Incorporating Liquidity Risk into Value-at-Risk Models

Jeremy Berkowitz

Graduate School of Management
University of California, Irvine

September 12, 2000

Abstract: This paper suggests a method of incorporating liquidity risk into VaR models. Such large scale risk models have been developed to measure daily firm-wide trading exposures and at many large banks form the basis of highly sophisticated systems for portfolio risk measurement and management. However, such models are overwhelmingly silent on liquidity risks. Our starting point is to posit a concrete definition of liquidity risk. In practice, operational definitions vary from volume-related measures to bid-ask spreads and to the elasticity of demand. We argue that elasticity based measures are of most relevance in that they incorporate the impact of the seller actions on prices. Armed with this definition, we build on existing theoretical results to derive methods to measure liquidity risk and incorporate such measures into standard market risk models.

Key Words: Liquidity, VaR, Specific Risk

Acknowledgements: I thank Jim O'Brien, Matt Pritsker and Pat White for helpful suggestions.

1. INTRODUCTION

It is well known that the Market Risk Amendment (MRA) to the Basle Accord allows qualifying banks to base their market risk capital requirements on a formal statistical model of portfolio returns. Currently, all major U.S. trading firms employ at least one large-scale model for measuring market risks. In many cases, banks have developed quite sophisticated risk models designed to measure trading exposures at the firm-wide level. While the specifics differ across firms, all models employ a standard risk metric, Value-at-Risk (VaR). Furthermore, large trading institutions now report summary data on their VaR estimates in annual reports to stockholders, as an indicator of market risk exposures.

However, far less attention has been paid to the 1997 modifications that allow for modeling of the specific risk component of portfolio risk.[1] Undoubtedly, these more recent modifications were enacted in part to address the fact that market risk models are almost without exception silent on liquidity risk. In particular, such models assume that the bank or firm in question is atomistic and can appropriately ignore its own effect on asset prices. In reality, however, this assumption has been violated -- at times in spectacular fashion. During the market turbulence of the summer of 1998, for example, well publicized liquidations and merely the threat of liquidation by some large players caused a substantial move in some asset prices.

The goal of this article is to suggest approaches to incorporating such liquidity risk into standard market risk models. The starting point of our approach is the recognition that liquidity risk is identically zero for assets that are held to maturity. Put another way, liquidity risk arises from the uncertain demand for funds which must be met by liquidating assets. While this definition may seem to be a tautology, formalizing the idea will turn out to have useful and far-reaching implications for regulatory supervision.

Since liquidating assets may depress prices, we interpret the additional risk as arising from (possibly time-varying) downward sloping demand curves. This approach to defining liquidity risk is consistent with Grossman and Miller (1988) and Kamara (1994) who discuss the risk that transactions in the treasury market occur at prices different than those currently quoted. Our approach is most similar to Duffee and DeMarzo (1999) who study liquidity in the context

---

[1] The 1997 textual changes to the *Amendment to the Capital Accord to Incorporate Market Risks* are available online at http://www.bis.org/publ/index.htm

of the security design problem, although our motivation and objectives are quite different.

Competing definitions of liquidity include the variation in trading costs such as bid-ask spreads (among others, Admati and Pfleiderer (1988), Foster and Viswanathan (1990), Bhushan (1991), Amihud and Mendelson (1991)) and the risk that a (solvent) borrower is unable to obtain funding as in Diamond (1991).

The bid-ask strand of liquidity research is largely driven by traders' time-varying inventory costs. Such issues, while undoubtedly related, are not of direct relevance to the present portfolio management context. Similarly, Diamond's definition appears ill-suited to agents with large, diversified portfolios of traded assets. Although some assets may suddenly have sharply reduced values, it is unrealistic to imagine that the agent is literally *unable* to sell any assets.

Finally, we note that intuitively appealing descriptions of the liquidity concept are discussed in Black (1971) and given more formal shape by Kyle (1985). In particular, Kyle suggests a tripartite definition of liquidity consisting of the cost of turning an asset around in a short time (tightness), the size of order flow needed to change prices a given amount (depth) and the recovery speed of prices after an uninformative shock (resiliency). While attractive, these transaction-based quantities are *measures* of liquidity -- they do not explain the underlying cause of differences in liquidity. As such, they are of limited use in modeling liquidity risk.

The remainder of the paper is organized as follows. Section 2 presents the background risk measurement framework in which we analyze liquidity risk. Section 3 suggests some computationally feasible approaches to quantifying the risk. Section 4 concludes.

2. FRAMEWORK

It will be useful to begin by stating the portfolio forecasting problem and then generalizing to incorporate liquidity problems. Suppose we are interested in approximating some given distribution of portfolio values, y_t.[2] The problem is typically tackled by combining a set of k market-wide factors, x_t , and pricing functions, $p(x_t)$. The risk factors are basic, economy-wide rates and prices such as interest rates and exchange rates. The pricing models are used to value all

[2]This is no more burdensome that VaR. To see why, note that if VaR is estimated by simulation, the entire distribution is approximated prior to taking the single percentile of interest. If VaR is estimated by variance-covariance methods, the entire distribution is known analytically.

assets as a function of the factors. For example, the delta approximation (first-order Taylor expansion) to Black-Scholes might be used to predict option prices. The valuation model can thus be written as,

$$(1) \qquad y_{t+1} = Q_t^{/} p(x_{t+1}) + e_t$$

where Q_t is an Nx1 vector of asset positions and e_t is idiosyncratic risk. The pricing model $p(x_t)$ is a mapping from R^k to R^N that produces an estimate of the prices of all assets in the portfolio for any given factor prices.

The most common method of approximating the CDF of y_{t+1} is to assume a factor distribution from which simulated values, $\hat{x}_{t+1}$, can be drawn. These $\hat{x}_{t+1}$ are then plugged into equation (1) to produce simulated portfolio values,

$$(2) \qquad \hat{y}_{t+1} = Q_t^{/} p(\hat{x}_{t+1})$$

After repeating this process many times, the simulated portfolio values can be sorted and tabulated to yield an estimate of the CDF of $\hat{y}_{t+1}$.

A key assumption in the standard approach is that asset prices are independent of the firm's own positions, Q_t. That is, the implicit assumption is that the firm is atomistic so that the changes in asset prices can be modeled *separately from portfolio size.*

This assumption has failed spectacularly in recent years. During the well publicized liquidity crisis in the summer of 1998, significant price moves in certain assets resulted from one large institution liquidating (or threatening to liquidate) its positions in those assets. Moreover, the failure of the assumption can arise even when the firm does not hold a large portion of the assets. In particular, the possibility that the seller has more information than buyers, asymmetric information, can give rise to precisely the same kind of downward sloping demand curve (Duffee and DeMarzo (1997, 1999)). This is a result of the classic lemons problem -- the more of the asset is sold, the more suspicious the market becomes that the motive for selling is privately known bad information.

When the possibility of downward sloping demand is considered, we can no longer separate quantities (positions) from prices. Since prices will now depend on the supply of assets and require consideration of more general specifications for portfolio valuation. We adopt the following to recognize the importance of "seller impact".

DEFINITION: *Liquidity risk is the uncertain change in portfolio value caused by liquidating assets to meet future cash requirements, above and beyond exogenous changes in factor prices.*

We are concerned with, for example, a portfolio manager or bank that must raise a given amount of cash for its clients or meet some debt obligation. The portfolio manager's problem is to sell some given quantity of assets, M_t, with a minimal reduction in the value of the remaining (retained) portfolio. We interpret these transactions as random shocks demanded by liquidity traders. Sales are transacted through a portfolio manager who must redeem the assets for her clients while simultaneously maximizing portfolio value.

We formulate the problem as in Bertsimas and Lo (1998).[3] The portfolio manager must sell a given number of shares, M_t, within a given horizon, T. The objective is to maximize the revenue received from the sales within the required time,

(3) $$\max_{\{q_t\}} E_t\left[\sum_t^T p_t q_t\right] \quad \text{subject to} \sum_t^T q_t = M_t$$

The law of motion for asset prices is assumed to be described by

(4) $$p_t = p_{t-1} - \theta q_t + x_t$$

where q_t is the amount of asset sold.[4]

[3]While their goal was to study minimizing execution costs when buying a large block of assets, their framework is equally applicable to the problem of maximizing the revenue from selling a block of assets.

[4]We choose this law of motion because it leads to particularly tractable results. The methods of this paper could accommodate the other curves considered in Bertsimas and Lo (1998).

In our context, market-wide changes in asset prices, x_t, are interpreted as rational reactions to information or preference shocks. As in Kyle (1985), we assume that the noise trades, q_t, are independent of informed trades. Under these conditions, Bertsimas and Lo (1998) show that the optimal solution to this problem is $q_t^* = M_t / T$. In equilibrium,

$$\hat{p}_{t+1} = p_t + x_{t+1} - \theta q_t^*. \tag{5}$$

Thus we are interested in forecasting

$$\begin{aligned} \hat{y}_{t+1} &= Q_t' \hat{p}_{t+1} \\ &= Q_t'(p_t + x_{t+1} - \theta q_t^*) \end{aligned} \tag{6}$$

The first term, $Q_t'(p_t + x_{t+1})$, represents the market risk component of portfolio risk. It measures the change in asset values due to market-wide factors movements, with the firm's own action having negligible effect on prices.

The additional $Q_t'(-\theta q_t^*)$ term accounts for possible price drops resulting from the market reaction to the firm's sales. The independence of q_t and x_t implies that the distribution of y_{t+1} can be decomposed into the separate effects implied by these two terms.

3. Measurement

Given the above outlined framework, it is now possible to quantify the additional risk due to the seller's own impact on asset prices. We first consider forecasting the expected value and conditional variance of y_{t+1}. Below we discuss evaluation of the entire distribution of y_{t+1}.

Mean-Variance Analysis

From equation (6), we have that

(7) $$E_t(y_{t+1}) = Q_t^{/}\left(p_t - \theta\, E\, q_t^* + E\, x_{t+1}\right)$$

Under our maintained assumptions the simple addition of variances is valid so that the *additional* variance from incorporating liquidity is given by

(8) $$Q_t^{/}\left(\mathrm{var}[\theta\, q_t]\right) Q_t \quad .$$

A regression equation of the form,

(9) $$p_{t+1} - p_t = \alpha - \theta\, q_t^* + x_{t+1} + \varepsilon_t$$

using historical observations on portfolio value and net flows delivers an estimate of θ.

Evaluating the Distribution Function

Using equation (6) and the independence assumption we can actually estimate the entire forecast distribution, preserving potentially important information in higher order moments. Using the inversion formula for characteristic functions (e.g., Billingsley (1986) p.357), we have that

(10) $$F(y_{t+1}) = \frac{1}{2\pi} \int_{-\infty}^{\infty} \frac{e^{-isk}}{is}\, E\, e^{is\, Q_t^{/}(p_t + \hat{x}_{t+1})} \bullet E\, e^{is\, -\theta Q_t^{/} q_t}\, ds$$

Equation (10) is easily approximated through numerical methods. In particular, we see that $E\, e^{is\, Q_t^{/}(p_t + \hat{x}_{t+1})} = \int e^{is\, Q_t^{/}(p_t + x_{t+1})} f(x_{t+1})\, d\, x_{t+1}$. At time t, Q_t and p_t are known. The distribution of factor prices, $f(x_{t+1})$, is typically assumed to come from a tractable family, such as the Normal or student-t, and then fit to past factor prices (see Berkowitz and O'Brien (2000)). If so, the expectation can actually be evaluated analytically. A second approach is to simulate factor prices by bootstrap or other methods so that $f(x_{t+1})$ is the empirical distribution of past prices. In this case, the integral must be evaluated numerically. Such methods are extremely fast and can be made arbitrarily accurate (e.g., Press (1997)).

Similarly, we have $E\, e^{is\text{-}Q_t'\theta q_t} = \int e^{is\text{-}\theta Q_t' q_t} g(q_t)\, d\, q_t$ where $g(q_t)$ is the distribution of liquidity trades. Given observations on q_t, we can either fit a Normal or other distribution to obtain a closed-form solution or use numerical methods. The unknown parameter, θ, is most simply estimated by fitting a linear regression.

An interesting point that emerges from this analysis is that net purchases are just as useful, from a statistical point of view, as net redemptions. Although risk managers are in reality concerned about sudden outflows rather than cash inflows, inflows can provide valuable information about the slope of the demand curve. To the extent that inflows raise prices, we can better ascertain the relationship between Q_t and p_t.

Identifying Quantities and Prices

A complication in practical implementation of these methods is that often datasets do not contain separate observations on quantities and prices. Rather, at the portfolio level net asset values (NAVs) and net flows are typically reported. Nevertheless, there are some simple ways of inferring the desired information from the available data.

Write the time t net asset value as $Q_t' p_t$ The dollar value of net new flows, q_t, is given by $(Q_t\text{-}Q_{t\text{-}1})' p_t$. This is the change in the number of shares times the new prevailing price per share. Therefore, the $NAV_{t+1} - q_{t+1} = Q_t' p_{t+1}$. Making use of one further transformation, we have

(11) $$\log\,(Q_t' p_{t+1} / Q_t' p_t) = \tilde{p}_{t+1} - \tilde{p}_t$$

where $\tilde{p}_t$ is the log price level. Equation (11) indicates that given observations on NAV and net flows, we can calculate $\tilde{p}_{t+1} - \tilde{p}_t$ as required for estimation of the demand equation (9).[5]

More generally, we note that the following accounting identity can be very useful in

[5]Actually equation (9) is written in terms of the level of prices p_t rather than the log level. This is not a serious complication and can be dealt with in several ways. For example, levels can be inferred by choosing a numeraire value for the initial price p_0.

inferring desired quantities from reported data:

$$\Delta NAV = \text{net flows} + \text{gross return} \quad (12)$$

where net flows exclude reinvested dividends (if any) and the gross return is given by $Q_t^{/}p_{t+1}-Q_t^{/}p_t$. Equation (12) states that the overall change in portfolio value can be decomposed into net flows in or out of the portfolio and the changes in asset prices. Given data on any two components of equation (12), the third can be inferred.

Liquidity Risk in Mutual Fund Complexes

In this section we illustrate our methods using data on mutual funds and net redemptions collected by Trim Tabs from February 1998 to January 2000. Trim Tabs collects daily data from 500 equity funds which are categorized by objective and aggregated into Investment Company Institute (ICI) Categories. We used data on the following four: Aggressive Growth, Growth, Growth & Income, Precious Metals. For market-wide factor prices, we took daily S&P500 and 30 year U.S. bond prices.

The first step was to estimate equation (9) which yielded the following estimates of θ,

Fund	Aggressive	Growth	Growth & Income	Precious Metals
θ	220.38	-3.37	33.21	18240.67
t-stat	(0.91)	(-0.01)	(0.09)	(4.51)

where t-statistics are shown in parentheses. Given these, we next estimate the one-step ahead mean and variance shown in (7) and (8) using a five day moving average of past fund and factor prices. Figure 1 displays the estimated overall variance $var_t(y_{t+1})$ (dashed line) and marginal variance due to liquidity risk (solid line) in units of $100 Billion.

For aggressive growth funds, the marginal variance due to liquidity shocks is small but noticeable, while for growth and growth & income funds liquidity risk is virtually zero. At the other extreme, liquidity risk makes up the vast majority of risk historically seen in precious metals funds.

Using the normal approximation for factor and fund price distributions, it is possible to calculate the 1-day ahead VaR as discussed above. Figure 2 shows estimates of the 99% VaR for each fund type with (dashed line) and without (solid line) incorporating liquidity risk. As one would expect from Figure 1, the VaR of aggressive growth funds are somewhat larger (more negative) when liquidity risk is accounted for. The VaRs of the growth and growth & income funds are unchanged when liquidity risk is measured. However, the estimated VaR for precious metals is on some days drastically larger when liquidity risk is incorporated.

4. Conclusions

In this paper we suggest a method of incorporating liquidity risk into VaR models. The importance of formalizing and measuring liquidity risk has recently taken on heightened recognition in the wake of the liquidity crisis of the financial turmoil following the Russian default episode of 1998. Nevertheless, standard market risk models currently in use at commercial banks are virtually without exception silent on liquidity risks.

The first step in tackling the question of how to measure and incorporate such risks into standard models is necessarily to formulate a rigorous definition of liquidity risk. Despite a very large literature on liquidity risk, economists have not converged on a universal definition. We argue that the salient risk is that asset prices at transaction will be substantially different than those quoted prior to transaction. As such, it is crucial that liquidity risk reflect the impact of the seller on prices so that elasticity-based measures are of particluar relevance.

Building on the work of Kyle (1985) and Bertsimas and Lo (1998), we show that under certain conditions the additional variance arising from seller impact can easily be quantified given observations on portfolio prices and net flows. It is also possible to estimate the entire distribution of portfolio risk through standard numerical methods.

The measurement of liquidity risk in the portfolios of banks, mutual funds, pension funds and other entities would substantially improve risk managers' ability to evaluate the firm's risk profile. Additionally, such measurement is likely to sharpen empirical tests of the forecast performance of internal risk models. Without incorporating liquidity risk, a model is unlikely to be able to accurately forecast tail behavior in the data. It is hoped that our approach will accelerate further progress in specific risk modeling and facilitate communication among financial institutions and regulatory bodies.

References

Admati, A. R. and P. Pfleiderer (1988), "A Theory of Intraday Patterns: Volume and Price Variability," *Review of Financial Studies*, 1, 3-40.

Amihud, Y. and H. Mendelson (1991), "Dealership Market: Market-Making with Inventory," *Journal of Financial Economics*, 8, 31-53.

Bertsimas, D. and Lo, A. W. (1998), "Optimal Control of Execution Costs," *Journal of Financial Markets*, 1-50.

Berkowitz, J. and J. O'Brien (2000), "The Performance of Large-Scale Porftolio Valuation Models," manuscript, Federal Reserve Board.

Bhushan, R. (1991), "Trading Costs, Liquidity, and Asset Holdings," *Review of Financial Studies*, 4, 343-360.

Black, F. (1971), "Towards a Fully Automated Exchange, Part I," *Financial Analysts Journal*, 27, 29-34.

Diamond, D. W. (1991), "Debt Maturity Structure and Liquidity Risk," *Quarterly Journal of Economics*, 2, 709-737.

Duffee, D. and P. DeMarzo (1999), "A Liquidity-Based Model of Security Design," *Econometrica*, 67, 65-99.

Foster, F. D. and S. Viswanathan (1990), "A Theory of the Interday Variations in Volume, Variance, and Trading Costs in Securities Markets," *Review of Fin. Studies*, 3, 593-24.

Grossman, S. J. and M. H. Miller (1988), "Liquidity and Market Structure," *Journal of Finance*, 43, 617-633.

Kamara, A. (1994), "Liquidity, Taxes, and Short-Term Treasury Yields," *Journal of Financial and Quantitative Analysis*, 29, 403-417.

Kyle, A. S. (1985), "Continuous Auctions and Insider Trading," *Econometrica*, 53, 1315-1335.

Press, W. H. (1997). *Numerical Recipes in C: The Art of Scientific Computing.* New York : Cambridge University Press.

Figure 1: Estimates of Overall Fund Variance and Variance due to Liquidity Shocks

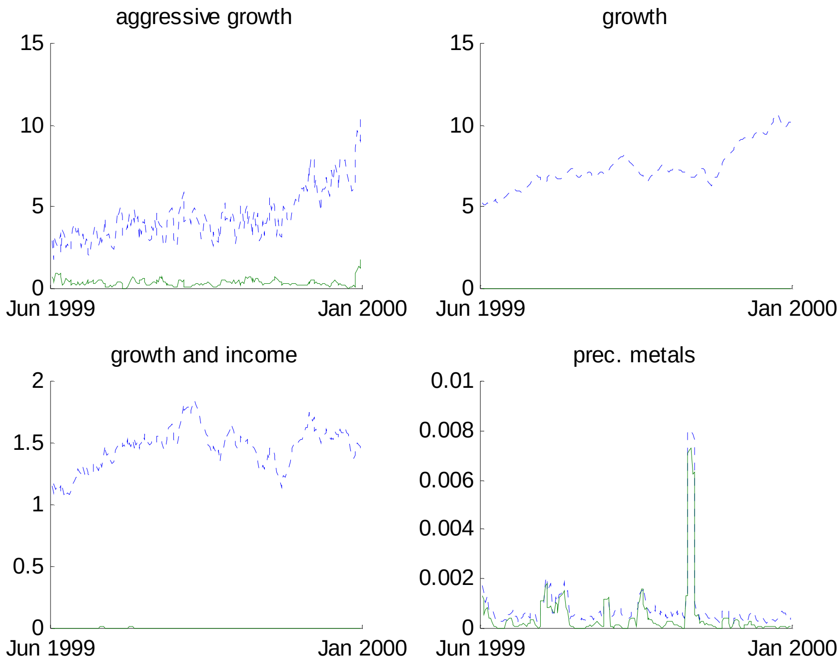

Figure 2: Estimates of VaR with and without Incorporating Liquidity Risk