Understanding and Monitoring the Liquidity Crisis Cycle

Richard Bookstaber

In the aftermath of the Long-Term Capital Management debacle, it is clear that a large hedge fund can have a systemic impact on the market. The high leverage, forced liquidations, declining liquidity, and cascade of widening spreads turned the Greenwich, Connecticut, hedge fund's losses into a global event. Without an infusion of capital to stem the need to liquidate to meet margin calls from its creditors, LTCM's demise—and the collateral loss to its creditors—appears to have been inevitable. The natural questions to ask are: What was the cause of the crisis? How can this type of crisis be prevented in the future?

Recent regulatory investigation into LTCM started with an ill-defined target, the community of "highly levered institutions." There are problems, however, with pointing to high leverage as the critical characteristic in the LTCM crisis and, for that matter, in the general description of hedge funds. The most immediate problem is defining what leverage means. After all, using conventional measures of leverage, the leverage of many hedge funds pales beside the 20+ times leverage of the broker/dealer community. Another problem is understanding why leverage should matter. A hedge fund that holds one-year U.S. T-bills with 10:1 leverage would not be considered in the same league as a fund that is levered 2:1 in riskier and less-liquid Russian Ministry of Finance bonds. Finally, not all hedge funds are highly levered by any definition. At any time, we can usually find some of the largest hedge funds with unencumbered cash—that is, not only with positions that are unlevered but also with free cash to spare.

The characteristics of hedge funds and other financial institutions that lead to potential crises do not rest entirely with their ability to take on leverage or with their ability to take large risks or to invest in illiquid markets. If a fund is highly levered but in instruments that have low risk and are highly liquid, the fund not only poses little risk to the market; it poses little risk to its investors. If it is in very risky instruments but unlevered, so that no creditors are involved and it has no risk of forced liquidation that could cascade into a problem for the markets, a fund's failure may be unfortunate for the investors but it does not have systemic implications. If it is in very illiquid instruments but not levered and has the stability of capital to allow a long holding period, a fund is no more of a concern than an insurance company that holds real estate in its portfolio.

What matters is the cycle that begins with the confluence of risk, leverage, and illiquidity—risk of loss coupled with leveraged positions, resulting in a need to liquidate into a market that cascades downward in price because of the rise in liquidation orders and the reduction in liquidity providers.

The Cycle

The liquidity crisis cycle consists of three stages. The first is a loss that acts as the triggering event. The second is a need by the fund to liquidate positions to meet the creditors' margin requirements—or have the positions liquidated for it. The third stage is a further drop in the fund's asset value as the market reacts to the fund's attempts to sell in too great a quantity or too quickly for market liquidity to bear. The drop in prices caused by the need to liquidate precipitates an additional decline in the fund's mark-to-market value, which leads, in turn, to yet more liquidations for margin or redemption purposes.

In principle, these stages can be modeled and the ultimate severity of the crisis can be measured. The likelihood of a triggering market shock can be determined by the distribution of the changes in market prices for the assets. For a fund that has a binding margin requirement, the amount it must liquidate because of the market shock will be related to the inverse of its leverage. For example,

Richard Bookstaber is director of risk management at Moore Capital Management, Inc. Moore Capital Management disclaims responsibility for any private publication or statement by any of its employees. The views expressed herein are solely those of the author and do not necessarily reflect the views of Moore Capital Management or the author's colleagues at Moore Capital Management.

if a 1 percent drop in the market occurs, the variation margin for a fund levered 10:1 will require a liquidation equal to 10 percent of its capital.

The next issue is how much the market price will drop as a result of this flow of liquidity demand into the market. This drop can be roughly estimated by looking at the impact of large trades on the market price, although in all likelihood this approach will underestimate the market impact of liquidations arising from a market shock. During a market shock, the liquidity suppliers—the market makers, broker/dealers, and bargain hunters in the trading community—will recognize the risk of a cascade of liquidity demand and be reticent to be the first ones to take on supply. Indeed, some of the liquidity suppliers may actually find themselves in the position of demanding liquidity.

If we could estimate the price impact from the flow of the margined fund's liquidations, we could chart the full course of the liquidity cycle for the fund. For example, suppose that when the fund sells 10 percent of its initial assets in response to the price shock, that sale lowers the market price by 1 percent. And suppose further that this relationship is linear, so for example, a sale equal to 5 percent of its initial asset size will drop the market price by 0.5 percent. The cycle will continue to work itself out by the subsequent 1 percent drop leading to a sale of 10 percent of the fund's remaining assets. The price will now drop by only 0.9 percent because now the fund's remaining assets are only 90 percent of its initial portfolio. This price drop will lead to a margin-induced sale by the fund of 9 percent of its still remaining assets, which in turn, will trigger further price drops by the market, inducing more margin sales, and so on. With the values we have assumed here, everything will finally converge with the fund down more than 30 percent and prices down by slightly more than 4 percent.

For lower levels of price elasticity of liquidity supply, we could end up without convergence. The cycle would become a market crisis as the drops in the market price led to similar needs to liquidate at other firms. The flood of liquidation would accentuate the price drop at every turn. The point of no return would come when the effect of liquidation elicited greater demand for margin for the remaining fund position than the amount of cash it would raise from the liquidation. That is, suppose that in order to raise $10 million to meet the 10 percent margin requirement, the fund must offer down the prices to an amount that causes the prices of its remaining positions to decline by more than $10 million. The fund then faces yet another call to meet the resulting mark-to-market loss. It is caught in an ever-widening downward spiral and cannot satisfy its creditors' or investors' demands no matter how aggressively it sells. Indeed, the very need for aggressiveness in liquidation becomes the root of the problem. In this situation, although poor investment decisions or adverse market events acted as a catalyst, once the losses moved beyond a critical point, the crisis became self-sustaining; it fed off the need for liquidity, and that need for liquidity did not come from the market but from the demands imposed by the fund's creditors and investors.

To understand the risk of crisis, we need to understand each stage of this cycle—the risk of a large loss, the risk that the loss will force liquidations, and the risk that a forced liquidation will adversely affect market prices.

Risk of Loss. In a simple world of static investments, profit or loss is equal to the position size times the change in asset price. If the position is static, then to assess the probability of loss, we need only measure the variability of the asset prices. We can view the assets' historical prices, estimate or assume correlations between the assets, and then measure the probability the portfolio of assets will change by a given amount. With the fund's position a given, we will then know the range of possible profit or loss that will occur with, say, a 95 percent probability.

The resulting measure is the portfolio's value at risk, a standard risk measure used by virtually all major investment firms. Unfortunately, VAR is particularly ill suited to assessing the risk of a market crisis. During a crisis, the correlation between the assets can change dramatically and unexpectedly, with the result that positions that were thought to be diversifying—or even hedging—end up compounding risk. The reason is that the data used to compute the variability and the correlation of the market prices for VAR analysis come from typical days, when prices and correlation are driven by economic relationships, but during a crisis, prices and correlation are driven by what various investors are holding in their portfolios and which of these investors need to liquidate. Two assets that are unrelated on a market or economic basis will exhibit high correlation in the course of a liquidity-driven market crisis if those who are being forced to liquidate a position in one asset also happen to have a position in the other asset. The correlation may be negative rather than positive if one of the assets is held long and the other short. Thus, in times of market crises, correlations between assets will tend toward 1 in absolute value, but we cannot predict ahead of time whether a correlation will be positive or negative.

Understanding the liquidity cycle helps put VAR into perspective as a risk-management tool and helps us understand its limitations in depicting "tail risk" (the risk from extreme events). VAR does not work for measuring risk in a situation like that of LTCM because VAR measures only an initial shock that leads to a crisis; it does not measure—and was not designed to measure—the effect a shock might have on the demands for liquidity in the market or the cascading effect of forced liquidation on markets being pushed toward illiquidity. VAR can measure the noise level, but it cannot gauge the stability of the snow pack in the face of that noise or predict the avalanche to follow.

Need to Liquidate. Losses may be painful for the manager and fund investors, but market losses alone do not transform a bad month into a market crisis. Indeed, if the investors are sophisticated and the fund is following its investment guidelines, market regulators can fairly stand back with a statement of "caveat emptor." The stage at which the broader market must start taking note of the losses is when those losses trigger a need to liquidate.

As we saw with the Tiger Fund, whose assets dropped by two-thirds by the time of its closing, investor redemptions can force liquidation just as surely as creditors' margin demands can. But many large funds have redemption windows and sources of long-term capital that reduce the impact redemptions can have on the need to liquidate. Redemptions, therefore, become a concern only in the cases of extremely large and illiquid positions.

Far more important than redemptions are the implications of leverage. If a fund is not levered and is not in a redemption window, then it cannot be forced into a noneconomic liquidation decision. In the short term, a fund crisis certainly will not be pleasant for the investors, and it may shorten the manager's career, but if the market environment is such that selling off the positions will be difficult, the manager of an unleveraged fund can stand fast. If a fund is levered 2:1, however, as losses approach 50 percent, it will be running out of capital and will need to sell assets to meet its obligations to its creditors and counterparties. If the leverage is 20:1, all that is needed to trigger a margin call and a forced liquidation of some positions is a drop of 5 percent in the asset holdings. If the managers do not liquidate to raise the cash, the lenders will do it for them.

The risk in leverage is thus not the leverage itself but the potential for that leverage to invoke a forced liquidation of positions. Being levered 20:1 in a spread between the yields of on-the-run (recently issued) and off-the-run (previously issued) Treasuries, where a move of more than 10 or 20 basis points (bps) is nearly unthinkable, may be less of a risk in this regard than being levered 2:1 on a stock. But in the case of LTCM, the unthinkable did occur, with the on-the-run / off-the-run Treasury spread, LIBOR spreads, and volatility spreads all moving far beyond their historical levels. So, the issue is not simply a move in the market and not simply the implications of that move leading to a need to liquidate; it is the feedback cycling between liquidation and market prices.

Market Liquidity. If an individual investor faces a margin call and cannot come up with the cash, the investor must quickly sell out the position. The margin call is usually made while a buffer still exists between the market price of the asset and the loan amount, so the lender is made whole by the sale. If a large hedge fund faces a margin call, however—or, for that matter, if many individual investors face margins calls at the same time—the flood of liquidity demand can lead to quite a different outcome.

If we are to sell an asset, someone must be on the other side of the trade to buy it, and the more of the asset we need to sell, the more capital must be at the ready on the other side. Economics 101 tells us that to attract more capital, we need to lower the price of the asset. In liquidation mode, moreover, quantity is not the only constraint; time is also a constraint: We need to sell, and we need to sell quickly. The only way to get capital quickly is to drop the price even more. But the number of investors who are ready to buy quickly is limited. Most investors make strategic plans for buying assets. They map out ideal holdings and do not monitor the market minute by minute, much less are they willing to jump into the market without careful consideration and analysis. Matters are made even worse during periods of liquidity crisis. The explosive changes in price and the increase in volatility make many investors second-guess their earlier investment plans because of a perception of higher risk, a reevaluation of asset fundamentals, or a shift in the return they are demanding.

The subset of market participants who have ready capital to deploy and who will act quickly with that capital are the liquidity providers—the market makers and the traders. Indeed, their economic value and profitability comes from keeping capital free to provide liquidity when prices signal a need for liquidity by other market participants. It is these liquidity providers who face the liquidation crisis in this analysis; in the crisis, the liquidity providers turn into liquidity demanders. With the front line of liquidity providers out of the picture,

a drop in prices—especially a drop that comes too quickly for the remaining potential sources of market liquidity to react—rather than eliciting supply, will signal to other market participants that the liquidity is gone from the markets. Thus, the rapid price drop actually reduces the willingness of others to come in and fill the breech.

The very ability to liquidate—clearly a desirable attribute of an investment portfolio—is, ironically, the root of the liquidity crisis cycle. The growing spiral of a crisis occurs for financial institutions but not for corporations facing bankruptcy because financial institutions are liquid and corporations are not. When the run-of-the-mill corporation cannot meet the terms of its creditors, no attempt is made to throw assets into the marketplace, because the assets have no immediate market. Indeed, the trigger of the liquidity cycle is unique to financial institutions because their assets are liquid and can be marked to market.

And whereas market risk is readily apparent, liquidity is not. Without actually doing the trade, it is difficult to gauge what the market implications of a large liquidation will be.

Measuring the Cycle

In a purely mechanistic physical universe, if we knew the location and velocity of every particle, we could apply the laws of Newtonian mechanics to project the future course of nature. In the world of finance, the elementary particles are the financial assets. In a purely mechanistic financial world, if we knew the position each fund had in each asset, the margin and leverage terms the fund had for each of those assets, and the ability and willingness of liquidity providers to take on those assets in the event of a forced liquidation, we would be able to fully understand the vulnerability of each fund.

It is tempting to get down to the elementary particles, to look at positions rather than dealing with aggregated measures such as liquidity, risk, and leverage. After all, if we knew the positions, we should know everything. The problem with stopping at elementary particles in the financial world is the same as in the physical world: the inherent randomness and complexity of the systems.

The push for position-level transparency is largely a reaction to the LTCM crisis. LTCM's creditors claim to have been caught unawares by the high leverage LTCM was using. Yet, position information was not needed to see LTCM's leverage and the resulting risk. That LTCM engaged in relative-value trading, principally in swap and mortgage spreads, was well known. The objective in those trades was to capture pricing differentials of 10–20 bps, just one-tenth to two-tenths of a percent over long holding periods. Yet, the return objective of the hedge fund—and its historical realized return for a number of years—was well above 20 percent. The only way to generate 20 percent returns from 20 bp mispricings is to lever 100:1. If the banks could not do this simple back-of-the-envelope calculation based on the most cursory analysis of LTCM's trading strategy, I cannot see what they would have gotten from poring over reams of position sheets.

Moreover, the difficulties with using position information come from many sources. Some instruments are complex and difficult to price. Some positions involve hedges for which a slight error in the transmission of one part or asynchronous pricing of the various legs of the strategy will grossly misstate the total exposure. The problems and inaccuracies in using position information to assess risk are apparent in the fact that, even with their huge resources, major investment banking firms choose to use summary statistics rather than position-by-position analyses for their firmwide risk management.

For the liquidity demander, transparency of positions is good. Liquidity demanders would like to announce their need to buy or sell as broadly as possible, even prearrange the time that they will trade and the amount they will trade. Doing so would lead to the largest group of liquidity suppliers bidding for their positions, and it would also indicate that a liquidity provider is not trading based on information; that is, the provider is not trying to pick anyone off.

For the liquidity supplier, transparency is bad. Unlike the liquidity demander, the liquidity provider does not intend to hold the position long term. Liquidity providers will come back to the market to sell off their positions—ideally, when other investors who need liquidity are on the other side of the market. If other traders know a provider's positions, they will also know that those positions are likely to be returned shortly into the market. The other traders will either be loath to be the first ones on the other sides of these trades or, knowing the overhang that remains in the market, will demand price concessions if they do trade. Thus, increased transparency will reduce the amount of liquidity provided for any given change in prices. A clear demonstration of the liquidity provider's desire for anonymity is the fact that in even the most liquid markets, the broker/dealer market makers frequently use brokers when they could easily enter their bids into the market directly.

The impact of transparency on the market when there is position overhang by liquidity pro-

viders can be seen in the case of LTCM.[1] Although still not widely recognized, one catalyst for the liquidity crisis of LTCM was the public announcement on July 6, 1998, of the closing of Salomon Smith Barney's U.S. fixed-income arbitrage unit, which created an overhang in the market for positions that LTCM shared. Market liquidity dropped precipitously, so when LTCM needed to sell off positions, it could find few willing to take the other side without a huge price concession.

Looking back past LTCM to other major crises, I am hard pressed to think of a case in which position information would have provided the information necessary to see the problems emerging. Much easier means of discovering the problems existed. In these cases, as in the case of LTCM, lack of information was not the difficulty; it was an inability to understand, interpret, or act on the information.

Instead of adding complexity and mass to the information, we should focus on the proper use of well-tested tools that can provide an assessment of a hedge fund's exposure to each of the three components of the liquidity cycle.

The first component, market risk, is the most widely analyzed component, and a number of tools are available for looking at this risk—most notably, VAR, stress testing, and scenario analysis.

The second component, leverage, can be analyzed as an adjunct to the market risk analysis. In the context of the cycle, leverage dictates the amount of loss that can be sustained before a fund is forced into liquidation mode. By itself, leverage is not very informative, however, because we need to know the likelihood that such a level of loss will occur. To analyze this likelihood, we need to couple leverage with two other measures—the risk of the portfolio and the aptitude of the managers for reducing risk in the face of loss. Can the manager reduce positions as losses mount; that is, are the holdings liquid? And, more importantly, *will* the manager reduce the holdings? In some strategies, the relative-value strategies of LTCM being a case in point, the instruments are liquid but the trading approach actually increases exposure in the face of loss. The thinking in these trades is that if a trade between two instruments loses money, the mispricing has increased. Therefore, if the trade was good before the loss, it is even better after the loss, and the position should, accordingly, be increased.

If we have a fund's trading history, we can construct a simple analysis to measure the leverage and the risk-reduction behavior of the fund in a VAR framework:

- *Measure the fund's risk.* The VAR already provides this measurement in a statistical sense, but restating this risk for a hedge fund in terms of a standard asset class is useful for providing an intuitive sense of the risk. This is especially useful when the decision-making investor is unaccustomed to statistical measures. For example, suppose we decide to use the U.S. equity market as represented by the S&P 500 Index for comparison and assume the fund has a one standard deviation move in daily profit and loss of $5 million. The first question to ask is: How much S&P 500 exposure would lead to the same level of risk? A typical move in the S&P 500 is about 1 percent, so it would take $500 million of equities to reach the same level of risk as the hedge fund portfolio. By restating the VAR in terms of some standard instrument's exposure, we no longer need to think in terms of abstract confidence intervals. We can base our thinking about risk on our market intuition.
- *Estimate the leverage.* The next step is to get an intuitive sense of leverage by looking at the asset exposure in the context of the total capital of the fund. If the fund in this example has $1 billion of capital, we can think of its risk leverage as being the same as if it were 50 percent invested in the S&P 500. With an S&P-equivalent risk-to-capital ratio of 0.5, the hedge fund does not appear to be levered. It may be taking positions on margin by holding futures contracts, or it might have a repurchase agreement position on Treasuries, but once the risk and size of these positions are taken into account, it is no more levered on a risk-related basis than if it invested its $1 billion to buy $500 million of stocks.
- *Measure the effective liquidity.* "Effective liquidity" means not only the liquidity of the assets but also the willingness of the managers to alter their exposure to those assets. This willingness can be seen in the variability in the S&P-equivalent exposure over time. If the exposure, adjusted for changes in asset size, is constant, then the fund tends to be illiquid, either because of the nature of the assets held or because of the risk-management approach of the manager. An exposure that varies widely and rapidly over time indicates that the assets are liquid and the manager is willing to take advantage of that liquidity.

The important aspects of the analysis of the fund's risk can be summarized by looking at the risk-to-capital ratio divided by the standard deviation of that ratio. The numerator is a qualitative measure of the likelihood a change in asset prices will require a liquidation of the fund's position. The denominator indicates how easily the market can absorb any liquidation and measures the fund's willingness to change and its experience in changing its exposure. Increasing risk-based leverage

will increase the numerator; a decline in effective liquidity will decrease the denominator. In either case, the ratio will increase, which implies a higher risk of crisis.

To see how this measure would work in practice, consider the implications this ratio would have had for LTCM. The fund's leverage would have led to a high risk-to-capital ratio. When the fund gave back capital early in 1998, the ratio would have doubled, signaling a major increase in the potential for crisis. The denominator would have been small because LTCM's long-term strategies maintained a constant level of exposure in spite of mark-to-market variations. One could argue that the stability of the risk level was not a sign of illiquidity—after all, LTCM traded in some of the most liquid markets in the world—but from an operational standpoint, this assertion remained untested. Whether because of market illiquidity or a management decision, LTCM was slow to vary positions. With a high numerator and low denominator, the fund would have had a high crisis risk ratio. This ratio could have been reduced in either of two ways: The numerator could have been reduced by reducing leverage, or the denominator could have been increased by taking shorter-term trading positions or by changing exposure in reaction to profitability.

The ratio gives an indication of whether a hedge fund is prone to a liquidity crisis, whether it will start an avalanche. Another application is to see what the likelihood is of being caught in that avalanche. If the manager of a hedge fund finds that other hedge funds with similar types of trading interests and tactics tend to have high ratios, the manager will know that there is a good chance that the markets in which the fund trades will be the center of a liquidity crisis and the fund will face an outlier event. The average value of the ratio across all funds, therefore, can be used to shed light on the component of risk that is the most difficult to measure but most critical—the tail risk that comes from market crises.

The risk of the assets being held is not the only thing that matters. A crisis can be triggered for low-volatility assets as well as for higher-risk assets. As long as the assets have some risk, leverage and a binding margin requirement can generate a crisis. Once the cycle of liquidations begins, the economics of the market and the inherent riskiness of the assets are replaced in importance by the extent of the margin and the ability of the market to absorb assets under conditions of market stress.

The ratio is a starting point in addressing the risk from a liquidity crisis. It approaches the problem qualitatively rather than trying the more pure but more difficult route of modeling the dynamics of the cycle. It faces the limitations of any statistical representation because it boils down a complex, multidimensional problem into a single number, with the objective of conveying as much of the relevant information as possible in that number.

Conclusion

A market crisis is not simply a "bad draw" from the distribution of day-to-day price moves. The genesis and dynamics of market crises have little to do with the information and the market flows that affect prices on typical days. A market crisis is a crisis of liquidity far more than it is a crisis born of information.

The role of liquidity in market crises has not been widely studied. One reason is that the role is difficult to investigate empirically. The liquidity that matters for a crisis is not the same as the liquidity that lies behind bid–ask spreads; it is not a phenomenon that can be readily observed. Market and economic information is widely available, but liquidity demand and supply are locked within the preferences and capital constraints of each trader. The role of liquidity in market crises, fraught as it is with the complexities inherent in any dynamic process, is also difficult to model analytically. And the market's behavior during a liquidity-based crisis can seem at odds with economic rationality. During the LTCM crisis, for example, U.K. swap spreads moved to historical highs as German rate volatility moved to near-historical lows. On-the-run versus off-the-run spreads in the U.S. Treasury market widened to levels that had never before been seen and seemed to leave huge potential profits untouched. These moves had nothing to do with information or economic relationships; they were the result of forced liquidations in the absence of liquidity providers.

Unfortunately, the lessons we can take from the LTCM story go only so far in illuminating future crises. The market dislocations of a liquidity-based crisis are both unpredictable and unstable; they will change every time a fund changes its positions.

Notes

1. See also Richard Bookstaber, "A Framework for Understanding Market Crisis," in *Risk Management: Principles and Practices* (Charlottesville, VA: AIMR, 1999).