

# Risk management lessons from Long-Term Capital Management

Philippe Jorion\*

*Graduate School of Management, University of California at Irvine, Irvine, CA 92697-3125, USA;*  
*e-mail: pjorion@uci.edu*

## 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 risk management lessons from LTCM. LTCM's strategies are analysed 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:** G11, G13, G14, G23.

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. (Chairman Alan Greenspan, 1998.)

The 1998 failure of Long-Term Capital Management (LTCM) is said to have nearly blown up the world's financial system. 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.

---

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

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.<sup>1</sup>

This issue is important as LTCM's failure has been widely ascribed to its use of Value at Risk (VAR).<sup>2</sup> If so, the disturbing implication is that the statistical risk-management methods spreading throughout the financial industry are woefully inadequate.

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 to ride out the turbulence of 1998. 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 of LTCM, which are analysed 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 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

Sources say Long Term Capital has warned investors to expect swings as high as 80% and as low as -20% in a given year. (*Business Week*, 29 August 1994.)

The seeds of Long-Term Capital Management (LTCM) can be traced to a highly profitable bond-arbitrage group at Salomon Brothers run by John Meriwether.<sup>3</sup> LTCM was founded by Meriwether in 1994, who left after the 1991 Salomon bond scandal.<sup>4</sup>

---

<sup>1</sup>This 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* (1999). The first systematic review of LTCM's downfall was by Dunbar (1998).

<sup>2</sup>See for instance *The Economist*, 14 November 1998 'Risk management: too clever by half'.

<sup>3</sup>The antics of the Salomon culture are vividly described in a colorful description of Wall Street, *Liar's Poker*, by Michael Lewis.

<sup>4</sup>Paul 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. This 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.

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. Formally, Long-Term Capital Management LP, a Delaware limited partnership, operated the fund, Long-Term Capital Portfolio LP (LTCP), which was organized as a Cayman Island partnership.

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 and can be quite risky.

The core strategy of LTCM can be described as 'relative-value', or 'convergence-arbitrage' trades, trying to take advantage of small differences in prices among closely related securities. Compare, for instance, an off-the-run Treasury bond yielding 6.1% versus 6.0% for the more recently issued on-the-run. The yield spread represents some compensation for liquidity risk. Over a year, a trade that is long the off-the-run and short the on-the-run would be expected to return 10bp for every dollar invested. The key is that eventually the two bonds must converge to the same value. This strategy was used in a variety of markets, long swap-government spreads, long mortgage-backed securities versus short government, long high-yielding versus short low-yielding European bonds, Japanese convertible bond arbitrage, equity pairs (stocks with different share classes) and so on. The firm also dabbled in non-arbitrage strategies, such as short positions in equity options, bets on takeover stocks, emerging market debt, and even catastrophe bonds.<sup>5</sup> Most of the time, these trades should be profitable—barring default or market disruption.

The problem with such strategies is that they generate tiny profits, so that leverage has to be used to create attractive returns. To control risk, the target ceiling risk level was set to the volatility of an unleveraged position in US equities (and the fund was advertised to investors as such). In essence, positions were obtained by an optimization with a constraint on volatility (with presumably some additional constraints such as the liquidity and concentration of positions). Thus, leverage had to be quite large, unusually so, as LTCM ended up with four times the asset size of the next largest hedge fund.

Initially, the new venture was eminently profitable. Capital grew from \$1 billion to more than \$7 billion by 1997. The firm was charging lofty 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's balance sheet was about \$125 billion. This represents the *total assets* of the fund, most of it borrowed.<sup>6</sup> Compared to equity of about \$5 billion only, this represents an astonishing leverage ratio of 25:1.

---

<sup>5</sup> *Institutional Investor*, December 1998.

<sup>6</sup> According to the President's Working Group report (1999), the fund had 60,000 trades on its books.

Even more astonishing was the off-balance-sheet position, including swaps, options, and other derivatives, that added up to a *notional principal* amount of \$1.25 trillion.<sup>7</sup> 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.<sup>8</sup> LTCM, however, was able to obtain unusually good financing conditions, with 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.

The swaps were subject to two-way marking-to-market, which means that their current market value was always close to zero. 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. Wall Street thus seems to have ignored its *potential exposure* to an LTCM default. In any event, LTCM was able to protect its short-term repo positions with term financing (e.g. 6 months) from Wall Street that carried no initial margin. LTCM also secured a \$900 million credit line from Chase Manhattan and other banks.<sup>9</sup>

On the asset side, LTCM was also 'bulletproofing' against a liquidity squeeze. LTCM had initially required investors to keep their money in the fund for a minimum of 3 years. The purpose of this 'lockup' clause, which was very unusual in the hedge fund industry, was to avoid forced sales in case of poor performance.<sup>10</sup>

This strategy worked excellently for LTCM in 1995 and 1996, with after-fees returns above 40%, as shown in Figure 1.<sup>11</sup> The fund placed large bets on convergence of European interest rates within the European Monetary System that paid off handsomely. Convergence had occurred in Europe, as the common currency, the

---

<sup>7</sup> As of December 1997, total swap positions amounted to \$697 billion, futures to \$471 billion, with options and other OTC derivatives accounting for the rest. To give an idea of the size of these positions, the BIS reports a total swap market of \$29 trillion on the same date. Hence, LTCM's swap positions accounted for 2.4% of the global swap market. The futures positions accounted for 6% of the \$7.8 trillion total. Only six banks had a notional derivatives amount above \$1 trillion at the time.

<sup>8</sup> Repo 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.

<sup>9</sup> *New York Times*, 23 October 1998. LTCM used \$500 million from that credit line, which was repaid with the proceeds from the rescue funds.

<sup>10</sup> This risk is very real for hedge funds. Indeed, 60% of hedge funds do not survive 3 years because investors remove their funds.

<sup>11</sup> Returns are derived from *Wall Street Journal* reports.

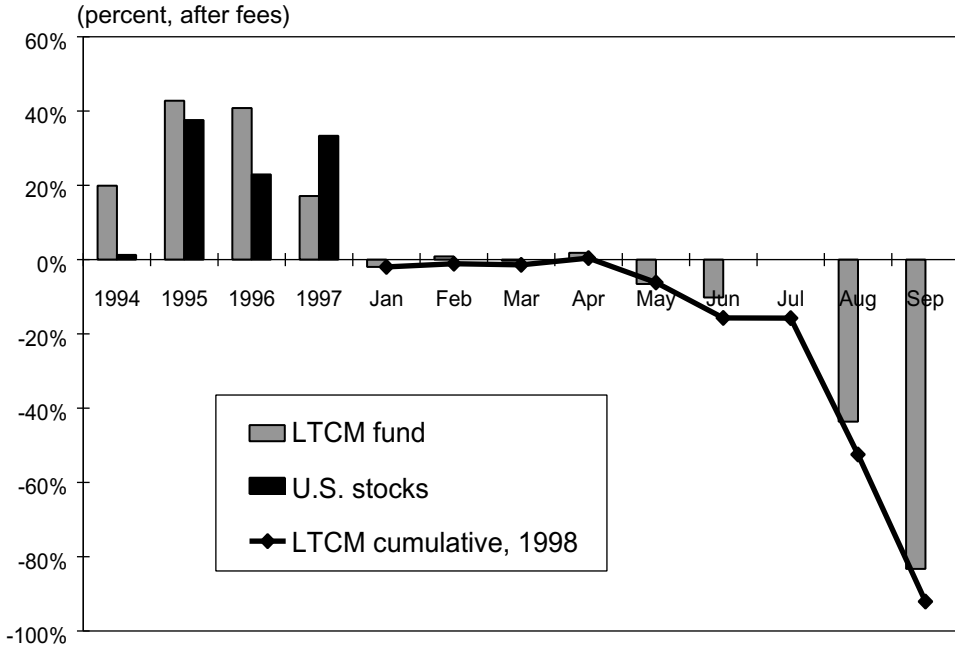


Fig. 1. LTCM's returns.

Euro, came into being on 1 January 1999.<sup>12</sup> By 1997, most spreads had narrowed substantially.

Figure 2 shows the evolution of relevant yield spreads since 1986.<sup>13</sup> 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. But this also meant that convergence trades had become less profitable. In 1997, the fund's return was down to 17% only.

This performance, in fact, was below that of US stocks which gained 33%. The leverage of the fund had also decreased from 25 to 18 due to the asset growth. 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 \$130 billion, explaining that 'investment opportunities were not large and attractive enough'. By shrinking the capital base to \$4.7 billion, the leverage ratio went

<sup>12</sup> Note 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.

<sup>13</sup> The US 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.

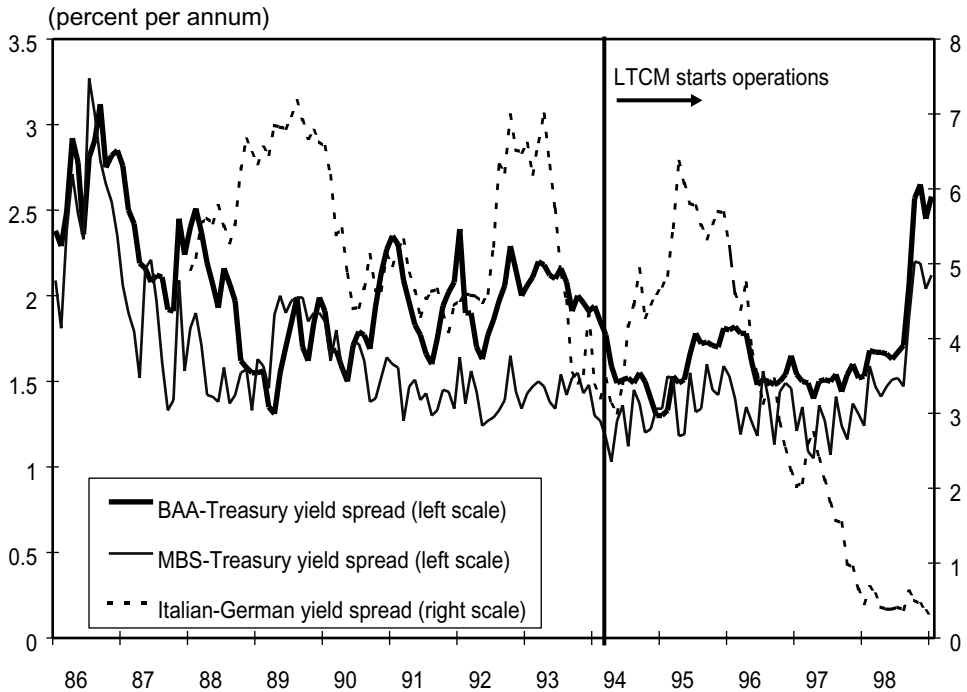


Fig. 2. Bond yield spreads.

back up to 28, amplifying returns to investors that remained in the fund, as shown in Figure 3. Some investors who were forced out were reported to be upset that the partners did not reduce their own equity.

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% loss in LTCM's capital, which dropped from \$4.7 to \$4.0 billion, thereby increasing leverage from 28 to 31.

Then came 17 August. 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, risk premia, and liquidity spreads jumped up sharply. Stock markets dived. LTCM lost \$550 million on 21 August alone. Swap spreads, which usually never moved by more than a couple of basis points every day, had moved by 21 basis points.

By August, the fund had lost 52% of its 31 December value. With assets still at \$126 billion, the leverage ratio had increased from 28 to 55:1. LTCM badly needed new capital. In his 2 September 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 21 September, the fund lost another \$550 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

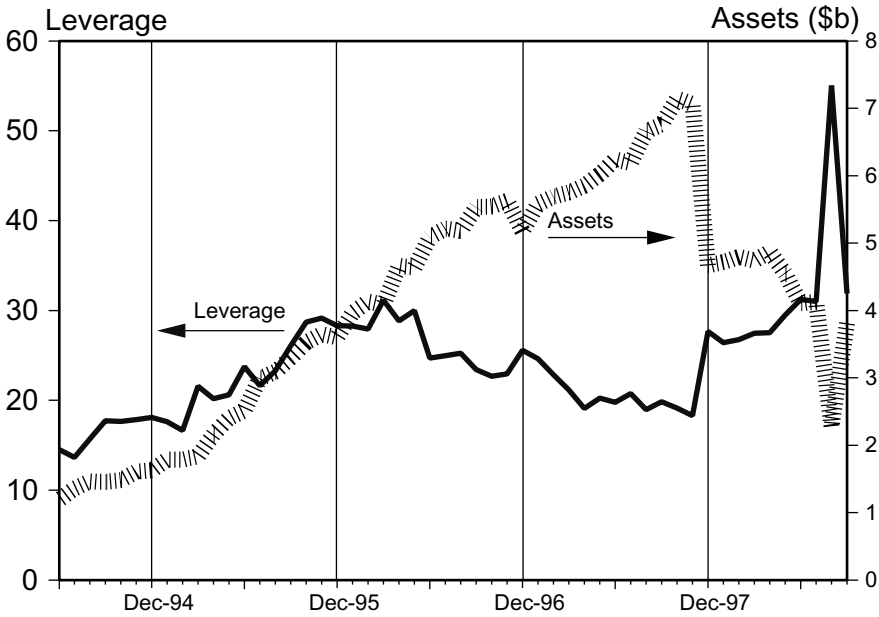


Fig. 3. LTCM leverage and asset growth.

position.<sup>14</sup> It then required 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 next-to-zero 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 been allowed to liquidate their collateral. In contrast, such liquidation is explicitly allowed under the US 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. One of the policy recommendations of the President's Working Group report (1999) was to clear up this uncertainty.

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

These fresh funds came just in time to avoid meltdown. By 28 September, 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. Out of the \$4.4

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

billion lost, a good \$3 billion came from the two main types of bets, interest rate swaps and equity volatility.<sup>15</sup>

After the rescue, LTCM operated 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. The portfolio was unwound over the following months. By the end of 1999, all of the money had been paid back to investors and John Meriwether had started a new hedge fund.

## 2. VAR and equity capital

LTCM now stands as a cautionary tale of the fallibility of some of the most sophisticated VAR models. (CTFC Chairperson Brooksley Born, 1999.)

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, banks now have the option to use their own VAR risk management model as the basis for required capital ratios.<sup>16</sup>

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.<sup>17</sup>

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 model misspecifications.

---

<sup>15</sup> See Lewis (1999). He reports that LTCM had sold 3- to 4-year US and European equity options at the end of 1997, when volatility stood at 20%. LTCM was also reported to have volatility positions on the CAC 40 French stock index that corresponded to 30% of the open interest (*New York Times*, 23 October 1998).

<sup>16</sup> These capital adequacy requirements apply to the trading operations of commercial banks and are detailed in the Basle Committee on Banking Supervision (1996) amendment to the 1988 Capital Accord.

<sup>17</sup> See for example, chapter 4 in Matten (1997).



For LTCM, the choice of the horizon is more delicate than the 10-day period required by 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 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.

Indeed, by the end of August, LTCM had \$2.3 billion of equity capital and \$1 billion excess liquidity. The firm could either reduce risk or raise fresh capital. Due to the enormous size of its positions, it could not cut its risk exposure fast enough. Similarly, it was unable to attract new investors. As the markets continued to slide in September, it became clear that the firm had grossly underestimated its capital needs.

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 derivatives that require no upfront investment. Consider for instance a bank with one investment only, a futures position with a notional value of \$100 million, a margin of \$10 million, and a payoff (or dollar return) of \$5 million.

How do we compute the *rate* of return on this \$5 million 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.<sup>18</sup>

Suppose, for instance, that VAR is estimated to be \$20 million at the 99% 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 4.

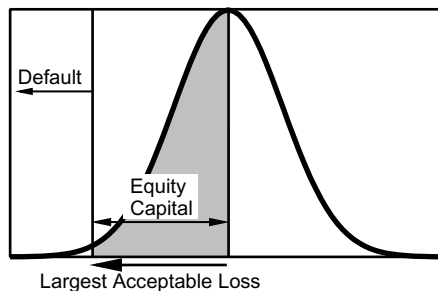


Fig. 4. Equity capital as a VAR measure.

<sup>18</sup> This 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.

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  $\$5/\$20 = 25\%$ . With a lower capital base, the rate of return would be even higher. The drawback, of course, is that the investment is more risky.

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 1-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 1-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 probably 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

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.

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, with six degrees of freedom, 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.

This framework also shows that the quantitative parameters are related. Cumulative default rates increase over longer horizons. Hence the same credit rating can be achieved with a higher confidence level over a short horizon or lower level over a longer horizon.

### 3. LTCM's risk management story

LTCM asserts that the portfolio was managed so that its target risk was no larger than the risk of an unleveraged position in the S&P 500. (Charles Smithson (1999), reporting LTCM information.)

How did the VAR concept 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 US equities.

The annual average volatility of the S&P 500 over 1978–97 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 stated that its target daily volatility was \$45 million around May 1998.<sup>19</sup> 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.

The \$45 million number, however, involves heroic assumptions. It assumes that volatility is constant, when in fact it can easily double in turbulent times. Also, focusing on the second moment is most useful when 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 *et al.* (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 forecasts.

---

<sup>19</sup> See Kolman (1999).

These biases should have showed up in the realized returns. Extrapolating the \$105 million 99% daily VAR to a monthly number yields \$480 million, using the square root of time and a 21-day month. Assuming an annual 18% expected return, the expected profit was \$71 million over a month. Thus, the worst dollar loss at the 99% level over a month should have been  $(\$480 - \$71) = \$409$  million. In May and June, the portfolio lost \$310 million and \$450 million. So, there must have been strong evidence that the models were starting to slip.

The firm reportedly tried to reduce its risk profile, but made a major mistake: instead of selling off less-liquid positions, or raising fresh capital, it eliminated its most liquid investments because they were less profitable. Indeed the 2 September 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 daily volatility should have 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 5 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 J.P. Morgan in 1994, applied to the daily change in credit spreads.<sup>20</sup> The figure shows that, first, the average volatility over 1995–97 was low by historical standards. Second, volatility creeps up on 25 August from 0.019 to 0.027; on 27 August, 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.

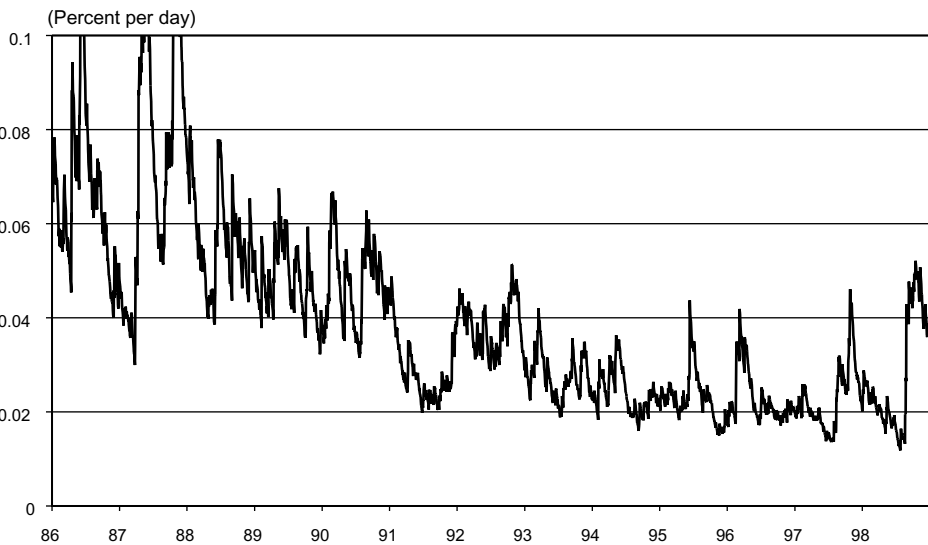


Fig. 5. Volatility of credit spreads.

<sup>20</sup> For the methodology, see J.P. Morgan’s (1994) RiskMetrics manual.

Even worse, on 21 August, the portfolio lost \$550 million. By 31 August, the portfolio had lost \$1,710 million in 1 month only. Using the presumed \$45 million daily (or \$206 million monthly) standard deviation, this translates into a 8.3 standard deviation event. Assuming a normal distribution, such an event would occur once every 800 trillion years, or 40,000 times the age of the universe. Surely this assumption was wrong.

While LTCM managers surely did not believe that returns were normally distributed, this line of reasoning goes to the heart of the explanation that the disaster was 'beyond the fund's capacity to anticipate'. In other words, they claimed this was a '100-year flood'.<sup>21</sup>

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.

In retrospect, the firm made major risk-management blunders. LTCM appears to have relied on recent history to estimate risk, assigning a low probability to events such as sovereign defaults 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.

### **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% 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% on 27 October 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 losses do occur once in a while, however. Most notably, the market lost 20.4% on 19 October 1987.

<sup>21</sup> *Wall Street Journal*, 20 May 1999.

LTCM is also reported to have done stress-testing of its portfolio. The firm had stress loss predictions that it could lose \$2.3 billion in a worst-case scenario. In fact, it lost \$4.4 billion during 1998. It also failed to appreciate that its very size made it impossible to manoeuvre once it had lost \$2.3 billion.

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 basically front-running the LTCM portfolio.<sup>22</sup> 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.

#### 4. Portfolio optimization and leverage

These models enable us to pursue a large position without a severe concern that our strategies are very exposed to interest rates or yield-curve changes. (Larry Hilibrand, August 1994.)

This section presents a stylized example of a portfolio optimization with two highly correlated assets. We use a government and a corporate bond portfolio series and find results that are quite consistent with LTCM's leverage. This exercise is representative of LTCM's relative-value trades using assets with high correlations. It also 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,  $\mu$  represents a vector of expected returns in excess of the risk-free rate, and  $\Sigma$  represents the covariance matrix of returns. The investor maximizes

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

where  $\lambda$  is the investor's risk tolerance, or inverse of risk aversion.<sup>23</sup>

The optimal positions in risky assets are given by

$$w = \lambda \Sigma^{-1} \mu, \quad (2)$$

with the remainder of the portfolio 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 1993–97 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.

<sup>22</sup> See Lewis (1999). Goldman Sachs later dismissed this argument as 'absurd'.

<sup>23</sup> We choose  $\lambda$  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 (1) 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		

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 optimization is illustrated in Figure 6.




Fig. 6. Portfolio optimization.

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?

### 5. The problem with portfolio optimization

In theory, LTCM will look anywhere for a correlation—provided the markets are liquid enough to construct trades and hedges. (*Institutional Investor*, November 1996.)

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 1 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 7 examines the volatility of the current optimized portfolio as the correlation coefficient decreases.




Fig. 7. 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 8. 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,  $\rho$  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  $\rho$  remained above 0.94. This is myopic, however, as a longer history shows that  $\rho$  was much lower five years earlier. It even reached a low of 0.75 in 1992.

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 1 month out of 10,000, or 1 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 1 month out of 306, or one year out of 26. Now the odds of failure are not so remote.

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.

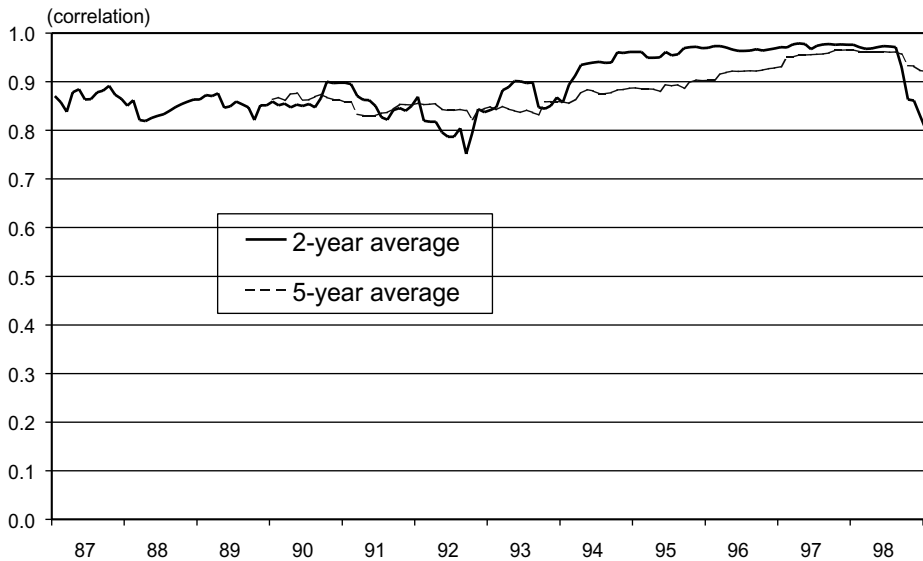


Fig. 8. Correlation between corporate and treasury yields.

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

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.<sup>24</sup>

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 (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.

<sup>24</sup>Michaud (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.

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.

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.

## 6. LTCM's risk profile

The lesson in this was ... Panic was good for business. (Michael Lewis, describing the lessons that Meriwether and his team had drawn from going through the October 1987 crisis.)

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.

To illustrate the driving factor behind LTCM's risks, Figure 9 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 5) 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 10 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

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

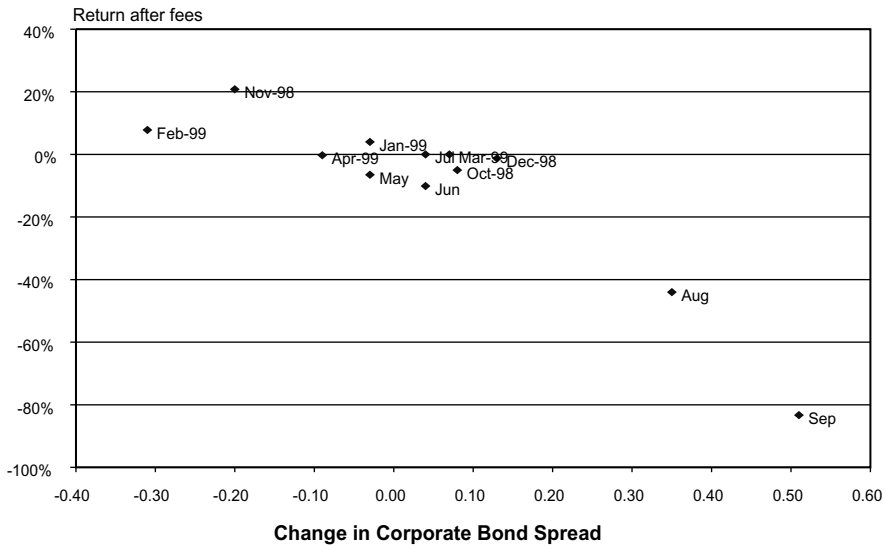


Fig. 9. Explaining LTCM returns.

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.

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




Fig. 10. Simulated daily volatility of the fund.

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 11, for instance, plots the distribution of at-the-money option volatilities over the period 1993–97. 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.<sup>25</sup>

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 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.<sup>26</sup> Overall, the portfolio returns had a distribution similar to a short position in an option, as shown in Figure 12.

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

<sup>25</sup> The 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-fifths of the funds survived through 1998. Similar rates of attrition are reported by Brown *et al.* (1999).

<sup>26</sup> Interesting progress has been made by Almgren and Chris (1998). Liquidity risks involves a tradeoff between *market impact*, which is the negative price effect of quick disposition, and *volatility*, which arises as the position is slowly eliminated. This leads to the notion of *liquidity-adjusted VAR*, which is the minimum VAR of any liquidation strategy for a given time horizon.




Fig. 11. Distribution of stock market volatility.

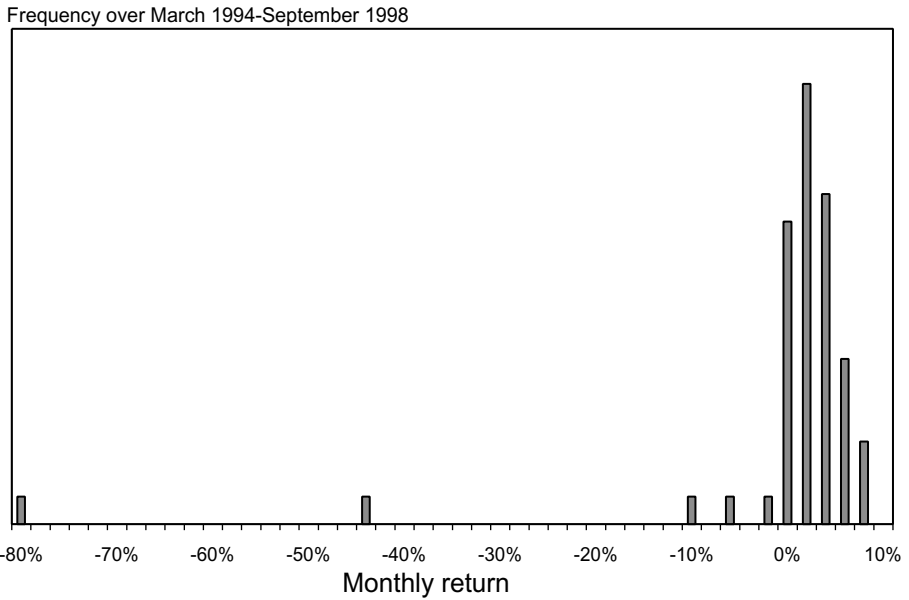


Fig. 12. Distribution of LTCM's monthly returns.

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

The fact that (modern financial engineering) failed does not mean that the science of risk management is wrong-headed; just that it is still an imperfect art in a world where the past holds lessons but provides few reliable precedents. (James Leach (1998), Chairman, House Banking and Financial Services Committee.)

LTCM's interpretation of its failure is that the events of 1998 were 'beyond the fund's capacity to anticipate'. In fact, it had severely underestimated its risk. The reported positions were singularly undiversified, as they were exposed to liquidity, credit, and volatility spreads, which are all similar risk factors. In addition, the payoff patterns of the investment strategy were akin to short positions in options. Even if it had measured its risk correctly, the firm failed to manage its risk properly.

The story of LTCM should not be taken as an indictment of VAR systems, which after all, performed reasonably well for the banking sector in 1998.<sup>27</sup> Instead, it provides a number of useful risk management lessons. First, it illustrates the danger of optimization biases, or 'traders gaming the system'. As this paper has shown, LTCM's strategy can be interpreted in terms of a constrained optimization, i.e., maximizing expected returns subject to a constraint on VAR. This strategy led the firm to its demise, as it created huge leverage and extreme sensitivity to instability in the correlations. We have also learned the danger of building measures of event risk based on very recent price data.

This is not to say that current systems could not be improved. Traditional risk-management models ignore asset and funding liquidity. When positions are relatively large and leveraged, it is important to account for the price impact of forced sales. We also need to make better use of stress-testing, not necessarily repeating past episodes but instead looking at worst-case scenarios for the portfolio at hand. Wall Street lenders also badly need to integrate market and credit risks. The Counterparty Risk Management Policy Group (1999) makes useful recommendations in that respect.

Finally, 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, like a short position in an option. 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

---

<sup>27</sup> The Basle Committee (1999) reviewed the performance of VAR model-based capital charges for market risk. During the second half of 1998, none of the institutions surveyed reported trading losses in excess of the capital requirements.

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

- Almgren, R. and Chriss, N., 'Optimal liquidation', *Working Paper* (University of Chicago, 1998).
- Basle Committee on Banking Supervision, 'Performance of models-based capital charges for market risk: 1 July–31 December 1998', (BIS, Basle, Switzerland, 1999).
- Basle Committee on Banking Supervision, *Amendment to the Capital Accord to Incorporate Market Risks* (BIS, Basle, Switzerland, 1996).
- Bollerslev, T., Chou, R. and Kroner, K., 'ARCH modelling in finance: a review of the theory and empirical evidence', *Journal of Econometrics*, Vol. 52, 1992, pp. 5–59.
- Born, B., *Testimony on Behalf of the Commodity Futures Trading Commission Before the Subcommittee on Capital Markets, Securities, and Government Sponsored Enterprises of the Committee on Banking and Financial Services* (US Congress: Washington, 3 March 1999).
- Brown, S., Goetzmann, W. and Ibbotson, R., 'Offshore hedge funds: survival and performance: 1989–1995', *Journal of Business* Vol. 72, January 1999, pp. 91–117.
- Counterparty Risk Management Policy Group, *Improving Counterparty Risk Management Practices* (New York, NY, 1999).
- Dunbar, N., 'Meriwether's meltdown', *Risk*, October 1998, pp. 32–36.
- Greenspan, A., 'Private-sector refinancing of the large hedge fund, Long-Term Capital Management', *Testimony before the Committee on Banking and Financial Services* (US Congress, Washington, 1 October 1998).
- J. P. Morgan, *Riskmetrics Technical Manual*, (J. P. Morgan, New York, 1995).
- Jorion, P., *Value at Risk: The New Benchmark for Controlling Market Risk* (Irwin-McGraw Hill, Illinois, 1997).
- Ju, X. and Pearson, N., 'Using value-at-risk to control risk taking: how wrong can you be?', *Journal of Risk*, Vol. 1, Winter 1999, pp. 5–36.
- Kolman, J., 'LTCM speaks', *Derivatives Strategy*, April 1999, pp. 12–17.
- Lewis, M., *Liar's Poker* (Penguin, New York, 1989).
- Lewis, M., 'How the eggheads cracked', *New York Times Magazine*, 24 January 1999, pp. 24–77.
- Matten, C., *Managing Bank Capital: Capital Allocation and Performance Measurement* (Wiley, Chichester, 1997).
- Merton, R., 'On the pricing of corporate debt: the Risk structure of interest rates', *Journal of Finance*, Vol. 29, 1974, pp. 449–70.
- Michaud, R., 'The Markowitz optimization enigma: is optimized optimal?', *Financial Analysts Journal*, Vol. 45, 1989, pp. 31–42.
- President's Working Group on Financial Markets, *Hedge Funds, Leverage, and the Lessons of Long-Term Capital Management* (Washington, DC, 1999).
- Pacelle, M., Smith, R. and Raghavan, A., 'Investors may see "LTCM, the sequel"', *Wall Street Journal*, 20 May 1999, C1.
- Smithson, C., *Managing Financial Risk 1999 Yearbook* (CIBC, New York, 1999).