

Risk Management Lessons from Long-Term Capital Management

Philippe Jorion

Draft: June 1999

Thanks are due to Neil Pearson and Neal Stoughton for useful comments.

Correspondence can be addressed to:
Philippe Jorion,
Graduate School of Management,
University of California at Irvine,
Irvine, CA 92697-3125,
(949) 824-5245, FAX: (949) 824-8469,
E-mail: pjorion@uci.edu

©1999 P. Jorion

Risk Management Lessons from Long-Term Capital Management

ABSTRACT

The 1998 failure of Long-Term Capital Management (LTCM) is said to have nearly blown up the world's financial system. For such a near-catastrophic event, the finance profession has precious little information to draw from. By piecing together publicly available information, this paper draws lessons from risk management practices at LTCM.

LTCM's strategies are analyzed in terms of the fund's Value at Risk (VAR) and the amount of capital necessary to support its risk profile. The paper shows that LTCM had severely underestimated its risk due to its reliance on short-term history and risk concentration. LTCM also provides a good example of risk management taken to the extreme. Using the same covariance matrix to measure risk and to optimize positions inevitably leads to biases in the measurement of risk. This approach also induces the strategy to take positions that appear to generate "arbitrage" profits based on recent history but also represent bets on extreme events, like selling options. Overall, LTCM's strategy exploited the intrinsic weaknesses of its risk management system.

JEL Classification Numbers: G-11, G-13, G-14, G-23

Risk Management Lessons from Long-Term Capital Management

The 1998 failure of Long-Term Capital Management (LTCM) is said to have nearly blown up the world's financial system.¹ Indeed the fund's woes threatened to create major losses for its Wall Street lenders. LTCM was so big that the Federal Reserve Bank of New York took the unprecedented step to facilitate a bailout of the private hedge fund, out of fear that a forced liquidation might ravage world markets.

For such a near-catastrophic event, the finance profession has precious little information to draw lessons from this failure. No doubt this is due to the secrecy of the hedge fund, which never revealed information about its positions, even to its own investors.

As LTCM is now in the process of rehabilitating itself, it is slowly disclosing information about its risk management practices. By piecing together publicly available information, this paper attempts to draw lessons from the LTCM failure.²

This issue is important as LTCM's failure has been widely ascribed to its use of Value at Risk (VAR).³ If so, the disturbing implication is that the method currently used to set capital adequacy requirements for the banking sector is woefully inadequate. Indeed commercial banks are now allowed, under the Basle framework, to use their own internal model as the basis for determining the amount of capital they should set aside against market risks. If these models are as flawed as LTCM's, the failure of the hedge fund can only presage calamities for the world financial system.

This paper is structured as follows. Section 1 first presents an overview of the LTCM saga. As the fund was highly leveraged, the key issue was the choice of the appropriate capital base. LTCM failed because it did not have enough equity capital. Section 2 reviews how Value at Risk can be used to assess the capital base needed to support a leveraged portfolio. This leads into the risk management practices at LTCM, which are analyzed in Section 3. In effect, LTCM used the tools of portfolio optimization to structure its portfolio, leveraging it by a factor of 25 to take advantage of so-called "arbitrage" trades.

This application, however, is fraught with danger. VAR has been primarily developed to measure and control risks. Optimizing a portfolio risk/return profile and using the resulting

¹Chairman Alan Greenspan (1998) testified that "Had the failure of LTCM triggered the seizing up of markets, substantial damage could have been inflicted on many market participants, including some not directly involved with the firm, and could have potentially impaired the economies of many nations, including our own."

²This paper expands on press reports based on interviews and presentations of the principals that appeared in the *Wall Street Journal*, the *New York Times*, and *Derivatives Strategy*.

³See for instance the *Economist*, November 14, 1998 "Risk Management: Too Clever by Half."

VAR to measure risk leads to serious optimization biases, which are illustrated in Section 4. Section 5 also shows that the nature of the bets taken by LTCM was singularly unidirectional. Finally, the last section provides some concluding comments.

1 How LTCM Lost its Capital

The seeds of Long-Term Capital Management (LTCM) can be traced to a highly profitable bond-arbitrage group at Salomon Brothers ran by John Meriwether.⁴ LTCM was founded by Meriwether in 1994, who left after the 1991 Salomon bond scandal.⁵

Meriwether took with him a group of traders and academics, who had been part of Salomon's bond-arbitrage group that had racked up billions of dollars in profits. Together, they set up a "hedge" fund using similar principles to those they had been using at Salomon.⁶

A *hedge fund* is a private investment partnership fund that can take long and short positions in various markets and is accessible only to large investors. As such, hedge funds are not regulated by the SEC. Of course, the term "hedge" is somewhat of a misnomer, if not misleading, since these investment vehicles are leveraged can be quite risky.

Initially, the new venture was eminently profitable. Capital grew from \$1 billion to more than \$7 billion by 1997. The firm was charging sky-high fees consisting of an annual charge of 2% of capital plus 25% of profits. By comparison, other hedge funds charge a 1% fixed fee and 20% of profits; the typical mutual fund fee is about 1.41%. By 1997, total fees had grown to about \$1.5 billion. LTCM's 16 partners had invested roughly \$1.9 billion of their own money in the fund.

Much has been said about LTCM's positions in the press. LTCM was supposed to have wagered \$125 billion. This represents the *total assets* of the fund, most of it borrowed.⁷ Compared to equity of about \$5 billion only, the size of assets appears ludicrous.

⁴The antics of the Salomon culture are vividly described in a colorful description of Wall Street, *Liar's Poker*, by Michael Lewis.

⁵Paul Mozer, a Salomon bond trader who reported to Meriwether, had tried to "corner" the primary Treasury auction market by submitting bids in excess of the firm's limit. When the Federal Reserve Bank learned that it had not been adequately informed of the infraction, it threatened to ban Salomon from participating in the primary Treasury market. These news led to a "run" on Salomon. Investors started to dump the firm's debt. As Salomon was highly leveraged, with \$150 billion in assets financed by only \$4 billion of equity and \$16 billion of debt, its very survival was in doubt. After Salomon's top management resigned and Warren Buffett took over as chairman, the ban was reversed, allowing Salomon to survive the crisis. Meriwether was fined \$50,000 by the SEC amid allegations that he failed to supervise traders properly.

⁶For a good review of LTCM's downfall, see Dunbar (1998).

⁷According to the President's Working Group report (1999), the fund had 60,000 trades on its books. The balance sheet consisted of over \$50 billion of long positions and short positions of an equivalent magnitude.

Even more scary was the off-balance-sheet position, including swaps, options, repurchase agreements and other derivatives, that added up to a *notional principal* amount of over \$1 trillion.⁸ Many of these trades, however, were offsetting each other, so that this notional amount is practically meaningless. What mattered was the total risk of the fund.

LTCM was able to leverage its balance sheet through sale-repurchase agreements (repos) with commercial and investment banks. Under “repo” agreements, the fund sold some of its assets in exchange for cash and a promise to repurchase them back at a fixed price on some future date. Normally, brokers require collateral that is worth slightly more than the cash loaned, by an amount known as a *haircut*, designed to provide a buffer against decreases in the collateral value.⁹ In the case of LTCM, however, the fund was able to obtain next-to-zero haircuts, as it was widely viewed as “safe” by its lenders. This must have been due to the fact that no counterparty had a complete picture of the extent of LTCM’s operations. But even when “fully collateralized”, these loans were exposed to the risk that LTCM could default at the same time as the collateral lost value.

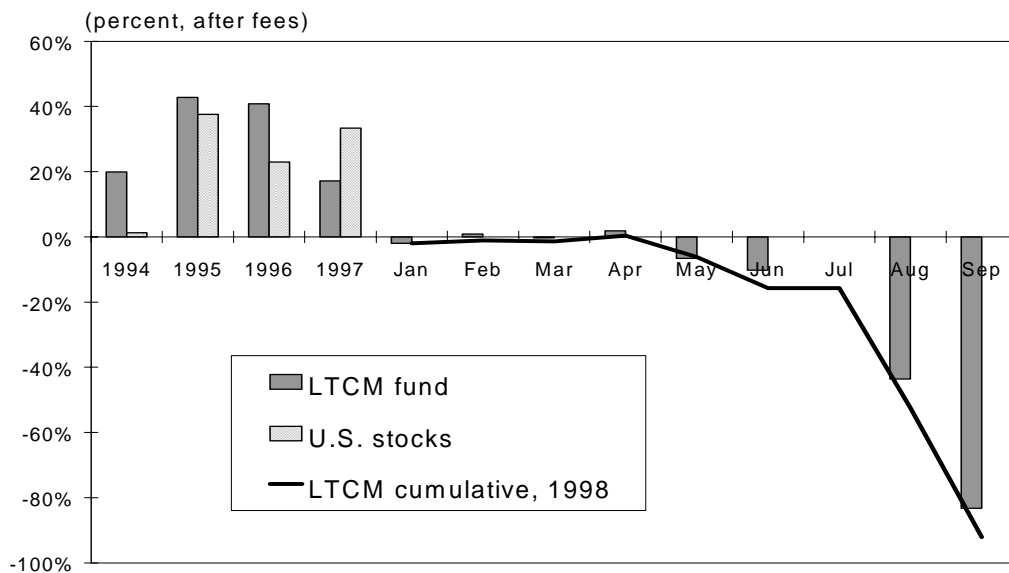
The core strategy of LTCM can be described as “convergence-arbitrage” trades, trying to take advantage of small differences in prices among near-identical bonds. Compare, for instance, a corporate bond yielding 10% and an otherwise identical Treasury bond with a yield of 7%. The yield spread of 3% represents some compensation for credit risk. If the corporate borrower does not default, a trade that is long the corporate bond and short the Treasury bond would be expected to return 3% for every dollar in the first bond. Short-term, the position will be even more profitable if the yield spread narrows further. The key is that eventually the two bonds must converge to the same value. Most of the time, this will happen—barring default or market disruption.

This strategy worked excellently for LTCM in 1995 and 1996, with after-fees returns above 40%, as shown in Figure 1.¹⁰ The fund placed large bets on convergence of European interest rates within the European Monetary System that paid off handsomely.

⁸Total futures positions amounted to over \$500 billion, swaps over \$750 billion, and options and other OTC derivatives over \$150 billion.

⁹Repo haircuts vary by instrument. Before LTCM faltered, there were no haircuts on Treasuries; haircuts for mortgage-backed securities were 3%; for dollar emerging market debt, 10%. The total haircut for LTCM’s \$110 billion in borrowing was reported to be around only \$500 million.

¹⁰Returns are derived from *Wall Street Journal* reports.

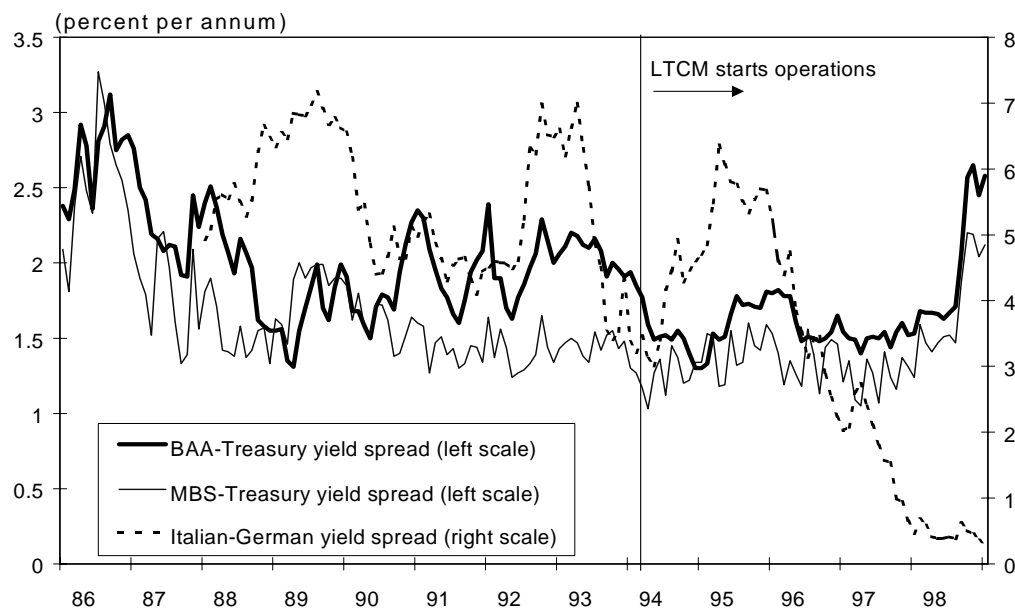
FIGURE 1 LTCM's Returns

By 1997, however, convergence had occurred in Europe, as the common currency, the Euro, came into being on January 1999.¹¹ Figure 2 shows the evolution of relevant yield spreads since 1986.¹² As the figure indicates, credit spreads were almost as narrow as they had ever been since 1986 and considerably lower than the average over the period 1986-93. Convergence trades had generally become less profitable. In 1997, the fund's return was down to 17% only. This performance, unfortunately, was trounced by U.S. stocks which gained 33%. This was embarrassing, as LTCM touted itself as having the same risk as equities. If it had lower returns, why would anybody invest in the fund? LTCM had to look for other opportunities.

To achieve the 40% returns it had become accustomed to, the firm had to assume greater leverage. So, LTCM returned \$2.7 billion of capital to investors in 1997 while keeping total assets at \$125 billion. By shrinking the capital base to \$4.7 billion, the leverage ratio went up, amplifying returns to investors that remained in the fund.

¹¹Note that there was a probability that the common currency would not occur for political reasons, which would have made these trades unprofitable. So, the historical data did not represent the true risks of the strategy. Even now, as currency risk has been eliminated in Europe, risk has not totally disappeared. As European governments have lost their ability to finance their deficit by printing money, the risk of default has increased. Thus currency risk has been transformed partially into credit risk, which is in fact harder to assess.

¹²The U.S. bond series are taken from the Federal Reserve's constant maturity 10-year yields, Moody's BAA yields, and fixed mortgage yields. The non-dollar yields are from the JP Morgan government bond indices.

FIGURE 2 Bond Yield Spreads

Unfortunately, this also increased the risks. Troubles began in May and June of 1998. A downturn in the mortgage-backed securities market led to a 16 percent loss in the value of equity. LTCM's capital had just dropped from \$4.7 to \$4.0 billion.

Then came August 17. Russia announced that it was “restructuring” its bond payments—de facto defaulting on its debt. This bombshell led to a reassessment of credit and sovereign risks across all financial markets. Credit spreads jumped up sharply. Stock markets dived. LTCM lost \$550 million on August 21 alone on its two main bets, long interest rate swap spreads and short stock market volatility.¹³

By August, the fund had lost 52% of its December 31 value. Its equity had dropped faster than its assets which still stood around \$110 billion. The resulting leverage ratio had increased from 27 to 50-to-1. LTCM badly needed new capital. In his September 2 letter to investors, Meriwether revealed the extent of the losses and wrote that “Since it is prudent to raise additional capital, the Fund is offering you the opportunity to invest on special terms related to LTCM fees. If you have an interest in investing, please contact...” There were no takers, though.

The portfolio's losses accelerated. On September 21, the fund lost another \$500 million, mostly due to increased volatility in equity markets. Bear Stearns, LTCM's prime broker, faced a large margin call from a losing LTCM T-bond futures position.¹⁴ It then required

¹³See Lewis (1999). He reports that LTCM had sold 3- to 4-year U.S. equity options at the end of 1997, when volatility stood at 20 percent.

¹⁴A prime broker provides centralized custodial, clearing, record-keeping, and financing for the hedge fund.

increased collateral, which depleted the funds's liquid resources. Counterparties feared that LTCM could not meet further margin calls, in which case they would have to liquidate their repo collateral.

A liquidation of the fund would have forced dealers to sell off tens of billions of dollars of securities and to cover their numerous derivatives trades with LTCM. Because lenders had required very low haircuts, there was a potential for losses to accrue while the collateral was being liquidated. In addition, as the fund was organized in the Cayman Islands, there was uncertainty as to whether the lenders could have liquidated their collateral. In contrast, such liquidation is explicitly allowed under the U.S. Bankruptcy Code. As it was believed that the fund could have sought bankruptcy protection under Cayman law, LTCM's lenders could have been exposed to major losses on their collateral.

The potential effect on financial markets was such that the New York Federal Reserve felt compelled to act. On September 23, it organized a bailout of LTCM, encouraging 14 banks to invest \$3.6 billion in return for a 90 percent stake in the firm.

These fresh funds came just in time to avoid meltdown. By September 28, the fund's value had dropped to \$400 million only. If August was bad, September was even worse, with a loss of 83%. Investors had lost a whopping 92% of their year-to-date investment. Of the \$4.4 billion lost, \$1.9 billion belonged to the partners, \$700 million to Union Bank of Switzerland, and \$1.8 billion to other investors.

LTCM is now operating under the control of the 14-member consortium, formally known as Oversight Partners I LLC. Helped by recovering financial markets, the portfolio gained 13% to December 1998. Since the bailout, the risk profile has been halved. As of this writing, though, the future of the fund is still in doubt.

2 VAR and Equity Capital

Value at Risk (VAR) has been heavily blamed for LTCM's failure. These criticisms are disturbing as VAR has emerged as a key component of capital adequacy measures for commercial banks and is spreading fast to other users. Indeed, the 1995 Basle Market Risk Capital Standards allow a VAR-based approach for setting regulatory capital for the market risks of commercial banks.¹⁵ It would be distressing to realize that VAR-based capital requirements are useless.

VAR can be defined as the worst loss that can happen under normal market conditions over a specified horizon at a specified confidence level. More formally, VAR measures the shortfall from the quantile of the distribution of trading revenues.

Initially, VAR was developed as a practical first-order gauge of financial risk. Its primary purpose was to communicate risks to stakeholders. Later, it developed into a common benchmark to compare and control risk across risk-taking units. More recently, VAR has been used to decide on the amount of equity capital necessary to support a trading activity.¹⁶

VAR measures are based on two quantitative parameters: the confidence level and the horizon. The choice of these parameters depends on the application. If VAR is simply used to report or compare risk, the parameters can be arbitrarily chosen, as long as they are consistent.

If, however, VAR is used as the basis for setting the amount of equity capital, the parameters must be chosen with extreme care. The confidence level must be high enough that the probability of exceeding VAR is very low. As to the horizon, it must be related to the liquidity of the assets, or the time necessary for an orderly liquidation (which depends on the size of the fund). Alternatively, the horizon should cover the time necessary to raise additional funds or for corrective action.

This explains why the Basle Committee has chosen a high confidence level of 99% and a 10-day horizon to determine the minimum capital level for the commercial bank. In addition, the resulting VAR is multiplied by a factor of three to account for non-normalities or model errors.

For LTCM, the choice of the horizon is more delicate than the 10-day period required the Basle Committee. Commercial banks are closely supervised by their regulator, who can step in at the first sign of trouble. Thus the 10-day period can be interpreted as the normal

¹⁵These capital adequacy requirements apply to the trading operations of commercial banks and are detailed in the Basle Committee on Banking Supervision (1995) amendment to the Internal Models.

¹⁶See for example, chapter 4 in Matten (1997).

reaction time should a commercial bank run into difficulties. For a hedge fund, the horizon should correspond to the period required to raise additional funds. This may be no easy matter, as additional capital will be needed precisely after the fund has suffered a large loss.

In the capital adequacy context, VAR can be viewed as a measure of “risk capital” necessary to support a financial activity. This resolves the paradox of how to calculate rates of return on investments that require no upfront investment, like futures. Consider for instance a bank with one investment only, a futures position with a notional value of \$100 million, a margin of \$10m, and a payoff (or dollar return) of \$5m.

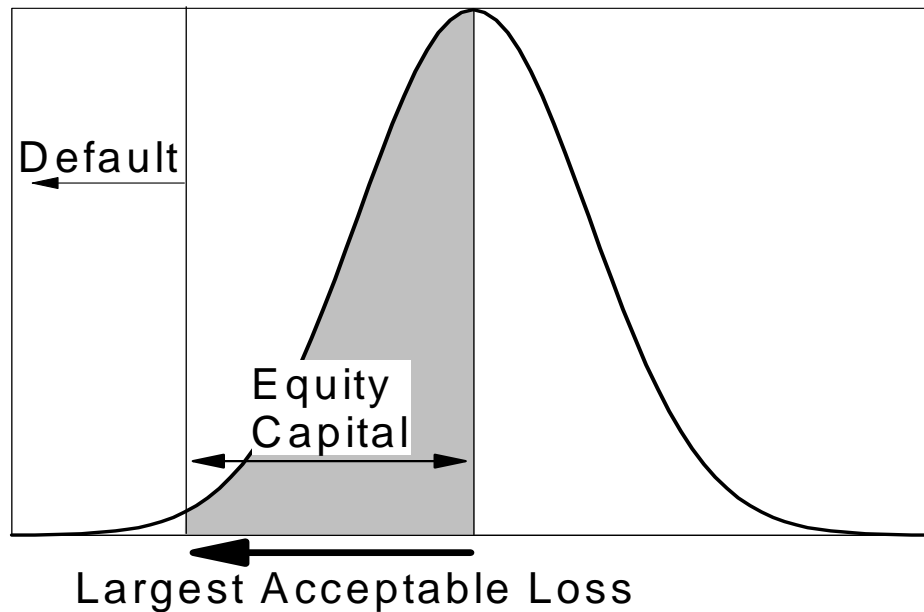
How do we compute the *rate* of return on this \$5m payoff, to compare it with other investments? In other words, what is the relevant denominator? The notional is not relevant, as it is never paid. Nor is the margin, which is a performance bond and typically pays interest anyway.

Instead, we could consider the amount of *equity capital* that needs to be set aside to cover most of the potential losses at a predetermined confidence level. Considering only market risks, this equity capital is basically a market VAR measure. Therefore, VAR measures the *economic capital*, defined as the aggregate capital required as a cushion against unexpected losses.¹⁷

Suppose, for instance, that VAR is estimated to be \$20m at the 99 percent level. As before, VAR is interpreted as the “largest” acceptable loss the bank is willing to suffer over a specified period. To cover this loss, the bank must maintain adequate equity capital. In other words, VAR is the amount of capital a firm allocates to self-insurance. This is illustrated in Figure 3.

Note that this capital is leveraged. That is, the equity is supported by borrowing. Ideally, the bank should select the amount of leverage to balance the marginal benefits of increased leverage with the marginal costs of increased default probability. In this case, the rate of return on equity is $\$5/\$20 = 25$ percent. With a lower capital base, the rate of return would be even higher. The drawback, of course, is that the investment is more risky.

¹⁷This ignores the time dimension or time value of money. If VAR is the loss on a fixed horizon, it should be discounted to find the required amount of current capital. More realistically, capital should be sufficient to cover a loss at any time over the horizon.

FIGURE 3 Equity Capital as a VAR Measure

This interpretation of VAR also helps to determine the appropriate confidence level. Consider a hypothetical bank for which all the risks are captured by the VAR measure. A one-year horizon is selected, assuming it corresponds to the horizon needed to adjust the level of capital, i.e. to raise additional equity. If the bank suffers a loss greater than the VAR within a year, its equity capital is wiped out and the bank defaults.

The bank can then set the VAR confidence level in relation to the desired credit rating. This is illustrated in Table 1, which reports the relationship between the credit ratings supplied by a rating agency, Moody's, and their historical one-year default probabilities. The last columns report the *equity coverage*, measured as the number of annual standard deviations necessary to achieve the desired credit rating.

As an example, assume the market risk is such that the daily standard deviation of profits and losses is \$100 million. As will be seen later, this was a good description of LTCM's risk profile. Assuming 252 trading days, this translates into an annual number of \$1.6 billion. How much capital should the bank set aside to achieve an Aa2 credit rating? The table shows that, if the distribution of losses is normally distributed, the required amount of equity is 3.5 times \$1.6 billion, or \$5.6 billion. If the bank estimates that the distribution of losses has fatter tails than the normal distribution, it could select the multiplier from a Student-t distribution, for example. In this case, the required amount of equity is higher, at \$12.5 billion. This is much more than LTCM's \$4.7 billion in capital.

TABLE 1 Choosing Equity Coverage from the Credit Rating
(Multiples of Annual Standard Deviation Assuming Various Distributions)

Desired Rating (Moody's)	1-Year Probability of Default	Equity Coverage (multiple of s.d.)		
		Normal	t(6)	t(4)
Aaa	0.009%	3.75	9.26	15.96
Aa1	0.015%	3.62	8.45	14.03
Aa2	0.022%	3.51	7.89	12.72
Aa3	0.035%	3.39	7.25	11.30
A1	0.05%	3.29	6.79	10.31
A2	0.06%	3.24	6.56	9.83
A3	0.07%	3.19	6.37	9.45
Baa1	0.13%	3.01	5.67	8.04
Baa2	0.16%	2.95	5.44	7.61
Baa3	0.70%	2.46	4.01	5.10
Ba1	1.25%	2.24	3.52	4.31
Ba2	1.79%	2.10	3.23	3.88
Ba3	3.96%	1.76	2.62	3.01
B1	6.14%	1.54	2.30	2.58
B2	8.31%	1.38	2.08	2.30
B3	15.08%	1.03	1.65	1.77

Notes: Equity coverage is defined as the number of annual standard deviations (s.d.) necessary to achieve desired default probability. Two distributions are used, the standard normal and Student-t with 6 and 4 degrees of freedom.

3 LTCM's Risk Management Story

How did this apply to the LTCM story? By itself, the \$1.2 trillion notional amount, or even \$125 billion in assets, is meaningless. What mattered was the overall volatility of the fund. LTCM aimed for a level of risk no greater than that of U.S. equities.

The annual average volatility of the S&P 500 over 1978 to 1997 was 15%. Applying this number to \$4.7 billion, we find a daily dollar volatility of $\$4,700 \times 0.15 / \sqrt{252} = \44 million. It was claimed that the portfolio was no more risky than an unleveraged investment in equities.

Indeed, LTCM admitted that its target daily volatility was \$45 million.¹⁸ In other words, the positions were allocated so as to maximize expected returns subject to the constraint that the fund's perceived risk was no greater than that of the stock market. While there may not have been a formal optimization process on total return and risk, the end result was the same: LTCM was maximizing returns while closely monitoring its volatility.

Let us see how this would translate into a capital requirement under the Basle framework for commercial banks. These rules specify that capital should be at least 3 times the 99% VAR over 10 business days. Assuming a normal distribution and the associated multiplier of 2.33, the daily 99% VAR is \$105 million. Applying the Basle rules, this translates into a minimum capital level of $3 \times \$105 \times \sqrt{10} = \$993m$, which seems adequately covered by the initial \$4.7 billion in equity. The 10-day period, however, is totally unrealistic for a hedge fund.

In addition, the \$45 million number involves heroic assumptions. It assumes that volatility is constant, when in fact it can easily double in turbulent times. One must also accept that the distribution of profit and losses is symmetrical, which is obviously incorrect when dealing with credit risk. Indeed, Merton (1974) has shown that long positions in credit-sensitive instruments can be interpreted as short positions in options, which have limited upside potential (the option premium) and large downside risk (due to the possibility of default). Even if the distributions were symmetrical, it is unlikely their tails could be well described by a normal distribution. Recent time-series models indicate that most financial series have fatter tails than the normal distribution, even when accounting for time-variation in risk. Bollerslev, Chou and Kroner (1992), for example, indicate in that, even after fitting a time-varying GARCH process, conditional standardized returns are better approximated by a Student t-distribution than by a normal distribution. Finally, this paper will show that using the same data to optimize a portfolio and compute its VAR results in biased risk

¹⁸See Kolman (1999).

forecasts.

These biases should have showed up in the realized returns. Extrapolating the \$105m daily VAR to a monthly number yields \$339m, using the square root of time and a 21-day month. Assuming an annual 18% expected return, the expected profit was \$71m over a month. Thus, the worst dollar loss at the 99% level over a month should have been (\$339-\$71)=\$268 million. In May and June, the portfolio lost \$310 million and \$450 million. So, it must have crashed through its daily VAR limit on a number of occasions.

Here is why *backtesting* becomes all-important. If the model is correctly calibrated, we should only observe 1 percent deviations, or one *exception* out of 100. The decision rule should take into account the variation in the data, though. To do this, we can use the Basle framework for determining whether the number of exceptions is acceptably low.¹⁹

Define N as the number of exceptions observed in the sample of size T . The confidence level is p . Using the normal approximation to the sequence of binomial processes, we can write the test statistic as:

$$t - test = \frac{[N - (1 - p)T]}{\sqrt{T(1 - p)p}}. \quad (1)$$

Two exceptions or more during these two months should have given strong indications that the model was invalid. For instance, with $p = 99\%$, $N = 2$, $T = 42$, we have $t - test = 2.45$, which is greater than the usual cutoff value of 2. The experience of May and June alone should have raised warnings flags that the model had gone awry.²⁰

The firm reportedly tried to reduce its risk profile, but made a major mistake: Instead of selling off less-liquid positions, it eliminated its most liquid investments because they were less profitable. Indeed the September 2 letter explained that “..steps have been taken to reduce our risks now, commensurate with our level of capital. We have raised the risk-return trade-off requirements for positions.” This made LTCM more vulnerable to subsequent margin calls.

After this adjustment, LTCM’s models indicated that its expected daily volatility had been indeed reduced to \$35 million. Something was wrong, however. Its actual daily volatility was closer to \$100 million. Either the market were getting more volatile or the models were seriously biased. Probably both.

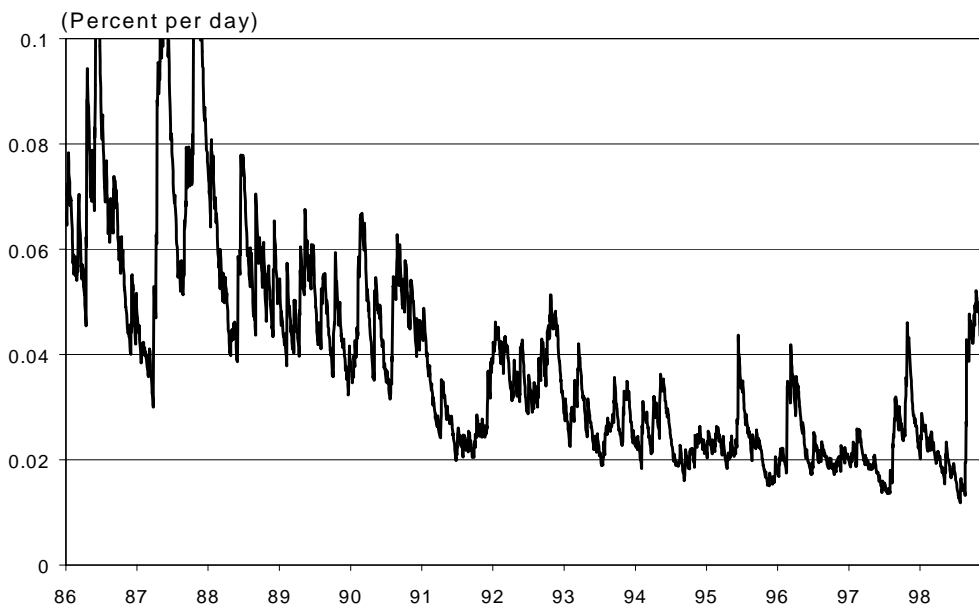
Figure 4 shows the predictable component of the daily volatility of changes in credit spreads. The forecast is based on exponential weighing, a standard method introduced by

¹⁹See the Basle (1996) document on backtesting.

²⁰More precisely, the exact p-value of observing 2 or more deviations under the binomial process is 0.066. In addition, observing too many deviations in a row indicates that the model does not capture an elevated market volatility.

J.P. Morgan in 1994, applied to the daily change in credit spreads.²¹ The figure shows that, first, the average volatility over 1995 to 1997 was low by historical standards. Second, volatility creeps up on August 25 from 0.019 to 0.027; on August 27, it moves and stays above 0.04, double its recent value. Thus standard risk management methods would have picked up the increase in market volatility.

FIGURE 4 Volatility of Credit Spreads



Even worse, on August 21, the portfolio lost \$500 million. By August 31, the portfolio had lost \$1,710 million in one month only. Using the presumed \$45 million daily (or \$206m monthly) standard deviation, this translates into a 8.3 standard deviation event. Assuming a normal distribution, such event would occur once every 800 trillion years, or 40,000 times the age of the universe. Surely the model was wrong.

If the true volatility was \$100 million, this was a 3.7 standard deviation event. Assuming now a distribution with fatter tails, such as the Student $t(4)$, such event should occur once every 8 years. This is probably not far off the mark.

Given a true standard deviation of \$100 million, the Basle minimum capital would increase to \$2.2 billion. The order of magnitude is now closer to the actual loss of \$1.7 billion in August. At the end of the month, LTCM was left with \$2.3 billion in equity, barely above the minimum.

²¹For the methodology, see J.P. Morgan's (1994) RiskMetrics manual.

In retrospect, the firm made major risk-management blunders. LTCM had relied on recent history to estimate risk, surely assigning a zero probability to events such as the Russian default and major market disruptions such as the 1987 crash, which also led to a flight to liquidity. Since credit-sensitive instruments are by nature affected by rare events, this approach missed the true risks of the portfolio. Going back to Figure 2, for example, it seems clear that the credit spreads of the 1995-97 period were low by historical standards. Any slowdown in the economy would widen these spreads. Box 1 gives another example of a famous hedge fund speculator who was caught short when the recent data failed him.

This type of strategy exposes the portfolio to catastrophic risk. Note how this interpretation fundamentally differs from LTCM's view of the events of 1998. LTCM claims the strategy is fundamentally sound, and that the events of 1998 were an aberration. Instead, the explanation here is that these "arbitrage" trades are doomed to take large losses once in a while and that the true risks cannot be picked up easily by VAR systems.

BOX 1 Victor Niederhoffer: The Education of a Speculator

Victor Niederhoffer outlined his investment philosophy in his book *Education of a Speculator*, which quickly became a best-seller. An eccentric and brilliant investor, he was a legend of the hedge fund business. Indeed he had compiled an outstanding track record—a 32% compound annual return since 1982.

Niederhoffer's mission was to "apply science" to the market. Although he had a PhD in business from the University of Chicago, he did not believe in efficient markets and traded on statistical anomalies. He believed, for instance, that the market would never drop by more than 5 percent in a single day. Putting this theory into practice, Niederhoffer sold naked out-of-the-money puts on stock index futures. When the stock market plummeted by 7 percent on October 27, 1997, he was unable to meet margin calls for some \$50 million. His brokers liquidated the positions, wiping out his funds. Comparing his downfall to LTCM's bailout, he later wished he had gotten "similar help."

Apparently, his views were narrowly based on recent history. It is true that the worst loss had been 3.1% in the previous 5-year period. Larger loss do occur once in a while, however. Most notably, the market lost 20.4% on October 19, 1987.

Admittedly, the losses were exacerbated due to the size of the positions and the drying up of liquidity in turbulent financial markets. LTCM also claims that the selling of securities was made worse by brokers "front-running" the LTCM portfolio.²² It is hard to verify these claims. More fundamentally, the firm had badly underestimated its risk and did not have enough long-term capital to ride out the turbulence of 1998.

²²See Lewis (1999).

4 Portfolio Optimization and Leverage

This section presents a stylized example of a portfolio optimization with two highly correlated assets. We use a government and a corporate bond series and find results that are quite consistent with LTCM's leverage. This exercise shows how leverage arises from such convergence trades and that results are extremely sensitive to errors in input parameters.

Assume the investor seeks to maximize a utility function that reflects the trade-off between expected portfolio returns $\mu_p = w'\mu$ and risk $\sigma_p^2 = w'\Sigma w$, where w represents a vector of portfolio weights, μ represents a vector of expected returns in excess of the risk-free rate, and Σ represents the covariance matrix of returns. The investor maximizes

$$U(w \mid \mu, \Sigma) = U(\mu_p, \sigma_p^2) = \mu_p - 1/(2\lambda)\sigma_p^2, \quad (2)$$

where λ is the investor's risk tolerance, or inverse of risk aversion.²³

The optimal positions in risky assets are given by

$$w = \lambda\Sigma^{-1}\mu, \quad (3)$$

with the remainder of the portfolio is invested in the riskless asset.

Table 2 presents a worked-out example of a portfolio optimization with two risky assets, a BAA-rated corporate bond series and a 10-year Treasury note series. The covariance matrix is constructed from the history of yield changes over the 5-year period 1993-1997 and from an assumed 7-year modified duration. The two bonds have similar volatility and high correlation, at 0.9654.

Expected returns are simply taken as yields on December 1997. With a credit yield spread of 1.53% and no allowance for defaults, the optimization should identify an "arbitrage" strategy.

Indeed, the portfolio takes a very large position in the corporate bond, \$19.66 for every \$1 of equity, offset by a large short position of \$15.60 in the Treasury bond. The sum of the bond positions is \$4.06, which implies a loan of \$3.06 at the risk-free rate. The expected return on the portfolio is 3.1% per month, which translates into 37% per annum.

This simple exercise captures the essence of LTCM's so-called arbitrage strategy. The leverage ratio is close to 20 and the annual expected return is close to 40%, as LTCM had hoped for. What is the problem with this strategy?

²³We choose λ so that an investor would be indifferent between investing in stocks and in cash. Assuming that the excess return on stocks is 8% with a volatility of 20%, we set (2) to zero, With $0.08 - 1/(2\lambda)0.04 = 0$, we find $\lambda = 0.25$.

TABLE 2 Portfolio Optimization with Two Assets

Input data:	Corporate Bond	Treasury Bond	Risk-free Asset
Expected return-yield (%pa)	7.28%	5.75%	5.36%
Volatility of return (%pm)	1.58%	1.90%	
Correlation	0.9654		
Output data:			
Position (for \$1 equity)	\$19.66	-\$15.60	-\$3.06

Optimal Portfolio	Return and Risk	
Initial Equity	\$1	\$4,700m
Expected return (monthly)	3.1%	\$145m
Volatility of return (monthly)	8.1%	\$382m
Expected return (annual)	37.0%	\$1,740m
Volatility of return (annual)	28.1%	\$1,322m
Volatility of return (daily)	1.8%	\$83m
Ratio of equity to SD	12.31	

5 The Problem with Portfolio Optimization

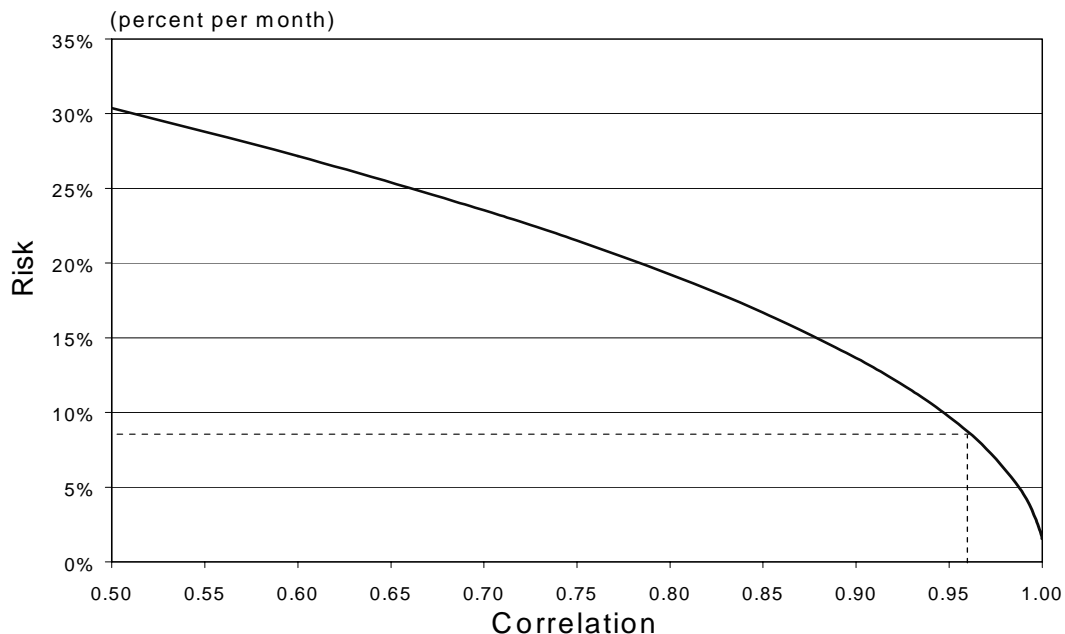
The crux of the problem is the presumed risk of the strategy. For every dollar invested, the monthly volatility is 8.1%. Thus to be wiped out over one month, the portfolio should suffer a loss of $\$1/\0.081 , which is 12.3 times the standard deviation. Define this ratio as the monthly equity coverage, or *safety factor*. As this loss appears unlikely, the amount of equity provides sufficient protection. Or, is this so unlikely?

This type of strategy hinges on the correlation staying at the high level of $\rho = 0.9654$. There are two problems with this. The first is the combined effect of portfolio optimization with estimation error. If the number happens to be high because of (positive) estimation error, the optimizer will tend to take large opposite positions in the two assets to exploit this illusory safety. The estimated correlation will be randomly distributed around the true value and any error will be amplified by the optimization.

The second problem is that the true correlation may change over time. Given that we start with a very high correlation, which is bounded by unity, there is probably more chance that it will drop rather than increase further.

This is why we need *stress-testing*, which involves changing the parameters to assess the sensitivity of the results. Figure 5 examines the volatility of the current optimized portfolio as the correlation coefficient decreases.

FIGURE 5 Effect of Correlation on Risk



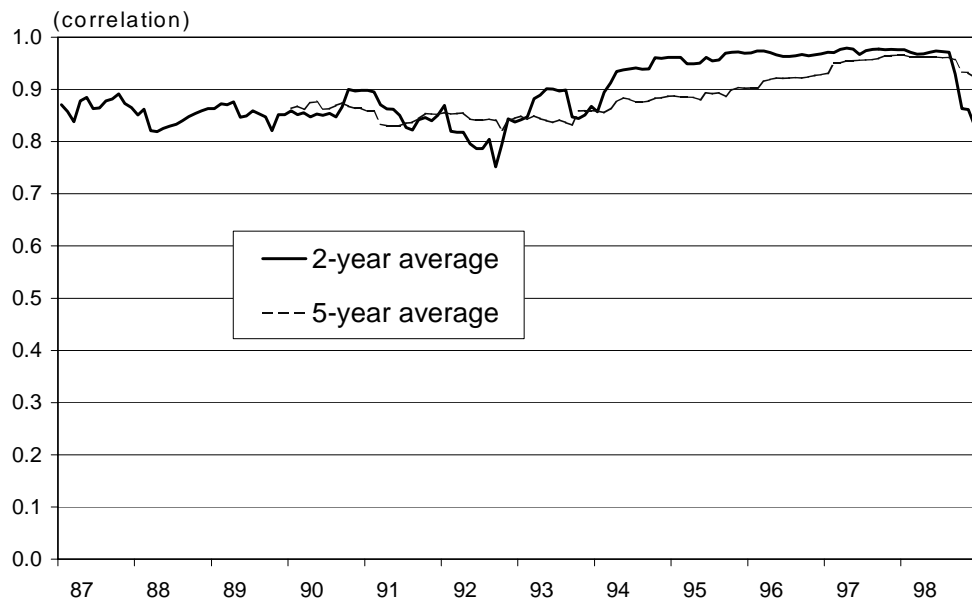
We start with a value of 8.1% at the estimated correlation of 0.965. If the correlation

decreases to 0.90, the portfolio volatility rises sharply, to 13.65%. For $\rho = 0.80$, the volatility increases further to 19.24%, which is more than twice the original number. At this level, the portfolio safety has fallen from 12.31 to 5.20. Here the amount of investor equity can only cover a 5 standard deviation move.

What actually happened to this correlation is described in Figure 6. The solid line represents the correlation estimated with a 2-year moving average (MA), while the dotted line shows a 5-year moving average.

In 1998, ρ drops sharply to 0.8, explaining why the convergence strategy suddenly went bad. Was this unexpected? Yes, if one focuses narrowly on the recent history of the fund, i.e. since 1994, when ρ remained above 0.94. This is myopic, however, as a longer history shows that ρ was much lower five years earlier. It even reached a low of 0.75 in 1992.

FIGURE 6 Correlation between Corporate and Treasury Yields



Armed with this information, we can now examine the effect of changing correlations on the safety factor and probability of ruin within a month. With $\rho = 0.97$, the safety factor is 13.19. Assuming that the distribution of losses is normal, Table 3 shows that the probability of monthly ruin is zero (within the precision of most numerical approximations). We do know, however, that financial time-series are “fat tailed”. Using a student t-distribution with 4 degrees of freedom, the probability of ruin is 0.000095, or one month out of 10,000, or one year out of 900. These odds seem low enough.

The problem is that as soon as the correlation drops to 0.80, the safety ratio decreases to 5.2. With a normal distribution, it is again highly unlikely that all of the equity could be wiped out. With a student $t(4)$, however, the odds move to one month out of 306, or one year out of 26. Now the odds of failure are not so remote.

TABLE 3 Safety Factor for Various Correlations

Rho	Std. Dev.	Safety Factor	Monthly Probability of Ruin		
			Normal	Student-6	Student-4
0.9999	1.56%	64.10	0.00000%	0.00000%	0.00002%
0.999	2.03%	49.26	0.00000%	0.00000%	0.00005%
0.990	4.55%	21.98	0.00000%	0.00003%	0.00127%
0.970	7.58%	13.19	0.00000%	0.00059%	0.00954%
0.950	9.71%	10.30	0.00000%	0.00245%	0.02507%
0.930	11.45%	8.73	0.00000%	0.00623%	0.04735%
0.900	13.65%	7.33	0.00000%	0.01650%	0.09227%
0.850	16.68%	6.00	0.00000%	0.04843%	0.19470%
0.800	19.24%	5.20	0.00001%	0.10099%	0.32637%
0.750	21.50%	4.65	0.00017%	0.17497%	0.48265%
0.700	23.54%	4.25	0.00108%	0.26957%	0.65891%
0.600	27.17%	3.68	0.01164%	0.51621%	1.05970%
0.500	30.36%	3.29	0.04943%	0.82675%	1.50539%

Admittedly, this example is simplistic. Can this be extended to the much larger number of positions assumed by LTCM? Yes, and the problem is probably even worse.

Estimation risk increases with the number of parameters. The more parameters are estimated, the greater the chance that errors will interact with each other and create a misleading picture of risk. This is a well-known problem in portfolio optimization. One cannot rely on the same covariance matrix (i.e. in-sample) to perform the portfolio optimization and to measure risk.²⁴

A similar problem occurs when traders try to “game” a VAR system. If a risk manager uses a VAR system for risk control or performance evaluation, traders may have an incentive to evade their risk limits. For instance, traders could move into markets or securities that appear to have low risk for the wrong reasons. Currency traders could have taken large positions in pegged currencies, which have low historical volatility but high devaluation risk. Or, traders exposed to a delta-normal VAR could take short straddles with zero delta

²⁴Michaud (1989), for instance, argues that mean variance optimizers are in effect “error maximizers” when using historical data. In other words, the optimizer will overweight assets with high recent average returns, whereas most of this performance may simply reflect luck.

(like Baring's Leeson). Such positions appear profitable, but only at the expense of future possible losses which may not be captured by VAR. More generally, a trader may be aware of measurement errors in the covariance matrix used to judge him. If so, he or she may overweight assets that have low estimated risk, knowing full well that this will result in a downward biased risk measure.

Ju and Pearson (1999) provide estimates of this potential bias. Table 4 shows the ratio of true VAR to the estimated VAR for various number of observations (T) and number of assets (N). For instance, with 100 days and 50 assets, the true VAR is 201%, or twice the estimated VAR. The bias increases as the number of assets increases relative to the number of observations, reflecting increased measurement error in the covariance matrix.

When time-variation in risk is modelled with an exponential model such as RiskMetrics, the number of effective observations is quite small. Table 4 shows that the bias can be quite substantial. With 50 assets, the true VAR is nearly 5 times the estimated VAR. The danger of relying on recent data explains why the Basle Committee disallows the use of such models and requires a long period for the estimation of VAR.

TABLE 4 Bias in VAR: Ratio of True to Estimated VAR

Model: Observations(T)	Number of Assets (N)			
	10	20	50	100
Moving-Average:				
50	123%	164%		
100	110%	124%	201%	
200	105%	111%	133%	199%
1000	101%	102%	105%	111%
Exponential:	135%	185%	485%	2174 %

Source: Adapted from Ju and Pearson (1999). The table reports the mean ratio of true VAR to estimated VAR assuming that the trader knows the true covariance matrix and maximizes expected return subject to a constraint on VAR. These numbers are independent of the true covariance matrix. Because the trader may not know the true matrix, these numbers represent worst-case estimates of the bias. The exponential model uses 100 data points and a decay factor of 0.94.

6 LTCM's Risk Profile

LTCM failed because of its inability to measure, control, and manage its risk. This was due in no small part to the fact that LTCM's trades were rather undiversified. LTCM was reported to have lost about \$1.5 billion from interest rate swaps positions, and a similar amount from equity options. Table 5 describes the exposure of various reported trades to fundamental risk factors. All the trades were exposed to increased market volatility. Most were exposed to liquidity risk (which is correlated with volatility). Many were exposed to default risk.

TABLE 5 Exposure of LTCM's Portfolio to Risk Factors

Trade	Loss if risk factor increases		
	Volatility	Default	Illiquidity
Long interest rate swap	Yes	Yes	Yes
Short equity options	Yes		
Long off-the-run/short on-the-run Treasuries	Yes		Yes
Long mortgage-backed securities (hedged)	Yes		Yes
Long sovereign debt	Yes	Yes	Yes

To illustrate the driving factor behind LTCM's risks, Figure 7 plots the monthly returns against monthly changes in credit spreads. The fit is remarkably good, indicating that a single risk factor would explain 90% of the variation up to the September bailout. Data points after the bailout are indicated with the year. Their slope is markedly lower, reflecting the decrease in the risk profile of the fund directed by the consortium.

We can combine the information on the exposure of the fund with the daily volatility of credit spreads (reported in Figure 4) to simulate the fund's volatility. For illustration purposes, we assume that the investment strategy in 1998 is constant throughout the life of the fund.

Figure 8 graphs the daily dollar volatility of the fund, which should be compared with the "presumed" value of \$45 million. The figure shows wide variations in the fund's volatility, which ranges from a low of \$55 million in July 1998 to a high of \$245 million in October 1998. It is interesting to note that, had such a historical analysis been performed, it would have shown several periods prior to 1998 when the volatility was close to \$200 million. Thus the target volatility of \$45 million was unrealistically low. This explains how the required capital cushion had been badly underestimated.

FIGURE 7 Explaining LTCM Returns

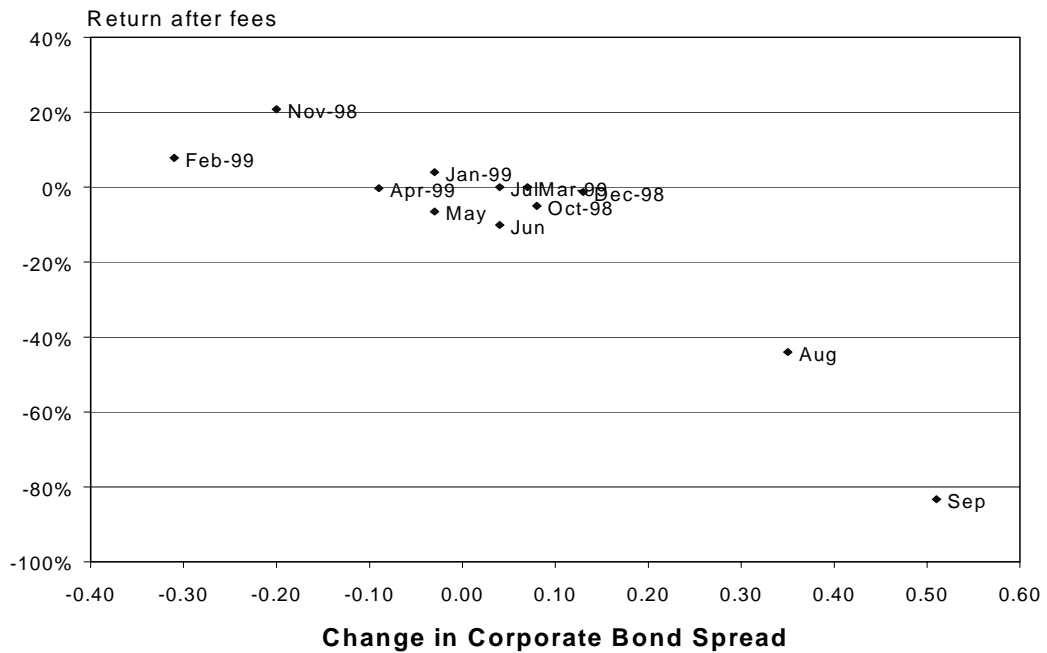
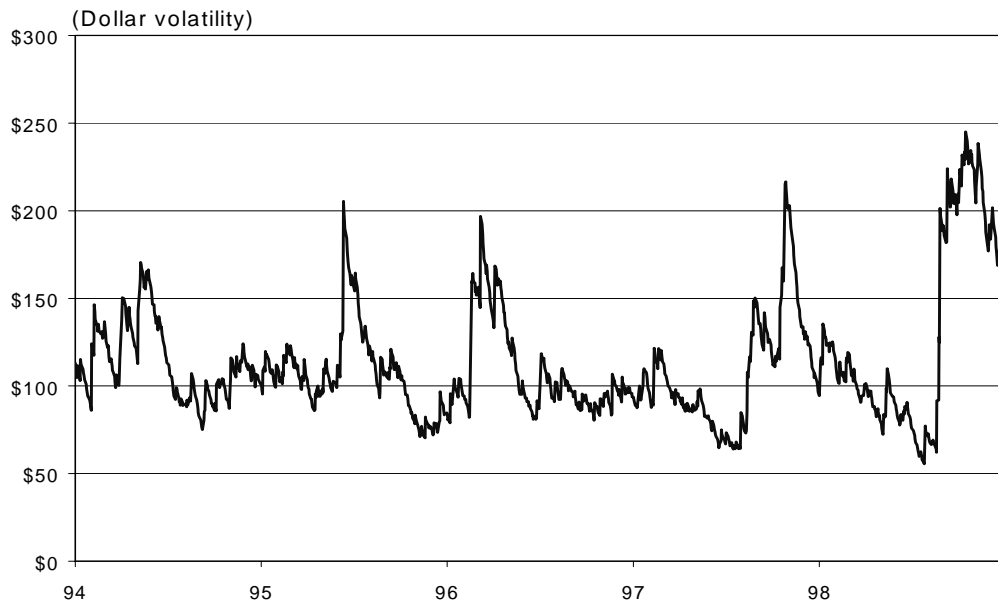


FIGURE 8 Simulated Daily Volatility of the Fund



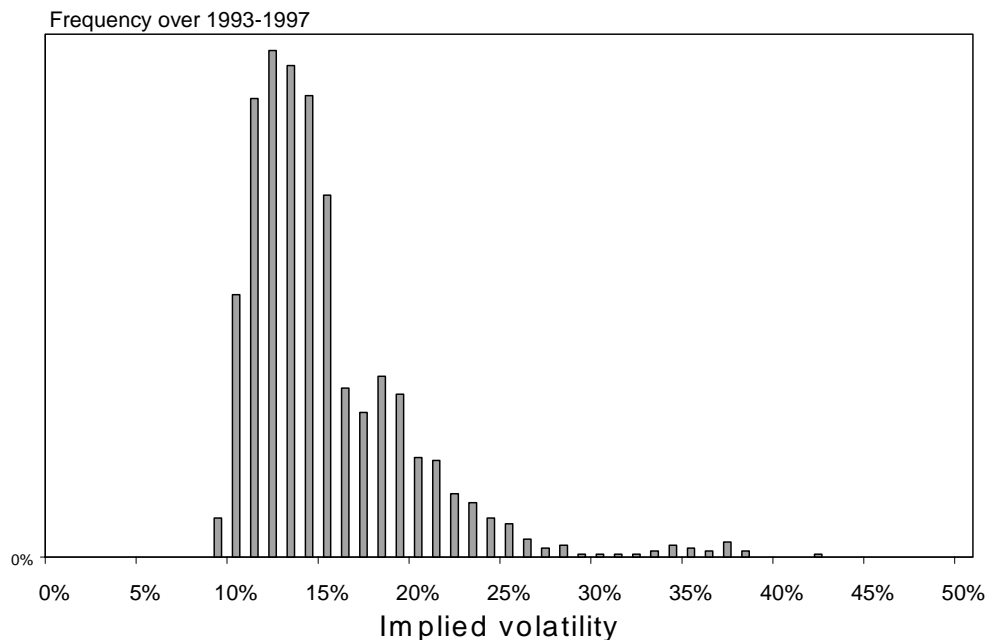
LTCM is a good example of the biases in portfolio optimization. Only a few of the “convergence” trades could truly qualify as arbitrage trades. One example is the long position in the off-the-run bond offset by a short position in the on-the-run equivalent. In this case, the probability of default is identical and eventually, the two bonds will converge to the same

value (in 30-some years). In the meantime, though, the spread could widen further, thus requiring the fund to carry enough capital to absorb temporary losses.

For most other trades, the portfolio was exposed to hidden risks. The typical profile of credit spread trades is to make small profits (which can be amplified through leverage) most of the time, but to take a big loss once in a while. This profile is akin to a short position in an option. To assess the possibility of catastrophic but rare events such as sovereign defaults, market or political disruptions, one cannot focus on recent history only.

Similarly, a short position in equity volatility has an asymmetric payoff profile. Figure 9, for instance, plots the distribution of at-the-money option volatilities over the period 1993 to 1997. The distribution is strongly asymmetric. Thus selling volatility exposes the portfolio to the risk of large losses due to increased volatility. Again, the position is akin to a short position in an option. This latent exposure to catastrophic risks explains why hedge funds can fail dramatically.²⁵

FIGURE 9 Distribution of Stock Market Volatility



To make it even worse, the LTCM portfolio was also exposed to liquidity risk, which causes losses when volatility increases. Traditional VAR methods assume that the fund is a

²⁵The President's Working Group report notes that the aggregate performance of hedge funds should be adjusted by their very high rates of disappearance. Starting in 1994, for instance, fewer than three-fifth of the funds survived through 1998. Similar rates of attrition are reported by Brown, Goetzmann, and Ibbotson (1999).

price-taker, i.e. is not large enough to affect prices. Risk managers now recognize that this assumption may not be appropriate for some users.

LTCM liked to compare itself to a Wall Street investment bank, like Salomon Brothers, which was leveraged by the same amount, about 25 times. This comparison is faulty for a number of reasons. First, bank portfolios are more diversified across risk factors. They deal with fixed income markets, currencies, equities, and commodities. Second, brokers' trading activities actually benefit from increased volatility, as volatile markets generate greater trading volume and revenues. This provides a natural hedge against increases in volatility and furthermore allows them to manage their inventory dynamically. Finally, investment banks can have access to fresh funds much faster than any hedge fund.

7 Conclusion

By relying on recent history, LTCM seems to have badly misjudged the possibility of a backup in credit yield spreads. It lulled itself into believing that its VAR was manageable. In fact, it had severely underestimated its risk due to its risk concentration and reliance on short-term history.

The story of LTCM should not be taken as an indictment of VAR. Instead, it serves to illustrate the danger of "traders gaming the system". As this paper has shown, LTCM's strategy exploited intrinsic weaknesses in its risk management system. Its strategy can be interpreted in terms of maximizing expected returns subject to a constraint on VAR. During this process, LTCM committed large-scale risk management blunders, ignoring warning signs given by VAR. In the end, LTCM failed because of it mismanaged its risk modelling.

The saga of LTCM has important lessons for convergence-arbitrage strategies. One view is that the strategy is fundamentally sound and that the events of 1998 were so unusual so as to be impossible to forecast. Hence these events are not likely to happen again.

Another view is that these strategies are actually designed to take a big loss once in a while. The fund claimed it was taking "non-directional" bets, i.e. that were not obviously dependent on the direction of markets. This strategy, however, induced the fund to make undiversified and highly leveraged bets on more subtle risks such as credit risks, political risks, or market disruptions. The implication is that such strategies are fundamentally dangerous and much riskier than measured by traditional risk management systems.

References

- Basle Committee on Banking Supervision, 1995, *An Internal Model-Based Approach to Market Risk Capital Requirements*, Basle, Switzerland: BIS.
- Basle Committee on Banking Supervision, 1996, *Supervisory Framework for the Use of 'Backtesting' in Conjunction with the Internal Models Approach to Market Risk Capital Requirements*, Basle, Switzerland: BIS.
- Bollerslev, T., R. Chou, and K. Kroner, 1992, "ARCH Modelling in Finance: A Review of the Theory and Empirical Evidence," *Journal of Econometrics* 52, 5–59.
- Brown, Stephen, William Goetzmann, and Roger Ibbotson, 1999, "Offshore Hedge Funds: Survival and Performance: 1989-1995," *Journal of Business* (January) 72: 91–117.
- Dunbar, Nicholas, 1998, "Meriwether's Meltdown," *Risk* (October) 32–36.
- Greenspan, Alan, 1998, "Private-sector refinancing of the large hedge fund, Long-Term Capital Management," *Testimony before the Committee on Banking and Financial Services*, October 1.
- J.P. Morgan, 1995, *Riskmetrics Technical Manual*, New York: J.P. Morgan.
- Jorion, P., 1997, *Value at Risk: The New Benchmark for Controlling Market Risk*, Irwin-McGraw Hill: Illinois.
- Ju, Xiongwei and Neil Pearson, 1999, "Using Value-at-Risk to Control Risk Taking: How Wrong Can You Be?" *Journal of Risk* 1 (Winter): 5–36.
- Kolman, Joe, 1999, "LTCM Speaks," *Derivatives Strategy* (April): 12–17.
- Lewis, Michael, 1989, *Liar's Poker*, Penguin Books: New York.
- Lewis, Michael, 1999, "How the Eggheads Cracked," *New York Times Magazine*, January 24, 24–77.
- Matten, Chris, 1997, *Managing Bank Capital: Capital Allocation and Performance Measurement*, Chichester, England: Wiley.
- Merton, Robert, 1974, "On the Pricing of Corporate Debt: The Risk Structure of Interest Rates," *Journal of Finance* 29, 449–70.
- Michaud, Richard, 1989, "The Markowitz Optimization Enigma: Is Optimized Optimal?," *Financial Analysts Journal* 45: 31–42.
- President's Working Group on Financial Markets, 1999, *Hedge Funds, Leverage, and the Lessons of Long-Term Capital Management*, Washington, DC.
- Pacelle, Mitchell, Randall Smith, and Anita Raghavan, 1999, "Investors May See 'LTCM, the Sequel'," *Wall Street Journal* (May 20): C1.