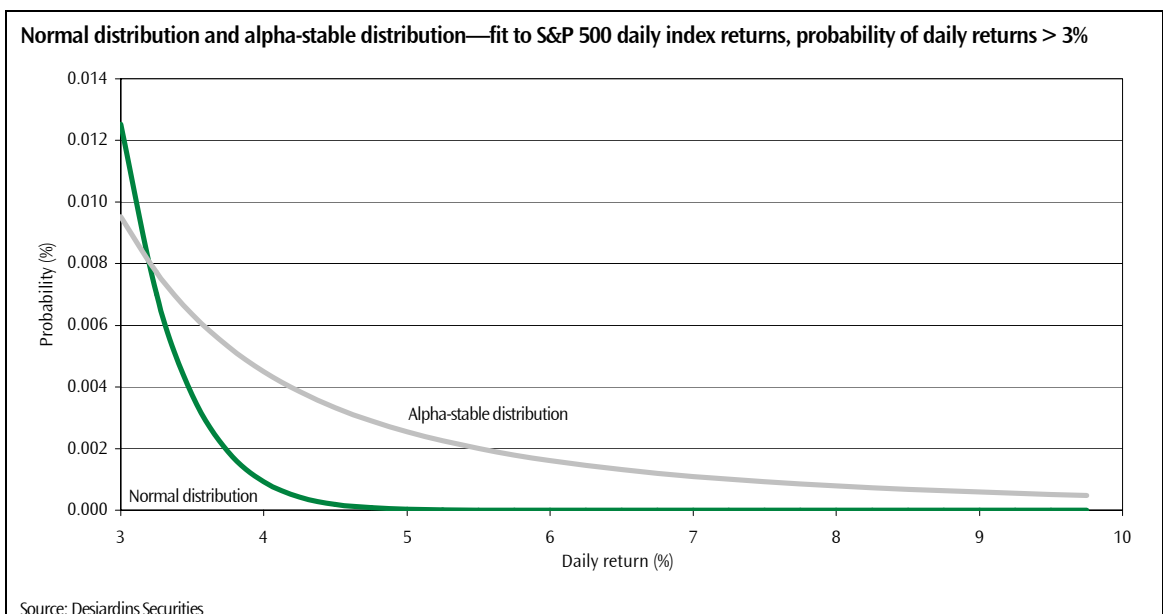


When a 10-sigma event occurs, the standard deviation increases to reflect the extreme event. However, even with an inflated standard deviation, the normal distribution will continue to under-predict the frequency of large moves! With only two parameters (mean and standard deviation), the normal distribution is insufficiently flexible to accommodate infrequent events.

In contrast, the alpha-stable distribution is more capable of ‘understanding’ extreme moves. As the chart above suggests, returns close to the average become more likely, as do extreme events in the distant tails. Events in the middle become less likely, which matches our observation that, for the S&P 500, extreme events are more likely than predicted, as are events within one standard deviation.

(To give an analogy, consider the difference between an off-the-shelf and a bespoke suit. The bespoke suit can be tailored to perfectly match the customer, creating an excellent fit. By contrast, off-the-shelf clothing must be designed to accommodate a wide range of body types, and cannot therefore produce the same ‘perfect fit’. Since the alpha-stable distribution has four parameters, compared with only two for the normal distribution, it can be ‘tailored’ to more closely match the observed data.)

It is difficult to observe tail behaviour, as both distributions fall to zero relatively quickly: as such, the following chart zooms in on the normal and alpha-stable distribution for daily returns of 3% or greater. As one can clearly see in the chart, the normal distribution very quickly falls to zero, while the alpha-stable distribution declines far less quickly.



Unfortunately, while the alpha-stable distribution makes tail events far more likely than the normal distribution, it actually **over-predicts the frequency of rare events** (at least compared with data actually observed in S&P 500 daily returns since 1928). For example, under the normal distribution, daily returns exceeding 9.25% are so unlikely that our computer is incapable of calculating the probability. In contrast, the alpha-stable distribution predicts that daily returns of 9.25% or greater should occur approximately once every 13 months—in other words, this should have occurred approximately 75 times since 1928, but has actually only been observed on 10 occasions.

Unfortunately, some arcane mathematics are necessary to describe this phenomenon (described in the footnote below²) In brief, when considering probability tails, the normal is ‘not big enough’, the alpha-stable is ‘too big’, and in actual practice, analysts must use a hybrid to find a match that is ‘just right’. One candidate proposed in the econophysics literature is the *truncated Levy distribution*. In that distribution, the central part of the distribution resembles an alpha-stable, but the extreme tails are either truncated entirely (eg the probability of very, very large events is set to zero), or the tails are replaced using an exponential function which decays to zero faster than the alpha-stable distribution but slower than the normal distribution. Ultimately, this allows for far more extreme daily returns, but when the distribution is ‘convolved’ to produce monthly or annual returns, the truncated tails ensure that these returns more closely resemble that of the normal distribution. Unfortunately, alpha-stable and truncated Levy distributions cannot be calculated without complicated mathematical techniques (inverse Fourier transforms, or alternatively, quadrature methods optimized for oscillatory functions), and have thus received only limited attention outside of the econophysics community.³

While the truncated Levy distribution provides a far more accurate fit than the normal distribution, a major flaw is that extreme events remain independent and uncorrelated. This is problematic, because actual stock returns tend to exhibit a ‘volatility clustering’ effect, where stock returns exhibit periods of extended calm, followed by periods of extreme up/down moves (such as the large 10% daily moves observed during the recent financial crisis).

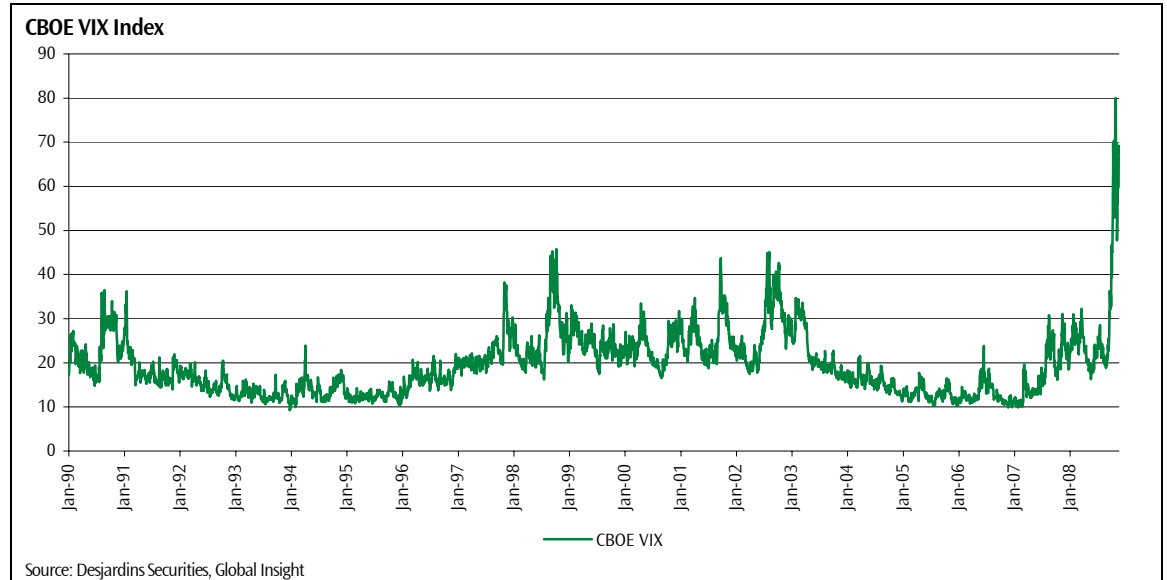
Ultimately, alpha-stable (or preferably, truncated Levy) distributions will far more accurately predict the **frequency** of daily stock returns vis-à-vis the normal distribution. However, none of these distributions will accurately replicate the **pattern** of returns actually experienced. **Simply put, fat tails are not sufficient to describe the behaviour of stock returns, and one must turn to more advanced models.**

² Under repeated convolution of the probability density, the convolved density will approach that of the normal. ‘Convolution’ can be thought of as a mathematical process which converts the distribution of daily stock returns into weekly or monthly stock returns—under the central limit theorem, stock returns should eventually resemble a normal distribution, but it might be necessary to use annual returns instead of daily returns. This is, in fact, what one observes in the data. The implication is that the normal distribution is not accurate for predicting daily returns, but is reasonable when applied to annual data. Unfortunately, if one uses the alpha-stable distribution, no amount of convolution will ever produce a normal distribution—in other words, annual returns in an alpha-stable world will show just as many extreme moves as daily returns.

³ While truncated Levy distributions have received limited attention, the Merton-style jump-diffusion models have received greater attention in the quantitative modeling community. In a jump diffusion, stock prices are assumed to follow a geometric Brownian motion most of the time (that is, daily returns are drawn from a normal distribution), but extreme ‘jumps’ are occasionally added to these returns at random. If one was unaware of the jump process, and assumed that daily returns came from a single distribution, it would resemble similar dynamics to a scenario where one simply drew from a truncated Levy distribution.

Stochastic volatility

As anyone who follows the CBOE VIX index will recognize, volatility is not constant; instead, volatility changes dramatically over time, with stock prices being extremely stable in some periods (a low VIX), while at other times being extraordinarily volatile.

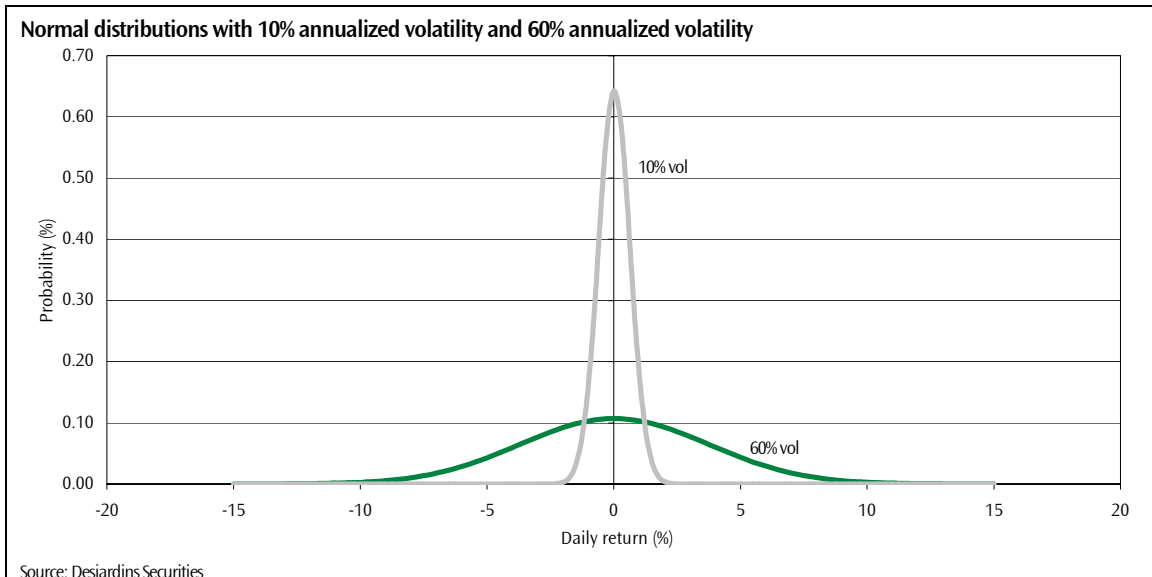


It is tempting to ignore short-term volatility fluctuations and instead calculate a ‘long-run’ volatility estimate, especially in contexts such as asset allocation, which are inherently ‘long term’ in nature. For example, using daily returns from 1928–2008, the S&P 500 index exhibited daily volatility of 1.157%. There are two flaws with this perspective: first, it assumes that stock returns are drawn from a distribution with constant volatility. If this is the case, stock returns clearly cannot be normally distributed, as there is simply too little probability in the tails to make these events possible.

Secondly, one would not expect extreme events to occur in ‘clusters’ in a constant-volatility environment—a single large move should occur in isolation, and it would be very, very unlikely to observe another large move in close proximity. **Yet, extreme stock returns tend to coincide with other extreme returns** (eg October/November 2008, where it was not uncommon to observe +6% days followed -6% days).

One resolution to this problem is stochastic volatility—in other words, volatility is assumed to vary randomly over time (oscillating around some long-term mean, but never settling down or becoming constant). In effect, we can assume once again that daily stock returns are normally distributed, but the volatility of this distribution changes slowly from day to day. When this volatility is high, returns will tend to be more extreme, and vice versa when volatility settles down.

The following chart plots two normal distributions, with the same average return but different standard deviations.



Clearly, the distribution calculated using a 60% annualized volatility has a far greater likelihood of extreme moves—in a 60% annualized volatility world, one would expect roughly one in every 275 days to exhibit returns of 10% or greater, whereas the probability would be almost zero if volatility was only 10% annualized.

Of course, this still understates the likelihood of tail events—from October 6, 2008–November 14, 2008, the actual frequency of ‘extreme tail events’ has been closer to one in 10 days. **In practice, one likely needs to combine both stochastic volatility and fat tails.**

The key point is that, if one erroneously assumes a single long-term volatility, the multitude of ‘tail events’ observed in the last month appear to be significant and abnormal outliers. Using the 80 years of daily data from 1928, the 11.58% return on October 13, 2008 would represent a 10+ sigma event, so unlikely in a ‘normal distribution world’ that such a day should not have occurred since the beginning of time. However, if volatility is stochastic, the same 11.58% return would only represent a 3.1-sigma event, and if fat tails are added, such ‘extreme’ moves become even more likely (if volatility is 60% annualized, as the VIX suggested, this implies an instantaneous daily volatility of 3.72%).

The Heston model

Stochastic volatility models are not new—one of the first such models was published in 1993 and became known as the Heston model. Under this model, stock returns (loosely speaking) continue to follow a geometric Brownian motion, with normally distributed day-to-day changes.

However, the distribution from which daily returns are drawn is allowed to change over time. Specifically, while the average daily return remains the same, the volatility fluctuates. Under the Heston approach, the ‘volatility process’ is modelled first: one assumes that volatility randomly changes over time, but it tends to revert toward a long-term average. This produces periods of high volatility, followed by periods of low volatility.

Once the volatility for a given day is determined, the actual return for that day is selected from a normal distribution having a volatility equal to the current value of the volatility process. A final and important feature of the Heston model is that volatility and asset returns are allowed to be correlated (or inversely correlated), matching another important feature observed in stock returns: historically, volatility of stock prices tends to be inversely correlated with their returns, such that as stock prices rise, volatility declines, and vice versa.

The Heston model has received considerable attention in the academic literature, and its popularity can be traced, in part, to the fact that closed-form analytic solutions exist (for instance, one can often calculate option prices under the Heston model without having to simulate the underlying dynamics—this is very attractive in high-frequency trading situations where traders often cannot wait for the seconds or minutes necessary to complete a Monte Carlo run). However, while a significant advance, Heston continues to improperly replicate the dynamics actually observed in asset prices, and yet more advanced stochastic volatility models have been proposed (such as SABR, among others).

Why is the structure of volatility relevant?

Volatility patterns are not simply of academic interest: risk management (value at risk and related variations), options pricing, and arbitrage and asset allocation techniques are directly influenced by 'fat tails' and stochastic volatility.

Firstly, consider asset allocation techniques. Asset allocation assumes that 'similar assets have similar performance', and that performance is driven predominantly by exposure to broad asset groups. If the behaviour of an asset class remains constant over time, investors can successfully target desired risk/return objectives through static allocations, rebalancing to fixed ratios as necessary. However, if volatility is stochastic, investors will only achieve the desired risk/return objectives in the (very) long run, resulting in sub-optimal portfolio holdings during interim periods.

For example, assume an investor desires a portfolio with 18% annualized volatility (approximately the 'long run' volatility of stocks, from 1928–2008). One option is simply to hold 100% stocks: with a sufficiently long holding period, the investor will likely experience this volatility 'on average' over time. However, this will include periods of extremely low volatility and extremely high volatility, such that *rolling volatility* in any specific sub-period is uncertain. Alternatively, one might dynamically adjust equity exposure, such that the portfolio is leveraged in periods of low volatility, but shifted toward cash in periods of high volatility—the result is a portfolio with a rolling volatility far closer to the desired 18% annualized target⁴. As a broad generalization, asset allocation models assuming constant volatility will tend to over-allocate to assets with high vol-of-vol: if investors are risk averse to unexpected developments in volatility, the static optimal portfolio should reduce its allocation to asset classes with a history of high vol-of-vol.

Options pricing and arbitrage are also particularly sensitive to assumptions about volatility patterns. For instance, fat tails alone cannot explain the behaviour observed in the volatility surface (implied volatility plotted against strike price and time to expiry), and stochastic volatility (and the correlation between volatility and the underlying) is generally necessary to explain volatility skews, at least in part. For instance, a naïve volatility arbitrage strategy is to sell out-of-the-money puts, as these options trade at significantly higher Black-Scholes implied volatilities vis-à-vis at-the-money options. It would seem reasonable that a properly constructed ratio trade would sell 'overpriced' OTM puts and hedge with at-the-moneys to capture an arbitrage profit.

However, if OTM options trade at higher Black-Scholes volatilities due to stochastic volatility (which makes extreme moves more likely) and correlation between volatility and the underlying (which makes it more likely to observe low prices in periods of high volatility), then it becomes quite possible that, despite their high Black-Scholes implied volatilities, OTM options may actually be *underpriced* if the market misprices vol-of-vol or asset/volatility correlation. Without structure imposed by a more advanced model, these options will almost always tend to appear 'rich', leading to improper arbitrages.

Finally, consider risk management. When properly structured, risk management has a similar function to that of asset allocation—investors decide on a desired risk/return level, and risk management ensures that the portfolio's performance tracks these objectives while reducing or eliminating exposure to unanticipated shocks to the greatest extent possible. The distinction is that asset allocation is typically 'static' (allocations do not change much over time), whereas risk management is an inherently dynamic activity. Value at risk (VaR) has become the industry standard, although a number of related risk measures are often considered alongside VaR (extreme shortfall, drawdown risk and so forth). Risk management typically has a short-term focus: current risk exposure is estimated and compared with the desired target, with exposure then adjusted accordingly. This leads to an ongoing process of short-term prediction/revision, and the underlying premise is typically to minimize forecast error.

When first implemented, VaR models relied on the normal distribution as a simplifying assumption, to create easily evaluated 'analytic formulas' that could be evaluated without extensive computational power. Unfortunately, when fat tails and stochastic volatility are present, it is far more challenging to find easily evaluated analytic formulas, and investors often must rely on simulation techniques. This involves the creation of thousands of 'simulated' stocks, where the statistical parameters of the simulated stock match those of an actual investment. Simulated holdings are then rolled up into simulated portfolios, and the returns on these portfolios are analyzed to determine risk exposure.

Of course, creating simulated stocks requires assumptions about asset price dynamics, and this is where alternatives to geometric Brownian motion are most helpful—for instance, if simulated stock returns are generated by a GBM process, VaR estimates will be significantly offside as they will fail to anticipate sudden, significant shocks. VaR models have received significant negative attention in recent months, especially from Nassim Taleb (who argues for the complete elimination of the

⁴ Whether or not this is useful, however, remains an open question. In principle, investors allocate assets because of risk aversion: the investor maximizes utility by holding a portfolio maximizing return for their desired risk tolerance. However, if volatility is stochastic, an investor cannot select the optimal portfolio, as there is inherent uncertainty about returns and volatility; it seems reasonable that risk-averse investors should also be risk averse with respect to portfolio allocation risk (eg the risk of selecting a portfolio which fails to meet expectations), and there should exist positive demand for products that transfer these risks to others.

VaR approach). In our view, this is somewhat extreme: the VaR approach remains extremely useful, especially for large portfolios holding several offsetting positions or products with complex payoffs (derivatives)—it is simply unrealistic to expect that such portfolios can be risk managed using non-quantitative rules of thumb.

With that said, however, the Basel VaR approach is fundamentally flawed. Under Basel II, financial institutions must calculate VaR using a 99% confidence interval. However, as Recardo Rebonato explains in thorough detail in his September 2007 book, *Plight of the Fortune Tellers: Why we need to manage financial risk differently*, it is virtually impossible to actually measure the 99th percentile accurately: for most portfolios, it would require several years (if not decades) of daily returns to properly measure the co-variance structure, yet due to parameter instability, it becomes a meaningless exercise: by the time risk managers have sufficient data to measure risk, the instantaneous risk exposure has changed. Ironically, Basel II imposes a one-year maximum lookback: even if parameters were stable over time, it would be next to impossible to reliably estimate the 99th percentile with such a short dataset.

Taleb suggests that this justifies eliminating VaR altogether—however, parallels should be drawn from engineering practices. For example, the Brooklyn Bridge is often referred to as being ‘notoriously over-engineered’: since the original engineers had no way of producing detailed quantitative estimates of expected live loads (or growth in load weights over time), they made a ‘best guess’ as to the required load-bearing capacity, and then multiplied that capacity several times to build in the necessary safety margin.

In the 1950s and 1960s, however (when post-war steel prices made over-engineering expensive), engineers began assuming that ‘sophisticated’ models could more accurately predict the loads which, in principle, allowed a reduction in the safety margin, and thus fewer materials and lower costs.⁵ What has ultimately been discovered (following the failure of several bridges from that era) is that forecasting loads is more uncertain than anticipated; by reducing the safety margin, engineers effectively discounted the possibility that forecasting models were flawed or subject to inherent uncertainty. Only in a world of perfect foresight can one reduce the safety margin to near zero.

In the VaR context, estimating risk at the 99th percentile is equivalent to building bridges without the necessary safety margin: risk managers are attempting to predict the unpredictable! Instead, it is far better to estimate risk at lower percentiles (the 85th–90th percentile can be far more accurately assessed) and gross up by an arbitrary safety margin. For example, if 90th-percentile VaR is 15% (implying that the portfolio has a 10% chance of falling by 15% or more during the investment horizon), the risk manager might gross up 90th-percentile VaR by 1.5x (eg 22.5%). If this is unacceptably high, various overlays would be employed to limit portfolio risk. In contrast, the risk manager might attempt to measure 99th-percentile VaR, finding it to be, say, 20%. VaR is not fundamentally flawed; rather, what is flawed is the practice of expecting VaR to be more precise than that of which it is capable. Quantitative risk measures are still necessary and crucial, but the industry must also recognize the inherent limitations and apply the appropriate margin of error.

⁵ “Generation of Bridges Was Built with Less Steel”, Eric Weiss, *Washington Post*, Sunday August 5, 2007: <http://www.washingtonpost.com/wp-dyn/content/article/2007/08/04/AR2007080401439.html>

STOCK RATING SYSTEM

Top Pick

Desjardins' best investment ideas—stocks that offer the best risk/reward ratio and that are expected to significantly outperform their respective peer group* over a 12-month period

Buy

Stocks that are expected to outperform their respective peer group* over a 12-month period

Hold

Stocks that are expected to perform in line with their respective peer group* over a 12-month period

Sell

Stocks that are expected to underperform their respective peer group* over a 12-month period

Not Rated

Stock is being covered exclusively on an informational basis

RISK QUALIFIERS

Average Risk

Risk represented by the stock is in line with its peer group* in terms of volatility, liquidity and earnings predictability

Above-average Risk

Risk represented by the stock is greater than that of its peer group* in terms of volatility, liquidity and earnings predictability

Speculative

High degree of risk represented by the stock, marked by an exceptionally low level of predictability

* Peer group refers to all of the companies that an analyst has under coverage and does not necessarily correspond to what would typically be considered an industry group. Where an analyst's coverage universe is such that 'relative' performance against a 'peer group' is not meaningful, the analyst will benchmark the rating against the most appropriate market index

Distribution of ratings

Rating category	Desjardins rating	Desjardins coverage universe (# of stocks)	% distribution	Desjardins Investment Banking (# of stocks)	% distribution
Buy	Top Pick/Buy	77	69	21	62
Hold	Hold	33	29	12	35
Sell	Sell	2	2	1	3
Total		112	100	34	100

DESJARDINS SECURITIES INC. LEGAL DISCLAIMERS

Dissemination of Research

Desjardins Securities makes all reasonable effort to provide research simultaneously to all eligible clients. Research is available to our institutional clients via FirstCall Research Direct, Reuters and Bloomberg. In addition, sales personnel distribute research to institutional clients via email, fax and regular mail.

Analyst Certification

Each Desjardins Securities research analyst named on the front page of this research publication, or at the beginning of any subsection hereof, hereby certifies that (i) the recommendations and opinions expressed herein accurately reflect such research analyst's personal views about the company and securities that are the subject of this publication and all other companies and securities mentioned in this publication that are covered by such research analyst, and (ii) no part of the research analyst's compensation was, is, or will be, directly or indirectly, related to the specific recommendations or views expressed by such research analyst in this publication.

Additional Disclosures

Desjardins Securities' equity research analysts are compensated from revenues generated by various Desjardins Securities businesses, including Desjardins Securities' Investment Banking Department. Desjardins Securities will, at any given time, have a long or short position or trade as principal in the securities discussed herein, related securities or options, futures, or other derivative instruments based thereon. The reader should not rely solely on this publication in evaluating whether or not to buy or sell the securities of the subject company. Desjardins Securities expects to receive or will seek compensation for investment banking services within the next three months from all issuers covered by Desjardins Securities Research.

Legal Matters

This publication is issued and approved for distribution by (i) in the US, Desjardins Securities International Inc., a member of the FINRA and SIPC, (ii) in Canada, Desjardins Securities Inc., a member of the IIROC and CIPF (collectively, "Desjardins Securities").

This publication is provided for informational purposes only, and does not constitute an offer or solicitation to buy or sell any securities discussed herein in any jurisdiction where such offer or solicitation would be prohibited. The securities mentioned in this publication may not be suitable for all types of investors; their prices, value and/or income they produce may fluctuate and/or be adversely affected by exchange rates. This publication does not take into account the investment objectives, financial situation or specific needs of any particular client of Desjardins Securities. Before making an investment decision on the basis of any recommendation made in this publication, the recipient should consider whether such recommendation is appropriate, given the recipient's particular investment needs, objectives and financial circumstances. Desjardins Securities suggests that, prior to acting on any of the recommendations herein, you contact one of our client advisors in your jurisdiction to discuss your particular circumstances. Since the levels and bases of taxation can change, any reference in this publication to the impact of taxation should not be construed as offering tax advice; as with any transaction having potential tax implications, clients should consult with their own tax advisors. Past performance is not a guarantee of future results.

This publication may contain statistical data cited from third party sources believed to be reliable, but Desjardins Securities does not represent that any such third party statistical information is accurate or complete, and it should not be relied upon as such. All estimates, opinions and recommendations expressed herein constitute judgments as of the date of this publication and are subject to change without notice.

Although each company issuing this publication is a wholly owned subsidiary of Desjardins Group, each is solely responsible for its contractual obligations and commitments, and any securities products offered or recommended to or purchased or sold in any client accounts (i) will not be insured by the Federal Deposit Insurance Corporation ("FDIC"), the Canada Deposit Insurance Corporation or other similar deposit insurance, (ii) will not be deposits or other obligations of Desjardins Group, (iii) will not be endorsed or guaranteed by Desjardins Group, and (iv) will be subject to investment risks, including possible loss of the principal invested.

The Desjardins trademark is used under licence.

© 2008 Desjardins Securities International Inc. and Desjardins Securities Inc. All rights reserved. Unauthorized use, distribution, duplication or disclosure without the prior written permission of Desjardins Securities is prohibited by law and may result in prosecution.

OFFICES

Montreal

1170 Peel Street
Suite 300
Montreal, Quebec H3B 0A9
(514) 987-1749

Toronto

145 King Street West
Suite 2750
Toronto, Ontario M5H 1J8
(416) 867-6000

Vancouver

200 Burrard Street
Suite 1510, Waterfront Centre
Vancouver, British Columbia
V6C 3L6
(604) 678-6141