Tail Risk Management

VINEER BHANSALI

VINEER BHANSALI
is a portfolio manager and the head of Analytics at PIMCO in Newport Beach, CA.
bhansali@pimco.com

The events unfolding as of the writing of this article show the importance of controlling portfolio exposures to fat-tail events. The obvious benefit of having the appropriate amount of disaster insurance is survival. When applied to financial market participants, a secondary benefit is observed—the survivors are better able to take advantage of the reduced liquidity that accompanies tail events and position themselves for attractive prospective returns. Thus, tail insurance is an offensive strategy for the long term, even though it may appear to be a defensive strategy in the short term.

The desire to outperform peers and indices over the last decade forced herds of portfolio managers to reduce or void insurance in their portfolios for improbable, yet possible, events. This trend culminated in a spiraling asset debacle whose end has yet to be realized. The emergence of structured products that were inefficient packaged sales of insurance was a natural consequence of the demand for high returns in a low-return world (Bhansali [2007a]). The very real quandary a portfolio manager now faces is one that requires the estimation of the optimal amount of insurance for a portfolio to protect against tail events. Tail events, by their very nature, are hard to quantify with traditional models, which pay little heed to rare events. Most investment processes rely on harvesting risk premia from systematic factors, whose variations are measured and exposures scaled based on history; by design, this statistical averaging process underweights the impact of tail events.

The purpose of this article is threefold. First, I will provide a characterization of tail events that is quantifiable; second, I will provide algorithms to quantify the probability and severity of these events; and third, I will propose a selection framework for tail hedges that may consist of securities, options, and strategies.

TAIL RISK AS MACRO RISK

Tail risk at the portfolio level is almost always a *systemic* risk. A systemic risk is one that puts pressure on the ability to fund levered holdings. In episodes of systemic risk, everyone desires liquidity, but no one is willing to provide it. According to Bookstaber [2007], the tight-coupling of financial markets requires that liquidity be available to the system at all times at some price, or the financial system will start to break down. Financing and liquidity are *macro* risks and their proper valuation requires macro models. Proper hedge construction requires a macro view and uses macro tools and markets.

Khandani and Lo [2007] have observed that new strategies in the market, especially those provided by hedge funds, are characterized by new betas (systematic selling of liquidity), which can be withdrawn very quickly. Many of these strategies are highly correlated; hence, a shock to one type of strategy causes ripples in other strategies. Further, the nature of the relationships is not static, so that a financial market dislocation in a strategy that is running many uncorrelated strategies can rapidly influence other strategies. Therefore, a tail-risk hedge portfolio that provides insurance against tail events should carry a "beta" to systemic and macro risks.

A statistical analysis of the returns of broad asset classes over the long term shows the presence of a handful of latent factors, or principal components, that explain the variation of most traditional investable assets. The striking observation from analyses such as these is that the presence of regimes follows a similar systematic behavior and loads most heavily on monetary policy. In other words, with suitable leads and lags, the quadrants of early and late recessions and early and late expansions in the developed markets (the U.S., in particular) can be mapped to early and late periods of Fed easing and tightening. In a finance-driven economy such as the U.S., it is not surprising that the tails which develop in these periods can be hedged with macro instruments which respond to central bank activity. Deleveraging risk is, to a large degree, a monetary policy risk; thus, tail events become macro events.

This observation has far-reaching consequences. The main consequence is that tail risk becomes a macro risk. To forecast and control against tail risk, it is necessary to step outside the world of historical estimates and calibration and to forecast structural changes (not easy!) and imagine improbable, high-severity scenarios. The immense benefit of thinking about tail risk in terms of a macro risk is that the construction of hedges is simplified. At any point in time, macro markets are the deepest markets. A macro market is a broad bond, stock, foreign exchange (forex), credit, or commodity market, or a combination of them. Despite pockets of illiquidity in the macro markets (e.g., the credit market could dry up with forex barely impacted), typically some sort of insurance remains attractively priced for a sufficiently long period due to the sheer mass of capital reallocation needed to align the markets.

When systemic crises happen, correlations rise in absolute value. This provides a free lunch because a completely disconnected macro asset class in normal times can become a good hedge against the tails in distressed times. In other words, if the cost of credit hedging is too high in the midst of a systemic crisis, then a hedge from the forex or equity option markets might actually be a more efficient credit hedge—over a long enough holding horizon—despite the apparent basis risk associated with it in normal times.

VALUATION OF TAIL-RISK HEDGES

Proper evaluation of the *cost* of insuring against tail risk is the second most important factor to evaluate. Credit market hedges were extremely cheap in 2007 because of the incessant selling of credit insurance through structured products and the demand driven by excess liquidity. Today, these same hedges by some measures are, at best, fairly priced. Thus, hedging credit with credit market instruments might no longer be an attractive solution. The estimation of option prices and sensitivities in general is based on the academic foundations of continuous trading and risk-neutral investors.

These ideal conditions are rarely met in turbulent markets. As a result, what are now considered traditional models of option pricing are flawed due to their over-simplifying assumptions (Bhansali [2007b]). A viable alternative approach is to run scenarios that investors are averse to and evaluate the expected value of portfolios under those scenarios with and without insurance assets. The difference in these two portfolio values is indicative of a fair price for the insurance from the investor's perspective, not the market price of the insurance. We all know that catastrophe insurance can trade too cheaply in the natural insurance markets (i.e., the best time to buy insurance is when it is not needed). Obviously, it is possible for the same to hold true in the world of finance, especially in a world of innovative financial engineering that ports risks from one type of market to another.

More than one method can be used to implement tail hedges, but our framework adopts a natural unity that would seem obvious to someone who is involved in insurance. The simplest way to insure portfolios is to buy high quality "insurance" securities. For example, credit market crises almost always occur with deleveraging and a grab for liquidity which usually results in Treasuries rallying in price terms, especially those with the shortest maturities. Therefore, short-term cash and Treasuries are natural asset-based hedges for tail risk. But these securities can become overpriced due to technical reasons.

The second way to insure portfolios is to buy contingent claims, or "option-like" securities. A year ago, out-of-the-money tranches on CDX and ITRAXX indices were literally being given away due to the demand for constant proportion debt obligations (CPDOs), levered super seniors, and other default-remote structured transactions. These option-like payoffs were priced below their theoretical expected value under a systemic risk outcome.

The third alternative to insure portfolios is to invest in strategies that are negatively correlated to tail risk. Among the traditional established strategies, Exhibit 1 demonstrates that only the systematic, trend-following managed futures strategy provides positive correlation to tail-risk indicators, such as the CBOE Volatility Index (VIX), while also being largely uncorrelated to the stock market. Copious amounts of research have been done to demonstrate that trend-following strategies empirically behave like a long position in lookback straddles, and hence are naturally long tail risk (Fung and Hsieh [2001]).

The fourth approach is to move the portfolio off the optimal frontier (i.e., to accept less return for the same amount of risk), which explicitly recognizes that the simplest mean–variance optimal frontier falsely assumes that risk can be measured by volatility (second moment) alone and that the investor has perfect forecasting ability, which is almost never true. One example of this is to reduce the exposure to spread products, such as corporate bonds, or low-quality mortgages, which have higher yield because of embedded default and illiquidity options. When it

EXHIBIT 1
Correlation of Typical Hedge Fund Strategies and VIX Index for Different Periods

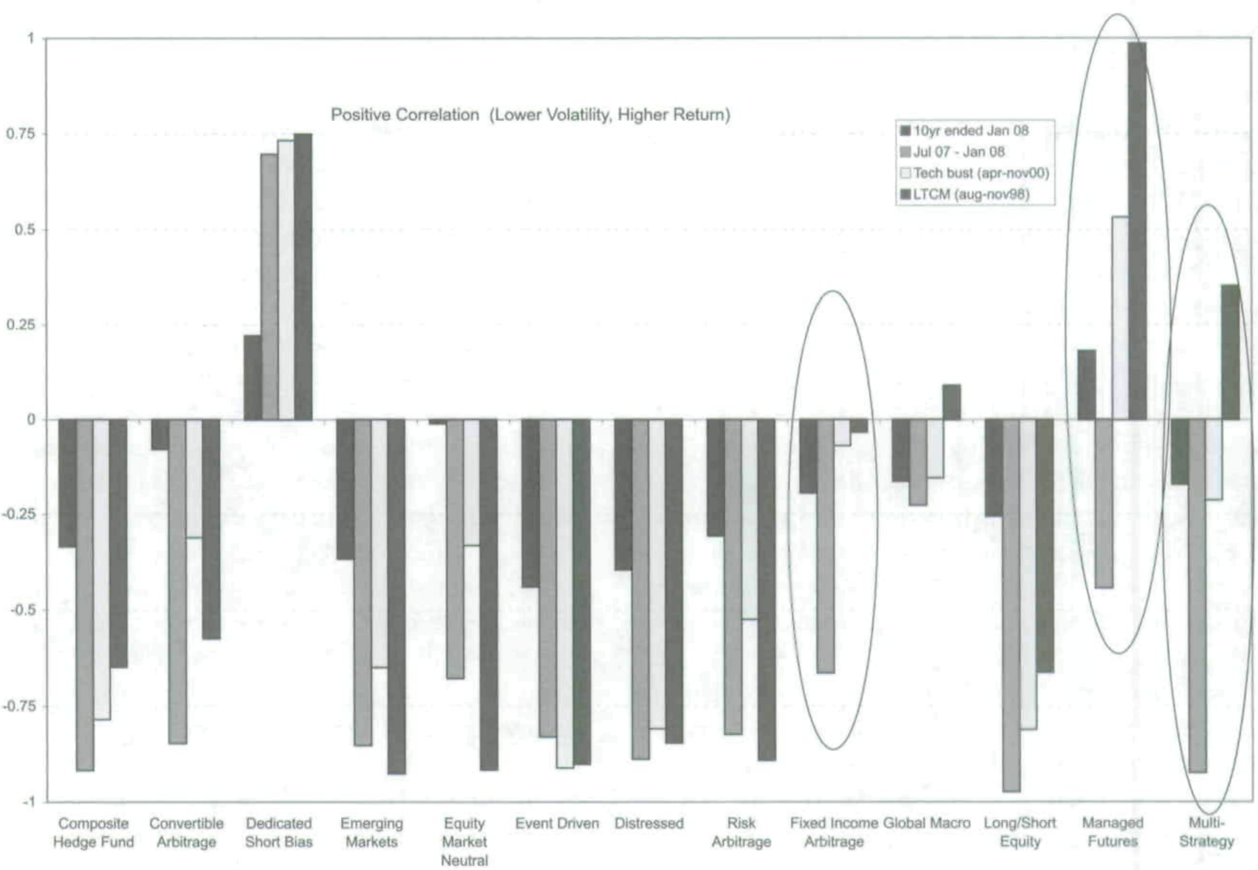

Source: PIMCO.

becomes evident that technicals have caused risk premia to compress to rich-enough levels, these assets promise a high probability of low-to-negative future returns. As shorts, these assets might actually become candidates for tail risk control.

To reiterate, I am not advocating portfolio insurance in which exposures are dynamically adjusted to changing market conditions. The problem with this type of zero-cost tail solution is that it assumes liquidity will be present in crises, yet liquidity typically all but evaporates during systemic shocks. In the approach discussed here, the portfolio is subjected to possible, but rare, supershocks, and insurance may be purchased using one of the four techniques just described. Clearly, there is an upfront, or running, cost to the tail hedge, but the role of the tail-risk portfolio manager is to reduce the cost of these hedges over the hedge horizon. This might not be as challenging as it sounds.

Frequently, when the horizon is long enough, hedges can almost be bought for "free." For instance, recently long-dated options on the forex, such as the dollar–yen exchange rate, had zero carry cost over a one-year horizon due to the roll-up of forward rates from the interest rate differential between dollar rates and yen rates. If, in periods of stress, one of the outcomes is deleveraging and a flight to low-yielding currencies such as the yen, the directionality of the forex movement and the associated increase in volatility makes the real-world expected value of the package much higher than the theoretical, risk-neutral value. Thus, the possibility of asymmetric payoffs in particular states-of-the-world makes the hedges worth more to the investor than their stand-alone theoretical value. For more discussion of how this strategy may result in excessive risk-taking and presence of cheap tail hedges, see Bhansali [2008].

To ascertain the value of a particular tail hedge, it is essential that the scenario analysis is performed with many variations of the inputs for the parameters. The inputs are more important than the level of refinement of the models. In the simple language of Black–Scholes options, this boils down to running the performance of option positions under various maturity, rate, volatility, and spot rate environments, using the option model simply as a crude nonlinear transformation machine between inputs and outputs. Special attention must be paid to consistently taking the volatilities and correlations to extreme values because the underlying assumption of joint lognormality undervalues the tails of the distribution (see Bhansali and Wise [2001]).

It is also possible to approach the problem by specifying other distributions with naturally fat tails, such as the Lévy distribution, but moving away from simple distributions, in practice, creates more problems with intuition than is compensated for by more accurate empirics. In any case, today's computational power makes it easy to substitute empirical distributions that can be evaluated numerically, especially on the tails, in lieu of theoretical, closed-form solvable models.

The extension of Black–Scholes to credit for this capability is the Gaussian copula model. While used and abused, this model—a simple and an effcient way to value options on portfolio losses—still has the ability to provide good intuition about the performance of tranches in periods of stress. Exhibit 2 shows the performance of a 7–100 tranche on the IG8 index for both instantaneous and one-year-delayed shocks. The first scenario in Exhibit 2 shows the percentage impact on a long protection tranche position for an instantaneous shock. The second scenario shows the percentage impact on a tranche position for a shock at a one-year horizon.

Tranche returns can be different as the implied base correlation moves around. As correlations rise, the senior tranches pick up additional return due to a higher probability of many names defaulting simultaneously, a hallmark of systemic risk. The one-year-horizon shock accurately captures the cost as the tranche—which is an option on losses—shortens in maturity and rolls down the credit curve, which typically is upward sloping.

Using credit derivatives as a reference, there is another way to look at systemic shocks. In Bhansali, Gingrich, and Longstaff [2008], index spreads for various indices were broken down using a three-jump approach, which was first used by Longstaff and Rajan. In Exhibit 3, the idiosyncratic component is the darker solid line, the sector risk is the dotted line, and the economy-wide risk is the light solid line. The economy-wide, or systemic, risk is currently at highly elevated levels. Because the economy-wide risk is carried by senior and super senior tranches, it provides a microscope for the valuation of liquidity and systemic risk in the market.

If we proxy systemic risk by the fraction of spread allocated to the systemic factor, or roughly speaking the fraction of expected losses accounted for by default-remote super senior tranches, we can construct a high-frequency measure for a systemic risk factor. Regression on this factor of other seemingly unrelated markets—such as municipals, asset-backed securities, and other spread

EXHIBIT 2
Tranche-Shock Scenario Analysis

PriceDate	6/26/08
index	IG8
tenor	10
funded	0
attPoint	7.00%
detPoint	100.00%
Coupon(bp)	82.8

Delay for shocks (years): 0

Spread Boost (bp)	(Attachment %) 64.69	97.03	32.34
	(Detachment %) 0.00	0	0
-25	-1.3	-0.35	-2.36
-10	-0.53	0.46	-1.58
0	0	0.99	-1.05
10	0.52	1.52	-0.51
25	1.32	2.32	0.31
50	2.66	3.65	1.71
100	5.32	6.24	4.49
200	10.38	11.12	9.83

Delay for shocks (years): 1

Spread Boost (bp)	(Attachment %) 64.69	97.03	32.34
	(Detachment %) 0.00	0	0
-25	-2.07	-1.17	-3.07
-10	-1.4	-0.47	-2.42
0	-0.94	0	-1.96
10	-0.49	0.46	-1.5
25	0.2	1.16	-0.79
50	1.38	2.33	0.42
100	3.71	4.63	2.87
200	8.24	9	7.64

Tranche-shock scenario analysis for a 7–100 senior tranche on the 10-year IG8 investment-grade index for various shocks to the underlying spread level and base correlations.
Source: PIMCO.

product—shows significant betas to systemic risk. In other words, many asset classes carry large amounts of systemic illiquidity risk. We found, with significant *t*-statistics, that from June 2006 through March 2008, for every 10 bps widening of the systemic component, without adjusting for duration, Lehman Brothers indices moved as follows: the ABS Floating Rate Index returned negative 53 bps, the U.S. High Yield Index returned negative 58 bps, the Municipal Bond Index returned negative 51 bps, the CMBS Index returned negative 63 bps, and the Securitized Index returned negative 38 bps. Each of these asset classes carries liquidity risk (which is different from credit

EXHIBIT 3

Allocation of CDX 5Y Index to Idiosyncratic, Sector-wide, and Systemic Risk Factors

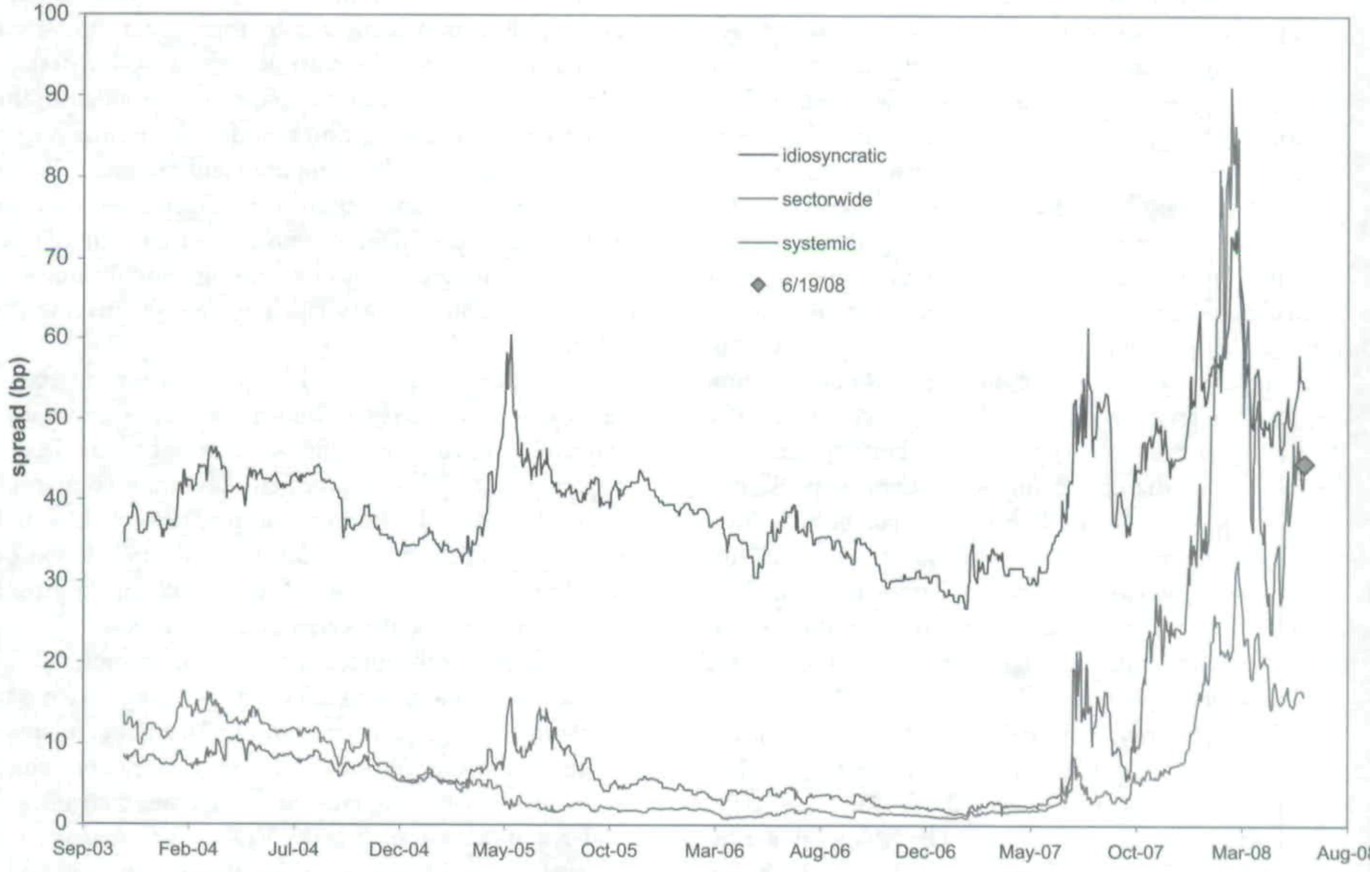

Source: Bhansali, Gingrich, and Longstaff [2008].

risk) and by stressing our liquidity risk measure we show how the asset classes can be expected to perform in periods of stress. It should be no surprise, therefore, that as of this writing, municipal bonds—which have one of the largest betas to this risk factor without much direct credit risk (especially general obligation bonds of natural AAA-rated states)—are trading at higher yield levels than Treasuries, despite the fact that they are tax exempt!

Thus, the construction of a tail portfolio hedge depends to a great extent on the following three factors: first, the scenario behavior of the portfolio; second, the scenario behavior of the hedges, net of cost; and third, the probability of the scenario occurring. Finding the best combination requires knowledge of how the portfolio to be hedged behaves under various shocks. The best combination of hedges is the one that produces the worst return to the portfolio under the stress scenario and controls the risk, at horizon, to the dangerous factors in the stress scenario.

This second point is critical because it enables the portfolio to survive to play in the multiperiod game. The goal is for your risks to adjust down in periods of stress so that you have liquidity while everyone else is clamoring for it. This might appear to suggest that tail hedges always have to be custom designed. Fortunately, because correlations between risk factors increase in periods of stress, this need is mitigated for all but the most exotic portfolios. For most investment portfolios, the underlying exposure to market factors, such as the Lehman Aggregate or S&P 500 indices, is all that is needed to optimize tail hedges.

The critical reader might wonder why cheap tail hedges can be found in almost all market environments. There are many potential explanations. For example, speculative demand of particular types and classes of assets may drive these asset prices very high. In low-return periods, as observed until the middle of 2007, "yield hogs" increased options sales as a source of carry. The belief of mean-reversion participants is that out-of-the-money options will never be exercised. This is generally true except when the leverage in the marketplace increases due to the majority of market participants simultaneously executing levered option sales, which forces the notional size to increase in order to generate the same carry.

At some point, this type of system becomes unstable to small noise, creating a domino effect of forced liquidation or of hedgers all trying to hedge at the same time. The following cost calculation can illustrate this point: In June 2007, to buy protection on the 7–100 tranche on the IG8 10-year index cost 15.63 bps per year. For a spread duration of 7.45 on the underlying index, the cost per unit of spread duration was 2 bps. Because the tranche rolls down to a shorter maturity, the annual roll-down cost of 2.67 bps is added to the previous number (2 bps) for a total 4.67 bps of cost per unit of spread duration. To buy protection on one billion notional, the total lifetime cost is about $12 million, which equals the annual cost times PV01 of 7.45.

The same tranche in February 2008 was trading at 91 bps per year. The new running cost per year of spread duration is 10.1 bps and the annual roll-down cost is 2.2 bps for a total of 12.3 bps per year. The new total lifetime cost is about $67 million. The delta of the tranche before the 2007 crisis was approximately 0.53. When multiplied by the spread duration of the 10-year index of 7.45, the result is an approximate sensitivity of 3.95% per 100 bps of index widening. In February 2008, the delta rose to 0.77 as the index widened and moved closer to being at the money, resulting in almost 5.66% of spread risk.

Evaluating these changes in the market and in risk, it is immediately clear why super senior levered notes, which essentially sold protection using these structures in the form of securities, are in severe distress today. First, the mark-to-market loss at five times is huge, equaling the net present value of premium change; second, the mark-to-market is more variable because the delta has increased; and third, the collateral that has to be posted to compensate for mark-to-market fluctuations is much more expensive due to lower T-bill rates and higher LIBOR rates. The fact that all of these events happened simultaneously is typical of systemic tail events.

Hedges are also available for long-term investors who are willing to commit capital to them because of the natural habitat formation of option market participants and the limitations of models used by short-horizon participants, such as option market makers. Most traditional option trading models are only appropriate for short-term trading, roughly for the time it takes a dealer desk to unload the package to another customer. Implicit in this framework is the assumption that deltas, gammas, vegas, and other greeks can be computed and executed. When the standard analysis is extended to longer-dated options, the unavailability of reference instruments, failure of risk-neutral pricing, continuous trading, and limitations imposed by counterparty risk controls all increase the likelihood of mispricing.

Neither is the probability question for tail events satisfactorily answered by reference to traded option prices. The reason is that the pricing of tail options in particular carries a significant amount of risk premium compensation to the seller, also known as lottery ticket risk, which alters the probability distribution. Simulation based on past observations is also not a totally satisfactory approach as each crisis differs in severity and magnitude.

The practical approach for estimating probabilities can follow a number of parallel paths. One approach is to sample from historical events with replacement and magnify the rare-event likelihood by some scale factor, which effectively reshapes the tails. Simultaneously, changes in a tail-risk indicator, such as the VIX, can be measured to indicate the likelihood of a crisis versus a normal environment. Correlations and volatilities from a distressed regime are used for this purpose. Regardless of how the probabilities are estimated, the important point is that because tail hedges are usually cheap in the context of long-lived portfolios, the probability calculation is relatively less critical than the knowledge that potential hedges exist at the right price. Recognizing this distinction can be the difference between survival and almost certain ruin.

OTHER IMPLEMENTATION CONCERNS

Once the macroportfolio tail hedges are put in place, much active management is still needed. As the cost of insurance across markets and maturities changes, a "dashboard" that shows the cost of hedges for different horizons in different markets can provide direction on rebalancing the tail hedge portfolio. Frequent opportunities arise to add *pseudo alpha* (value that reduces the cost of the hedge), particularly during times of rolls for indices or due to short-term imbalances in pricing. The portfolio manager faces the choice of keeping the hedges static (i.e., letting them age) or keeping them current in the most on-the-run versions by evaluating the trade-off between basis risk and liquidity/transactions costs.

Also key are the issues of monetization and financing. When the rare event does happen, can hedges efficiently be converted into cash? In our experience, because tail risks are accompanied by deleveraging, transaction efficiency is skewed to the part that is long the hedges; in other words, in a demand surge, the seller of reinsurance is the price setter for liquidity. Financing is also a concern because hedges are usually implemented through derivatives. Payment for out-of-the-money equity options and forex options is usually made upfront, where as payment for credit derivatives is as-you-go. The structure of the current crisis has revealed that counterparty risk is very important. Many participants found that insurance written by monoline insurers was not as solid as they had thought.

CONCLUSION

In this article, I have discussed why tail risks in portfolio construction should be thought of as systemic risks. Because systemic risks are macro risks, a proper tail-risk hedging program should take into account the relative pricing of broad macro markets and strategies. A tail-risk hedging program should also evaluate combinations of hedging alternatives to find the best combination for immunizing the portfolio against improbable, but not impossible, shocks. Lest it appear that tail-risk hedging is only about crisis events, note that the current low yields on long-term Treasuries may indicate that although credit insurance is fully priced, tail inflation risk is still cheap. That is, the tails to watch out for might arise from an uncontrollable rise in inflation that erodes the value of the dollar and the fixed coupons of long-term bonds. Indeed, stochastic simulations show that in a world where the central bank is simultaneously pursuing inflation-targeting and an asymmetric policy against deflation, yield curves will be flatter when the economy is robust and steeper when the central bank is forced to ease aggressively (Bhansali, Dorsten, and Wise [2007]). While only time will show the truth or falseness of this prognostication, one thing is certain—the hedge for this type of tail risk will not be a micro hedge.

ENDNOTE

I wish to thank clients and PIMCO colleagues who paid attention to the need to hedge tail risk before the subprime crisis of 2007. I would also like to thank Bob Gingrich and Francis Longstaff for their collaboration on recent research.

REFERENCES

Bhansali, Vineer. "Markowitz Bites Back: The Failure of CAPM, Compression of Risky Asset Spreads, and the Path Back to Normalcy." *PIMCO Viewpoints*, 2007a.

——. "Putting Economics (Back) into Quantitative Models." *Journal of Portfolio Management* (Spring 2007b), pp. 63–76.

——. "Correlation Risk: What the Market Is Telling Us and Does It Make Sense?" in *Credit Risk: Models, Development, and Management*, ed. Niklas Wagner. Chapman and Hall/CRC Press, 2007c.

——. "Volatility and the Carry Trade." *Journal of Fixed Income*, Vol. 17, No. 3 (Winter 2007), pp. 72–84.

Bhansali, Vineer, Matthew Dorsten, and Mark Wise. "Asymmetric Monetary Policy Rules and the Yield Curve." Working paper, 2007. Available at www.ssrn.com.

Bhansali, Vineer, Robert Gingrich, and Francis Longstaff. "Systemic Risk: What Is the Market Telling Us?" Forthcoming in *Financial Analysts Journal*, 2008.

Bhansali, Vineer, and Mark Wise. "Forecasting Portfolio Risk in Normal and Stressed Markets." *Journal of Risk*, Vol. 4, No. 1 (Fall 2001).

Bookstaber, Richard. *A Demon of Our Own Design: Markets, Hedge Funds, and the Perils of Financial Innovation*. Hoboken, NJ: Wiley, 2007.

Fung, William, and David Hsieh. "The Risk of Hedge Fund Strategies: Theory and Evidence from Trend Followers." *The Review of Financial Studies*, Vol. 14, No. 2 (2001), pp. 313–341.

Khandani, Amir, and Andrew Lo. "What Happened to the Quants in August 2007?" Working paper, MIT (November 2007).

Longstaff Francis A, and Arvind Rajan. "An Emperical Analyzis of the Pricing of Collateralized Debt Obligations." *Journal of Finance*, Vol. 63, No. 2 (April 2008) pp. 529.