

Adjusting value-at-risk for market liquidity

Liquidity remains a key risk factor in many portfolios, but quantifying it remains an open question. Here, David Cosandey offers a new macroeconomic approach to quantifying liquidity risk based upon trading volume, and incorporates it into VAR, testing his model against empirical data

The Asian and Russian crises of the late 1990s, and the Long-Term Capital Management (LTCM) debacle, demonstrated the need to better understand market liquidity risk. Several methodologies have been suggested that include market liquidity effects in value-at-risk. Some of them rely on bid-ask spreads (Bangia et al, 1999, and Monkkonen, 2000). These approaches raise the difficulty of gathering long time series of bid/ask spreads for different portfolio sizes and securities. Moreover, they are derived from a micro view of the market, whereas liquidity is typically a system-wide effect, crying for a macro view. Another approach (Krakovsky, 1999) has the downside of relying on an empirical parameter, L , which has to be assessed by the traders themselves – no sure basis for risk control.

This paper aims to introduce a new and simple methodology for calculating a market VAR that includes liquidity effects. It relies on trading volume as the central parameter. The scheme has the threefold advantage of being intuitive, relying on accessible information and deriving from a macro view of the market.

In this paper, liquidity risk means market liquidity risk, as opposed to balance-sheet liquidity risk – that is, we shall talk of the danger that selling a large position may induce a price drop, due to insufficient absorbance capacity in the market.

The liquidity issue

In many markets outside of the very large stock exchanges, liquidity is an issue – whether for western European non-blue-chip corporate bonds or Asian convertible bonds. It is even more of an issue in very thin markets, such as Russian equity, where trading volumes typically are a thousand times smaller than at the New York Stock Exchange. But even high-volume markets, such as the NYSE, are affected by liquidity risk in times of rare but intense crises – the realm of the risk manager. It hurts then all the more since it is unexpected. The giant yen/dollar market suffered sharp drops in liquidity during short, but intense, bouts of panic, such as October 1998, which left investors helpless and spreads yawning.

In table A, we give a few indicative figures, as a reminder of

how varied volumes can be around the world and in different periods. The Russian figures are particularly low, even when account is taken of the fact that only about a quarter of all transactions are reported to the Russian Trading System (RTS). Quickly unwinding a \$50 million portfolio of shares would trigger a larger adverse effect on prices at RTS than at the NYSE. The aim is to quantify this difference.

Assessing the potential impact of selling

In standard VAR, one records past daily price fluctuations, measures their theoretical effect on portfolio value and selects the 99% worst case. This is done with effective historical price changes, ie, when the portfolio was not sold. To adjust market VAR for liquidity effects, one has instead to estimate the daily price changes that would have occurred, given market depth of the day, if one had sold the portfolio. These corrected returns should then include both components: the ‘normal’ price move of the day and the liquidity-driven, self-induced price move. The 99th percentile should be selected from these potential price changes.

To assess the impact on price of selling a portfolio is best done by working with supply and demand curves. By knowing these curves exactly in real time, then aggregating them over the time horizon (say, one day), the risk controller could fathom the impact that selling a portfolio would have, within any given day. They would just need to displace the supply curve by the number of securities that they intend to sell (in figure 1, shifting up the supply curve by a distance ΔN along the vertical axis).

However, getting to know the supply and demand curves can be difficult. Subscribers to some electronic exchanges, such as the Swiss Stock Exchange, can read real-time outstanding bids and offers, complete with price and volume, for the equity markets. But this cannot be expected to be the case for all instruments and locations. In the currency markets, no data is generally available. Also, one would need to know the evolution over time of the demand and supply curves’ shapes, as for yield curves, in order to derive VAR figures.

In other words, the pragmatic VAR specialist cannot rely on exact supply and demand curves to assess liquidity risk. An approximation is needed. The approximation that we suggest here is to assume that the amount of money available on the counterparty (buying) side over the trading day remains the same in case of selling, and that the amount of securities offered on the supply side rises by exactly the portfolio sold.

Transaction exchange triangle

The situation can be expressed geometrically with a triangle (see figure 2). The total amount, A , paid for the securities in the given trading day and the total number, N , of securities sold on that day are the right-angled sides of the triangle. The trade-weighted average

price, P , of the share on this day is the slope of the triangle's hypotenuse. This describes the situation when the portfolio is not sold.

If a supplementary portfolio containing ΔN shares is sold within the same day, a new exchange triangle forms, with A and $N + \Delta N$ as the right-angle sides. The new mean price of the day is given by the slope of the new hypotenuse $P' = A/(N + \Delta N)$. As a consequence of the simplifying assumptions, only ΔN contributes to the price change triggered by the portfolio sale (no more money comes out on the demand side, ie, $\Delta A = 0$).

Supply and demand curves

In figure 1, these simplifying assumptions are translated graphically by drawing the demand curve as a function inversely proportional to price P . The number N of demanded securities is given by the formula $N = (N_0 \times P_0)/P$, where $N_0 \times P_0$ is the observed trading volume in currency units, without the investor selling. The supply curve is approximated as a horizontal line, shifted by the portfolio size ΔN when the investor sells.

The supply and demand curves in figure 1 are aggregated daily curves. Their intersection point has co-ordinates P_0 , the trade-weighted mean price of the given day, and N_0 , the daily trading volume expressed in number of securities.

These assumptions about the supply and demand curves are conservative. They should not underestimate liquidity-induced price effects because the demand curve is actually expected to rise more, and the supply curve is expected to go down, for falling asset prices. They are practical, because the two figures (quantity of money and securities exchanged without the investor selling) are known: they are the trading volume of the day, in currency units and in number of securities. Trading volumes can be found for many asset classes. Daily US, European and Asian equity and bond figures are publicly released. The foreign exchange markets are trickier, but figures are released occasionally (in particular, the Bank for International Settlements' surveys).

Quantifying liquidity-driven returns

With these assumptions, we can now calculate the liquidity-driven price falls that selling a portfolio within any day would cause, given market depth. These liquidity-driven returns can be used for deriving a pure liquidity VAR. We shall then include both effects, liquidity-driven price moves and general market moves (those used in standard VAR), in order to get a complete, liquidity-adjusted, market VAR.

For the first step, the equations are:

$$P = \frac{A}{N} \quad P' = P + \Delta P = \frac{A}{N + \Delta N} \Rightarrow \quad (1)$$


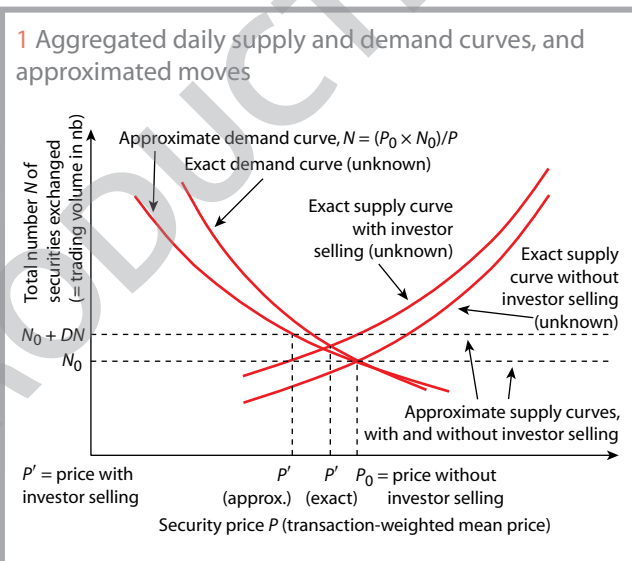
$$\Delta P = \frac{-A \times \Delta N}{N \times (N + \Delta N)} \quad \frac{\Delta P}{P} = \frac{-\Delta N}{N + \Delta N}$$

where A is the trading volume on a given day, in currency units, without investor selling; N is the trading volume on the same day, in number of securities, without investor selling; ΔN is the size of the investor portfolio (in number of securities) to be sold; P is the trade-weighted security price on a given day without investor selling; P' is the trade-weighted security price on the same day with investor selling; and ΔP is the liquidity adverse effect on price, ie, the difference between P and P' .

Equation (1) implies that liquidity effects are felt when the portfolio size, ΔN , starts to be non-negligible compared with trading volume, N . In extreme cases, when the portfolio to be unwound is much larger than market depth, as for LTCM in September 1998, the equation shows, as expected, that the asset price

A. Mean daily traded volumes, averages over indicated periods

Market	Daily volume (\$ million)
Russian Stock Exchange (RTS), February 1999	5
Russian Stock Exchange (RTS), 2000	23
Poland Stock Market, February 1999	50
Switzerland Stock Exchange (SWX), 1999	2,200
Taiwan Stock Exchange (TSEC), 1999	3,400
New York Stock Exchange (NYSE), 2000	44,000
US Treasury bonds, December 1999	186,000



could sharply drop to almost zero, effectively forbidding selling. Within this formalism, one could also show the inverse effect, ie, letting prices rise by massively entering a shallow market. An investor suddenly entering into Russian equities with millions to invest could have a terrible time, multiplying the prices, but it would be difficult to get out unscathed. Indeed, buying a dollar amount ΔA under a trading volume A would, according to the same formalism, raise prices by an amount of about $\Delta P/P = +\Delta A/A$.

Aggregating liquidity and market returns

Now we combine liquidity-induced negative returns of equation (1) with the daily 'general market moves', ie, the ones used for

B. Liquidity-adjusted 99% value-at-risk and 99% shortfall for different portfolio sizes on the Swiss and Russian stock markets

Measuring liquidity risk	Standard without liquidity effects*	Liquidity-adjusted	
		Portfolio of Sfr100,000	Portfolio of Sfr1 million
CSGN shares: 99% VAR	-6.8%	-7.3%	-7.4%
CSGN shares: 99% shortfall	-8.4%	-9.2%	-9.4%
LKOH shares: 99% VAR	-15.6%	-30.2%	-80.8%
LKOH shares: 99% shortfall	-20.4%	-49.8%	-88.3%

All figures given in relation to portfolio size
 * Equivalent to very small position with liquidity effects

standard market VAR, the historical returns. The calculation is straightforward. On a given day, we have an average price P_0 and a volume N_0 , so:

$$P_0 = \frac{A_0}{N_0} \quad P_1' = P_0 + \Delta P = \frac{(N_0 P_0) + N_0 \times \Delta P_{market}}{N_0 + \Delta N} \Rightarrow$$

$$\frac{\Delta P}{P_0} = \frac{N_0^2 \times \Delta P_{market} - (N_0 P_0) \times \Delta N}{(N_0 P_0) \times (N_0 + \Delta N)} \quad (2)$$

where A_0 is trading volume on initial day (in currency units) $A_0 = P_0 N_0$; N_0 is trading volume on the initial day (in number of securities); ΔN is the size of the portfolio to be sold by the investor (in number of securities); P_0 is the trade-weighted security price on the initial day; P_1' is the trade-weighted security price on the next day (when the investor has sold); ΔP is the security price change between the initial day and the next day after the investor has sold; and ΔP_{market} is the security price change between the initial day and the next day if the investor does not sell.

When the general market return is set to zero, one sees that equation (2) reverts to equation (1), as desired. Thus, we now have an expression giving both components of market risk: the return occurring when the portfolio is not sold, and the liquidity-driven return occurring when the portfolio is sold.

This approach is not too dissimilar to Krakovsky's (1999). A close inspection of Krakovsky's equation (2) and extensions reveals that the parameter L can be identified as the trading volume – expressed in number of securities, with the to-be-sold portfolio included, ie, $L = N + \Delta N$. If this interpretation is correct, there is no need for the empirical parameter L to be assessed by traders. In his paper, Krakovsky allows his parameter L to vary with portfolio size, but without showing how to model this dependence. This dependence is clear in our formalism.

Liquidity-adjusted VAR: empirical results

From the new, liquidity-adjusted returns, it is straightforward to derive a liquidity-adjusted VAR. As an illustration, we calculate VAR for two portfolios, one in a liquid market and the other in an illiquid market, using historical VAR simulations.

We synthetically built two single-stock portfolios, one of Swiss stocks, denominated in Swiss francs, and the second of Russian stocks, denominated in roubles. Both portfolios were worth Sfr100,000 on December 1, 2000. The first portfolio contained 331 Zurich-traded shares of the financial services company Credit Suisse Group (CSGN). The second portfolio contained 6,306 Moscow-traded shares of Russian oil concern OAO Lukoil (LKOH). We calculated liquidity-adjusted historical returns from equation (2) and derived the 99th percentile of the return

distribution (the VAR) and the corresponding shortfall (tail expectation). We then compared these figures with standard VAR to assess the impact of liquidity. In a second step, we repeated the calculation with portfolios 10 times larger (worth Sfr1 million on December 1, 2000), ie, comprising 3,310 CSGN shares and 63,060 LKOH shares, to make clear that liquidity-adjusted VAR is dependent on portfolio size.

We used equity price and trading volume data from July 11, 1997 to November 20, 2000, ie, 882 days. The prices in the time series were close values – an approximation for daily mean prices. There was no need to test multi-instrument portfolios since VAR aggregates the same way with liquidity-adjustment.

We selected stocks with very different market liquidities. CSGN stocks traded a median daily volume of Sfr270 million during our time interval. Extreme values were Sfr50 million and Sfr2,700 million. Median daily trading volume for Lukoil shares was Sfr4 million, with extreme values of Sfr0.01 million and Sfr60 million.

The Swiss and Russian portfolios, equal in value at the end of the period (December 2000), need not have been so over the three-year period, of course. The number of shares was kept constant. The aim was just to build portfolios of roughly comparable magnitudes.

The VAR figures we calculated here are local-currency based, ie, they are Swiss franc- and rouble-based. The results are shown in table B. The different profit/loss distributions, with and without liquidity effects, constructed to extract the VAR and shortfall figures are displayed in figures 3 and 4.

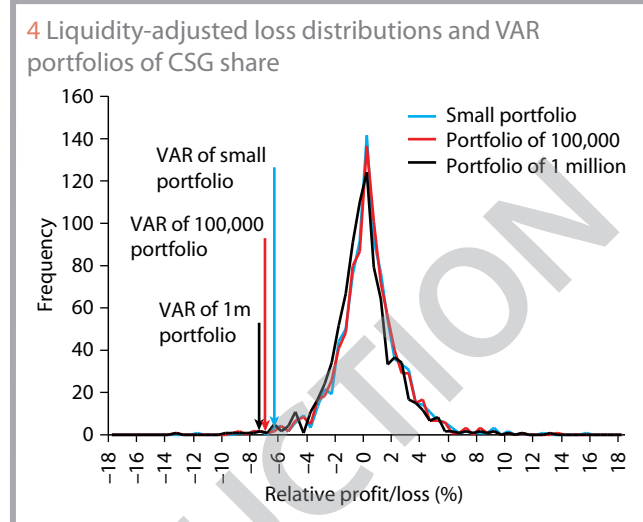

As expected, liquidity yields much more impact with larger portfolios and thinner markets. It is felt more heavily on the Russian Stock Exchange than on the Swiss one, and is measurably larger for bigger portfolios. The sale of the Sfr100,000 and Sfr1 million portfolio has a negligible effect on prices on the Credit Suisse shares market, as is clear when bearing in mind the typical trading volume of 270 million. But liquidity effects cannot be neglected for the same portfolio sizes on the Lukoil shares market with 4 million volume. Liquidity risk dramatically compounds Russian VAR figures. With standard VAR, 16% of the second portfolio value is at risk, whereas with liquidity considered, 81% of the value is at risk!

Not only is the liquidity effect heavier in the shallower market, it also upsets the structure in the profit/loss distribution tail. In the Russian market, worse returns in the tail of the distribution were not the same with and without liquidity effects. In the Swiss market, the relative order of the price returns in the tail stayed unchanged. This implies that, at least in thin markets, liquidity has its own, autonomous dynamic. It does not merely pull down all returns by amounts proportional to market moves. On some days, liquidity falls when prices do not fall so much, and, inversely, liquidity might fall much less than prices on certain times. Therefore, a separate modelling of liquidity makes sense.

The VAR of the small portfolio corresponds to a VAR without liquidity effects, ie, to standard VAR.

Monte Carlo versus historical simulations

Historical simulations seem to offer some advantages over Monte Carlo simulations, as far as liquidity-adjusted VAR is concerned. For Monte Carlo simulations, a mean value and a standard deviation of trading volume has to be calculated, as for other risk factors. An assumption then has to be made about the distribution function of the trading volume. This is difficult because that function is unknown and needs not be similar to that of the returns. The correlation of volume with other risk factors has to be derived as well. To avoid the well-known unstable correlations problem (the danger



that correlations are higher in crisis times), a safe choice could be to assume correlations of one between price and liquidity.

Historical simulations offer the clear advantage of correctly mirroring correlation between market risk and liquidity risk, and the distribution function of the trading volume.

Highly illiquid instruments

For relatively liquid instruments, such as the main currencies and equities, the one-day time horizon works fine to obtain a VAR figure. For relatively illiquid instruments, such as non-blue-chip corporate bonds, which change hands a few times a month at most, one should divide the monthly volume by 20 to obtain a proxy for daily data. Of course, this reduces the observed volume volatility. In historical simulations, this leads several successive daily returns to be shifted by the same amount, thus smoothing out extreme moves and hence shortening the profit/loss distribution tail. On the other hand, using effective daily volumes would make no sense. Zero-volume days would be modelled as 100% negative returns in case of portfolio selling, which would incorrectly inflate the VAR figures. In any case, the problem of zero-volume days is not specific to liquidity-adjusted VAR; since they lack a price, these days represent a problem in standard VAR as well.

As a better fix, one may expand the VAR time horizon so that each time window always contains at least one trade.

Time scaling of VAR

The conversion of one-day into 10-day VAR figures is mostly done at banks through multiplying daily VAR figures by the square root of 10, following the rules of the random walk model. Liquidity-adjusted VAR scales differently in time. The purely liquidity component gets smaller for larger time intervals, in relative terms, because it is inversely proportional to time.

To calculate 10-day liquidity-adjusted VAR, one should combine the empirically observed 10-day standard returns with the 10-day liquidity returns (obtained by adding the trading volumes of 10 succeeding days). This would be the proper, direct way, but would require much more data, extending further back into the past.

As an approximation, one might multiply the daily market returns, $\Delta P_{\text{market}}/P$, by \sqrt{t} (in our case, $t = 10$) and the trading volumes, N , by t in equation (2). When these changes are done, it is seen in equation (2) that the liquidity component of the return is proportional to $1/t$. As the market component of the returns is proportional to \sqrt{t} for

larger time horizons, liquidity-adjusted VAR tends towards standard VAR. This makes sense: the longer the time available to unwind a position, the less the loss suffered through volume effects.

Further refinements

Further refinements could include smarter modelling of the supply and demand curves. Typically, one could assume that the demand curve rises faster for lower prices, in order to obtain smaller liquidity adjustments to returns. Instead of the function $N = N_0 \times (P_0/P)$ one could set, eg, $N = N_0 \times (P_0/P)^a$ where the parameter a would be some positive constant. Such an improvement would be difficult to calibrate, though. The constant a should then be determined through empirical studies, in a way similar to what is done in Persaud (2000). It might change with time and might be different for each instrument.

Concluding remarks

The strange thing about liquidity risk is that, when disaster strikes and trading volume shrinks, market theories no longer apply since the market itself vanishes. Market specialists understandably dislike contemplating market disappearance, but it must be modelled all the same.

The approach presented here has the clear advantages of being intuitive, easy to implement and conservative. It relies on the trading volume, which is a macro measure of the market. It is, of course, a first approximation but, in new territory, first approximations are what is needed by risk managers. ■

This article first appeared in *Risk* in October 2001. At the time, David Cosandey was head of asset/liability management at Zurich Cantonal Bank. He is now a risk manager at Bank Morgan Stanley in Zurich. Email: david.cosandey@morganstanley.com

References

- | | |
|---|---|
| Bangia A, F Diebold, T Schuermann and J Stroughair, 1999
<i>Liquidity on the outside</i>
<i>Risk</i> June 1999, pages 68–73 | Monkkonen H, 2000
<i>Margining the spread</i>
<i>Risk</i> October 2000, pages 109–112 |
| Kravovsky A, 1999
<i>Gap risk in credit trading</i>
<i>Risk</i> March 1999, pages 65–67 | Persaud A, 2000
<i>The liquidity puzzle</i>
<i>Risk</i> June 2000, pages 64–66 |