

Modeling Liquidity Risk

With Implications for Traditional Market Risk Measurement and Management

Anil Bangia

Oliver, Wyman & Company

Francis X. Diebold

University of Pennsylvania,
Stern School, NYU
and Oliver Wyman Institute

Til Schuermann

Oliver, Wyman & Company

John D. Stroughair

Oliver, Wyman & Company

November 1998

This draft/print: December 21, 1998

Abstract

Market risk management under normal conditions traditionally has focussed on the distribution of portfolio value changes resulting from moves in the mid-price. Hence the market risk is really in a "pure" form: risk in an idealized market with no "friction" in obtaining the fair price. However, many markets possess an additional liquidity component that arises from a trader *not* realizing the mid-price when liquidating her position, but rather the mid-price minus the bid-ask spread. We argue that 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 risk and is therefore an important component to model.

We develop a simple liquidity risk methodology that can be easily and seamlessly integrated into standard value-at-risk models, and we show that ignoring the liquidity effect can produce underestimates of market risk in emerging markets by as much as 25-30%. Furthermore, we show that the BIS inadvertently is already monitoring liquidity risk, and that by not modeling it explicitly and therefore capitalizing against it, banks will be experiencing surprisingly many violations of capital requirements, particularly if their portfolios are concentrated in emerging markets.

We thank Steve Cecchetti and Edward Smith 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 behaviour.”

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

I. INTRODUCTION

The recent turmoil in the capital markets has led experts and laymen alike to cast liquidity risk in the role of the culprit. Inexperienced and sophisticated players were all caught by surprise when markets dried up. Unsurprisingly, the first to go were the emerging markets in Asia and more recently in Russia. Then it spilled over into the US corporate debt market which was indeed much more surprising. Its most famous victim has been Long Term Capital Management (LTCM). Spreads appeared to widen out of the blue; but this could have been predicted.

More generally, it is a well acknowledged fact that the standard Value-at-Risk (VaR) concept used for measuring both market and credit risk for tradable securities lacks a rigorous treatment of liquidity risk. At best, the risk for large illiquid positions is adjusted upwards in an *ad hoc* fashion by utilizing a longer time horizon in the calculation of VaR that at best is a subjective estimate of the likely liquidation time of the position. This holding-period adjustment is usually carried out using the square root of time scaling of the variances and covariances rather than a recalculation of variances and covariances for the longer time horizon.¹

However, the combination of the recent rapid expansion of emerging market trading activities and the recurring turbulence in those markets has propelled liquidity risk to the forefront of market risk management research. New work in asset pricing has demonstrated how liquidity plays a key role in security valuation and optimal portfolio choice by effectively imposing endogenous borrowing and short-selling constraints, as argued by Longstaff (1998) and the references therein. Liquidity also plays a major role in transaction costs as trades of large illiquid positions typically execute at a price away from the mid-price. BARRA’s Market Impact Model and other such models quantify the market impact cost, defined as the cost of immediate execution, for establishing and liquidating large positions. Jarrow and Subramanian (1997) consider the effect of trade size and execution lag on the liquidation value of the portfolio. They propose a liquidity adjusted VaR measure that incorporates the liquidity discount, volatility of liquidity discount, and the volatility of time horizon to liquidation. Although conceptually attractive, there is no available data or procedure to measure the model parameters such as mean and variances for quantity discounts or execution lags for trading large blocks.

In this article, we present a framework for treating liquidity risk using readily available data and integrating it with VaR calculations for market risk measurement of tradable securities. 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 which the market is unable to absorb easily. In a nutshell, we address a fundamental concern being raised currently in the markets, that current models ignore valuable information contained in the distribution of bid-ask spreads.

¹ To see why this is inappropriate, see Diebold, Hickman, Inoue and Schuermann (1998).


The rest of the paper is organized as follows. In section 2 we present our conceptual framework for understanding market risk, liquidity risk, and their interaction. In section 3 we describe the various components of overall market value uncertainty and techniques for their measurement, with emphasis on the neglected liquidity risk component and our approach to modeling it, and we also include worked examples for FX instruments. 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

We generically define risk as *uncertainty about future outcomes*. Market risk is primarily concerned with describing uncertainty about prices or returns due to market movements. Measurement of market risk therefore involves describing and modeling the return distribution of the relevant risk factors or instruments. Traditional market risk management under normal conditions usually deals exclusively with the distribution of portfolio value changes via the distribution of trading returns. These trading returns are based on mid-price, and hence the market risk is really in a “pure” form: risk in an idealized market with no “friction” in obtaining the fair price. However, many markets possess an additional liquidity component that arises from traders *not* realizing the mid-price when liquidating a position either quickly or when the market is moving against them, but rather that they realize the mid-price minus some spread. Marking to market 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. We argue that the deviation of this liquidation price from the mid-price, also referred to as the market impact or liquidation cost, and the volatility of this cost are important components to model in order to capture the true level of overall risk.

We conceptually split uncertainty in market value of an asset, i.e. its overall market risk, into two parts: uncertainty that arises from asset returns, which can be thought of as a pure market risk component, and uncertainty due to liquidity risk. Conventional VaR approaches such as JP Morgan’s (1996) RiskMetrics™ focus on capturing risk due to uncertainty in asset returns but ignore uncertainty due to liquidity risk. The liquidity risk component is concerned with the uncertainty of liquidation costs. Figure 1 summarizes our market risk taxonomy.


Figure 1: Taxonomy of Market Risk



Conceptually, we can express these ideas as a market / liquidity risk plane or *Risk Cross* (Figure 2 below) which considers the joint impact of the two types of risk. Most markets and trading situations fall into regions I and III; we observe that market risk and liquidity risk components are correlated in most cases.

For instance, FX derivative products in emerging markets have high market and liquidity risks and therefore fall into region I. The spot markets for most G-7 currencies on the other hand will fall into region III due to the relatively low market and liquidity risks involved. Most normal trading activity occurs in these two regions and is subject to *exogenous* liquidity risk, which refers to liquidity fluctuations driven by factors beyond individual traders' control. We distinguish this from *endogenous* liquidity risk, which refers to liquidity fluctuations driven by individual actions, such as an attempt to unwind a very large position. A trader holding a very large position in an otherwise stable market, for example, may find herself in region II.

Figure 2: Normal Trading on the Market Risk Cross




The risk cross is of course a simplification of a more complex relationship between markets and position sizes, involving both exogenous and endogenous components of liquidity risk. Moving along the liquidity axis, in particular, can be done either by moving across markets (i.e., increasing exogenous illiquidity), say from G-7 to emerging, or within a market by simply increasing one's position size (i.e., increasing endogenous illiquidity).

Exogenous illiquidity is the result of market characteristics; it is common to all market players and unaffected by the actions of any one participant (although it can be affected by the joint action of all or almost all market participants as happened in several markets in the summer of 1998). The market for liquid securities, such as G7 currencies, is typically characterized by heavy trading volumes, stable and small bid-ask spreads, stable and high levels of quote depth.² Liquidity costs may be negligible for such positions when marking to market provides a proper liquidation value. In contrast, markets in emerging currencies or thinly traded junk bonds are illiquid and are characterized by high volatilities of spread, quote depth and trading volume.

Endogenous illiquidity, in contrast, is specific to one's position in the market, varies across market participants, and the exposure of any one participant is affected by her actions. It is mainly driven by the size of the position: 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. This relationship is qualitatively depicted in Figure 3 below.

² Quote depth is defined as the volume of shares available at the market maker's quoted price (bid or ask).

Figure 3: Effect of position size on liquidation value



Market impact models such as one developed by BARRA quantify the above relationship between the transaction price and trade size. If the market order to buy/sell is smaller than the volume available in the market at the quote, then the order transacts at the quote. In this case the market impact cost, defined as the cost of immediate execution, will be half of the bid-ask spread. In our framework, such a position only possesses exogenous liquidity risk and no endogenous risk. If the size of the order exceeds the quote depth, the cost of market impact will be higher than the half-spread. The difference between the total market impact and half-spread is called the incremental market cost, and constitutes the endogenous liquidity component in our framework. Endogenous liquidity risk can be particularly important in situations when normally fungible market positions cease to be fungible; a good example would be when the cheapest to deliver bond of a futures contract switches.

Quantitative methods for modeling endogenous liquidity risk have recently been proposed by Jarrow and Subramanian (1997), Chriss and Almgren (1997), and Bertsimas and Lo (1998), among others. Jarrow and Subramanian, for example, consider optimal liquidation of an investment portfolio over a fixed horizon. They characterize the costs and benefits of block sale vs. slow liquidation, and they propose a liquidity adjustment to the standard VaR measure. The adjustment, however, requires knowledge of the relationship between the trade size and both the quantity discount and the execution lag. Clearly, there is no readily available data source for quantifying those relationships, which forces one to rely on subjective estimates.

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. This is because exogenous liquidity risk is characterized by the volatility of the *observed* spread with no reference to the relationship of the *realized* spread to trade size. The upshot is that we can incorporate liquidity risk into VaR calculations in a simple and powerful way.

III. MODELING UNCERTAINTY IN MARKET VALUE

A. Capturing asset return uncertainty using basic VaR

We begin by defining one-day asset returns at time t , r_t , to be the log difference of mid-prices:

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

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

$$P_{99\%} = P_t e^{\lfloor E[r_t] - 2.33 \sigma_t \rfloor} \quad (2)$$

where $E[r_t]$ and σ_t^2 are the first two moments of the distribution of the asset returns, and the multiple of 2.33 for the standard deviation arises from the assumption of normality. Without any loss of generality, we assume that the expected value of daily returns $E[r_t]$ is zero. The standard parametric VaR (henceforth P-VaR) is

$$P - VaR = P_t \left[1 - e^{\lfloor -2.33 \sigma_t \rfloor} \right] \quad (3)$$

We have assumed normality to simplify the exposition, but as will be made clear later, our methods do not rely crucially on normality. It is well known in empirical finance that asset returns exhibit volatility clustering: variances are not constant but change over time, as large changes tend to follow large changes and small changes tend to follow small. In order to account for this effect, we calculate volatility as an exponentially weighted moving average (EWMA) of past squared returns, as in the industry standard VaR model in JP Morgan's (1996) RiskMetrics.TM Another class of volatility models that has been developed to deal with volatility clustering is GARCH (Generalized Autoregressive Conditional Heteroskedasticity).³

B. Incorporating Exogenous Liquidity Risk

The above expressions for worst case price ($P_{99\%}$) and potential loss (P-VaR) only consider the volatility of the mid-price, whereas on average we would expect the bid-price to be $\frac{1}{2} \cdot \bar{S}$ (i.e. $\frac{1}{2}$ times the average spread) below that. Moreover, we are interested not in average circumstances but in unusual, tail-event circumstances, albeit due to overall market conditions (exogenous liquidity risk) rather than as a result of an individual trader attempting to dispose of an unusually large position (endogenous liquidity risk). We define exogenous liquidity risk measurement in terms of a confidence interval or a tail probability. We define the exogenous cost of liquidity (COL) 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.

$$COL = \frac{1}{2} \left[P_t \left[\bar{S} + a \tilde{\sigma} \right] \right] \quad (4)$$

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 $[\text{Ask-Bid}]/\text{Mid}$), $\tilde{\sigma}$ is the volatility of relative spread, a is the scaling factor such that we achieve roughly a 99% probability coverage.

Because spread distributions are very far from normal, we cannot rely on Gaussian distribution theory for guidance on the value of the scaling factor a . Our empirical analysis indicates that a ranges from 2.0 to 4.5 depending on the instrument and market in question. This range is determined largely by the shape of


³ While GARCH models have become quite popular in asset *pricing* (especially for options), they have yet to overtake EWMA as the model of choice for risk management. This is largely because high-dimensional GARCH models, necessary for portfolio-level risk management, are nigh impossible to estimate, whereas the grossly simplified EWMA model adapts readily to high (500+) dimensions.

the spread distribution — generally speaking, the greater the departure from normality, the larger a . However, since the empirical distributions of spread do not fit well into a standard parametric family for skewed distributions such as the log-normal, any scaling of the volatility to achieve a desired tail probability will necessarily be an approximation.


For illustration, consider the plots of densities of a few spreads. In Figure 4 we plot the density of relative spread (in basis points) with the best-fitting normal (same mean and standard deviation as the data) superimposed for illustration. Notice that 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. Notice especially that the scale of the Indian rupee spread density is nearly an order of magnitude larger than that of the yen. We seem to be seeing different regimes, some of high liquidity, others of low liquidity.

Figure 4: Distribution of spreads for three currency pairs.


**Daily Spread Distribution for Japanese Yen
5/95 - 5/97**



**Daily Spread Distribution of Thai Baht
5/95 - 5/97**




**Daily Spread Distribution for Indian Rupee
5/95 - 5/97**



One of the central functions of a market maker is liquidity provision. Campbell, Lo and MacKinlay (1997, section 3.2.2 and references therein) survey a literature which seeks to explain the bid-ask spread specifically from a market maker’s perspective, and this literature seems to break up this process into two components. The market making component is deemed to cover the order-processing and inventory costs to the market maker. This component, also termed the effective spread, is the market maker’s true compensation for providing liquidity. The other component of the spread compensates the market maker for taking the other side of a potentially information-based trade. Since market makers cannot distinguish the informed from the uninformed, they are forced to engage in potentially losing trades and should be accordingly rewarded.

In order to treat the liquidity risk and market risk jointly, we make a simplifying but reasonable assumption that in adverse market environments extreme events in returns and extreme events in spreads happen concurrently. Loosening this assumption merely complicates the algebra without bringing anything conceptually new. The correlation between mid-price movements and spread is not perfect, but it is nevertheless strong enough during extreme market conditions to enable and encourage us to treat market and liquidity risk as experiencing extreme movements simultaneously. Hence in calculating liquidity-risk adjusted VaR we incorporate *both a 99th percentile movement in the underlying and a 99th percentile movement in the spread*, as graphically summarized in Figure 5.

Figure 5: Combining Market and Liquidity Risk



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


$$P' = P_t e^{\lfloor -2.33 \sigma_r \rfloor} - \frac{1}{2} [P_t (\bar{S} + a\tilde{\sigma})] \quad (5a)$$

$$LAdj - VaR = P_t [1 - e^{\lfloor -2.33 \sigma_r \rfloor}] + \frac{1}{2} [P_t (\bar{S} + a\tilde{\sigma})] \quad (5b)$$

C. Adjustment for Fat Tails

For many instruments, especially in emerging markets, we find that the assumption of normality is violated: market returns are not normally distributed, instead they possess “fat-tails”; more formally, they are leptokurtic. Figure 6 displays the return distribution of two emerging market instruments, Mexican Peso/US Dollar exchange rate and 28-day cetes (Mexican government bond) rate, along with cumulative normal returns. It can be clearly seen that the return distributions of both the assets deviate significantly from normality; we see far too many extreme (both positive and negative) returns relative to those predicted by a normal distribution.

Figure 6: Comparison of Actual Returns with Cumulative Normal Returns



Sample Period: Jan'94-April'96

When asset returns deviate significantly from normality, the use of standard deviation multiples based on a normality assumption (such as 2.33 for 1% worst return) leads to an underestimation of risk. Therefore, we designed a correction factor θ for P-VaR to take account of leptokurtic or “fat-tailed” distributions, such that:

$$P - VaR = P_r \left[1 - e^{[-2.33\theta\sigma, 1]} \right] \quad (6a)$$

The correction factor is such that $\theta = 1$ if the asset return distribution is normal, and $\theta > 1$ is an increasing function of the “fat-tailedness” causing the deviation from normality. A commonly used statistic to characterize the fatness of tails of distributions is the kurtosis. Table 1 displays the kurtosis of spot exchange rates of selected currencies against the US dollar. Clearly, return distributions with higher kurtosis have fatter tails and therefore should have higher θ . In fact, since the unconditional kurtosis of a return distribution is a measurable characteristic, it seems to be a natural yardstick to estimate the correction factor θ . By expressing θ as an explicit function of the unconditional kurtosis of the return distribution, in essence we are improving our estimate of VaR (or a fixed percentile-level return) in equation (6a) by utilizing two distributional characteristics of the return distribution, namely its second and fourth moments.

Table 1: Kurtosis of Selected Spot rates against USD⁴

<i>CURRENCY</i>	<i>KURTOSIS</i>
Thai Baht	4.9
British Pound	5.4
German Deutschmark	6.3
French Franc	6.7
Japanese Yen	7.0
Malaysian Ringgit	7.5
Singapore Dollar	8.4
Indonesian Rupiah	15.8
Brazilian Real	18.1

⁴ Source: Reuters. 1/95 - 5/97 daily.

Mexican Peso	20.5
Indian Rupee	23.2

Sample period: Jan. '95 to May '97

In order to derive an explicit relationship between kurtosis κ and the correction factor θ , we first considered the relationship between standard deviation, kurtosis and the empirical tail probability for the t -distributions with various degrees of freedom. The family of t -distributions present themselves as a convenient control sample where the kurtosis (and fatness of tails) can be adjusted by changing the number of degrees of freedom. Recall that for large sample sizes — large degrees of freedom —, the t converges to the normal distribution. The relationship between kurtosis κ and the correction factor for t -distributions is well captured by the following empirically derived relationship:

$$\theta = 1.0 + \phi \ln(\kappa/3) \quad (6b)$$

where ϕ is a constant whose value depends on the tail probability (1%, 2.5% etc.). Clearly, if the distribution is normal, then $\kappa = 3$ and $\theta = 1$: no adjustment is needed. We estimated the value of the constant ϕ by regressing the right hand side of equation (6a) with historical VaR (H-VaR) for 14 FX currencies. For 1% tail probability VaR measures, the value of ϕ is 0.4.

Clearly historical simulation (HVaR) is also designed to capture the non-normalities of market return data, since it makes no distributional assumption of any kind about the market return process. Thus, one could argue that the “fat-tail” adjustment proposed here is redundant. However, the purpose of the adjustment is not to replace HVaR, rather to complement it by leveraging a relatively small data set to improve the accuracy of the PVaR. In fact, HVaR’s advantage is achieved at a price: by dropping the parametric constraint, one needs much more data to learn about the process.⁵ Moreover, during periods of high volatility or market turmoil, the fat-tail adjustment will have a clear advantage. Since HVaR by construction is based on an *unweighted* average of the past while the parametric methods use weighted averages, recent information carries the same weight as market information from several months or even years ago.⁶ The asymmetric updating scheme of the parametric methods precludes this.

D. Conditional vs. Unconditional Measures

Conventional VaR, particularly parametric VaR, is a conditional measure, because it takes as input (usually) a conditional variance model such as EWMA or GARCH. When the portfolio contains a lot of nonlinear instruments, i.e. options and derivatives whose value varies nonlinearly with movements in their underlyings, one may condition out all states of the world by conducting a Monte Carlo computation using, usually, conditional volatilities and covariances. In contrast, standard historical simulation (H-VaR) is also an *unconditional* VaR metric since one computes the tail probability non-parametrically from the unconditional distribution of market returns. An intriguing suggestion toward moving to a more conditional approach was recently put forth by Boudoukh, Richardson, and Whitelaw (1998).

Our liquidity risk enhanced VaR is a cocktail of these conditional metrics plus unconditional spread risk plus a fat-tail adjustment using the unconditional kurtosis. Clearly a next step is to enrich the model by considering conditional (i.e. time-varying) spread risk models as well as time-varying kurtosis. We stress that especially the latter is a difficult exercise since it is already quite difficult, i.e. requires large samples, to obtain reliable *unconditional* estimates of higher moments without imposing the additional demands of conditioning on a specific model.

⁵ A study by the Federal Reserve Bank of New York indicates that the real advantages of H-VaR require approximately four years of data. See Hendricks (1996).

⁶ A recent method proposed by Boudoukh, Richardson, and Whitelaw (1998) outlines a weighted H-VaR scheme that seems promising.

E. Calculation Examples

We will look at two currency exchange rates against the US dollar, a highly liquid one (the Japanese Yen, JPY) and a more thinly traded one (the Thai Baht, THB). Our standard measure of market risk involves the volatility of returns. However, we will see that this can be a deceptive indicator of *liquidity* risk. We will do this step-by-step by first looking at basic market risk, then adding the fat-tail adjustment, and finally considering the liquidity component. We break up the period of analysis into pre and “post” Asia crisis (where post means specifically after the beginning of the crisis) by incorporating a breakpoint in May 1997. Note that the Thai Baht went from pegged to free floating over this period, with a dramatic impact on its return and spread volatility.

a) Including the Fat Tail Adjustment

For the former period, σ_{JPY} was 1.12% and θ_{JPY} was computed to be 1.34, while for the THB, $\sigma_{THB} = 0.19\%$ and $\theta_{THB} = 1.2$. The 99% worst case returns for JPY and THB are

$$\begin{aligned} r'_{JPY} &= (1.34) \cdot (1.12\%) \cdot 2.33 = 3.50\% \\ r'_{THB} &= (1.2) \cdot (0.19\%) \cdot 2.33 = 0.53\% \end{aligned}$$

To translate this back into prices from returns we simply define the 1% worst price P^* as:

$$\begin{aligned} P^*_{JPY} &= \text{¥}126.735 \cdot e^{-3.50\%} = \text{¥}122.38 \\ P^*_{THB} &= \text{B}26.105 \cdot e^{-0.53\%} = \text{B}25.97 \end{aligned}$$

In other words, on May 2, 1997, we would expect the probability of a 1-day price drop in the Yen from ¥126.735 to ¥122.38 or less to be no more than 1%. For the latter period, we obtain $\sigma_{JPY} = 2.00\%$, $\theta_{JPY} = 1.4$ and $\sigma_{THB} = 5.48\%$, $\theta_{THB} = 1.7$. The 99% worst case returns and prices can similarly be calculated as:

$$\begin{aligned} r'_{JPY} &= (1.4) \cdot (2.00\%) \cdot 2.33 = 6.50\% \\ r'_{THB} &= (1.7) \cdot (5.48\%) \cdot 2.33 = 21.7\% \\ P^*_{JPY} &= \text{¥}127.17 \cdot e^{-6.50\%} = \text{¥}119.14 \\ P^*_{THB} &= \text{B}53.55 \cdot e^{-21.7\%} = \text{B}43.10 \end{aligned}$$

b) Adding Liquidity Risk

For the period leading up to May 2, 1997 we obtain (including the kurtosis adjustment as above):

$$\begin{aligned} \text{Bid}_{JPY} &= \text{¥}122.38 - \frac{1}{2}[\text{¥}122.38 \cdot (6.6\text{bp} + a \cdot 1.7\text{bp})] \\ \text{Bid}_{THB} &= \text{B}25.97 - \frac{1}{2}[\text{B}25.97 \cdot (6.3\text{bp} + a \cdot 4.1\text{bp})] \end{aligned}$$

and for the period from May 2, 1997 to January 22, 1998 we obtain:

$$\begin{aligned} \text{Bid}_{JPY} &= \text{¥}119.14 - \frac{1}{2}[\text{¥}119.14 \cdot (7.1\text{bp} + a \cdot 2.7\text{bp})] \\ \text{Bid}_{THB} &= \text{B}43.10 - \frac{1}{2}[\text{B}43.10 \cdot (76.4\text{bp} + a \cdot 47.4\text{bp})] \end{aligned}$$

We measured the liquidity risk scaling factor a for very high liquidity markets like the JPY to be 2.5 and for moderately high liquidity markets like the THB to be 3.5, so we have, for the first period:

$$\begin{aligned} \text{Bid}_{JPY} &= \text{¥}122.31 \\ \text{Bid}_{THB} &= \text{B}25.94 \end{aligned}$$

and for the second period:

$$\text{Bid}_{\text{JPY}} = \text{¥}119.06$$

$$\text{Bid}_{\text{THB}} = \text{B}42.58$$

Again, this presumes that the 1% tail event in market and liquidity risk *perfectly* correlated, which is a conservative assumption, and one which seems reasonable given our empirical analysis.

c) Summary

Let's summarize the progression of market risk calculations for our currency trader by taking into account both the market and liquidity components to total market risk. Of principal interest here is the marginal impact of the liquidity component, heretofore not explicitly measured, as revealed in Tables 2a and 2b.

Table 2a: FX Market & Liquidity Risk Summarised “Pre Crisis” Period

	JPY	THB
Price on May 2, 1997 (P_t)	¥126.735	B26.105
Return volatility (σ_t)	1.12%	0.19%
Fat tail factor (θ)	1.34	1.2
Market component of VaR, $P_t \left[1 - e^{\theta - 2.33 \theta \sigma_t} \right]$	¥4.35	B0.14
Liquidity component of VaR, $\frac{1}{2} \left[P_t \left(\bar{S} + a \tilde{S} \right) \right]$	¥0.07	B0.03
Total Adjusted Value at Risk	¥4.42	B0.17
% liquidity component	1.5%	16%

Table 2b: FX Market & Liquidity Risk Summarised “Post Crisis” Period

	JPY	THB
Price on January 22, 1998 (P_t)	¥127.17	B53.55
Return volatility (σ_t)	2.0%	5.48%
Fat tail factor (θ)	1.4	1.7
Market component of VaR	¥8.03	B10.45
Liquidity component of VaR	¥0.08	B0.52
Total Adjusted Value at Risk	¥8.11	B10.97
% liquidity component	1.0%	5%

What is particularly revealing is the difference in risk contribution of the market and liquidity components for these two currencies. Intuitively we would expect that the Yen be much more liquid than the Baht; arguably the JPY/USD spot market is one of the most liquid markets in the world. By looking only at the conventional measure of market risk, i.e. based only on the return volatility, we miss almost a fifth of the total market risk for the Thai Baht for the initial period, whereas one could safely ignore the liquidity component for the Japanese Yen. Moreover, the comparisons across the East Asian crisis reveal what is suspected, namely that the floating of the Thai Baht and subsequent heavy trading has increased the pure market risk by raising its return volatility substantially, thereby decreasing the proportion of liquidity risk from 16% of total market risk to 5%. One interpretation for the drop in liquidity risk is the pressure released by depegging the Baht. The markets heretofore could express their sentiment (i.e. excess demand for USD) only via the spread distribution before the devaluation. Both returns and spreads became more volatile thereafter, but “returns risk” now dominates “spread risk” in the Baht currency market.

IV. A NEW PORTFOLIO VAR

A. Ex-Post and Ex-Ante Measures

We have discussed univariate or instrument-level market and liquidity risk. The covariance matrix in VaR allows us to make a seamless transition from a univariate to a multivariate or portfolio-level market risk measure. The covariance matrix measures the correlation structure between the different market returns as well as volatility. One approach for a full portfolio level treatment for liquidity risk is to estimate the correlation structure of the market spreads in an analogous fashion. However, spread distributions are not nearly as well behaved as return distributions making this approach quite difficult.

Another approach is to compute a portfolio level bid and ask series by simply taking the weighted sum of the bids and asks of the instruments. One can then use the univariate methodology above to adjust the pure market risk for liquidity risk. This is straightforward, and we will call this the *ex-post* liquidity measure. Alternatively one can redefine the price at the instrument repricing level and incorporate the liquidity risk there. This output serves as input to the covariance matrix VaR calculator. We thereby leverage the existing covariance structure at the cost of having to rewrite the repricing engines. We call this the *ex-ante* liquidity risk metric.

B. Backtesting some Portfolios

Using the ex-post liquidity measure, we examined the effect of our liquidity risk and fat-tail adjustment add-on to standard VaR for the case of four sample portfolios over a sample period from January 1995 to March 1997. Table 3 contains the exact description of the portfolios. Specifically we considered the BIS backtesting framework over the course of one year, but we considered the liquidation price to be the bid or ask price, not the mid-price. We then considered the impact on regulatory capital.

Table 3: Compositions of four equally-weighted portfolios for backtesting analysis.

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 violations an institution's market risk model experiences over the course of a trading year (specifically, 252 trading days). The institution is then assigned a regulatory color (green, yellow or red) with a corresponding capital multiplier. From a bank's perspective, both the over- and underestimation of VaR results in tied up capital. Overestimation will yield too few violations; this indicates inefficiently used capital. On the other hand, underestimation results in an uncomfortably high number of violations; BIS regulations penalize this with capital charges as shown in Table 4.

Table 4: BIS Capital multiplier based on number of backtesting violations

# OF EXCEPTIONS IN 250 DAYS	PR (EXCEPTIONS) IF TRUE 99% VAR	BIS ZONES	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

Figure 7 displays the results of backtesting for the four portfolios considered. The numbers of backtesting violations during 240 trading days are shown for two models: a standard VaR model, and the same model with our liquidity risk add-on. While the liquidity risk add-on allows for fewer violations in emerging market portfolios, the gains, as expected, are small in more developed markets. If we then include a fat-tail adjustment (see Figure 8), the results are even stronger. In fact when both liquidity and kurtosis adjustments are made to the standard VaR model, all four portfolios fall under the green regulatory zone.

Figure 7: Backtesting results for sample portfolios with and without the add-on for liquidity risk





Figure 8: Backtesting results for sample portfolios with and without kurtosis-based adjustment for fat tails



V. CONCLUDING REMARKS

Traditional Value-at-Risk measures are obtained from the distribution of portfolio returns computed at the bid-ask average prices. That method underestimates risk by neglecting the fact that liquidation occurs not at the bid-ask average price, but rather at the bid-ask average 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 liquidity risk can produce substantial underestimates of overall risk, particularly in emerging market securities.

What does this mean in terms of capitalization of a trading operation? Put differently, what are the regulatory implications of ignoring liquidity risk? BIS regulations stipulate only that the number of VaR violations be monitored, not the way that VaR is computed. Neglecting liquidity risk will lead to underestimation of overall risk, under-capitalization, and too many violations. Hence the BIS, whether intentionally or not, *is* quite appropriately monitoring liquidity risk. The regulated need to do so as well.

Finally there is a compensation implication for trading rooms: 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 market risk, but also for liquidity risk. Otherwise, performance will be incorrectly assessed and dealer compensation will be distorted.

REFERENCES

- Boudoukh, J., M. Richardson and R. Whitelaw (1998), "The Best of Both Worlds," *Risk*, 11, 64-67.
- Bertsimas, D. and A.W. Lo (1998), "Optimal Control of Execution Costs," *Journal of Financial Markets*, 1, 1-50.
- Campbell, J.Y., A.W. Lo and A.C. MacKinlay (1997), *The Econometrics of Financial Markets*. Princeton: Princeton University Press.
- Chriss, N. and Almgren, R. (1998), "Optimal Liquidation," Manuscript, Department of Mathematics, University of Chicago, Goldman Sachs & Co., and Courant institute of Mathematical Sciences.
- Diebold, F.X., A. Hickman, A. Inoue and T. Schuermann (1998), "Scale Models," *Risk*, 11 (January), 104-107.
- Dunbar, Nicholas (1998), "Meriwether's Meltdown," *Risk*, 11 (October), 32-36.
- JP Morgan (1996), "RiskMetrics -- Technical Document," Fourth Edition, New York.
- 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.
- Hendricks, D. (1996), "Evaluation of Value-at-Risk Models Using Historical Data," Federal Reserve Bank of New York, *Economic Policy Review*, April, 39-69.