



Max Darnell
Partner
Chief Investment Officer



Ken Ferguson
Partner
Co-Director of Global Derivatives



Dori Levanoni
Partner
Co-Director of Global Derivatives



Anna Chefter
Director
Risk Management

Risk Shouldn't Be the Elephant in the Room

Risk shouldn't be the elephant in the room. Investment managers and their clients spend a majority of their time together talking about the return side of the equation, while risk often receives little or no attention.

Together they discuss at length the philosophy that underlies the return generation process, they talk about the specific drivers of return, and they may comb over the historical attribution of return to those drivers.

The return objective is always discussed, but little discussion is held about how risk is to be managed, and even less about what the client's preferences or objectives are regarding risk. One reason for this is that risk is more complex than return, and it is more challenging to talk about the myriad of different risks that may be faced and how the manager or client would or should respond to them. This note is an attempt to motivate more discussion about risk with our clients.

We're often asked why a firm like ours would want to create a separate function within our walls that focuses resources exclusively on matters of

risk (our independent Risk Office was created in 2004). Because we're known for being "quant," it is assumed that when it comes to risk, we must have it fully quantified, captured, and contained. The central problem with this assumption is that

For risk management to be truly effective, we must have a clear understanding about what our clients' individual objectives are with respect to risk.

for risk management to be truly effective, we must have a clear understanding about what our clients' individual objectives are with respect to risk, and to achieve that, we felt a need to dedicate

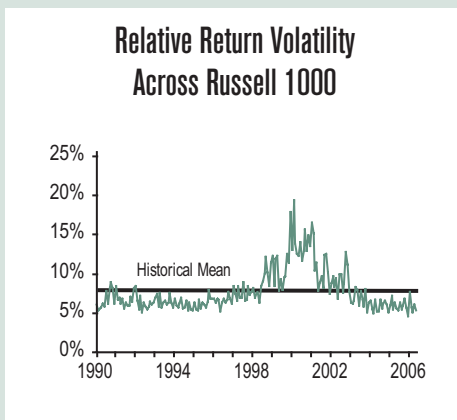
resources that would focus exclusively on risk research and on engaging our clients in discussions about risk.

Take recent experience, for example. Many (not all) investment managers have been delivering less realized risk than their clients have expected them to deliver. As a result, there are many managers who have delivered positive results, but have, to the disappointment of their clients, fallen short of the return objective. Should clients fault their managers for not taking enough risk? Should investment managers increase the risk in the portfolios in response to the low volatility that they've seen in the markets? The answers to these questions are complex and do require that the investment managers and their clients establish a clear understanding about the options, their implications, and the client's objectives with regard to matters of risk. We'll sort through these issues in the comments that follow.

Volatility Has Been Low

Relative return volatility is what matters to active management, and the volatility of relative returns has been low by historical standards. Most focused upon perhaps is the relative return volatility across the US equity market. Relative return volatility tells us how differentiated the returns to stocks have been relative to each other. This is important to active management for the obvious reason that the larger the return

differentials, the larger the potential gain to be captured through active stock selection.



At the end of 2005, relative return volatility reached the lowest level ever seen since 1990. In December, 2005, relative return volatility (cross-sectional volatility) in the Russell 1000 fell to 4.6%, which is 40% lower than its historical mean of 7.8%. Over the first six

months of 2006, we've continued to see this volatility reach into the lowest historical decile.

With two important exceptions, volatility elsewhere has also been low by historical standards. Commodity prices have, of course, been volatile, as have global bond prices. Equity market indices such as the S&P 500, EURO STOXX, and MSCI World have displayed volatility levels between 50%-70% of what is normal, while currency volatility and TOPIX volatility has tended to be only marginally lower than what is normal.

Index	Recent Volatility*	Normal Volatility**	Recent Relative to Normal (%)
MCSI World	8.08%	14.45%	55.92%
S&P 500	9.72%	18.32%	53.06%
EURO STOXX	13.79%	20.39%	67.60%
TOPIX	17.05%	19.60%	86.99%
EUR vs. USD	8.75%	9.91%	88.38%
GSCI	24.34%	24.27%	100.29%
Citigroup World Government Bond Index	8.48%	8.65%	98.03%
US 10 Year Treasury	5.93%	7.89%	75.18%
Russell 1000 Cross-sectional	5.29%	7.78%	68.08%
MSCI Euro Cross-sectional	5.47%	9.36%	58.49%
Country Selection Stock Cross-sectional	1.90%	3.89%	48.78%
Country Selection Bond Cross-sectional	0.66%	1.59%	41.30%
GSCI Cross-sectional	8.11%	7.83%	103.62%

* Recent Volatility is annualized using daily volatility from 06/01/2005 until 05/31/2006

** Normal Volatility is the average of annualized daily volatilities from 01/01/1997 until 05/31/2006

What To Do in a Low Volatility Environment

Most investors are familiar with the fact that risk – as measured by the narrowest definition of risk, price volatility – varies across time from high to low. Risk is commonly thought of, therefore, as a *variable* quantity. While this is fair when considering measures of risk individually and in isolation, it is a potentially dangerous mindset to apply to risk management broadly where risk is properly understood to be multidimensional.

A better mindset would be to regard risk as a *fixed*, multidimensional quantity.

A better mindset would be to regard risk as a fixed, multidimensional quantity. One can seek to change its location or change its shape, but the fact is that it is hard to reduce its overall magnitude if we're talking about risk in the broadest sense of the term. Think of risk as though it were the volume of air contained in a balloon. If you squeeze one end of a balloon, the air simply gets pushed elsewhere and the balloon expands at the other end. The volume in the balloon doesn't change, it is merely shifted or reshaped.

Thinking this way leads us to look for the consequences of a change in the magnitude of one type of risk, or a change in the shape of another type of risk. When relative price volatility is low, as it has been in the US equity market recently, for example, this mindset leads us to look for the offsetting consequences of the fact that risk has fallen in this respect. One consequence of this is a heightened risk that active managers may fail to deliver the level of returns that investors are in need of. Risk measured as price volatility is the source of their return, so as that risk falls, the risk that they may fall short of their return goals rises!

How do we then manage the consequences that accompany low relative return volatility? There are three alternative responses that the active managers may take. They may choose to stay the course, taking no action at all, they may choose to scale the

positions to raise the level of active risk embedded in the active portfolio, or, they may choose to change the composition of active risks.

Do nothing. Doing nothing at all in response to below normal levels of volatility is a perfectly appropriate response under certain conditions. First of all, there may be costs – trading costs, leverage costs, or undesired risks – that come with taking action. Doing nothing at all may simply be better than the alternatives, therefore. The best answer may be to take a long-term view of the risk and return objective and to allow risk to vary around the long-term target, above and below, in such a way that the long-term objective is met on average over time rather than at each point in time. This would be particularly favorable if the opportunity set and/or the quality of the risk adjusted return varied in proportion to the level of market risk.

Scale positions. This is best described with an example. Let's say, for example, that the current set of active positions is estimated at today's level of market volatility to hold only 50% of the desired active risk level. One way to increase the active risk to reach the desired active risk level is to simply double the size of each active position. This means, for example, that each 1% underweight position would become a 2% underweight position and each 1% overweight position would become a 2% overweight position. Such scaling of positions would scale risk proportionately. Such an example of the scaling of positions is shown in the exhibit below.

One of the notable advantages of derivatives portfolios lies in the fact that you can typically scale positions this way without any material, incremental cost. Such is not the case with portfolios constructed from cash securities. In an equity portfolio, for example, scaling may result in net short positions, and there would be a financing cost associated with shorting stocks (the shortfall between the short rebate and the return on cash). Furthermore, scaling may require "leverage" in the sense that the sum of the absolute value of the longs and the shorts exceed 100% of the portfolio value, in which case prime brokers will charge a leverage, or borrowing cost. These costs may offset the benefit of scaling and lead one to do nothing instead.

There is another important consequence of scaling that we will give special attention to later.

Change the composition of risks. Another way to raise the risk level is to shift the concentration of active risk into those assets or positions that carry a higher level of active risk. This may mean, for example, focusing more risk in higher beta stocks in an equity portfolio, or taking more stock specific risk and less style, sector or industry active risk since stock specific risk is measurably higher than the latter types of risk. The downside of this approach is that it may mean forming a less diversified allocation of active risk, or it may result in a portfolio with the right risk level, but a lower expected risk adjusted return.

All three of these responses may be appropriate under different circumstances.

The Unintended Costs of Scaling

Beyond the explicit costs that may come with scaling certain types of assets (e.g., the leverage cost associated with equities as we've described above), there are other costs that may be less obvious that may also make scaling less appealing. For example, the costs of continuously scaling positions to maintain a constant targeted risk level may not be commensurate with the benefit. Volatility levels are constantly changing, and risk estimates are very noisy, so treating them with a


Treating risk estimates with a high degree of precision is likely to generate high transactions costs with no better a fit between realized and desired risk levels.

high degree of precision and responding to small changes in them on a short-term basis is likely to generate high transactions costs with no better a fit between realized and desired risk levels.

While there are no *universal* statements that can be made about the relationship between estimated risk (the risk that is "expected" based upon current positions and current market volatility) and realized risk (that which actually occurs), we've provided just one example to drive home the point about how noisy this relationship can be. In the

Scaling Positions

50% OF DESIRED ACTIVE RISK LEVEL



To double the risk
 →
 Double the positions

100% OF DESIRED ACTIVE RISK LEVEL




chart below, we have used a Monte Carlo approach to forming Global Tactical Asset Allocation portfolios (1000 portfolios at each point in time), where the expected risk (parameterizing the estimate on 260 days of trailing daily data at each point in time, and weighting these observations to produce a half-life of 130 days) is held precisely to 1% for all 1000 portfolios formed at each point in time. The solid line displays the average of the risk realized for each of the 1000 portfolios over the subsequent twenty four months, while the bars show the average range of realized risk (these show the +/- one standard deviation of the realized risk for the 1000 portfolios at each point in time).

The point of all of this is to show that realized risk will display a wide band around the target, or “expected” risk. It should be immediately evident in this example that paying transactions costs to constrain expected risk within a very narrow range is not going to yield a more accurate alignment of expected and realized risk than if one allows a relatively loose band within which expected risk may fall. In our Global Macro work, for example, we find that transactions costs may rise as much as 15% to 25% when forcing expected risk to be pegged at the long-term target at all points in time rather than allowing expected risk to fall within a range of +/-20% of the target risk level (e.g., and range of 8%-12% would be applied to a long-term risk target of 10%).

To avoid errors in this vein, one should take the investment horizon into account. It is easier to allow positions that turn over quickly to be more tightly constrained around a risk target than positions that turn over more slowly, for the simple reason that the risk is more likely to vary, the longer the holding period. Continuously rescaling long-term positions to meet short-term risk objectives may do nothing more than generate additional turnover costs. Simply holding the position as risk varies between high and low, averaging out to the long-term target is likely to result in a superior net-of-transactions costs risk-adjusted return. The best answer, therefore, may often be to allow a range around the targeted risk level, and to allow the expected risk to vary anywhere within that range.


Possibly even more important is the impact that scaling has on tail risk.

Possibly even more important is the impact that scaling has on tail risk. Let’s say you have a strategy that displays historically a 5% volatility and a \$10m VaR at the 99% level for each


\$100m invested. During a period of low market volatility, you may estimate that the near-term volatility of the strategy has fallen to 2.5%. Should you double the size of your positions, you’d raise the risk level to the targeted 5% volatility, but you would bear what is likely to be an unintended doubling of your VaR to \$20m. A doubling of your downside risk is the consequence of scaling your positions to maintain a constant level of risk in this example.

For those accustomed to thinking of risk as a variable quantity, this is not an intuitive outcome. The reason we’re doubling the size of our positions is because risk levels are low. The assumption naturally occurs that because volatility levels are low, the size of any potential outliers – those picked up by VaR – should be suppressed as well. The problem with this assumption is that tail events occur *independently* from yesterday’s risk levels. A \$10m hit to a set of positions is just as likely to occur in a low risk environment where near-term risk is measured to be lower as it is to occur in a high-risk environment where near-term risk is measured to be higher. *Risk “events” are events because they reflect a sudden shift out of whatever the recent risk context has been.*

The Loose Relationship Between Expected and Realized Risk

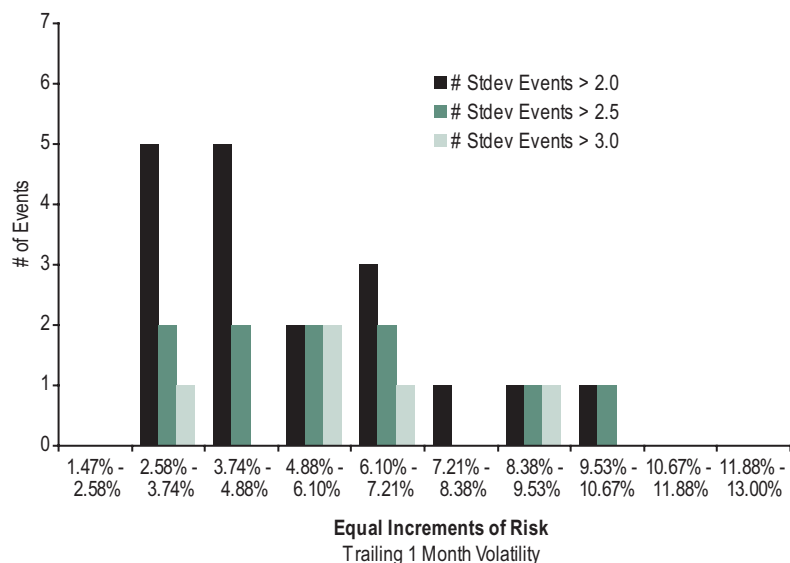


Does the distribution change shape when volatility is low?



An analogy may be helpful here. Investors, we believe, tend erroneously to think of risk in the same way they think about dynamite. Dynamite grows less stable and more vulnerable to an unintended explosion as the temperature rises. Investors intuitively think that market shocks are more likely to occur when volatility is high. The truth is that market volatility is just as subject to shocks, or “explosions,” when market conditions are “cool” as when they are “hot.” Unlike dynamite, the “instability” of market returns is ever present and independent of current market volatility conditions. In the chart below, we show the frequency of large (i.e. 2.0 sigma, 2.5 sigma and 3.0 sigma) events in the S&P 500, broken down by recent volatility. It is quite apparent that the frequency of the larger (2.5 sigma and 3.0 sigma) events bears little, if any, relationship to the recent level of measured volatility.

Occurrence of Large Return Events is Independent of Recent Risk Level



On the subject of risk “events”, we should draw attention to another possible mistake that can be made by managing too tightly to a risk target when the estimate of risk relies heavily or exclusively on trailing measures of risk (tomorrow’s risk typically *is* estimated based upon trailing risk levels, so this problem is not uncommon). The problem occurs when we have very recently experienced a volatility shock. Subsequent to a shock, risk systems will tend to both (a) overestimate future risk, and (b) compel investors to reduce risk

Reducing risk subsequent to a spike in risk may be exactly the wrong response for a fundamentally-oriented investment strategy.

at what, for fundamentally oriented investment approaches, may be exactly the wrong time. Market

inefficiencies occur for many reasons, but perhaps the most compelling reason for them to exist is that investors will behave less rationally in the context of sudden market events, particularly risk events. These may be the best time to take risk, not the time to trim risk back!

Moving the Dialogue Forward

The starting point for this dialogue is to focus on those dimensions of risk that have the most direct impact upon the return objectives. Here are four questions that may be useful as a starting point. Note that the issues they address are overlapping.

1. How sensitive is the client to tail risk? The more sensitive the client is to tail risk, the more the manager should target long-term risk characteristics over short-term risk characteristics. This means the risk level, as measured over shorter time spans, is going to vary more across time.
2. Over what time period is it important for the client to hit the return goal (or risk goal if that’s how the client focuses on the objective)? The shorter the time period over which the goal must be reached, the more the manager needs to manage to short-term measures of risk. Allowing the risk level to fall with falling market volatility, for example, will put the return goal out of reach.
3. What role does the portfolio play in the client’s portfolio of portfolios? If the portfolio is highly differentiated from the client’s other portfolios in terms of the type of strategy and the assets or asset classes involved, then the client should have less risk sensitivity at the individual portfolio level. Alternatively, the manager’s strategy may share more in common with other managers in the client’s portfolio of portfolios. In such circumstances it will be important to understand whether the manager plays a key role as a diversifier or not. Those who provide a higher level of diversification within the pool will more likely need to maintain a more constant risk level to provide a more continuous risk offset to the other managers.
4. Can constraints be lifted to allow managers more room to manage risk optimally? Fewer and fewer constraints are being placed on investment managers. Because constraints may serve as pseudo-objectives (defining what the manager

cannot do), it is best if the dropping of constraints is accompanied by a greater clarity of the investment objectives. By removing constraints and replacing them with a better definition of the objective, managers will be freer to deliver on those objectives and should deliver better results. The broadening use of so-called “extended mandates,” e.g., 120/20 and 130/30 equity mandates is a good example of such. It was rarely the client’s intention to cause managers to place most of their risk in the largest 30 names of the S&P 500 by placing long-only constraints on the portfolio, but that was the outcome. By releasing the manager from the long-only constraint, replacing it with clarity that the objective is to minimize the impact that a small number of shorts can have on the portfolio, managers can add more value with less risk.

For example, derivatives can be very effective tools in constraining and reshaping risk. This can be particularly important in dealing with tail risk or in simply being very specific about which risks managers want to take and which they don’t.

These issues are only the beginning of what can potentially be addressed. We look forward to continuing this dialogue with our clients.

First Quadrant, L.P.
800 E. Colorado Boulevard, Suite 900, Pasadena, California 91101
626.683.4223
Marketing Services: FQ_Updates@firstquadrant.com
www.firstquadrant.com

Copyright © by First Quadrant, L.P. 2006, all rights reserved.