Risk Management for Event-Driven Funds

Philippe Jorion

Many portfolio strategies are "event driven" (i.e., designed to benefit from price movements caused by corporate events, such as a merger). These strategies involve payoffs with discontinuous and skewed distributions that conventional risk methods do not measure well. This article develops methods to measure the forward-looking risk, based on current positions, of portfolios exposed to such discrete events. The method is applied to independent events and to the more realistic case of events that are not independent. For mergers and acquisitions, empirical estimates of deal-break correlations are positive but low, which implies that most of this event risk is idiosyncratic and diversifiable. The methodology allows assessment of the risk and return of various portfolio structures for event-driven funds.

Many hedge fund strategies are "event driven"; that is, they are intended to benefit from price movements caused by such corporate events as restructurings, bankruptcies, mergers, acquisitions, or other special situations. Such trading strategies involve payoffs that have discontinuous distributions. Either the event happens or it does not, which is a binary distribution. These distributions pose a particular challenge to risk management because of their discontinuous nature, as well as the asymmetry of their payoffs. In addition, the *history* of price movements may not be relevant for measuring risk in event-driven strategies.

Consider, for example, a portfolio manager positioned to take advantage of a proposed acquisition. The acquirer makes an offer that exceeds the current market price by an amount known as the "merger premium." Offers can take the form of cash or stock of the bidding company. Right after the deal announcement, the target company's price jumps up in proportion to this premium. There is still uncertainty, however, as to the success of the deal, which explains why the price of the target company generally trades at a discount to the offer price, known as the "arbitrage spread." The typical position for merger and acquisition (M&A) risk arbitrageurs is to buy the target company and, in the case

of stock deals, short the acquirer. If a deal succeeds, the position generally creates a modest gain.[1] In contrast, in the case of deal failure, the target company's price may fall back to its preannouncement level, which is a big loss.[2] Thus, the position is exposed to an asymmetrical "deal risk." And this risk cannot be measured from historical data. My goal in this article is to develop methods to measure the risk of portfolios exposed to such discrete events.

Current industry methods for managing the risk of such positions are rather elementary. Moore, Lai, and Oppenheimer (2006) described a recent survey of risk management methods used by 21 risk arbitrageurs in mergers and acquisitions. The primary method used to control event risk is position limits. The first type is a limit on the fraction of the portfolio invested in any deal, which is typically 10 percent of the portfolio value. This percentage implies a minimum of 10 positions. In the survey sample, the number of deals in a portfolio varied between 20 and 100, with a median of 30. The second type is a limit on the maximum loss of any position (e.g., 5 percent of the portfolio value if the deal fails). This loss represents a fall in the stock price back to its preannouncement level. These two types of constraints are largely overlapping.

Such single-name limits help diversify the portfolio but are totally ad hoc. The reason is that, as revealed in the Moore survey, M&A arbitrageurs do not measure their *total* portfolio risk. Portfolio risk can be summarized with a single measure of downside risk, such as value at risk (VAR).[3] Risk measures for event-driven funds, however, cannot rely on traditional risk models. Such models do not take into account the uncertainty generated by the currently unfolding events; in addition, the history

Philippe Jorion is Chancellor's Professor of Finance at the University of California, Irvine, and managing director at Pacific Alternative Asset Management Company, Irvine, California.

Note: Pacific Alternative Asset Management Company is a fund-of-hedge-funds investment firm offering strategic alternative investment solutions.

of price movements is not directly relevant for measuring risk. Indeed, Schachter (2006, p. 69) stated that VAR "is not well suited" to these cases. The lack of suitability is certainly the case for conventional VAR measures based on recent historical data. Nevertheless, this article shows that it is perfectly feasible to construct a forward-looking distribution of portfolio profits and losses based on current positions and to use a VAR measure to summarize risk.

Such a measure would allow risk arbitrageurs to gauge optimum diversification, which depends on the number of deals, the expected profit on each deal, and the contribution of each deal to portfolio risk. Also, VAR can be used to determine the optimal amount of economic capital required to support the portfolio or, equivalently, the optimal amount of leverage. Moore et al. (2006) reported that about two-thirds of arbitrageurs use leverage.

To my knowledge, this article is the first to provide quantitative tools to measure the portfolio risk of event-driven funds. The methodology proposed here expands on methods developed for credit portfolios, which also involve discrete events (i.e., defaults). It relies on position information to construct the forward-looking distribution of profits and losses. The approach is illustrated with the use of empirical estimates of deal-break probabilities and correlations for a sample of mergers and acquisitions.

Drawbacks of Conventional Risk Measures

Typically, risk measures are return based. For example, for an M&A fund, one could take the volatility of historical returns to the fund over a recent window. This approach is simple but has severe shortcomings. It does not allow the arbitrageur to react quickly to changes in the portfolio composition or in the trading strategy. In addition, it obscures the structural drivers of risk and, hence, makes it difficult to manage the risk profile of the portfolio.

Recently, position-based risk measures, such as VAR, have become widely used. Conventional VAR measures combine current position information with the recent history of risk factors, which assumes that this history can be applied to the future. Current positions are taken as fixed over the risk horizon. This approach could be termed "conventional position based." Generally, it is a vast improvement over return-based methods. It provides a structural, bottom-up risk model because it can handle changes in portfolio composition and because it can be used to manage portfolio risk. As discussed by Jorion (2007) in the case

of hedge funds, this methodology demonstrates that position information is essential to understanding and managing risks.

In the case of event-driven funds, however, the conventional application of position-based VAR methods can be misleading for a number of reasons. First, the event itself changes the nature of the stochastic process, from a random walk to a discrete-outcome process. In the case of cash deals, a position in the target company is similar to an unsecured bond issued by the acquirer that matures on the closing date. Thus, historical data cannot be used in the usual fashion.

For example, consider the data provided by Mitchell and Pulvino (2001). They reported that over the 1963–98 period, a value-weighted portfolio of M&A deals had an annual realized volatility of 9.3 percent. This number represents the actual risk of a portfolio invested in an average of 31 M&A positions taken after the announcements. The authors did not provide the volatility of the same portfolio before announcements, but surely that amount should be greater than the volatility of a value-weighted market index because the index is better diversified. Over the same period, the market index had a volatility of 15.1 percent, which is a lower limit for the position-based risk of the M&A portfolio. Still, take this number as a conventional risk measure for the M&A portfolio. Then, the actual risk is (15.1 − 9.3)/15.1 = 38 percent lower than the conventional position-based risk. The difference is substantial.

An additional problem is that the recent history of an event includes the announcement of the event itself. This news creates a large jump in the price of the target, which artificially increases the risk measure. The shorter the risk measurement window, the larger the effect. For example, consider an actual M&A hedge fund. The fund is not leveraged and contains more than 150 positions. Over the short period from May 2005 through April 2007, this fund had a return-based volatility of 2.2 percent per year. Conventional risk forecasts were also measured for this fund from monthly position data combined with a four-year window for the risk factors. Over the same period, these volatility forecasts averaged to 9.8 percent. Here again, the bias is substantial. The actual risk turned out to be (9.8 − 2.2)/9.8 = 78 percent less than the conventional position-based risk measure.

Perhaps simple rules of thumb could be developed for correcting conventional position-based risk measures. Nevertheless, such adjustments could not be universal because the relationship between a conventional position-based risk measure and its return-based counterpart depends on

the structure of the portfolio. And the structure involves the number of positions, the distribution of deal payoffs, and historical risk measures. Note, for instance, that the actual risk of the first example M&A portfolio was 9.3 percent, which is much higher than the 2.2 percent risk of the second M&A portfolio (albeit over a much shorter and different period). This difference can be explained by the greater diversification of the second portfolio, which had 150 positions against 31 for the first portfolio.

On the other hand, position-based information can be used directly to construct forward-looking measures of risk that account for the discrete-outcome process that characterizes event-driven funds. This construction is the main purpose of this article.

Measuring Risk from Position Information

Suppose an investor is investing in a stock that is currently in play as a takeover target at a fixed offer price. The distribution of payoffs is characterized by the two states of the world shown in **Table 1**. It is a cash deal in which the stock of the target currently, after the announcement, trades at $100. This analysis ignores directional market risk and focuses on event risk only. The event indicator is b. In the case of failure, or deal break, $b = 1$ with probability $p = 0.15$ and the dollar return, R, is an absolute loss, AL, of –$15. In the case of success, $b = 0$ and the absolute profit is $AP = +$5.

Table 1. Single-Deal Distribution

Measure	Failure (b = 1)	Success (b = 0)
Probability	15%	85%
Payoff (R)	–$15	+$5

This example assumes a probability of failure, $E(b)$, of 15 percent, which is typical of empirical studies of M&A (e.g., Baker and Savasoglu 2002). More generally, the portfolio manager is supposed to assess the probabilities and payoffs of each deal and to add value through skillful selection of deals. In this example, the expected profit is positive at

$$E(R) = 15\% \times (-\$15.00) + 85\% \times (\$5.00) = \$2.00.$$

For simplicity, this $2 is assumed to be above the risk-free return.

The volatility for this Bernoulli distribution is

$$\sigma(b) = \sqrt{p(1-p)}. \qquad (1)$$

The distribution, however, is strongly skewed in the direction of losses.

Here, the horizon is assumed to be the average time needed for the resolution of a deal. Mitchell and Pulvino (2001) indicated that the average horizon of a deal is 59 trading days, or about three calendar months. In practice, the horizon is uncertain. The distribution of payoffs could be extended to include an intermediate case in which the deal is still pending. With the quarterly horizon extrapolated to a year, the expected excess rate of return is[4]

$$\frac{\$2.00 \times 4}{\$100.00} = 8\% \text{ per year.}$$

This example is calibrated to realistic numbers: Mitchell and Pulvino (2001) found that simulated excess returns from risk arbitrage strategies average 9.25 percent per year, which is equivalent to the 2 percent per quarter assumed here.[5] These numbers represent pure alpha because they are adjusted for the risk-free rate and the market beta. The methodology presented here does not, however, depend on these numbers.

The next issue is aggregation at the portfolio level. This type of problem is similar to building a distribution of portfolio credit losses. To simplify, we define L as a positive number representing the loss relative to the success case [e.g., –(–$15 – $5) = $20 for Table 1]. We assume first a homogeneous portfolio: All deals have equal-sized loss L and equal failure probability p.

If the events are independent, building this distribution is straightforward. The total number of deals is N. The total number of failures, k, then follows the binomial distribution, f:

$$f(k; N, p) = \binom{N}{k} p^k (1-p)^{N-k}. \qquad (2)$$

The distribution of total portfolio losses, TL, which is k multiplied by L, is given directly by Equation 2. It is also a binomial distribution, which can be summarized by a dispersion measure, such as VAR, or the lowest quantile at a specified confidence level.

We can now explore the effect of changing the number of deals, N, on the shape of the distribution. To maintain comparability across portfolios, we keep the total dollar exposure fixed for the total portfolio at, say, $100. Hence, the size of each deal, L, is inversely related to N among various portfolio structures. The distribution can be characterized by its VAR at, say, the 95 percent confidence level. (Because the distribution is discrete, I will report the loss associated with a confidence level at least equal to 95 percent.) **Table 2** reports the VAR together with the associated actual confidence level and number of failures. VAR is measured as the negative of the dollar loss over the period and is generally a positive number. The distribution of losses is described in **Figure 1** for $N = 5$, 10, and 100.

**Table 2. Distribution of Independent Deals**
(constant exposure of $100)

No. of Deals	Expected Profit	VAR ≥ 95%			VAR ≥ 99.9%		
		VAR	Actual Confidence Level	No. of Failures	VAR	Actual Confidence Level	No. of Failures
1	$2.00	$15.00	100.0%	1	$15.00	100.0%	1
5	2.00	3.00	97.3	2	11.00	99.99	4
10	2.00	1.00	95.0	3	7.00	99.99	6
20	2.00	1.00	97.8	6	4.00	99.98	9
30	2.00	0.33	97.2	8	2.33	99.92	11
40	2.00	0.00	97.1	10	2.00	99.96	14
50	2.00	−0.20	97.0	12	1.40	99.93	16
100	2.00	−0.80	96.1	21	0.40	99.94	27

Table 2 shows that the expected dollar profit is independent of the number of deals, which simply reflects the fixed total exposure of the portfolio. VAR, however, does change with the number of deals. With a single deal, the minimum 95 percent VAR implies a loss of $15.00, which actually corresponds to a 100 percent confidence level, or 1 failure. With 10 independent deals, the 95 percent VAR goes down to $1.00. With 40 deals, the 95 percent VAR is $0. In other words, there is a low probability (less than 5 percent) of losing any money on this portfolio. The reason is that the expected profit is positive

and the dispersion of payoffs shrinks rapidly as N increases. As the number of deals increases to 100, VAR becomes −$0.80, which means that the probability of making a profit less than $0.80 is low. Thus, increasing the number of independent deals shrinks the spread of the distribution rapidly, as illustrated in Figure 1. As N increases, the binomial distribution converges to a Poisson distribution.

Note also that, because of the discrete nature of outcomes, VAR does not always change smoothly when parameters change. Consider, for instance, VAR at the minimum 95 percent confidence level. It

Figure 1. Distribution of Portfolio Losses with Independence
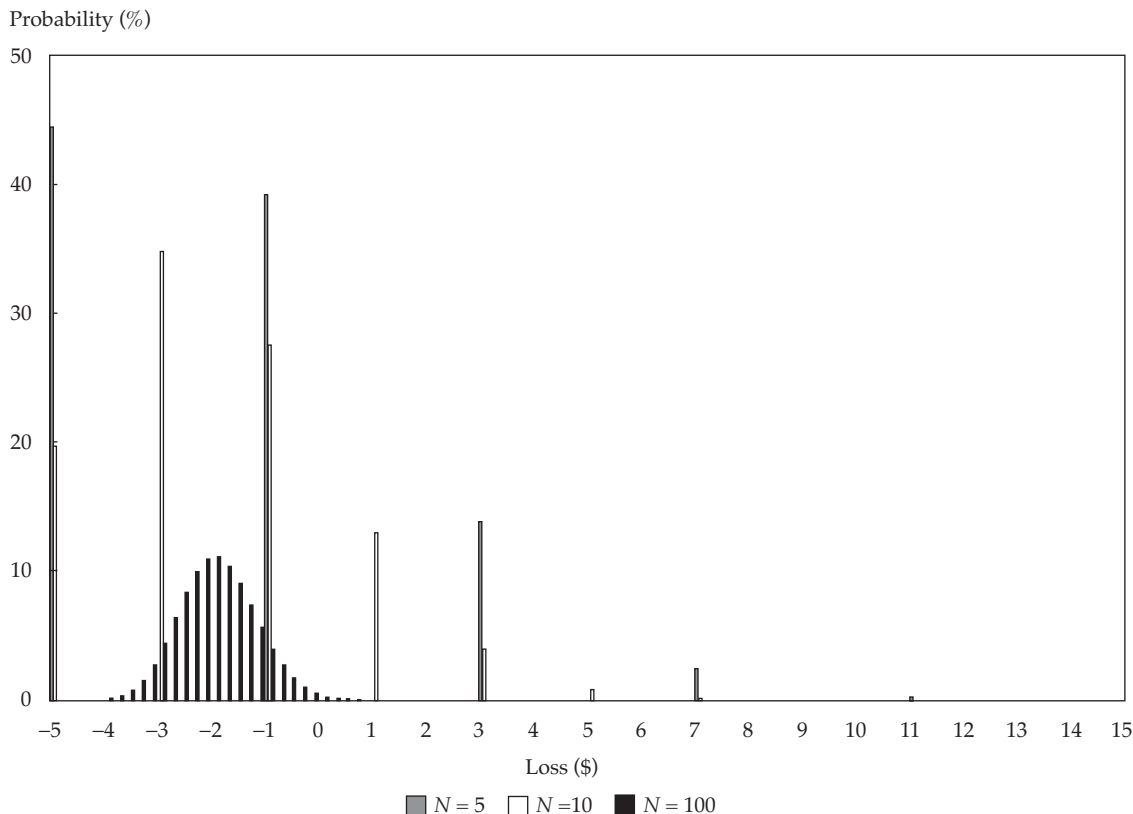
stays at \$1 for $N = 10$ and 20. For $N = 10$, the actual confidence level is just at its required minimum of 95.0 percent. At $N = 20$, the actual confidence level increases to 97.8 percent, which corresponds to six deal failures. The confidence level for five deal failures is still below 95 percent. As N increases further, VAR eventually drops, but in discrete steps. Thus, VAR is not a smooth function of the number of deals.

This distribution could be used to infer the maximum amount of leverage a fund could safely take on. If the fund wants to maintain a single-A credit rating, the confidence level can be taken as 1 minus the default probability over a year. Credit-rating agencies report a historical default rate of approximately 0.1 percent for A rated companies, which implies a confidence level of 99.9 percent.[6] This VAR measure, which can be interpreted as economic capital, is given in the right-hand columns of Table 2.

Table 2 shows that for a single deal, the 99.9 percent VAR is \$15.00. So, if the notional value of the position is \$100.00, the portfolio can support a maximum leverage of \$100.00/\$15.00, or 6.7 times. For 10 deals, economic capital shrinks to \$7.00. Thus, leverage could go to \$100.00/\$7.00, or 14.3 times. This would generate a rate of return on economic capital of \$2.00/\$7.00, or 28.6 percent per quarter.[7] With 100 deals, economic capital is \$0.40. Because this number is so low, leverage can increase even more, generating a rate of return of 500 percent. Of course, this very high number is unrealistic because it assumes independent deals. I examine the effect of dependencies next.

Empirical Estimation of Break Distributions

In practice, the events in the portfolio may not be independent. A large fall in the market may lead to many deals breaking during the same period. Mitchell and Pulvino (2001) showed, for example, that a 20 percent decrease in the market return increases the probability of deal failure by about 9 percent. This contemporaneous dependency of deal breaks on the market is akin to positive correlation across events.

This section provides empirical estimates of the deal-break probability and the average break correlation. For the study reported here, I considered a large sample of M&A deals, 1,765, involving at least one North American company and concluded between 1997 and 2006. The sample includes only deals with an announced total value more than \$500 million. The data were collected from Bloomberg

and are summarized in **Table 3**. Deals were tabulated by their outcome every quarter. The empirical analysis does not consider actual arbitrage spreads. In practice, the portfolio manager would also need to evaluate the size of each spread in relation to the break probability for individual deals.

Table 3. Description of M&A Deals, 1997–2006

Outcome	No. of Deals Completed per Quarter	Average Value of Deals (\$ millions)	Mean Completion Days for Deals
Failure	5.3	5,377	138
Success	38.8	3,385	137
Total	44.1		

In this 40-quarter period, 212 deal breaks and 1,553 completed deals occurred, which gives a 12 percent failure rate. The average number of deals completed in a quarter is 44.1, with a mean completion time of 138 calendar days, or about 4.5 months.

To estimate correlations between deal failures, I used a standard methodology developed for credit events.[8] Initially, each deal was assumed to have the same unconditional probability of failure p and all pairwise correlation coefficients ρ were assumed to be equal. These assumptions can be relaxed easily.

Suppose that we observe a number of failures F_t during a period t out of a total of N_t observations. The probability of failure can then be estimated by

$$p = E(b) = \sum_{t=1}^{T} w_t \frac{F_t}{N_t}, \qquad (3)$$

where w_t is the weight assigned to each observation, here taken as $(1/T)$. The probability of a joint deal break, p_{12}, can be estimated from the ratio of the total number of pairs breaking during the period, or $F_t(F_t - 1)/2$, to the total number of pairs combinations, or $N_t(N_t - 1)/2$. De Servigny and Renault (2002) showed via simulations that small-sample properties of the joint default probability estimator are slightly better for

$$p_{12} = E(b_1 b_2) = \sum_{t=1}^{T} w_t \frac{(F_t)^2}{(N_t)^2}. \qquad (4)$$

The deal-break correlation can be constructed from the covariance between the two events:

$$
\begin{aligned}
\mathrm{cov}(b_1 b_2) &= \rho[\sigma(b)\sigma(b)] \\
&= E(b_1 - p)(b_2 - p) \\
&= E(b_1 b_2) - p^2 \\
&= p_{12} - p^2.
\end{aligned}
\qquad (5)
$$

Replacing $\sigma(b)$ by the standard deviation of a Bernoulli variable, we find the correlation can be written as

$$\rho = \frac{p_{12} - p^2}{\left[\sqrt{p(1-p)}\right]\left[\sqrt{p(1-p)}\right]}. \tag{6}$$

Note that the deal-break correlation is also related to the variance of the time series of fraction of breaks during each period. We define this fraction as

$$x_t = \frac{F_t}{N_t}. \tag{7}$$

A sharply fluctuating x_t implies clustering of failures, or positive correlation within quarters. Using Equations 3, 4, and 5, and ignoring the degrees-of-freedom adjustment, we find that the variance of x can be written as

$$\begin{aligned}
\sigma^2 &= \frac{1}{T} \sum_{t=1}^{T} (x_t - \bar{x})^2 \\
&= \left[\frac{1}{T} \sum_{t=1}^{T} x_t^2\right] - \bar{x}^2 \\
&= p_{12} - p^2 \\
&= \rho\left[p(1-p)\right].
\end{aligned} \tag{8}$$

So, a high variance is equivalent to a high failure correlation.

Figure 2 shows the fraction of deals failing each quarter (Q). From Equation 3, the average break probability is estimated to be 12 percent. There is little first-order autocorrelation in the series, which can be assumed to be independent across quarters. Figure 2 reveals substantial variation, however, in the fraction of deals failing over time. It reached a peak of 27 percent in Q1 2003. The portion was zero for two quarters—Q1 1997 and Q1 2004. This variation can be interpreted in terms of clustering of failures, or correlations across deal breaks within quarters.

Using Equation 4 to estimate p_{12} and using Equation 6 to calculate the correlation, we find the break correlation is estimated to be 0.03. The failure probability is 0.12. To assess the reliability of these estimates, I simulated their distribution by bootstrapping the sample of 40 quarterly fractions. Based on 1,000 replications, the 95 percent significance interval for the break probability is (0.102, 0.137) and the interval for the break correlation is (0.017, 0.045). Therefore, the break correlation is positive and significantly different from zero. This finding justifies the use of an average break correlation of 0.03 in the following section.

More generally, the portfolio manager could model the probability of failure as a function of a number of deal variables.[9] Here, the focus is on the aggregate fraction of breaks, not individual deal data, and I have used stock market returns and credit spreads as independent variables.

A distinction can be made between "predictive regressions," which model predictable time variation in the failure probability, and "explanatory regressions," which model correlations through the influence of common risk factors. The results are presented in **Table 4**.

Figure 2.   Fraction of Deals Failing, Quarterly 1997–2006
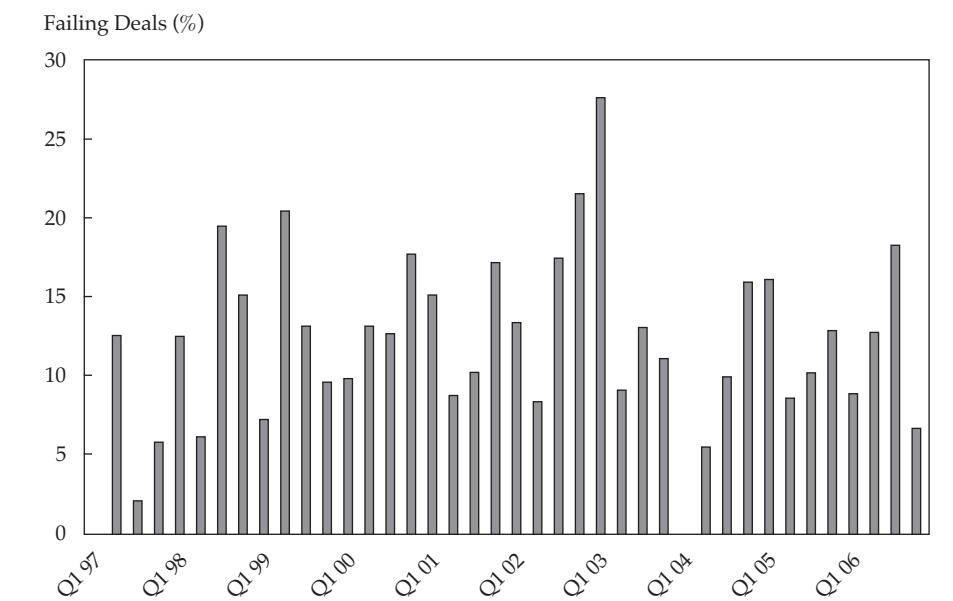
Table 4. Explaining and Predicting Fraction of Breaks for M&A Deals, 1997–2006

(standard errors in parentheses)

Independent Variable	Explanatory Regressions		Predictive Regressions		
	Model 1	Model 2	Model 1	Model 2	Model 3
Constant	0.12	0.10	0.13	0.05	0.07
	(0.01)	(0.01)	(0.01)	(0.02)	(0.03)
$RM(t)$	–0.08				
	(0.10)				
$RM(t) > 0$		0.13			
		(0.18)			
$RM(t) < 0$		–0.36*			
		(0.22)			
$RM(t-1)$			–0.26**		–0.14
			(0.10)		(0.11)
$CS(t-1)$				1.22**	0.94**
				(0.37)	(0.43)
R^2 (%)	1.7	7.1	15.5	22.0	25.3

Note: *RM* is the total return on the S&P 500; *CS* is the credit spread between the Lehman U.S. High-Yield and Treasury indices.

  *Significant at the 10% level.
**Significant at the 5% level.

The first column in Table 4 reveals a negative coefficient on the stock market return during the same quarter, $RM(t)$. That is, the fraction of breaks increases when the stock market falls during the same quarter. Splitting up market returns into positive and negative values shows that the effect is asymmetrical. Positive returns are basically uncorrelated with deal breaks, presumably because acquirers always have the option to revise their offer price upward. Large negative market returns, however, are significantly associated with a greater proportion of deal breaks. As expected, a large drop in the stock market reduces the value of the target and makes it more likely that the acquirer will walk away from the deal.

The –0.36 coefficient on $RM(t) < 0$ means that when the stock market drops by 20 percent, the break probability increases by 7 percent. This conditional dependence on the market explains the positive correlations across deal breaks.[10] After conditioning on this common factor, the break correlation becomes unimportant. In practice, using a simple measure of unconditional positive correlation is sufficient because market downturns cannot be predicted reliably.

Focusing next on the predictive regressions in Table 4, we see that the coefficient on the stock market return during the previous quarter, $RM(t-1)$, is negative and strongly significant, indicating that falling stock prices increase the fraction of breaks the next quarter. In addition, the coefficient on the credit spread at the start of the quarter, $CS(t-1)$, is positive and highly significant. As expected, higher financing costs lead to a greater fraction of deals failing. Even though these two variables are negatively correlated, both appear with the correct sign in the multiple regression in the last column of Table 4. The result is a predicted break probability that varies between 9 percent and 18 percent over this period.

Note that this calibration exercise is limited to large North American deals. To lower the failure correlation, a portfolio manager could include in the portfolio M&A deals from smaller companies and companies outside North America. More generally, the methodology described next can be used to assess the benefits from searching for lower correlations.

Building Portfolio Distributions with Correlated Events

The dispersion of the distribution of losses shrinks rapidly as the number of independent deals increases. In practice, however, positive correlations temper this reduction in risk. The question is how to adjust the measure of portfolio risk. Because M&A events are similar in structure to credit events, the methodology of credit risk models can be used to measure event risk.

Various methods have been proposed to measure the risk of credit portfolios. One approach is the "binomial expansion technique" (BET) advocated

by Moody's Investors Service (1996) to evaluate collateralized debt obligations. This technique approximates the actual distribution of N correlated events with a simpler binomial distribution for another number, D, of uncorrelated events. D is called the "diversity score." It is computed for each deal i from event probabilities p, deal size L, and bivariate correlations ρ:

$$D = \frac{\left(\sum_{i=1}^{N} p_i L_i\right)\left(\sum_{i=1}^{N} (1 - p_i) L_i\right)}{\sum_{i=1}^{N} \sum_{j=1}^{N} \rho_{ij} \sqrt{p_i \left(1 - p_i\right) p_j \left(1 - p_j\right)} L_i L_j}. \quad (9)$$

Equation 9 is obtained by matching the first and second moments of the true distribution with the binomial distribution, $f(k; D, p)$. It is, therefore, an approximation only. As previously, VAR can be computed for this new distribution.

For illustration, we assume that all probabilities, all losses, and all correlations are equal, even though this need not be the case. The diversity score then reduces to

$$\begin{aligned} D &= \frac{(NpL)\left[N(1-p)L\right]}{\left[N + \rho(N-1)N\right]\sqrt{p(1-p)p(1-p)}LL} \\ &= \frac{N^2}{\left[N + \rho(N-1)N\right]}. \end{aligned} \quad (10)$$

If the events are independent, all correlations are zero and Equation 9 reduces to $D = N$, as one would expect. More generally, a positive correlation leads to $D < N$, which implies that the distribution has longer tails than with independent events.

Table 5 compares the distributions of a portfolio of $N = 30$ deals for various correlation coefficients. The case of zero correlation is the same as previously, with a 95 percent VAR of $0.33. The diversity score decreases with the correlation. For example, when the average correlation increases to 0.03, the diversity score falls from 30 to 16. Accordingly, the 95 percent VAR increases sharply—from $0.33 to $1.25. With a correlation of 0.10, the 95 percent VAR moves to $3.57. As expected, higher correlations

lead to greater risk. When the correlation is 0.50 or above, VAR moves to $15.00, which is the same as for a portfolio of a single deal. In this case, the portfolio receives no diversification benefit.

Figure 3 compares the distribution of losses with $N = 30$ in two cases, $\rho = 0$ and $\rho = 0.05$. The distributions are heavily skewed. In each case, the expected return is $2. With zero correlation, the right tail falls off quickly. With positive correlation, the tail is much longer. Thus, a modest amount of correlation has a substantial impact on portfolio risk.

Methods more general than the BET approach can be used. For instance, the distribution of events can be constructed from Monte Carlo simulations. In this framework, each event is driven by a latent random variable such that a drop below some cutoff point simulates a break with a specified probability. Correlations can be induced between these random variables so as to generate specified default correlations.[11] Such a bottom-up approach, which is similar to structural models of credit risk, can be used to derive a more general distribution of gains and losses for the event-driven portfolio. This approach will be used in the next section.

As an illustration of the effect of number of deals on the risk profile of the portfolio, we can perform a simple experiment based on the actual number of deal breaks and successes in the 1997–2006 period. The average number of deals per quarter is 44. The first portfolio is equally invested in all deals. Consistent with the previous analysis, a success is assigned a return of +5 percent; a failure has a return of –15 percent. **Figure 4** displays the hypothetical time series of returns on this portfolio. During the first quarter, for example, the failure rate is exactly 0 percent. The return is then computed as

$$\begin{aligned} R &= 0\%(-15\%) + (1 - 0\%)(+5\%) \\ &= 5\%, \end{aligned}$$

which is the highest feasible value. For the entire sample, the average return is close to 2 percent. This all-deal portfolio has low volatility. In this experiment, it loses money in only 1 quarter out of 40.

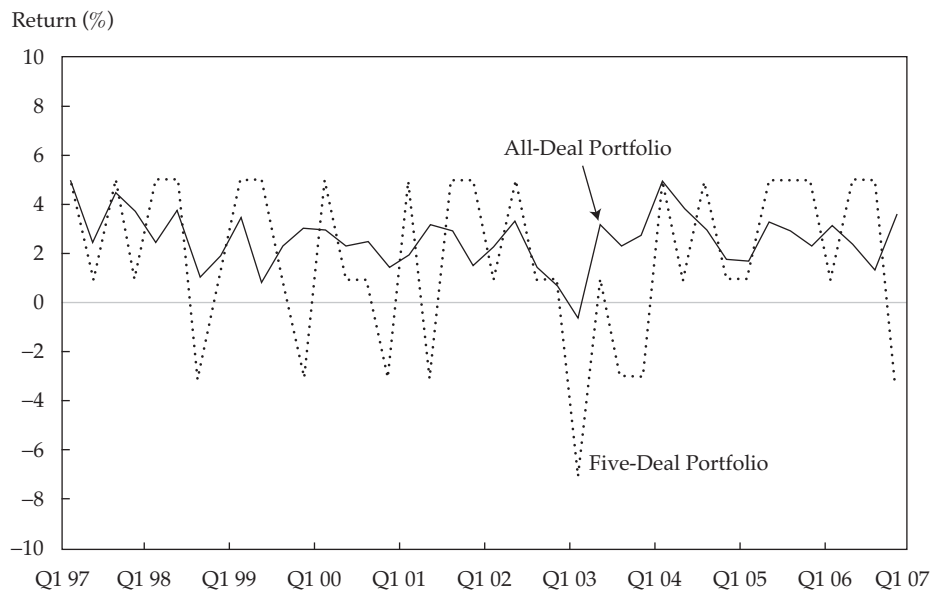**Table 5. Distribution of Dependent Deals in the BET Approach** ($N = 30$)

Correlation	Diversity Score	Expected Profit	VAR ≥ 95%	VAR ≥ 99.9%	
				VAR	Actual Confidence Level
0.00	30.0	$2.00	$ 0.33	$ 2.33	99.92%
0.03	16.0	2.00	1.25	5.00	99.98
0.05	12.2	2.00	1.67	5.00	99.93
0.10	7.7	2.00	3.57	9.29	99.99
0.20	4.4	2.00	5.00	10.00	99.95
0.50	1.9	2.00	15.00	15.00	100.00

Figure 3. Distribution of Portfolio Losses with Various Correlations
($N = 30$)

Distribution (%)

VAR = $5.0

VAR = $2.3

Loss ($)

■ Correlation = 0   □ Correlation = 0.05
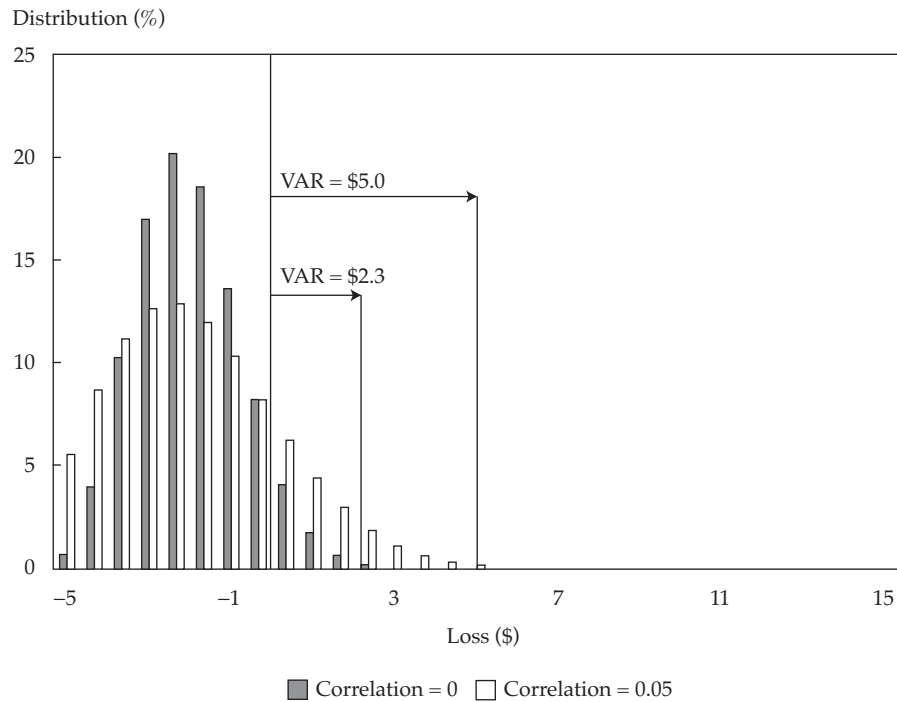
Figure 4. Time Series of Portfolio Returns for Two Portfolio Sizes

Return (%)

All-Deal Portfolio

Five-Deal Portfolio

Q1 97   Q1 98   Q1 99   Q1 00   Q1 01   Q1 02   Q1 03   Q1 04   Q1 05   Q1 06   Q1 07

The second portfolio in Figure 4 consists of five deals. Each quarter, the outcome was selected randomly by using the deal-break fraction in that same quarter as the probability of failure. This approach preserved the pattern of correlations in the data. Figure 4 shows a typical outcome. This five-deal portfolio is much more volatile than the all-deal one, even though it has the same mean. It suffers losses

during eight quarters. The worst loss is –7 percent. Clearly, the portfolio with a greater number of deals is the safer one. Another way of looking at it is that the all-deal portfolio could be leveraged three times and have the same level of risk as the five-deal portfolio while delivering three times the expected return.

In summary, to construct the distribution of returns for an actual fund, the portfolio manager

needs to estimate the distribution of returns for each deal. This distribution rests on (in addition to the deal-break correlation) the size of the position, the probability of the deal breaking, and the payoffs in the two states. This information can be used to construct Monte Carlo simulations of portfolio returns.

Assessing Economic Capital and Leverage

This methodology can be used to infer economic capital levels or, conversely, the number of deals required to keep a fixed level of risk for event-driven portfolios. Leverage can be assessed from economic capital. Recall that with one deal only, economic capital is $15, implying a maximum leverage of $100/$15 = 6.7 times, at which point, the rate of return on economic capital is $2/$15 = 13.3 percent per quarter.

For a 10-deal portfolio with correlation of 0.03, VAR is $9 at a minimum 99.9 percent confidence level. As a result, the portfolio could be leveraged by a factor of $100/$9, or about 10 times the original equity investment, generating an expected return of $2/$9 = 22 percent.

This analysis shows that increasing the number of deals substantially lowers risk because deal risk is largely idiosyncratic. When N increases to 30 and 100, economic capital drops to, respectively, $5.00 and $3.80. **Figure 5** illustrates how economic capital (interpreted as VAR) decreases

with the number of deals. Lower risk allows greater leverage, increasing expected returns.

This powerful diversification effect is driven by the low correlation, taken as 0.03 here. Next, **Figure 6** shows how correlation affects economic capital (as VAR) for various portfolio sizes. The effect is rather dramatic. For $N = 100$, increasing the correlation from 0 to 0.10 pushes VAR from $0.40 to $9.00. Beyond a correlation of 0.30, there is basically no diversification effect: VAR converges to the worst loss of $15.00. Based on historical experience, however, average correlations are very unlikely to reach such values. The previously computed 95 percent confidence interval for ρ ranges from 0.017 to 0.045. This translates into a confidence interval of $2.30 to $5.20 for VAR with $N = 100$. So, even for this break correlation range, there is still a fair amount of diversification.

Economic capital is sensitive also to the break probability but less so than to correlations. **Figure 7** shows that if the probability changes from the base number of 15 percent to 20 percent, VAR (or economic capital) increases from $3.80 to $5.00 when $N = 100$. The 95 percent confidence interval for the estimated break probability is (10.2 percent, 13.7 percent) for this sample. So, the worst increase from the midpoint probability is only 2 percent. Increases in the break probability are important, however, because they directly reduce the expected return on the portfolio, keeping the arbitrage spread fixed.

Figure 5.   Economic Capital (VAR) and Portfolio Size
(break correlation = 0.03)
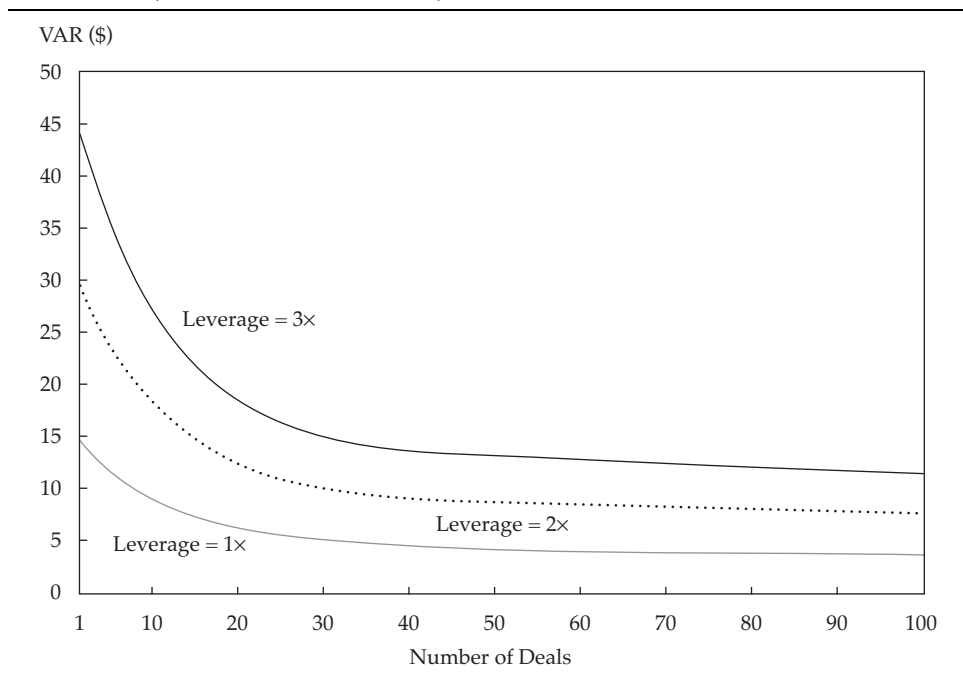
Figure 6.   Economic Capital (VAR) and Break Correlation
(break probability = 15 percent)
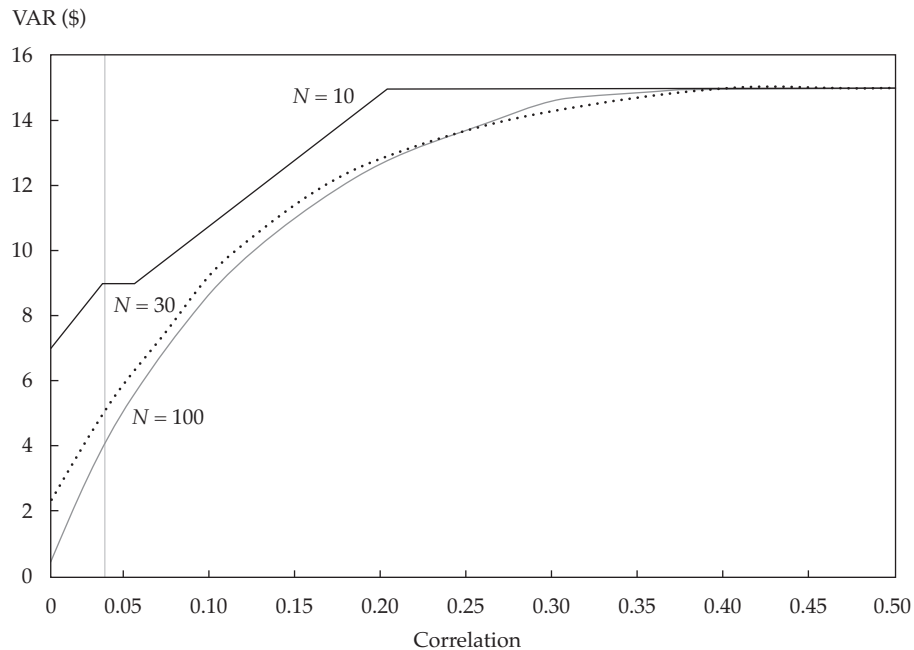
Figure 7.   Economic Capital (VAR) and Break Probability
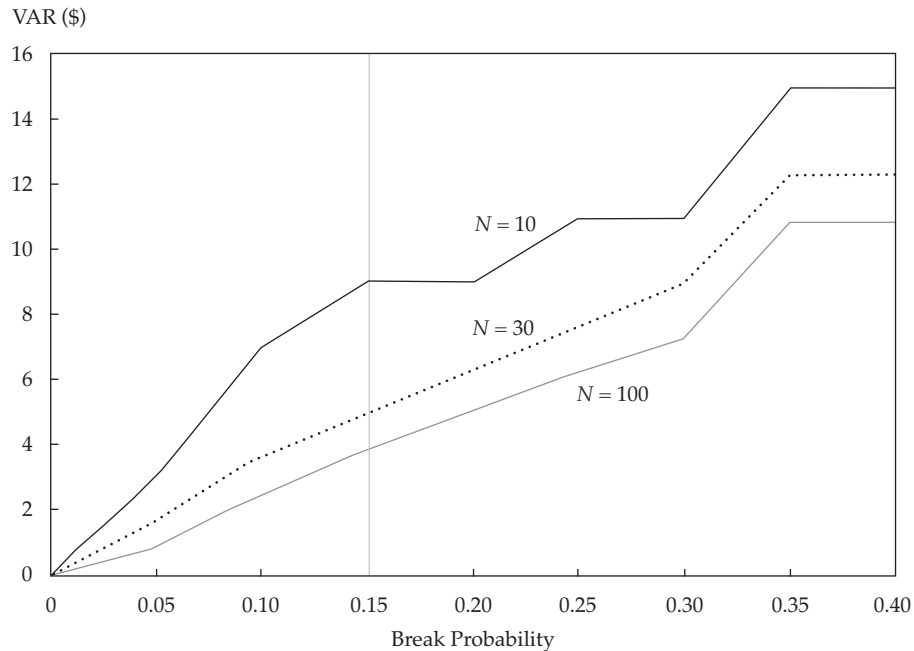(break correlation = 0.03)

Figure 8 illustrates the trade-off between expected return and risk for various levels of leverage and portfolio size under the assumption of a baseline correlation of 0.03. Different levels of leverage are represented by horizontal lines. For a given portfolio size N, expected (excess) returns and risk increase linearly with leverage. A portfolio with $N = 100$ can be three times leveraged and still have a level of risk similar to a portfolio with $N = 10$ while delivering three times the expected return.

This analysis seems to lead to the conclusion that having a larger number of deals is always an improvement. This inference arises because we assumed that the expected profit per deal remains

Figure 8.   Trade-Off between Expected Return and Economic Capital
(break correlation = 0.03)
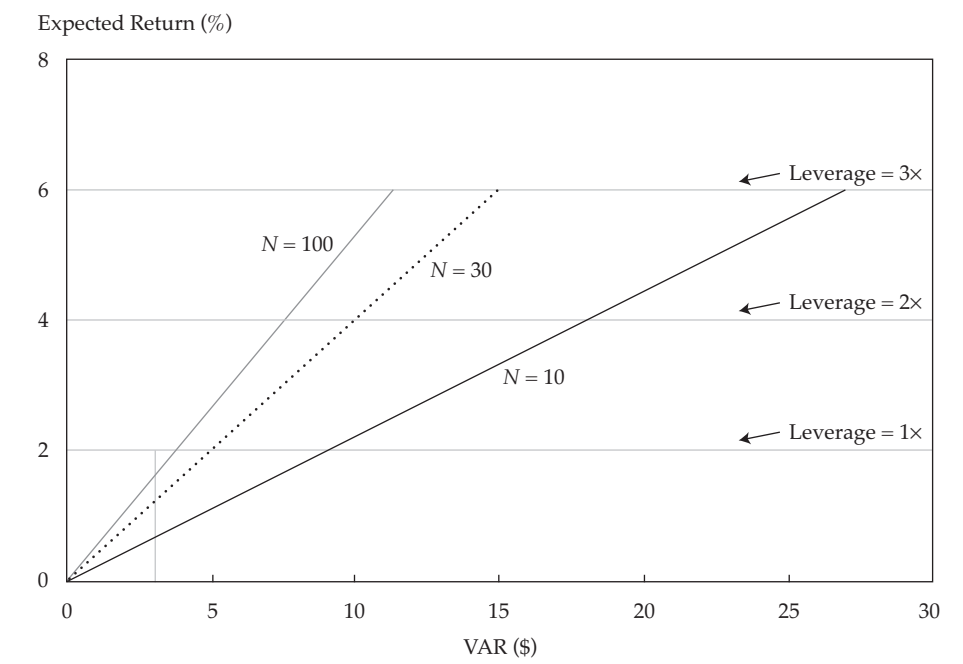
constant at $2. In practice, portfolio managers have expertise in evaluating some types of deals more than others. Thus, when the number of deals to evaluate is increased, the profit per deal should decrease to reflect information acquisition costs. Therefore, some number of deals should be optimal as a reflection of the trade-off between risk and profits. Without a measure of portfolio risk, however, such assessment is not even possible.

Finally, note that the analysis of economic capital should also account for nonevent risks, including market risk and some idiosyncratic risk.[12] These additional risks can be incorporated in an enlarged model. Stress tests should also be performed to evaluate the sensitivity of the results to changes in input parameters, including break probabilities and correlations.

Conclusion

Event-driven portfolio strategies face discrete deal risk. In contrast to other types of market risks, where distributions are continuous, deal risk is represented by a binary variable, which is the success or failure of a deal. This type of risk cannot be captured by conventional position-based risk measures.

Current positions, however, can be used to construct the distribution of portfolio returns. To do so, the risk manager needs to estimate the probability of success for each deal, the payoffs from success and failure, and the joint correlations across deals. I have described new methods developed to construct quantitative measures of portfolio risk for event-driven portfolios.

Once constructed, this distribution can be summarized by the usual risk measures, such as VAR, which can then be used to assess acceptable levels of leverage. I found that in the case of independent deals, the economic capital required to support a portfolio of 30 independent deals is six times lower than that needed for a single-deal portfolio. Thus, the diversified portfolio could be levered six times more than a single-deal position and still maintain the same level of risk. At the same time, this higher leverage implies a higher return on capital.

In practice, diversification benefits are limited because of contemporaneous clustering of deal breaks. This article provided empirical estimates indicating that the average deal-break correlation is around 0.03 for large North American deals. This positive correlation decreases the benefits from diversification and increases the required economic capital. Even so, a portfolio of 100 deals can be leveraged three times and still have a level of risk similar to that of a 10-deal portfolio with no leverage. Thus, diversifying among deals reduces risk sharply.

This methodology can be used to improve current risk management practices for event-driven portfolios. Without such quantitative methods, assessing the optimal risk–return profile

of the portfolio is difficult. On the one hand, increasing the number of deals provides diversification benefits. On the other hand, with limited expertise or resources, increasing the number of deals is costly because it decreases the average expected return per deal. The optimal portfolio represents the point at which the marginal cost equals the marginal benefit from increasing deals. Without a formal measure of the marginal benefit, this optimal position cannot be identified. Armed with these portfolio measurement tools the port-folio manager can search for deals that are profitable yet not too correlated with each other.

In short, this article provides the means to assess the costs and benefits of different structures for event-driven portfolios.

This article qualifies for 1 CE credit.

Notes

1. Success, however, does not simply mean that the original bidder takes control of the target. Other bidders can intervene in the meantime. Success may be defined as a situation in which the price of the target goes to or above the offer price as a result of any acquisition. Equivalently, failure is when the target remains independent.
2. The drop can be more than the initial run-up in price if the deal reveals negative information about the target. In the case of positions that involve both the target and the acquirer, the payoff on the strategy depends also on the price movement in the acquirer's stock.
3. See Jorion (2006) for an exposition of VAR.
4. Various explanations have been advanced for this "risk arbitrage premium." One is that risk arbitrageurs need compensation for bearing event risk. The information-rich environment of risk arbitrage requires specialized expertise (e.g., trading and legal), which needs compensation. Another explanation is that some investors are forced to sell because of investment guidelines that would be breached after a takeover. For instance, a required minimum level of portfolio yield could be lacking if the merged company pays no dividend. Finally, these strategies have nonlinear exposures to the market because the break probability increases when the market falls. This characteristic makes the strategy akin to a short put position, which should command a risk premium.
5. After estimated (and conservative) costs, however, they reported an average return of about 3.5 percent.
6. Under the newly established Basel II rules, commercial banks need to maintain regulatory capital in excess of their 99.9 percent credit VAR over a horizon of one year. In the example here, the horizon is one quarter, so the annual default rate is even lower than that for A rated credits.
7. Normally, financing costs should be taken into account. The estimate of 2 percent expected return, however, is already in excess of the risk-free rate, so there is no need for further adjustment. The only portion of financing costs that is ignored is the spread between the borrowing and deposit rate.
8. See, for example, Bahar and Regal (2001); De Servigny and Renault (2002).
9. Mitchell and Pulvino (2001) estimated the probability of deal breaks from individual deal characteristics and general market factors. For example, "hostility," the most important indicator, increased the probability of failure by 12.8 percent. They also found that the beta to the S&P 500 Index is greater on the downside than on the upside.
10. Note that a similar factor structure is used for credit portfolio models, where correlations between defaults are typically induced by one-factor models.
11. In our example, with $p = 15$ percent, we would need a correlation of latent variables calibrated at 0.068 to obtain the default correlation of 0.03 observed when multivariate normal random variables are used.
12. Horizon risk also exists; it is the risk of a lengthening of the horizon for the deal resolution. This lengthening happened in 2004, for instance, when the effects of the Sarbanes–Oxley Act temporarily slowed down deal activity.

References

Bahar, Reza, and Krishnan Regal. 2001. "Measuring Default Correlation." *Risk*, vol. 14, no. 3 (March):129–132.

Baker, Malcolm, and Serkan Savasoglu. 2002. "Limited Arbitrage in Mergers and Acquisitions." *Journal of Financial Economics*, vol. 64, no. 1 (April):91–115.

de Servigny, Arnaud, and Olivier Renault. 2002. "Default Correlation: Empirical Evidence." Working paper, Standard & Poor's.

Jorion, Philippe. 2006. *Value at Risk: The New Benchmark for Managing Financial Risk.* New York: McGraw-Hill.

———. 2007. "Risk Management for Hedge Funds with Position Information." *Journal of Portfolio Management*, vol. 34, no. 1 (Fall):127–134.

Mitchell, Mark, and Todd Pulvino. 2001. "Characteristics of Risk and Return in Risk Arbitrage." *Journal of Finance*, vol. 56, no. 6 (December):2135–2175.

Moody's Investors Service. 1996. "The Binomial Expansion Technique Applied to CBO/CLO Analysis." *Moody's Special Report* (December).

Moore, Keith, Gene Lai, and Henry Oppenheimer. 2006. "The Behavior of Risk Arbitrageurs in Mergers and Acquisitions." *Journal of Alternative Investments*, vol. 9, no. 1 (Summer):19–29.

Schachter, Barry. 2006. "Limits on Far VAR." *Risk*, vol. 19, no. 3 (March):69.