

Modeling Liquidity Risk

With Implications for Traditional Market Risk Measurement and Management

Anil Bangia

Oliver, Wyman & Company

Francis X. Diebold

Stern School, NYU
University of Pennsylvania
and Oliver Wyman Institute

Til Schuermann

Oliver, Wyman & Company

John D. Stroughair

Oliver, Wyman & Company

This draft/print: April 26, 1999

Abstract

Market risk management has traditionally focussed on the distribution of portfolio value changes produced by changes in the midpoint of bid and ask prices. Hence market risk is traditionally assessed under the assumption of an idealized market with a negligible bid-ask spread. In reality, however, spreads can be both wide and variable; hence a superior approach would recognize that positions will *not* be liquidated at the mid-price, but rather at the mid-price less the uncertain bid-ask spread. Liquidity risk associated with the uncertainty of the spread, particularly for thinly traded or emerging market securities under adverse market conditions, is an important part of overall market risk and is therefore important to model. We do so, proposing a simple liquidity risk methodology that can be easily and seamlessly integrated into standard value-at-risk models. We show that ignoring the liquidity effect can produce underestimates of market risk in emerging markets by as much as thirty percent. Furthermore, we show that because the BIS is already implicitly monitoring liquidity risk, banks that fail to model liquidity risk explicitly and capitalize against it will likely experience surprisingly many violations of capital requirements, particularly if their portfolios are concentrated in emerging markets.

We thank Steve Cecchetti, Edward Smith and two anonymous referees for helpful comments and suggestions. All remaining errors are ours.

“Portfolios are usually marked to market at the middle of the bid-offer spread, and many hedge funds used models that incorporated this assumption. In late August, there was only one realistic value for the portfolio: the bid price. Amid such massive sell-offs, only the first seller obtains a reasonable price for its security; the rest lose a fortune by having to pay a liquidity premium if they want a sale. ...Models should be revised to include bid-offer behavior.”

Nicholas Dunbar (“Meriwether’s Meltdown,” *Risk*, October 1998, 32-36)

I. INTRODUCTION

Experts and laymen alike often point to liquidity risk as the culprit responsible for the recent turmoil in world capital markets. Both sophisticated and inexperienced players were surprised when markets evaporated and spreads widened drastically. More generally, it is widely acknowledged that the standard Value-at-Risk (VaR) concept used for measuring market and credit risk for tradable securities suffers from failure to account for liquidity risk. At best, VaR for large illiquid positions is adjusted upwards in an *ad hoc* fashion by using a longer time horizon that reflects a subjective estimate of liquidation time.

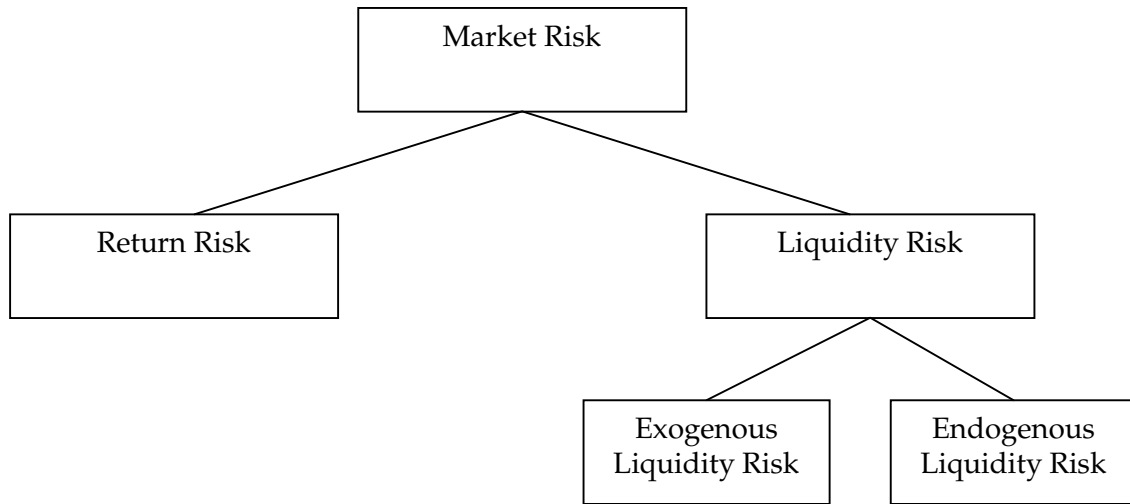
Motivated by the defects of VaR as traditionally implemented, we present in this article a framework for incorporating liquidity risk in market risk measurement and management. In section 2, we present our conceptual framework for understanding market risk and its components, one of which is liquidity risk. We make an important distinction between *exogenous* liquidity risk, which is outside the control of the market maker or trader, and *endogenous* liquidity risk, which is in the trader’s control and usually the result of sudden unloading of large positions that the market cannot absorb easily. In section 3, we describe the various components of overall market risk and techniques for their measurement, with emphasis on the neglected liquidity risk component and our approach to modeling it. Our liquidity risk adjustment uses readily available data and is easily integrated with traditional VaR calculations, and we provide several worked examples. In section 4, we broaden the analysis from one instrument to an entire portfolio, and we display backtesting examples that reveal the very different results that can be obtained depending on whether liquidity risk is or is not incorporated. We conclude in section 5 with additional discussion of selected issues.

II. CONCEPTUAL FRAMEWORK

“Risk” refers to uncertainty about future outcomes. Traditional market risk management deals almost exclusively with portfolio value changes driven by trading returns. Trading returns are calculated from mid-price, and hence the assessed market risk corresponds to an idealized market with no friction in obtaining the fair price. However, risk in many markets possesses an additional liquidity component: traders do *not* realize the mid-price when liquidating a position quickly, or when the market is moving against them. Instead, they realize the mid-price less the bid-ask spread. Marking to market mid-price therefore yields an underestimation of the true risk in such markets, because the realized value upon liquidation can deviate significantly from the market mid-price. If one wants an accurate assessment of overall market risk, it is important to model the distribution of the deviation of the liquidation price from the mid-price, which is sometimes called the cost of liquidity.

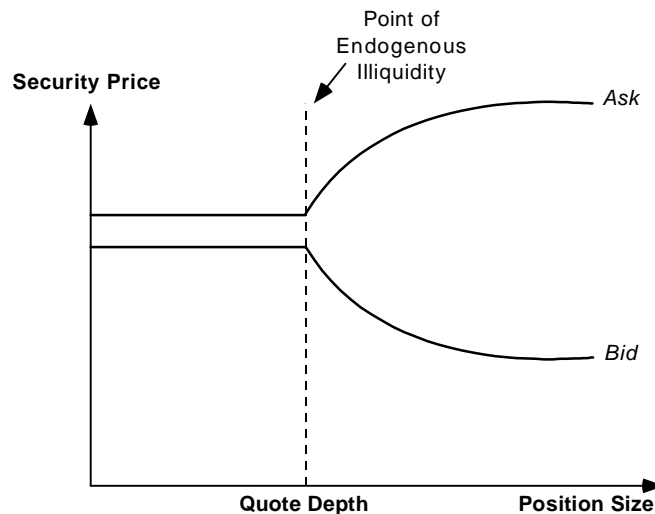
More precisely, we conceptually split market risk into two parts: return risk, which can be thought of as a pure market risk component, and liquidity risk. We further split liquidity risk into endogenous and exogenous components, as summarized in Figure 1.

Figure 1: Taxonomy of Market Risk



Exogenous illiquidity is the result of market characteristics; it is common to all market players and unaffected by the actions of any one participant. Markets with very low and stable levels of exogenous illiquidity, such as G7 currencies, are typically characterized by heavy trading volumes, stable and small bid-ask spreads, and stable and high levels of quote depth. *Endogenous* illiquidity, in contrast, is specific to one’s position in the market, depends on one’s actions, and varies across market participants. It is mainly driven by position size: the larger the size, the greater the endogenous illiquidity. A good way to understand the implications of the position size is to consider the relationship between the liquidation price and the total position size held, which we depict qualitatively in Figure 2.

Figure 2: Effect of Position Size on Liquidation Value



If the size of the order is smaller than the quote depth, the order transacts at the quote, and the cost of immediate execution is half the bid-ask spread. Such a position possesses only exogenous liquidity risk. If the size of the order exceeds the quote depth, the cost of immediate execution is greater than half the spread. The excess over the half-spread reflects endogenous liquidity risk, which can be particularly important in situations when normally fungible market positions cease to be fungible. Quantitative methods for modeling endogenous liquidity risk have been proposed recently by Jarrow and Subramanian (1997), Chriss and Almgren (1997), Bertsimas and Lo (1998), Longstaff (1998) and Kravovskiy(1999),

among others. Such methods, however, typically rely on models whose key parameters (e.g., relevant elasticities) are unknown and extremely difficult to gauge, due to a lack of available data.

In this paper we approach the liquidity risk problem from the other side, focusing on methods for quantifying *exogenous* rather than endogenous liquidity risk. Our approach is motivated by two key facts. First, fluctuations in exogenous liquidity risk are often large and important, as will become clear from our empirical examples, and they are relevant for all market players, whether large or small. Second, in sharp contrast to the situation for endogenous liquidity risk, the data needed to quantify exogenous liquidity risk are widely available. The upshot is that we can incorporate exogenous liquidity risk into VaR calculations in a simple and powerful way.

III. MODELING UNCERTAINTY IN MARKET VALUE

A. Incorporating Exogenous Liquidity Risk in VaR Calculations

Asset returns are the log difference of mid-prices:

$$r_t = \ln[P_t] - \ln[P_{t-1}] = \ln \left[\frac{P_t}{P_{t-1}} \right],$$

Taking a one-day horizon over which the change in asset value is considered, and assuming that one-day returns are Gaussian, the Value-at-Risk at 99% confidence level is

$$VaR_t = P_t \left[\Phi^{-1} \left(e^{-b_{E_t[r_t]} - 2.33 \sigma_t} \right) \right],$$

where $E_t[r_t]$ and σ_t^2 are the first two conditional moments of returns. With no loss of generality, we set $E_t[r_t]$ to zero.

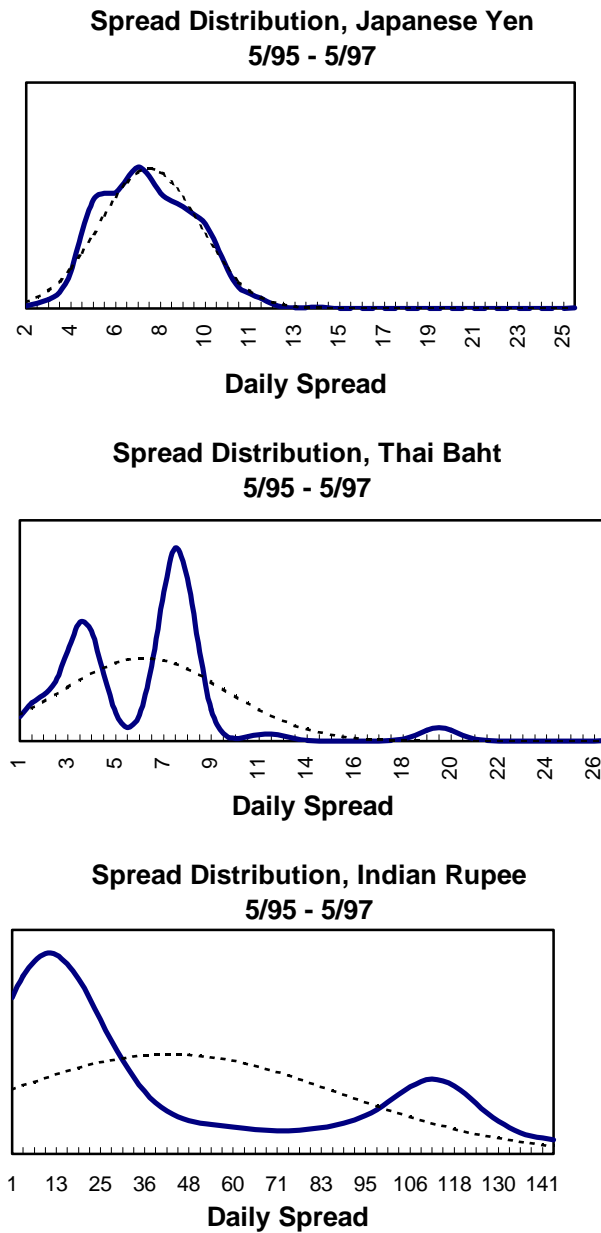
The above expressions for VaR consider only the distribution of the mid-price, whereas on average we would expect the bid to be less than the mid-price by half the average spread, $\frac{1}{2} \bar{S}$. Moreover, risk managers are interested not in average spreads but rather in extreme spreads associated with extreme market conditions. We define the (exogenous) cost of liquidity based on a certain average spread, \bar{S} , plus a multiple of the spread volatility, $a \cdot \tilde{\sigma}$, to cover most (say 99%) of the spread situations:

$$COL_t = \frac{1}{2} \left[P_t \bar{S} + a \tilde{\sigma} \right],$$

where P_t is today's mid-price for the asset or instrument, \bar{S} is the average *relative spread* (where relative spread, a normalizing device which allows for easy comparison across different instruments, is defined as ([Ask-Bid]/Mid)), $\tilde{\sigma}$ is the volatility of relative spread, and a is a scaling factor that produces 99% probability coverage.


Because spread distributions are typically far from normal, we cannot rely on Gaussian distribution theory to help determine the scaling factor a . Empirical explorations indicate that a ranges from 2.0 to 4.5 depending on the instrument and market in question. For illustration, we show in Figure 3 the densities of relative spreads (in basis points) for three currency returns, with the best-fitting normal superimposed. As we move to anecdotally less liquid markets, the densities become less normal. Moreover, spread distributions are not nearly as well behaved as return distributions; sometimes, for example, they appear multi-modal. We seem to be seeing different regimes, some of high liquidity, others of low liquidity.

Figure 3: Distribution of Spreads for Three Currencies



In order to treat the return risk and liquidity risk jointly, we make the conservative simplifying assumption that extreme return events and extreme spread events happen concurrently. The correlation between mid-price movements and spreads is not perfect, but it is nevertheless strong enough during extreme market conditions to enable and encourage us to view return risk and exogenous liquidity risk as experiencing extreme movements simultaneously. Hence when calculating liquidity-risk adjusted VaR we incorporate both a 99% mid-price return on the underlying and a 99% spread, as we summarize graphically in Figure 4.

Figure 4: Combining Return Risk and Liquidity Risk



To translate from returns back to prices, we simply define the 1% worst-case price (P') and parametric liquidity-adjusted VaR for a single asset as

$$P'_t = P_t e^{b \cdot 2.33 \sigma_t} \left[-\frac{1}{2} [P_t \bar{C} + a \tilde{\sigma} \eta] \right]$$

$$L-VaR_t = VaR_t + COL_t = P_t e^{b \cdot 2.33 \sigma_t} \left[-\frac{1}{2} [P_t \bar{C} + a \tilde{\sigma} \eta] \right].$$

B. A Detailed Example

We examine two exchange rates against the U.S. dollar, one highly liquid (the Japanese Yen, JPY) and the other much less so (the Thai Baht, THB). We first examine basic asset-return risk, and then we consider the exogenous liquidity component. We split the sample period into pre and post May 1997, which corresponds to the onset of the Asian crisis.

For the pre-crisis period we calculate the risk measures as of May 2, 1997. The U.S. dollar exchange rates for the Yen and Baht were 126.74 and 26.11, and the volatilities were 1.12%, and 0.19%. Hence we calculate the 99% worst case prices as:

$$P_{JPY}^* = ¥126.74 \cdot e^{-2.33 \cdot 1.12\%} = ¥123.47$$

$$P_{THB}^* = B26.11 \cdot e^{-2.33 \cdot 0.19\%} = B25.99.$$

In other words, we calculate that on May 2, 1997, the probability of a one-day movement in the Yen from ¥126.74 to ¥123.47 is 1%. Next, we incorporate exogenous liquidity risk by adding the 99% bid values, computed using our estimates of spread means, volatilities, and scaling factors, yielding

$$Bid_{JPY} = ¥123.47 - \frac{1}{2} [¥126.74 \cdot (6.6bp + 2.5 \cdot 1.7bp)] = ¥123.40$$

$$Bid_{THB} = B25.99 - \frac{1}{2} [B26.11 \cdot (6.3bp + 3.5 \cdot 4.1bp)] = B25.96.$$

In Table 1a we summarize the progression of market risk calculations, which account for both the return risk and exogenous liquidity risk components of total market risk. The difference in liquidity risk contribution is strikingly different for the two currencies: the marginal impact of the liquidity component is only 2% for the Yen, but 19% for the Baht. This makes good sense, because the JPY/USD spot market is arguably one of the most liquid in the world, whereas the THB/USD market most definitely is not. Hence, at least during the period at hand, one could have safely ignored Yen liquidity risk, whereas ignoring Baht liquidity risk would have resulted in substantial underestimation of overall market risk.

Table 1a: FX Return Risk and Liquidity Risk, Pre-Crisis

	JPY	THB
Price on May 2, 1997 (P_t)	¥126.74	B26.11
Average return volatility (σ_t)	1.12%	0.19%
Mean and volatility of relative spread $\bar{C}, \tilde{\sigma}_h$	6.6 bp, 1.7 bp	6.3 bp, 4.1 bp
Return risk component of VaR, $P_t \theta - e^{b - 2.33 \theta \sigma_t} g_j$	¥3.26	B0.12
Liquidity risk component of VaR, $\frac{1}{2} [P_t \bar{C} + a \tilde{\sigma}_h]$	¥0.07	B0.03
Total liquidity adjusted VaR	¥3.33	B0.14
Liquidity component %	2.1%	19%

In Table 1b we present the analogous calculations for the post-crisis period. Comparison of the results for the two periods reveals that the proportion of exogenous liquidity risk in total market risk dropped from 19% to 9%, as the floating of the Baht and subsequent heavy trading increased return risk substantially. Thus, although both return and liquidity risk increased with the crisis, the share of return risk increased.

Table 1b: FX Return Risk and Liquidity Risk, Post-Crisis

	JPY	THB
Price on January 22, 1998 (P_t)	¥127.17	B53.55
Average return volatility (σ_t)	2.0%	5.48%
Mean and volatility of relative spread $\bar{C}, \tilde{\sigma}_h$	7.1 bp, 2.7 bp	76.4 bp, 47.4 bp
Return risk component of VaR	¥5.79	B6.42
Liquidity risk component of VaR	¥0.09	B0.65
Total liquidity adjusted VaR	¥5.88	B7.07
Liquidity component %	1.5%	9%

IV. ADJUSTING PORTFOLIO VaR FOR EXOGENOUS LIQUIDITY RISK

We have discussed return risk and exogenous liquidity risk, and their combination, at the level of a single financial instrument. Now we move to the level of an entire portfolio. In traditional portfolio VaR analysis, which ignores liquidity risk, the covariance matrix of asset returns is the key bridge from single-instrument to portfolio risk. A liquidity-adjusted portfolio VaR analysis could proceed in similar fashion, but it would require an assumption of multivariate normality of the spreads and estimation of the spread covariance matrix. Unfortunately, as we have seen, spread distributions are not nearly as well behaved as return distributions, which renders the approach of dubious value.

An alternative and more credible approach is to compute a portfolio-level bid and ask series by simply taking a weighted average of the individual instrument bids and asks, and then to use the instrument-level L-VaR methodology to adjust standard portfolio VaR for exogenous liquidity risk. We will now illustrate this approach using four portfolios, which range from Asian telecom equities to G7 currencies, from January 1995 to March 1997, as detailed in Table 2.

Table 2: Compositions of Four Equally-Weighted Portfolios

S.E. Asian Telecoms	Emerging Mkts. Mixed	US Equities / Bonds	G7 Currencies
Singapore Telecom	Indonesian rupiah	IBM 5yr issue (6.375%)	Japanese yen
Telecom Asia (Thai)	Polish zloty	IBM 10yr issue (7.25%)	British pound
Telekom Malaysia	Telecom Asia (Thai)	AT&T	French Franc
Singapore dollar	Genting (Malaysia)	DuPont	Deutschemark
Thai baht	City Devs (Singapore)	GE	
Malaysian ringgit			

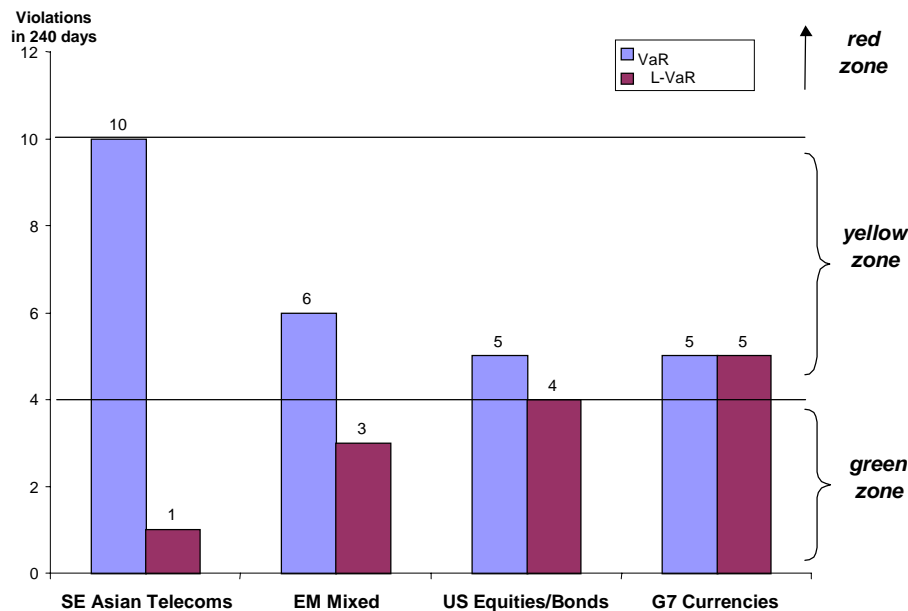
Recall that BIS assigns regulatory capital according to the number of VaR violations experienced over the course of a trading year, as illustrated in Table 3. Based on violations, banks are assigned regulatory colors with corresponding capital multipliers. From a bank's perspective, both over- and underestimation of VaR results in misallocation of capital. Overestimation produces too few violations, which indicates inefficient capital use, and underestimation produces too many violations, which result in capital charges.

Table 3: BIS Capital Multiplier

# OF VaR VIOLATIONS IN A TRADING YEAR	PR (VIOLATION) IF TRUE 99% VaR	BIS ZONES	CAPITAL MULTIPLIER
0	8.1%	Green Zone	3.0
1	20.5%		3.0
2	25.7%		3.0
3	21.5%		3.0
4	13.4%		3.0
5	6.7%	Yellow Zone	3.4
6	2.7%		3.5
7	1.0%		3.65
8	0.3%		3.75
9	0.08%		3.85
10 or more	0.01%	Red Zone	4.0

In Figure 5 we display VaR backtesting results for the four portfolios considered. We show the number of violations during a trading year for two models: a standard VaR model, and the same model with our adjustment for exogenous liquidity risk. The liquidity risk adjustment never worsens performance, and it typically produces substantial improvement. As one moves from right to left in Figure 5, the liquidity of the underlying assets decreases, and the risk assessment gains from incorporating a liquidity adjustment increase dramatically.

Figure 5: VaR Violations with and without Liquidity Risk Adjustment



V. CONCLUDING REMARKS AND DIRECTIONS FOR FUTURE RESEARCH

Traditional VaR measures are obtained from the distribution of portfolio returns computed at mid-market prices, which underestimates risk by neglecting the fact that liquidation occurs not at the mid-market price, but rather at mid-market less half the spread, and the spread may fluctuate widely. Hence we developed and illustrated a simple measure of exogenous liquidity risk, computed using the distribution of observed bid-ask spreads. Our results suggest that ignoring exogenous liquidity risk can produce substantial underestimates of overall market risk, particularly for portfolios containing emerging market securities.

The insight that a simple liquidity adjustment can produce substantial accuracy gains in assessed VaR has wide-ranging implications for best-practice financial risk management, and it also has a number of more specialized implications. An immediate implication for trader compensation, for example, is that performance evaluation should be based on returns adjusted for risk – *including* liquidity risk. In recent years, many financial institutions have seen growth in their emerging markets trading activity due to higher margins. A risk-adjusted view of performance in those markets should account not only for return risk, but also for liquidity risk. Otherwise, performance will be incorrectly assessed and dealer compensation will be distorted.

A second implication involves regulatory capital guidelines. BIS regulations stipulate that the number of VaR violations be monitored, but not the way that VaR is computed. Neglecting liquidity risk will produce spuriously low estimates of overall market risk, under-capitalization, and too many violations. Hence the regulator, whether intentionally or not, *is* quite appropriately monitoring liquidity risk. The regulated need to do so as well.

REFERENCES

- Bertsimas, D. and A.W. Lo (1998), “Optimal Control of Execution Costs,” *Journal of Financial Markets*, 1, 1-50.
- Chriss, N. and R. Almgren (1998), "Optimal Liquidation," Manuscript, Department of Mathematics, University of Chicago, Goldman Sachs & Co., and Courant Institute of Mathematical Sciences.
- Dunbar, N. (1998), “Meriwether’s Meltdown,” *Risk*, 11 (October), 32-36.
- Jarrow, R. and A. Subramanian (1997), “Mopping up Liquidity,” *Risk*, 10 (December), 170–173.
- Longstaff, F. (1998), “Optimal Portfolio Choice and Valuation of Illiquid Securities,” Manuscript, Department of Finance, UCLA.
- Krakovsky, A (1999), “Gap Risk in Credit Trading,” *Risk*, March, 65-67.