Forced Liquidations, Fire Sales, and the Cost of Illiquidity

RICHARD R. LINDSEY AND ANDREW B. WEISMAN

RICHARD R. LINDSEY is the chief investment strategist for liquid alternatives at Janus Capital in Denver, CO, and a fellow at the Courant Institute of New York University in New York, NY.
richard.lindsey@janus.com

ANDREW B. WEISMAN is the chief investment officer for liquid alternatives at Janus Capital in Denver, CO.
andy.weisman@janus.com

Institutional investors seeking diversification often build portfolios using collections of securities with widely varying characteristics.[1] To help achieve diversification, investors generally use the common currencies of reported return, volatility, and correlation to construct or optimize their portfolios. As a result, investors using this approach are often drawn to investment opportunities that appear to exhibit diversifying properties simply because of the limited price discovery associated with those investments. Such opportunities are often relatively illiquid when compared with traditional investments such as large-cap stocks or sovereign debt, and investors frequently take for granted that they receive a liquidity premium that compensates them fairly for the lack of liquidity.[2] A variety of approaches have been proposed to incorporate liquidity into the portfolio optimization process: Seigel [2008] and Leibowitz and Bova [2009] developed methods for institutional investors to explicitly take liquidity into account when determining optimal asset weights; Ang, Papanikolaou, and Westerfield [2011] characterized an investor's optimal liquidity policy when there are frictions in the market; Lo, Petrov, and Wierzbicki [2003] added liquidity as an additional constraint in a mean–variance optimization; and Kinlaw, Kritzman, and Turkington [2013] incorporated liquidity as a shadow allocation to the portfolio.

The most common way to measure illiquidity in investments such as hedge funds or private equity is serial correlation in the investment's reported return series,[3] because such serial correlation is frequently viewed as the result of price smoothing caused by exposure to less liquid securities or investments. More sophisticated investors may adjust the return data by taking into consideration observed serial correlation in order to decode the portfolio's true volatility; they thereby correct both the volatility and the investment's risk-adjusted performance (Scholes and Williams [1977], Geltner [1993], Getmansky, Lo, and Makarov [2004], Bollen and Pool [2008], Anson [2010 and 2013]). Simply adjusting for serial correlation, however, fails to measure or capture the core risk and cost of illiquidity: forced liquidations and fire sales.

Forced liquidations typically occur when illiquid portfolios become overvalued relative to their true market value (a relevant, timely valuation) and the reported valuation is no longer credible. Fire sales, or rapid asset sales that depress prices, result when managers attempt to sell illiquid instruments or investments quickly. When these sales occur during significant adverse movements in the broader market, with an associated high demand for general liquidity, the price depression is exacerbated. Most importantly, forced liquidations and fire sales often occur without warning, because they are precipitated by factors outside the investor's control.

Many investors do not understand the true risk or cost of illiquidity until a forced liquidation or fire sale actually occurs—unfortunately, too late to help them. But by applying the barrier option–pricing framework presented in this article to an illiquid investment's expected return, investors can often know the probable cost of illiquidity in advance. The method described here lets investors use a combination of market data and experience in a consistent, analytically rigorous framework to derive a fair-value estimate of the cost of illiquidity.

CAUSES OF ILLIQUIDITY

When it comes to liquidity, not all securities are created equal. Ask any experienced portfolio manager to describe the agony associated with exiting a losing investment in the face of an illiquid, declining market—there is a qualitative difference and a real cost to exiting an illiquid portfolio, compared with exiting a liquid portfolio. Yet, in spite of the lessons learned during the global financial crisis of 2007 to 2008, many institutional investments remain illiquid, and this illiquidity may be exacerbated during an extended low-interest-rate environment with low volatility, because asset managers are seeking higher returns.

One can think of the primary cause of illiquidity as a mismatch between the underlying investment's funding and the horizon over which the investment can be sold. Further, the greater the leverage employed in the investment, the more likely it is that illiquidity will have a deleterious effect on the investment's value in a declining market. Positions in exchange-traded securities can generally be sold very quickly, although the seller of a large holding may experience significant price decline. In contrast, for an investment in real estate, even if the investor is willing to sell at a steep discount, it may still take a long time to find a buyer. When investments are supported through the use of any form of short-term leverage[4] or transactions that have embedded liquidity puts,[5] there exists the potential for a funding mismatch and therefore illiquidity.

Additional factors can increase any particular investment's illiquidity. Contractual terms, such as redemption notice periods, lockups, or gates, have liquidity costs that Ang and Bollen [2010] explored. They estimate that a three-month redemption notice period, combined with a two-year lockup for hedge funds, costs investors 1.5% of their initial investment. If a gate is imposed, there is an additional cost that can exceed 10.0%.

Other, so-called network factors may not be as readily apparent. The use of common service providers (custodians, prime brokers, securities-lending counterparties, or pricing providers); common investors, such as funds-of-funds or large institutions (Bhattacharya, Lee, and Pool [2013]); or strategies that turn out to be correlated in unanticipated ways (Boyson, Stahel, and Stulz [2010]) can create unforeseen illiquidity. In general, factors that cause implicit linkages may serve to create or increase illiquidity for a particular investment or portfolio.

Investors in collective investment vehicles (such as hedge funds or private equity) are also subject to the actions of other investors in the same (or similar) portfolios. In some circumstances, if even a single large investor decides to exit an investment, it can cause managers to sell assets to meet redemptions (Gennaioli, Shleifer, and Vishny [2012]). A high degree of leverage in the portfolio can also result in a rapid decrease in the investment vehicle's value and thereby cause other investors to react. The economic advantage to being an early redeemer if a portfolio or asset is under stress is well known (see Mitchell, Pedersen, and Pulvino [2007] or Chen, Goldstein, and Jiang [2010], among other examples), because slower investors end up holding shares of an increasingly less liquid portfolio (Manconi, Massa, and Yasuda [2012]). For that reason, research has shown, sophisticated investors withdraw much more quickly when there are questions associated with an investment's liquidity (Schmidt, Timmermann, and Werners [2013]).[6]

LIQUIDITY AND REALITY

Consider first the nature of returns associated with trading and pricing a liquid security (or portfolio of securities) versus those of an illiquid security. For liquid securities, there exists virtually continuous, objective, and reliable price discovery; generally even large transactions do not suffer a significant increase in cost upon liquidation. In contrast, illiquid securities may trade only by appointment, at infrequent intervals, and without reliable, objective, public reporting.[7] Sellers of illiquid securities cannot be certain about the market value of their holdings, but sellers know (or should know) that a security's sale price is likely to vary depending on the amount to be sold, the need to effect the transaction quickly, and price pressure associated with other sellers in the market.[8] All else being equal, if a seller needs to sell a large amount in a short time, the price can decline dramatically.

This lack of a liquid, transparent pricing mechanism tends to produce relatively predictable behavior from the managers of less liquid portfolios. For example, it has been documented that portfolios of less liquid securities exhibit a high degree of positive serial correlation.[9] In these cases, a significant proportion of the current period's return can be statistically explained by the prior period's return. This serial correlation is, of course, less likely to be the product of a wide-scale exploitable market anomaly than the result of the valuation practices of the managers of such less liquid portfolios. (It may be useful to think of private equity or certain hedge fund investments.) Absent objective pricing, these portfolios tend to get priced conservatively. In other words, the manager adjusts a stock's prices by some proportion of the perceived difference between the point where the stock was marked during the prior period and the point at which the stock is thought to be tradable today. This approach does not necessarily imply nefarious behavior on the part of the manager,[10] but it can represent a simple Bayesian updating rule, using a partial adjustment or adaptive expectations approach, which is rational in that it minimizes the mean squared difference between the estimated value and the market value (Quan and Quigley [1991]).

Over time, assuming markets move both up and down, the illiquid asset's reported value (reported asset value or R) will be a relatively unbiased predictor of the true,[11] tradable value (true asset value or N).[12] However, to the extent that a portfolio is conservatively marked, the reported returns will be autocorrelated; in other words, the current or observed return r_t^o will be determined, in part, by the prior period reported return $r_{t-1}{}^o$.

Why is this so? Consider the following equation, which describes a typical conservative pricing strategy. The reported return for the illiquid portfolio in the current period r_t^o is determined by adjusting the change in portfolio value by some proportion, λ (the reporting adjustment), of the difference between the prior period reported value R_{t-1} and N_t, the true, tradable value today.

$$r_t^o = \frac{\lambda\left(N_t - R_{t-1}\right)}{R_{t-1}}$$

We assume that the portfolio N_t, follows a discrete Brownian motion:

$$N_t - N_{t-1} = N_{t-1}\mu\Delta t + N_{t-1}\sigma\varepsilon_t\sqrt{\Delta t}$$

Here, the portfolio's true value in the current period N_t is equal to its true value in the prior period N_{t-1} multiplied by the trend rate of return μ and the one-period time step Δt, plus the assumed volatility of the process σ multiplied by the product of ε_t (a standard normal random variable) and the square root of the one-period time step (Δt). Thus, the change in value in any given period is the result of a combination of a trend and a random shock.

Substituting the Brownian process into the equation for r_t^o and rearranging results in the following, an expanded representation of the observed rate of return over a period of length (Δt):

$$r_t^o = \frac{\lambda\left(N_{t-1}\left(\mu\Delta t + 1\right) - R_{t-1}\right) + \lambda\left(N_{t-1}\sigma\varepsilon_t\sqrt{\Delta t}\right)}{R_{t-1}}$$

Finally, taking the expectation of r_t^o yields[13]

$$E\left[r_t^o\right] = \frac{\lambda\left(N_{t-1} - R_{t-1}\right) + \lambda N_{t-1} m\Delta t}{R_{t-1}}$$

As a consequence, some proportion of the expected value for period Δt will be explained by the actual change in portfolio value and some proportion will be explained by the difference between N_{t-1} and R_{t-1}. However, for smaller values of λ (where the portfolio manager incorporates less of the change in the portfolio's value through time), we would expect the absolute value of $(N_{t-1} - R_{t-1})$ to grow larger through time. The effect of this partial adjustment is to induce first-order serial correlation $\rho_{s(1)}{}^o$ in the observed return series as the expected return in each period becomes increasingly dominated by the difference between N_{t-1} and R_{t-1}. This first-order serial correlation is proportional to $(1 - \lambda)$, even though the error term that helps drive the portfolio's true return may be independent through time. For a quick and dirty estimate of the extent to which a portfolio manager is conservatively marking a portfolio (understating the change in its real value), one can simply calculate the first-order serial correlation and subtract it from 1 to yield an estimate of the proportion λ of the true change in portfolio value that the manager is reporting.[14]

As noted, even with a process that systematically under-adjusts for changes in valuation from time period to time period, the reported asset value may still be a relatively unbiased representation of the true value. Given that, the question becomes whether this misreporting is merely a benign understatement of the portfolio's true volatility, which we can ignore, or has a real cost to the investor.

THE BARRIER OPTION FRAMEWORK

To answer this question, we begin by examining the dynamics of how the two related processes N_t and R_t evolve over time. Exhibit 1 illustrates how R_t might track N_t over a randomly generated period of 60 time intervals (for this simulation, assume that $\mu = 0.05$, $\sigma = 0.25$, and $\lambda = 0.25$). The conservative reporting process tends to smooth portfolio valuation through time and, as expected, exhibits less price volatility than does the actual underlying series.

EXHIBIT 1
Hypothetical Growth of $100

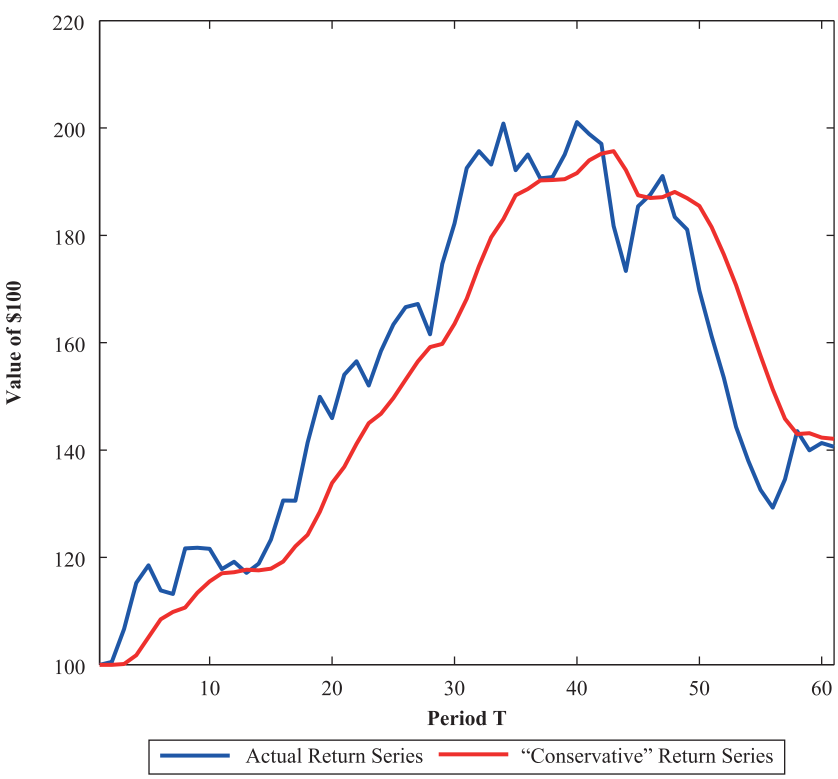

Exhibit 2 depicts the individual period differences between R_t and N_t, where values above zero correspond to periods when the manager overvalues the portfolio relative to its true value. The manager's conservative approach to portfolio valuation results in periods of both significant under- and overvaluation. From practical standpoint, undervaluation tends to be less important than overvaluation to an investor, because it's rarely harmful to own an investment that is worth more than its stated value. Undervaluation is also rarely a concern for third parties providing financial support to a portfolio, as those parties are overcollateralized. But parties providing financing are usually quite interested in overvaluation. For example, a prime broker extending credit to finance portfolio positions will want to ensure that the manager's appraisal of portfolio value doesn't exceed some rational tradable value by more than a reasonable margin, because that value serves as collateral for the financing.

This reasonable margin of overvaluation may be referred to as the credibility threshold L. Exhibit 2 sets the threshold at 15.0%. When a manager exceeds the credibility threshold, interested parties often respond; prime brokers tend to take action promptly, but the response can be slower from larger, more bureaucratic organizations, such as institutional investors, particularly when financial reporting is delayed (as is the case for hedge funds). Regardless of where it arises, a breach of the credibility threshold is likely to trigger forced behavior from the manager. In other words, the manager will be required to sell some or all of the illiquid portfolio in a relatively short time, typically in a descending (or thin) market. The single-period loss that occurs thus consists of two components. The first is a loss governed by $(R_t - N_t)$, the extent to which the portfolio was overvalued. The second is a liquidation penalty P associated with a fire sale of the illiquid portfolio in a (typically) descending market. This liquidation penalty increases when the portfolio contains significant leverage, because it is likely that more of the portfolio will need to be liquidated. Large, single-period losses of this type are relatively common in financial markets and

EXHIBIT 2

"Conservative" Less Actual Returns

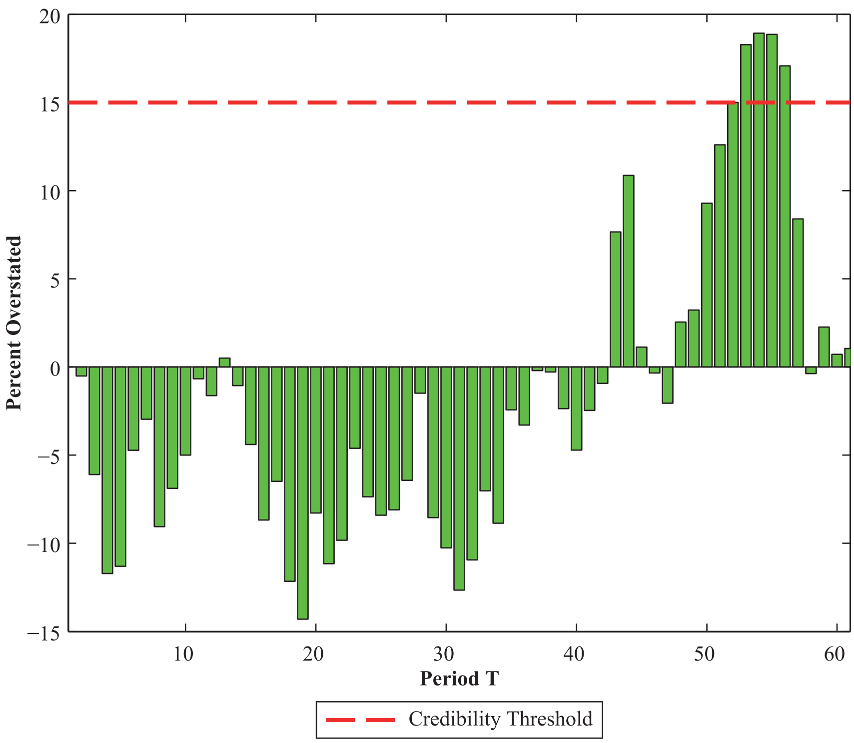

tend to be larger than losses estimated through the use of conventional and highly data-dependent methodologies such as value at risk (VaR) or expected shortfall (CVaR). However, this article suggests that these losses are somewhat predictable and that, by formalizing the basic structural dynamics described previously, it is possible to develop an objective framework for analyzing the cost of illiquidity.

When the credibility barrier is breached and a manager is required to liquidate positions as outlined, we can model the associated cost as an up-and-in barrier option on the path of reported portfolio valuation.[15] When the path of reported value is overvalued and exceeds the threshold (barrier), the option pays (the payment is negative and represents a loss to the investor) the sum of two components, the amount by which the portfolio's value was overstated and the additional loss associated with the forced liquidation. If the investor has a set of prior beliefs about a) the portfolio's return and volatility characteristics, based on the observed mean, standard deviation, and serial correlation;[16] b) the conditions that will elicit a forced sale of the portfolio, i.e., a realistic estimate of the credibility threshold; and c) the liquidation cost of being forced to sell a relatively illiquid portfolio during a stressful market period (an estimate of the liquidation penalty), then pricing the option is relatively straightforward using Monte Carlo techniques.[17]

First, simulate the portfolio's true value using a discrete Brownian motion, a function of the observed volatility, and trend rate of return plus its associated estimated valuation lag, λ (which is based on 1 minus the observed first-order serial correlation, $1 - \rho_{s(1)}^{o}$).[18] Then calculate the individual period differences between the two processes and, when the difference exceeds the assumed credibility threshold, apply a payout equal to the difference between the two series at the time of the breach plus the assumed transaction penalty associated with a forced liquidation. Do this 100,000 times, calculating the net present value of whatever payout occurs for every one-year path and noting that many paths will have no associated penalty. Finally, calculate the mean of the resulting payoffs, including the zero-valued payoffs.

The option model is structured as a one-year option, so that the price translates as a haircut to the reported annualized rate of return associated with the investment. The price of the option is the de facto price of the risk assumed by investing in the less liquid portfolio, because the conservatively valued portfolio may become significantly overvalued and thereby force a sudden, expensive liquidation. To arrive at the portfolio's liquidity-adjusted expected return, we simply subtract the option's dollar value from the portfolio's expected return. That is, if the option is valued at $2 and the investment's expected return is 10%, the adjusted return would be 8%.

AN EXAMPLE

To illustrate this approach, our base case is a portfolio with mean expected return, μ, of 6%; volatility, σ, of 12%; with a riskless rate in the market of 2%, all on an annualized basis.[19] The one-year barrier option is valued on an initial $100 portfolio priced every week, or 52 times per year. When varying other parameters, we fix λ (the reporting adjustment) at 25%, *L* (the credibility threshold) at 15%, and *P* (the liquidation penalty) at 25%.[20] The evaluation of the option price is always based on 100,000 Monte Carlo simulations.[21]

Exhibit 3 illustrates how the cost of the de facto option varies as a function of the reporting adjustment λ and credibility threshold *L* when the liquidation penalty is held constant at 25%. As λ decreases, i.e., as serial correlation in the portfolio increases, the option value increases. Because we have used a one-year option on a $100 portfolio, we can directly interpret the option value as an annual percentage cost to the investor for the illiquidity in the portfolio. For example, with a λ of 25%, a credibility threshold of 15%, and a liquidation penalty of 25% (our base case, approximately left of center in Exhibit 3), the option has a value of $15.54. This means that the investor should adjust the expected return for an investment in the portfolio by −15.54%. Because we have assumed an expected return of 6% for this portfolio, the illiquidity option that the investor is providing to the manager consumes all of the expected investment return, leaving an undesirable −9.54% liquidity-adjusted expected return.[22] It is important to recognize that, for all credibility thresholds less than 20%, the option's value is significant over a range of reporting adjustments (serial correlations).

Turning our attention to the influence of the credibility threshold, Exhibit 3 shows that as *L* increases (as the manager is given more latitude to overstate performance), the cost of the illiquidity option decreases. This result is expected, because the likelihood that the credibility threshold will be breached goes down as the threshold increases. Significantly, no single investor controls this threshold. Although a given investor may be very lax—i.e., have a very high credibility threshold—it is the market's credibility threshold that matters.[23]

Exhibit 4 illustrates the relationship between the reporting adjustment and the liquidation penalty when the credibility threshold is constant at 15%. For higher reporting adjustment factors (lower serial correlations), the liquidity option's cost is low. But for lower reporting adjustment factors (higher serial correlations), the liquidity option's cost can increase dramatically. As might be expected, the option's cost monotonically increases in the size of the liquidation penalty, *P*. The cost function's steepness as the serial correlation increases underscores the cost of vanishing liquidity in a portfolio. As discussed previously, when a portfolio is under stress, investors can be left with fewer and fewer liquid positions, resulting in higher and higher serial correlation. Exhibit 4 shows that, in those cases, the illiquidity option's cost can easily overwhelm the portfolio's expected return.

Exhibit 5 shows the relationship between the liquidation penalty and the credibility threshold when the reporting adjustment is held constant at 25%. As previously discussed, when the credibility threshold is high and supervision is lax, the option's cost is low. Otherwise, as the liquidation penalty increases, the option's cost also increases. For sensible ranges of the credibility threshold (5% to 15%) and relatively high serial correlations (75%, for example), the entire range of liquidation penalties results in significant cost associated with the illiquidity option. Although the value of the illiquidity option diminishes with lower

EXHIBIT 3
Option Price as Function of λ and *L*

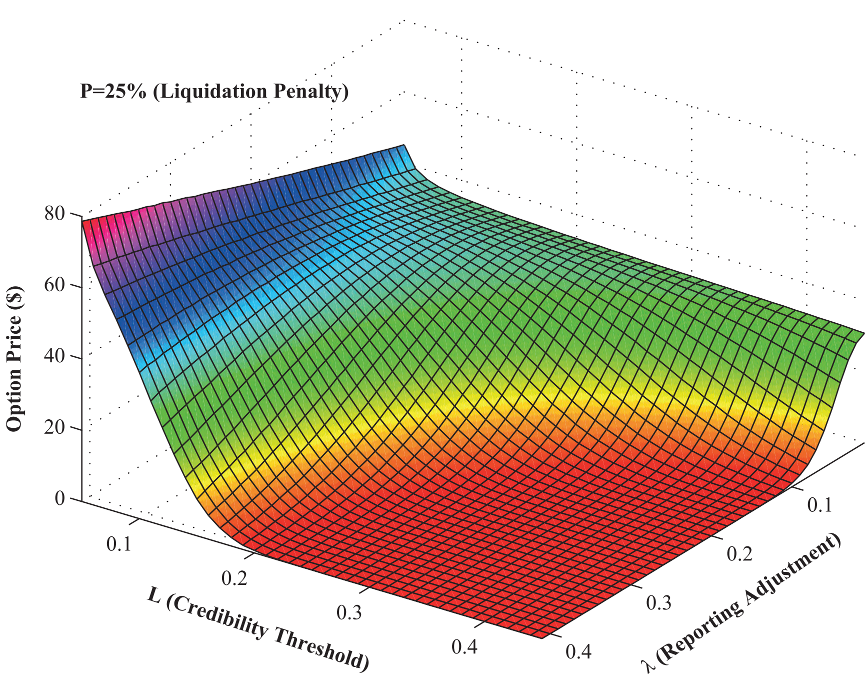

EXHIBIT 4
Option Price as Function of λ and *P*

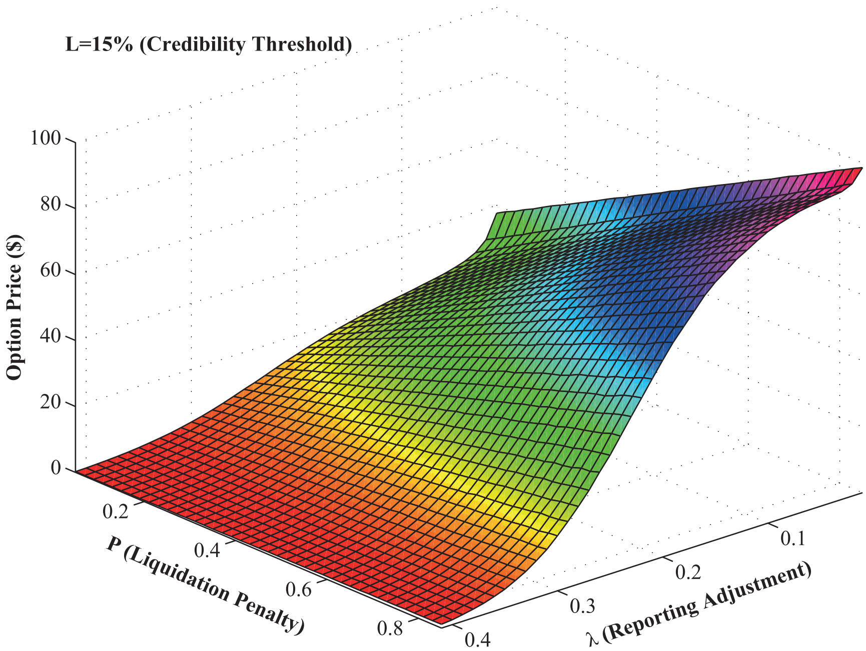

EXHIBIT 5
Option Price as Function of *P* and *L*

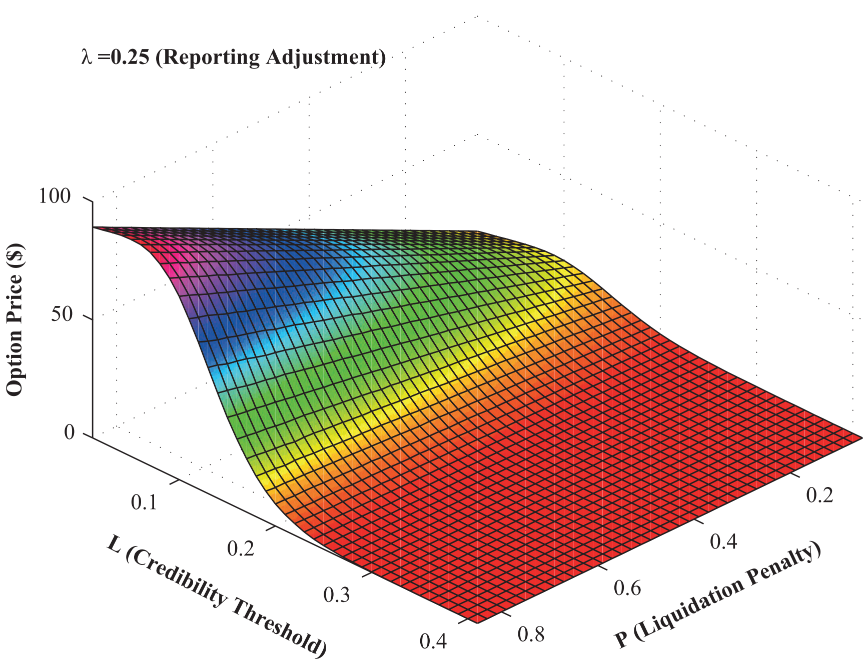

serial correlation, the associated cost can still be important enough to investors to warrant estimation and consideration.

It is also useful to consider the effect of the other parameters on the illiquidity option's cost. When the riskless interest rate increases, the option's cost goes down due to discounting—of course, the investment's relative value should also be subject to the same discounting. When the number of pricing periods (i.e., the frequency at which the portfolio value is priced or marked) goes up, the cost of the option also goes down, because the market can more quickly identify any overvaluation and therefore act earlier on any breach of the credibility threshold. For portfolios with higher volatility, the cost of the option increases because there is a greater likelihood of overvaluation, and for portfolios with higher expected returns, the cost of the option is lower because a greater expected return tends to offset the illiquidity option's cost.

For the liquidation penalty, we have been assuming a base case value of 25%. Research by Ramadorai [2008], who analyzed transactions in the secondary market for hedge funds, found that, for those transactions involving fraud or collapse, the average discount to reported NAV was 49.6%, almost twice our base case value. As discussed earlier, as the liquidation penalty increases, the illiquidity option's cost increases monotonically.

If we attempt to interpret the illiquidity option as the liquidity premium embedded in an investment, it appears to be much too high: approximately 15% in our base case. This is because the option approach used here does not price the liquidity premium investors usually think of, but rather it represents the cost associated with price-smoothing an illiquid investment. When combined with a triggering event, this results in an abrupt sale into a deteriorating market. The illiquidity option's size is a function of the magnitude of manager mispricing and the cost of liquidating at an unfavorable price. The greater the degree of leverage employed in the underlying investment, the larger the cost associated with any fire sale liquidation. Clearly, manager actions can mitigate the illiquidity option's value. For example, the manager could vary the reporting adjustment (λ) dynamically, rather than using the static approach modeled here.[24] This would allow the manager to control the degree of over- or undervaluation associated with the investment. In practice, we expect that managers would use such an approach.

A manager could also liquidate a portion of the portfolio as the credibility threshold is approached. However, to modify behavior as the barrier is approached, the manager would need a well-formed expectation about the level of the credibility threshold, which may be difficult to obtain because the barrier is not pre-established and is set by actors outside the manager's control. In addition, the incentives to cheat may increase as the barrier is approached. Many managers get into trouble when they decide to hide bad results in an attempt to limit the scope of investor withdrawals. They often hope that the market will turn and bail them out of the situation. But most managers lack the flexibility to alter their books. Attempting to institute a portfolio insurance strategy or sell off part of the book will exacerbate the situation, because the positions will now have an accurate mark. Managers and investors can generally live with a bit of variation around an illiquid investment's true value, but if a manager decides that the credibility threshold is close, trying to modify holdings or positions may actually trigger a fire sale. Nonetheless, because we are not modeling an environment where a manager is responding dynamically, we believe that the option value represents an upper bound on the cost of illiquidity.

Finally, turning to the assumptions associated with expected returns and volatility, J.P. Morgan Asset Management [2012] projects hedge fund expected returns in the range of 5% to 7% per year, with volatilities in the range of 7% to 13% per year—consistent with the assumptions of our base case. Importantly, J.P. Morgan estimates private equity's expected returns at about 9%, with a volatility of 34.25%. Because volatility serves only to increase the cost of the illiquidity option, only a high expected return (with reasonable levels for the other parameters) can serve to offset illiquidity's cost in the portfolio. We could question whether middle, single-digit returns are truly sufficient for investors to bear the cost of illiquidity that many hedge fund and private equity portfolios contain.

PRICING LIQUIDITY IN ALTERNATIVE INVESTMENTS

We are now in a position to consider several real-world applications of the model and the implications for investors in less liquid portfolios. We begin by considering five well-known hedge fund indices constructed by HFRI: a) fixed income—convertible arbitrage, b) distressed/restructuring, c) multi-strategy, d) fixed income—corporate, and e) emerging markets—Russia/Eastern Europe. Exhibit 6 contains estimates for the first-order serial correlation from each index's inception, along with the reporting adjustment implied in each case. The values shown represent index-level estimates, so it is reasonable to assume a potentially significant degree of variability in the underlying funds that constitute the indices.

As shown in Exhibit 6, mean index serial correlations for most of these strategies lie in the 50% to 60% range, implying an adjustment factor of less than 0.50, i.e., where managers, on average, reflect less than 50% of the true period-to-period change in portfolio value. Depending on an investor's assumptions concerning the liquidation penalty and the market credibility threshold, as well as estimates of expected return and volatility, the adjustment to observed returns (based on the cost of the illiquidity option) could be quite significant.

We would expect, however, for indices constructed with a large number of underlying funds, that at the index or composite level, various managers' over- and undervaluation would be diversified away[25] and that the liquidity option on an index should therefore be zero

EXHIBIT 6
First-Order Serial Correlation of Select HFRI Hedge Fund Indices

HFRI Index	Serial Correlation	"Estimated" λ
Fixed Income—Convertible Arbitrage	58%	0.42
Distressed/Restructuring	53%	0.47
Multi-Strategy	50%	0.50
Fixed Income—Corporate	48%	0.52
Emerging Markets—Russia/Eastern Europe	38%	0.62

or close to zero. In fact, for the first four indices in Exhibit 6, that is exactly what we find: the estimated option value is zero in each case. However, the Russia/Eastern Europe emerging-market index, which reflects the lowest serial correlation in Exhibit 6, has a monthly reported mean of 1.44% (17.30% annualized) and a monthly reported standard deviation of 7.69% (26.64% annualized). The liquidity option is priced at 13.52,[26] which reduces the seemingly large 17.30% return to only 3.78% per year and demonstrates that serial correlation alone is not sufficient to determine whether or not there is a cost associated with lack of liquidity.

The next application uses the Morningstar-CISDM hedge fund and CTA database, which contains data for both live and dead funds.[27] Eliminating CTAs and funds-of-funds, the initial sample contained data for 13,540 hedge funds. Selecting hedge funds with a serial correlation greater than 0.01 and for which there were at least 24 months of returns results in a final sample of 3,654 hedge funds. Using this final sample, we computed the value of the liquidity option using a mean, standard deviation, and serial correlation estimated from the return series of each fund, less the last three months (e.g., if there were 30 months of data, we used the first 27 months to estimate the parameters, to avoid including the period in which the fund might fail). In all cases, we used a risk-free rate of 5%,[28] a credibility threshold of 15%, and a liquidation penalty of 25%. The mean annualized return for this sample of funds was 11.79%, the mean annualized standard deviation was 13.88%, and the mean serial correlation was 0.2032. The mean option value for these 3,654 hedge funds was 5.52, implying an actual liquidity-adjusted mean annual return of 6.27%.[29]

To focus on funds for which the liquidity option represented a real reduction in return, we selected hedge funds with positive average returns and an option value greater than 1 (that is, where the option reduces the expected return by 1% or more). This resulted in a final sample of 1,031 funds. For this sample, the mean annualized return was 18.75%, the mean annualized standard deviation was 25.67%, and the mean serial correlation was 0.1959. The mean option value for these 1,031 hedge funds was 15.82, implying an actual liquidity-adjusted mean return of 2.93% on an annualized basis—a sharp contrast to the supposed 18.75%.[30] Narrowing the sample lets us further explore the liquidity option's application to real data.

Evaluating the relationship between the various attributes of the funds in the sample, we make the following observations:[31] a) as we saw with the HFRI indices, serial correlation alone is not a proxy for the liquidity option's value; b) serial correlation does not have a particularly strong relationship to maximum drawdown; c) a fund's mean return has a weak relationship to the liquidity option's value for that fund; d) a fund's mean return has no relationship to the maximum drawdown associated with that fund.

As we would expect from standard option-pricing theory, the liquidity option's value has an extremely strong relationship with fund volatility. A regression of option value on volatility has an adjusted R-squared of 0.8475 with a *t*-statistic of 75.66 for the regression coefficient. However, we find little relationship between the maximum drawdown experienced in the hedge fund sample and volatility, because the vast majority of drawdowns occur with volatilities (monthly standard deviations) of less than 12.5%. So while volatility of returns is an important determinant of the liquidity option's value, it appears to be unrelated to a fund's maximum drawdown. We find, in contrast, a reasonably strong relationship between the drawdown and the option value. A regression of maximum drawdown on option value has an adjusted R-square of 0.3353 with a *t*-statistic of 22.81 for the regression coefficient. Although not a perfect predictor, it is clear that higher drawdowns are generally associated with higher option values and that the computed option value can serve as a useful screening tool for investors.

It is also informative to look at individual hedge funds and the pricing of their associated options. Exhibit 7 shows the results for 16 hedge funds drawn from our final sample, ordered by annualized return. These results underscore the general results already described. For example, the two funds with highest serial correlation, Alta Partners and Bear Stearns Asset Backed Securities, have relatively low-valued liquidity options, whereas the liquidity options for Bridgewater Partners and Okumus Opportunity, which have relatively low serial correlation, represent a significant reduction in the funds' expected returns.

As a final example, we consider the Bear Stearns High-Grade Structured Credit Strategies fund. Its marketing materials state that the "fund seeks to generate total annual returns through 'cash and carry' transactions and capital markets arbitrage."[32] It further states that "the Fund generally invests in high-quality floating-rate-structured finance securities." The High-Grade Struc-

Exhibit 7

Select Hedge Fund Returns and Liquidity-Adjusted Returns

Fund Name	Annualized Return	Serial Correlation	Option Value	Adjusted Annual Return
Deephaven Credit Opportunities	1.88%	0.42	$1.32	0.56%
Bridgewater Partners	3.58%	0.20	$8.21	–4.62%
Thames River European A (EUR)	6.76%	0.03	$1.28	5.47%
Thames River Property Growth & Income (EUR)	7.24%	0.43	$1.64	5.60%
Glenrock Global Partners (BVI), Inc.	7.42%	0.11	$1.01	6.40%
Everest Capital Intl.	8.31%	0.27	$9.34	–1.03%
Rocker Partners	9.42%	0.09	$7.46	1.96%
Alta Partners L.P. (Onshore)	11.13%	0.64	$1.67	9.46%
Rainbow Global High Yield (USD)	16.03%	0.28	$14.31	1.71%
Marathon Emerging Markets	16.20%	0.23	$20.05	–3.85%
Bear Stearns Asset Backed Securities	24.86%	0.60	$2.24	22.61%
Okumus Opportunity A	27.53%	0.04	$36.48	–8.95%
Viaticus	31.22%	0.05	$33.03	–1.81%
Lancer Offshore	33.63%	0.16	$9.68	23.95%
Galleon Omni Technology (B)	40.25%	0.13	$45.76	–5.50%
Infinity Emerging Opportunities	74.16%	0.12	$88.80	–14.64%

tured Credit Strategies fund is a poster child for a fire sale and rapid liquidation. As subprime mortgage delinquencies grew, the value of CDOs the fund held dropped. The fund's prime brokers asked for more cash collateral. The fund attempted to meet those collateral calls by liquidating assets in a rapidly deteriorating market, but values fell quickly and collateral requirements rose rapidly, leading to eventual collapse. The fund failed in spite of an attempt to stabilize fund value through Bear Stearns's substantial injection of capital.

Before the fund collapsed, the annualized mean return was 12.40% with an annualized standard deviation of 1.50% and a serial correlation of 0.6365. Straightforwardly pricing the liquidity option produces a value of zero,[33] seemingly at odds with the approach advocated here. But notice that the reported standard deviation is extremely low, much lower for example than the standard deviation for the comparable HFRI index (the Fixed-Income Asset-Backed Index), which is 4.03%.

Exhibit 8 plots the value of the liquidity option for the Bear Stearns fund when the fund volatility is assumed to be a varying percentage of the volatility of the HFRI Fixed-Income Asset-Backed Index. The Bear Stearns fund's reported volatility is approximately 37% of the index's volatility, which results in an option value essentially equal to zero. But as the fund's assumed volatility approaches that of the index, the value of the liquidity

Exhibit 8

Option Price as Function of Index Volatility

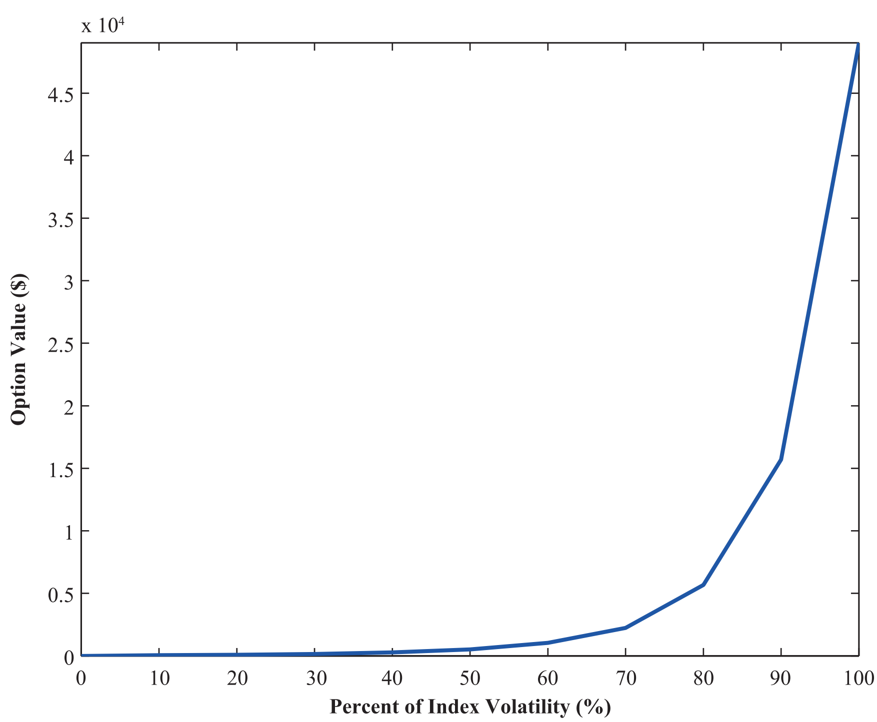

option explodes exponentially. (Note that the option value on the vertical axis is in tens of thousands). This example underscores the importance of appropriate estimates for the parameters of the liquidity option model. In the case of the Bear Stearns fund, the reported volatility was significantly lower than what the market should have expected given the fund's assets. For this reason, it is important to have a sanity check on a fund's reported values when computing the value of the liquidity option—artificially low returns volatility or artificially high expected returns can make the analysis superfluous.

CONCLUSION

Illiquidity, even in its milder forms, is not a benign condition that results merely in lower reported portfolio volatility. Underestimating the cost of illiquidity during the portfolio construction process can produce significant distortions: understated volatility, understated correlations, overstated expected returns, and an overallocation to less liquid securities. One need only consider J.P. Morgan's predicament as it was forcibly extricated from its overvalued London Whale trades during an abbreviated, career-defining, one-month period, or for the investors in quant hedge funds in August 2008 to appreciate how large these costs can be. To adjust properly for illiquidity's embedded risks, we propose a straightforward barrier option-pricing model that provides an objective framework for incorporating market data and investor experience and thereby lets institutions evaluate the true cost of illiquidity in their investments.

It is also important for investors to recognize that an investment's liquidity is not necessarily constant through time. When a portfolio comes under funding stress, the easiest and quickest way to alleviate that stress is to sell the most liquid assets in the portfolio. But as liquid assets are sold, the portfolio becomes even more illiquid, generally over short time horizons—driving up the cost of the embedded illiquidity option, sometimes rapidly.

As this article has shown, illiquidity in the market has discernible costs. The barrier option-pricing methodology described here can provide both a screen and a rigorous adjustment mechanism for evaluating investments that appear attractive, but may actually be unduly risky. Thus, in the case of the Bear Stearns fund previously discussed, a modest upward adjustment to the fund's reported volatility revealed an embedded option cost that dwarfed the reported returns, because of the fund's highly serially correlated returns. This article's rigorous framework for assessing liquidity gives an investor the means to make an annualized adjustment to reported returns that corrects for illiquidity.

ENDNOTES

[1]The authors thank Yakov Amihud, Mark Anson, Ashwin Alankar, Mark Baumgartner, Peter Carr, Mark Kamstra, Mark Kritzman, Masao Matsuda, Fabio Savoldelli, Myron Scholes, the participants of the 2014 Columbia MAFM Practitioner's Seminar, and participants of the Fall 2015 Q-group conference the participants of the January 2015 IAQF-Thalesian seminar series for useful comments and discussion. Errors remain our sole responsibility.

[2]Sparrow and Ilijanic [2010] quantified the value of liquidity in a trading context, and Amihud, Mendelson, and Pedersen [2005] provided a literature survey of the theoretical ways liquidity effects asset price.

[3]It went up (down) last period, so odds are it will go up (down) again this period—also referred to as autocorrelation. See Scholes and Williams [1977], for example.

[4]Examples of short-term leverage may include margin borrowing (see Garleanu and Pedersen [2009] or Brunermeier and Pedersen [2009], for example) or the use of futures, options, or swaps (see e.g., Office of Financial Research [2013]).

[5]Transactions that contain contractual obligations requiring liquidity upon demand, such as securities lending or repo transactions (see e.g., Keane [2013] or Office of Financial Research [2013]).

[6]Hill [2009] argued that options with long volatility exposures hedge some liquidity risk; Bhaduri, Meissner, and Youn [2007] applied this concept in a number of hedging strategies; and Golts and Kritzman [2010] suggested that investors can protect for unscheduled capital calls by purchasing liquidity options that pay the investor in those parts of the world when such calls are likely.

[7]This is often the reason for low measured correlation and volatility.

[8]Acerbi and Scandolo [2008] formalized this observation.

[9]See, e.g., Weisman [2002] or Getmansky, Lo, and Makarov. [2004].

[10]Nor does it preclude such behavior.

[11]"True" should be thought of in a generic sense as being some prudent central value for the price that would be obtained during a given period if one wished to buy or sell the illiquid asset or portfolio.

[12]This article uses the terms "asset" and "portfolio" interchangeably—one either has a single illiquid asset (e.g., an investment in private equity or a hedge fund) or a port-

folio that contains illiquid assets (e.g., the even more illiquid portion of a hedge fund).

[13]The expectation is obtained by applying Ito's lemma, where $m = \mu - \sigma^2/2$ and μ is the arithmetic mean return.

[14]This relationship is explored in Online Appendix B.

[15]We are not the first researchers to apply option theory to the issue of liquidity. Chaffe [1993] used Black–Scholes option pricing to value illiquidity in private company valuations. Longstaff [1995] used risk-neutral valuation to develop an upper bound on the value of marketability, and Golts and Kritzman [2010] used a cliquet option to hedge unexpected capital calls.

[16]Given that these are derived from returns reported by the manager, it is important to also use additional, independent sources to validate these parameters, or the model may produce misleading results. See the discussion regarding the Bear Stearns fund in the section "Pricing Liquidity in Alternative Investments."

[17]Technically, the investor gives the option to the manager of the illiquid portfolio, because the manager generally benefits from holding illiquid assets through an asymmetric compensation structure (i.e., the manager shares gains on the upside but does not share investors' losses on the downside). See Keane [2013] and Huang et al. [2011].

[18]To generate paths for the portfolio's true value, it is necessary to adjust the observed volatility for the serial correlation. We use the methodology developed by Geltner [1993]:

$$\sigma = \sqrt{\left(\frac{1-\rho_{s(1)}^{o2}}{\left(1-\rho_{s(1)}^{o}\right)^2}\right)\sigma_o^2}$$

where σ is the adjusted volatility, σ_o is the observed volatility, and $\rho_{s(1)}^{o}$ is the observed serial correlation.

[19]These are approximate market values at the time of writing.

[20]We discuss the reasonableness of these values later in this section.

[21]Matlab code to price this option is included in Online Appendix A.

[22]This has not been adjusted for the cost associated with notification periods, lockups, or gates as described in the section "Causes of Illiquidity."

[23]More precisely, it is the lowest credibility threshold of any investor or third party who has the ability to trigger a fire sale.

[24]This type of dynamic or state-dependent choice could theoretically be embedded in a general equilibrium setting.

[25]Geltner [1993] demonstrated a similar result in the case of real estate appraisals.

[26]Based on the estimated mean, standard deviation, and serial correlation, a riskless rate of 2.00%, a credibility threshold of 15%, and a liquidation penalty of 25%.

[27]Using both live and dead funds minimizes survivorship bias in the analysis.

[28]This is approximately the average risk-free rate for the period.

[29]1,357 of the funds had a zero option value.

[30]Of course, there is a purposeful selection bias here.

[31]Exhibits that illustrate each of these observations are found in the online material accompanying this article.

[32]There were two Bear Stearns funds following similar strategies. We use only one as an example.

[33]Again, assuming a credibility threshold of 15%, a liquidation penalty of 25%, and a risk-free rate of 5.0%.

REFERENCES

Anson, M. "Measuring a Premium for Liquidity Risk." *The Journal of Private Equity*, Spring 2010, pp. 6-16.

——. "Risk Management and Risk Budgeting in Real Estate and Other Illiquid Asset Classes." *CFA Institute*, March 2013, pp. 1-11.

Acerbi, C., and G. Scandolo. "Liquidity Risk Theory and Coherent Measures of Risk." *Quantitative Finance*, Vol. 8, No. 7 (October 2008), pp. 681-692.

Amihud, Y., H. Mendelson, and L. Pedersen. "Liquidity and Asset Prices." *Foundations and Trends in Finance*, Vol. 1, No. 4 (2005), pp. 269-364.

Ang, A., D. Papanikolaou, and M. Westerfield. "Portfolio Choice with Illiquid Assets." Working paper, 2011, SSRN 1697784.

Ang, A., and N. Bollen. "Locked Up by a Lockup: Valuing Liquidity as a Real Option." *Financial Management*, No. 39 (2010), pp. 1069-1095.

Bhaduri, R., G. Meissner, and J. Youn. "Hedging Liquidity Risk." *The Journal of Alternative Investments*, Vol. 10, No. 3 (Winter 2007), pp. 80-90.

Bhattacharya, U., J. Lee, and V. Pool. "Conflicting Family Values in Mutual Fund Families." *Journal of Finance*, Vol. 68, No. 1 (2013), pp. 173-200.

Bollen, N., and V. Pool. "Conditional Return Smoothing in the Hedge Fund Industry." *Journal of Financial and Quantitative Analysis*, Vol. 43 (June 2008), p. 267-298.

Boyson, N., C. Stahel, and R. Stulz. "Hedge Fund Contagion and Liquidity Shocks." *Journal of Finance*, Vol. 65, No. 5 (2010), pp. 1789-1816.

Brunnermeier, M., and L. Pedersen. "Market Liquidity and Funding Liquidity." *Review of Financial Studies*, Vol. 22, No. 6 (2009), pp. 2201-2238.

Chen, Q., I. Goldstein, and W. Jiang. "Payoff Complementarities and Financial Fragility: Evidence from Mutual Fund Outflows." *Journal of Financial Economics*, Vol. 97 (2010), pp. 239-262.

Garleanu, N., and L. Pedersen. "Margin-Based Asset Pricing and Deviations from the Law of One Price." Working paper, New York University, 2009.

Gennaioli, N., A. Shleifer, and R. Vishny. "Neglected Risks, Financial Innovation, and Financial Fragility." *Journal of Financial Economics*, Vol. 104, No. 3 (2012), pp. 452-468.

Getmansky, M., A. Lo, and I. Makarov. "An Econometric Model of Serial Correlation and Illiquidity in Hedge Fund Returns." *Journal of Financial Economics*, Vol. 74, No. 3 (December 2004), pp. 529-609.

Golts, M., and M. Kritzman. "Liquidity Options." *The Journal of Derivatives*, Fall 2010, pp. 80-89.

Hill, J. "A Perspective on Liquidity Risk and Horizon Uncertainty." *The Journal of Portfolio Management*, Vol. 35, No. 4 (Summer 2009), pp. 60-68.

J.P. Morgan Asset Management. *Long-Term Capital Market Return Assumptions*. 2012.

Keane, F. "Securities Loans Collateralized by Cash: Reinvestment Risk, Run Risk, and Incentive Issues." *Current Issues in Economics and Finance*, Federal Reserve Bank of New York, Vol. 19, No. 3 (2013).

Kinlaw, W., M. Kritzman, and D. Turkington. "Liquidity and Portfolio Choice: A Unified Approach." *The Journal of Portfolio Management*, Vol. 39, No. 2 (2013), pp. 19-27.

Leibowitz, L., and A. Bova. "Portfolio Liquidity." Morgan Stanley Research, 2009.

Lo, A., C. Petrov, and M. Wierzbicki. "It's 11 pm—Do You Know Where Your Liquidity Is? The Mean-Variance Liquidity Frontier." *Journal of Investment Management*, Vol. 1 (2003), pp. 55-93.

Longstaff, F. "How Much Can Marketability Affect Security Values?" *Journal of Finance*, Vol. 50 (1995), pp. 1767-1774.

Manconi, A., M. Massa, and A. Yasuda. "The Role of Institutional Investors in Propagating the Crisis of 2007-2008." *Journal of Financial Economics*, Vol. 104, No. 3 (2012), pp. 491-518.

Mitchell, M., L. Pedersen, and T. Pulvino. "Slow Moving Capital." *American Economic Review*, Vol. 97, No. 2 (2007), pp. 215-220.

Office of Financial Research. "Asset Management and Financial Stability." U.S. Department of Treasury, September 2013.

Quan, D., and J. Quigley. "Price Formation and the Appraisal Function in Real Estate Markets." *Journal of Real Estate Finance and Economics*, Vol. 4, No. 2 (1991), pp. 127-146.

Ramadorai, T. "The Secondary Market for Hedge Funds and the Closed-Hedge Fund Premium." Working paper, University of Oxford, 2008.

Schmidt, L., A. Timmermann, and R. Wermers. "Runs on Money Market Mutual Funds." Working paper, 2013, SSRN-2024282.

Scholes, M., and J. Williams. "Estimating Betas from Nonsynchronous Data." *Journal of Financial Economics*, Vol. 5, No. 3 (December 1977), pp. 309-327.

Siegel, J. "Alternatives and Liquidity: Will Spending and Capital Calls Eat Your Modern Portfolio?" *The Journal of Portfolio Management*, Fall 2008, pp. 103-114.

Sparrow, C., and D. Ilijanic. "The Value of Liquidity." *The Journal of Trading*, Winter 2010, pp. 10-15.

Weisman, A. "Informationless Investing and Hedge Fund Performance Measurement Bias." *The Journal of Portfolio Management*, Summer 2002, pp. 80-91.