STANDARD &POOR'S

RATINGSDIRECT®

April 15, 2008

Trading Losses At Financial Institutions Underscore Need For Greater Market Risk Capital

Primary Credit Analysts:
Thierry Grunspan, Paris (33) 1-4420-6739; thierry_grunspan@standardandpoors.com
Nick Hill, London (44) 20-7176-7216; nick_hill@standardandpoors.com
Tanya Azarchs, New York (1) 212-438-7365; tanya_azarchs@standardandpoors.com

Secondary Credit Analysts:
Bernard de Longevialle, Paris (33) 1-4420-7334; bernard_delongevialle@standardandpoors.com
Charles D Rauch, New York (1) 212-438-7401; charles_rauch@standardandpoors.com
Craig Bennett, Melbourne (61) 3-9631-2098; craig_bennett@standardandpoors.com

Additional Secondary Credit Analysts:
Elie Heriard Dubreuil, Paris (33) 1 4420 6748; elie_heriard_dubreuil@standardandpoors.com
Yuri Yoshida, Tokyo (81) 3-4550-8581; yuri_yoshida@standardandpoors.com
Carina Lopez, Buenos Aires (54) 11-4891-2118; carina_lopez@standardandpoors.com

Table Of Contents

642798 | 300004403

Table Of Contents (cont.)

Trading Losses At Financial Institutions Underscore Need For Greater Market Risk Capital

The securities markets changed dramatically in 2007, shaking the trading businesses of banks and showing up in their risk measurements. The main metric, the aptly named value at risk (VAR), was rising in conjunction with soaring market volatility. VAR estimates maximum loss for a certain time period--for instance, one-day--to a given confidence interval--such as 99%. However, many banks posted losses much higher than VAR and even greater than their regulatory requirements for the capital they need to hold against market risks.

This situation illustrates the shortcomings of VAR models. Most notably, they are designed to predict losses under normal trading conditions. In addition, they ignore or underestimate certain risks, notably the increasing amounts of idiosyncratic risk arising from new and complex financial instruments that are a feature of today's trading desks. However, this doesn't mean that VAR is useless. It provides valuable information under most market conditions. Understood properly and with the appropriate caveats, Standard & Poor's Ratings Services believes VAR remains an important risk management tool.

To better reflect the magnitude of trading portfolios' underlying risks, we envisage making a series of upward adjustments to capital requirements under Basel II as part of the calculation of our proposed risk-adjusted capital ratio (RAC) (see "Criteria: Request For Comment: Standard & Poor's Risk-Adjusted Capital Framework For Financial Institutions," published April 15, 2008, on RatingsDirect). Our approach to assigning capital to traded market risk is twofold. First, we scale up regulatory VAR capital by a factor of three, to capture a longer holding period and a higher confidence interval as well as the "fat tail" outlying risks that characterize assets traded on the securities markets. Secondly, we use a higher charge for operational risk than under Basel II's standardized approach, which helps compensate for "model risk" for capital based on VAR models and the risk of fraud on trading desks. We believe the resulting capital charge will more closely approximate stress losses and banks' own estimates of economic capital.

The wide variation in banks' trading losses in recent months highlights the need for greater disclosure regarding their trading portfolios, which are often very large in nominal terms but are relatively opaque. Under the RAC's specific charge adjustments (SCA) for estimating the capital needs of a trading business, we will seek more details from financial institutions about regulatory capital requirements and specific risk add-ons, the extent of high credit risk components in trading books, asset liquidity, and trading strategy concentrations. We will adjust our charges accordingly.

Banks Post Huge Losses In Comparison With VAR Estimates

In the second half of 2007, several banks posted trading losses that were many times greater than their VAR estimates. To be sure, VAR had soared due to high historic volatility, but still grossly underestimated market risks. As a result, banks reported a large number of "backtesting exceptions"--the number of instances when the actual loss exceeded VAR for that day (see chart 1). That number statistically shouldn't exceed between two and three times a year for VAR at the 99% confidence level (see appendix 1 for details and terminology about VAR). The large number of VAR excesses in the second half casts doubt on the robustness of VAR models in stress conditions.

Moreover, banks generally express backtesting exceptions in real profit and loss (P&L). But if expressed in "hypothetical P&L," they would have been much higher. In our view, hypothetical P&L is the best way to assess the robustness of VAR models.

Chart 1

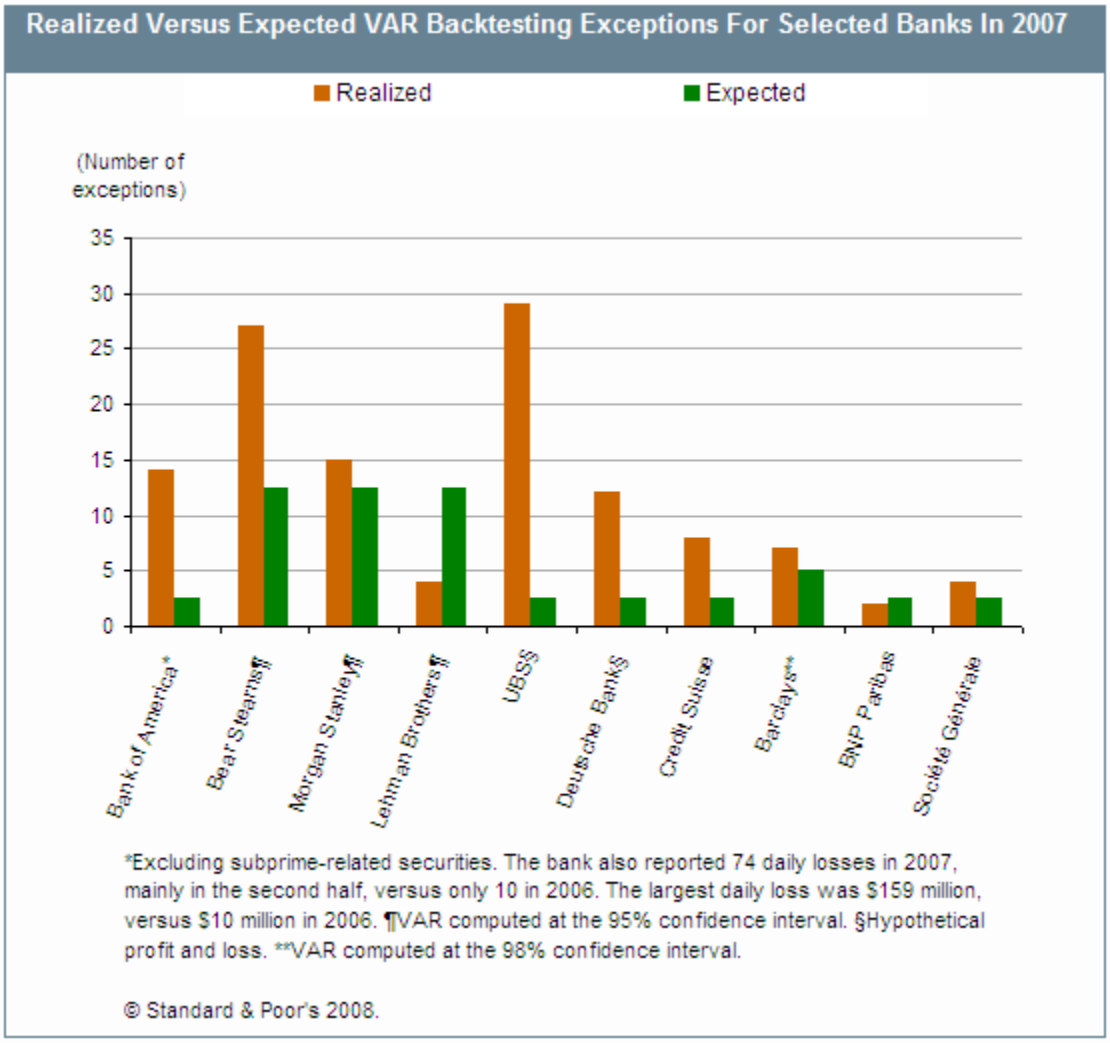

Even after excluding subprime-related securities in the calculation of VAR and the trading business' P&L, Bank of America N.A. reported 14 backtesting exceptions for the one-day 99%. In the third quarter of 2007, Deutsche Bank AG recorded trading losses of €1.6 billion, of which about €900 million was imputable to designated proprietary trading and other trading positions excluding collateralized debt obligations (CDOs) and residential mortgage-backed securities (RMBS). In addition to the €1.6 billion, the bank recorded mark-to-market losses of €759 million on leveraged loans and loan commitments. These examples show that VAR excesses in the third and fourth quarter 2007 are not solely the result of subprime exposures--the metric underestimates market risk for a broad range of assets.

Banks' own estimates of economic capital requirements for market risk support our approach for scaling up regulatory capital requirements for market risk. For example, Deutsche Bank's published estimate of the capital

required for traded market risk was €1.8 billion at the end of 2007. Meanwhile, its regulatory market risk-weighted assets (RWA) were €14 billion, implying a regulatory charge of €1.1 billion--40% less than its own economic capital estimate.

We've Adopted A Hybrid Approach For Determining Capital Adequacy For Market Risk In the Trading Book

Assigning capital to market risk is a difficult task. Ideally, one would construct a common set of plausible but severe scenarios and use these as a common framework for determining stress losses. This is theoretically possible, but to ensure consistency, a great number of assumptions would have to be fixed, and financial institutions would need to conduct extensive bespoke work. Even then it would be hard to achieve true comparability due to the complexity of the relationships among different traded assets.

Other methods are barely any more satisfactory. One option would be to use the absolute size of the trading book's balance sheet as a proxy for risk. However, size generally bears little relation to risk because it doesn't distinguish among different instruments. More importantly, the definition of the trading book varies widely among institutions and accounting regimes. For example, Deutsche Bank had trading assets under U.S. Generally Agreed Accounting Principles (GAAP) of €448 billion on Jan. 1, 2006; by contrast, under International Financial Reporting Standards (IFRS) it had trading assets at fair value of €1.009 trillion, and total balance sheet assets were 44% higher. A large part of the difference stems from the gross-up presentation of derivatives and repurchase agreements under IFRS, versus their net presentation under U.S. GAAP. Furthermore, banks hold a considerable amount of market risk off balance sheet. Given these issues, we've adopted a hybrid approach for determining capital adequacy for market risk in the trading book. We combine a VAR-based measure with a revenue-based charge for operational risk and some specific add-ons to reflect other risks not well captured--such as specific risk or trading asset liquidity risk.

Traded risk

Despite the inconsistencies, problems, and incompleteness of the measure, we believe that VAR validated by regulators remains one of the most valuable pieces of public information that banks disclose to the market. The validation process ensures a certain level of quality and consistency. In addition, VAR is based on current trading positions rather than past exposures. Plus, banks use it for daily risk management purposes. And some of the problems with VAR are partly captured by upward regulatory adjustments--such as the regulatory scaling factor and the backtesting add-on--as part of the validation process. That's why our approach is to take the regulatory VAR capital charge as our starting point, consistent with our practice of using regulatory Basel II data as our starting point for credit risk estimation. Then, to be consistent with the one-year, 99.9% confidence interval that we use in our RAC analysis, we apply a scaling factor of 3 to the VAR-based regulatory capital charge. Expressed differently, the regulatory approach itself is 3 x 10-day 99% VAR (assuming a scaling factor of 3), which is roughly equivalent to three-month 99% VAR (using the square root of time rule). By multiplying the regulatory capital charge by 3, our method can be said to be broadly equivalent to 33 x one-day 99% VAR or to one-year 99.9% VAR (the passage from 99% to 99.9% VAR is done assuming fatter tails than typical log-normal distribution implies).

Operational and model risk

The RAC also envisages a higher operational risk charge for trading activities and investment banking--set at 25% of revenues versus 18% for the Basel II regulatory standardized approach. Furthermore, we base our capital charge for operational risk on the maximum of the last three years of gross revenues, versus the average of the last three

years for the standardized approach. This reflects our view that operational issues, such as fraud, can be material in investment banking, as recent cases have shown. In addition, this charge also represents a proxy for model risk and associated shortcomings of VAR for our estimates of market risk capital.

The combined trading risk and operational risk charges are a little more than twice the estimated regulatory capital requirement (see table 1).

Table 1

Estimated Regulatory Capital Versus Standard & Poor's Risk-Adjusted Capital For Selected Banks

(Bil. $)	2007 revenues	2006 revenues	2005 revenues	10-day 99% VAR for 2007 (mil. $)	Estimated regulatory capital charge	Standard & Poor's risk-adjusted capital
UBS Investment Bank	(0.4)	18.7	15.0	462.9	3.8	10.2
Merrill Lynch Global Markets & Investment Banking	(7.6)	14.3	10.2	372.9	2.5	8.0
Citigroup Markets and Banking	2.7	21.2	19.0	449.0	4.4	10.7
Bear Stearns Capital Markets Inc.	3.9	7.3	5.7	148.7	1.6	3.6
Credit Suisse Investment Banking	16.4	17.7	13.4	310.8	4.1	8.1
Morgan Stanley Institutional Securities	16.1	21.1	15.5	390.9	4.7	10.0
Société Générale Corporate & Investment Banking	6.4	10.0	8.2	193.4	2.2	4.8
Deutsche Bank Corporate and Investment Bank	23.6	23.7	22.7	387.1	5.8	10.6
Bank of America Capital Markets and Advisory Services	0.3	8.5	6.8	166.3	1.6	4.1
Totals	61.4	142.5	116.5	2,882.1	30.7	70.2

Excluding capital requirements for counterparty risks in the trading book. §Assuming a scaling factor of 3 and that the operational risk charge is computed according to the standardized approach (18% of the average of the last three years of revenues). VAR--Value at risk.

Bank losses in second-half 2007 varied considerably, which tells us that more factors are at play than financial institutions currently disclose. For this reason, we'll be asking them for more information about their trading risk for use under the RAC's specific charge adjustments (SCA), for example:

- A breakdown of the regulatory capital charge for market risk and the nature of add-ons,
- The nature and extent of credit risk in the trading book,
- The liquidity of trading assets,
- Indicators of trading strategy concentration, and
- Other market risk metrics such as stress test results and expected shortfalls.

For reasons of transparency, we encourage banks to improve their public disclosure of such risk factors to enhance market discipline--as called for under Pillar 3 of Basel II.

VAR Is Not A Measure Of Losses In Times Of Stress, But Rather Of Large Losses Under Normal Trading Conditions

Because it looks into the rear-view mirror, VAR can't predict the consequences of the stress that might lie ahead

Parametric, historical, and Monte Carlo VAR all are based on historical data on risk factors. They therefore may provide limited insight into risks that aren't present in the model's "time window." In particular, if the time window doesn't encompass any periods of illiquidity, VAR will fail to produce a relevant measure of risk on some positions. Let's take the example of positions taken against the TED spread--here, the difference between the three-month London interbank offered rate (LIBOR) and the interest rate on the three-month U.S. Treasury bill. Although the spread didn't vary much before mid-2007, it then widened considerably because of the liquidity crunch (see chart 2). During the first weeks of August 2007, a bank exposed to a widening of the spread would have incurred losses that VAR models couldn't have predicted.

Positions on pegged currencies are another example of when historical data may fail to deliver a good VAR estimate. If a bank holds speculative positions betting that a peg between two currencies isn't going to be relaxed, VAR can't estimate the market risk satisfactorily, since the relevant risk factor--the exchange rate between the two currencies--shows no sign of movement at all during the time window. However, if the peg is actually relaxed, losses could be significant. Typically stress tests can estimate these risks but VAR estimates can't.

Chart 2

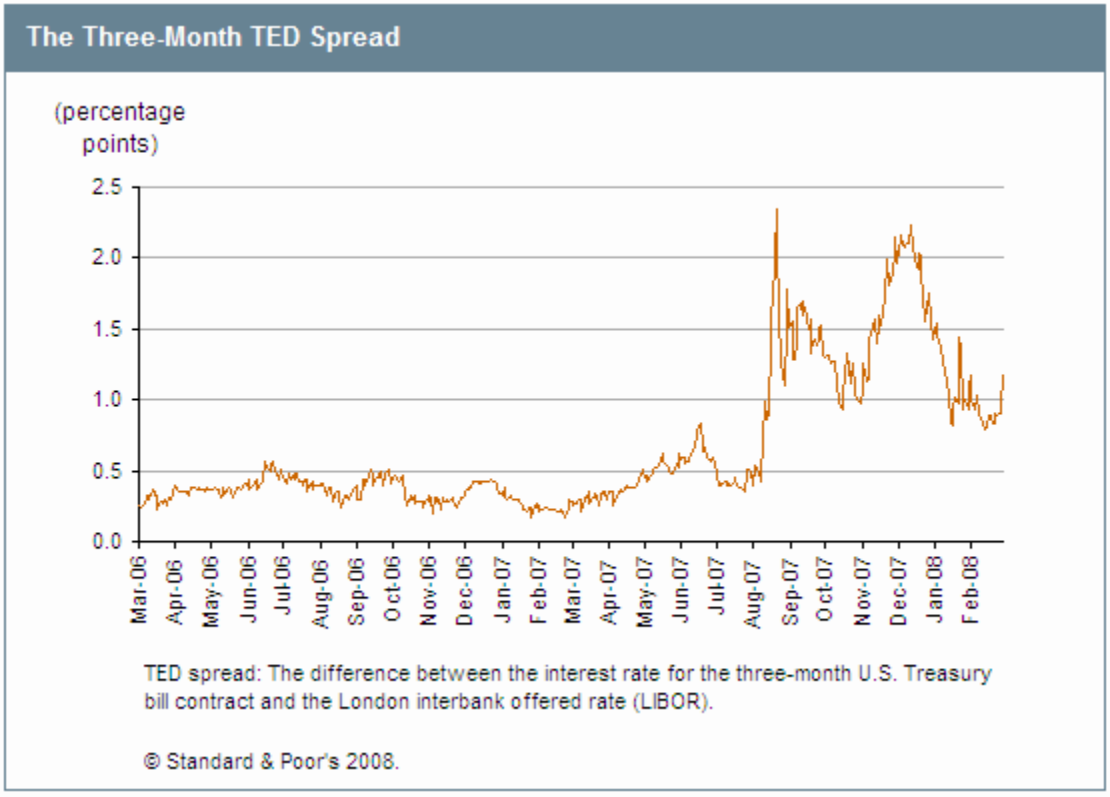

In a stress scenario, the 10-day horizon is inadequate

In some circumstances and because of illiquidity problems, some positions can take several days or weeks to unwind, compared with only a couple of hours in a normal market. When a bank faces pockets of illiquidity for some of its assets--such as certain structured products today--the 10-day horizon is inadequate.

In times of stress, asset returns exhibit possible autocorrelations, and correlations among risk factors change

In times of stress, bad days tend to cluster. However, the square root of time rule, which banks use to scale up one-day VAR into regulatory 10-day VAR, assumes that daily returns are independent across time. In this respect, disclosures by UBS AG provide a useful example. The bank publicly reports two VAR measures for the same confidence level: the first at a one-day horizon and the second at a 10-day horizon. The second is not derived from the one-day VAR, but computed separately using 10-day shocks for all the relevant risk factors. In the fourth quarter of 2007, 10-day VAR was 40% higher than what one-day VAR and the square root of time would suggest. Based on the same time window of five years of historical data, one-day VAR at the 99% confidence level was Swiss francs (CHF) 149 million and 10-day 99% VAR was CHF665 million. Independent daily returns across time would have resulted in 10-day VAR of CHF471 million (the square root of 10 x 149). Note that on average in the second quarter of 2007, one-day VAR and 10-day VAR figures were perfectly compatible with the independent return hypothesis.

Moreover, most VAR models use correlations among risk factors that aren't stressed. In a crisis however, correlations change and the benefits of asset diversification in the trading portfolio may be overestimated. In other words, when market conditions deteriorate, extreme movements can occur in all risk categories simultaneously.

By not looking beyond the 99th percentile, VAR, and hence the regulatory capital charge, fail to capture risks on some products

For "gap risk" products (like constant proportion portfolio insurance, or CPPI), the allocation rule between nonrisky and risky assets is such that there is no risk for the bank unless the risky asset--sometimes a fund of funds--suffers a drop in value between two reallocations by more than the cushion (see appendix 2). In most cases, the size of the cushion is sufficient for 10-day 99% VAR--and hence the regulatory capital charge--to be negligible. However, 10-day VAR at a much higher confidence level (say 99.9%) can be very large. The risk for the bank stems mainly from the capital guarantee attached to the product. The regulatory capital charge doesn't capture this extreme event, which could materialize in the case of an overall collapse of the hedge fund market. In terms of stress tests or of expected shortfalls, gap risk products represent a much higher share of the risk than what VAR typically demonstrates.

Even Under Normal Trading Conditions, VAR Can Underestimate Market Risk

A notable source of risk is often not chosen as a risk factor in VAR

VAR models can't capture "basis risk" between two market parameters if both are "mapped" or linked to the same risk factor. Consider a simple negative basis trade strategy where a bank has a long position on a corporate bond, hedged by buying credit protection on the bond issuer through a credit default swap (CDS). Most of the time, CDS margins and corporate credit spreads move in tandem so that these strategies are almost risk-free. When concerns over the possible default of a given corporation grow, the bank loses money on the bond but makes up the loss on the CDS. But when the market is more concerned about liquidity than about defaults, CDS margins increase less than credit spreads and the bank loses money. That's because banks need cash to buy a bond, and cash is hard to

come by in a liquidity squeeze. On the contrary, banks don't need cash to enter into a swap. This is what has happened during recent months when some investment banks took hits as a result of these strategies. However, for most banks, VAR models couldn't capture this risk since CDS margins and credit spreads for a given issuer--and a given maturity--were mapped onto the same risk factor and the strategies were risk-free from a VAR point of view. In general, in most VAR models, implicit volatility shocks for in-the-money and out-of-the-money equity options are assumed to match at-the-money (ATM) volatility shocks--for a given underlying and a given maturity. A market risk that VAR doesn't predict could materialize if there's a change in the "volatility smile." Likewise, ATM long-term volatility shocks are often derived from ATM short-term volatility shocks using some simplifying assumptions, and it's possible to construct a portfolio of risky optional positions with zero VAR.

Finally, some market parameters are assumed to remain constant over the VAR horizon. This includes parameters whose levels are essential for determining the mark-to-model value of positions--such as recovery rates for RMBS, CDOs, or equity correlations. Typically, through its business of selling equity basket options to customers, an investment bank loses money when correlation increases. This risk materializes during a crisis, but VAR won't capture it if correlation isn't modeled as a distinctive risk factor. For credit derivatives, most VAR models assume that the "correlation smile" remains constant. (A correlation smile materializes when the market prices the default correlation between the bonds backing a CDO at different levels depending on the attachment point of the tranche.) Everything else being constant, a long position in an equity tranche, hedged by a short position in a mezzanine tranche, will appear to be risk-free, whereas in reality these strategies have proved to be risky.

VAR pricers are sometimes too simplified to quantify risks properly

For computational reasons, VAR models use pricers to revalue positions that are often more simplified than the complex ones used to compute the daily trading book P&L. Under certain market conditions, this can underestimate the risk. For example, when Taylor expansions at order one are used to price exotic options in the VAR model, whereas full revaluation is used for vanilla options, the market risk on exotic desks would be grossly underestimated. VAR wouldn't capture the negative gamma effect on the exotic options sold to customers while it would get the benefit of the positive gamma effect on the hedges.

VAR underestimates specific risk

Some specific VAR models map bonds to different risk factors depending on the sector, maturity, and current level of the credit spread. Within the 10-day horizon of VAR, the credit spread could increase so that the bond should in theory migrate to a new risk factor, corresponding to a higher level of risk. However, most of the time VAR doesn't factor in this potential migration risk--let alone the default risk--over the 10-day horizon. Everything else being equal, VAR treats interest rate-specific risks as if issuer risk remained constant over the 10-day horizon.

Credit spreads can also widen for reasons unrelated to changes in the credit quality of the issuer. Lack of liquidity is one reason. Downgrades of guarantors, or concerns about the possible downgrade of a guarantor, are another. In our view, the Basel II regulatory "incremental default charge," which banks are expecting to implement in 2010, should cover not only default risk but also event risk, such as unexpected rating downgrades over the horizon of the VAR or credit spread changes that are unrelated to the credit quality of the issuer.

The 10-day horizon is irrelevant for nonliquid positions, even in normal times

The 10-day horizon is inappropriate for nonliquid positions like stakes in hedge funds, which can only be redeemed at certain dates provided the fund manager is notified in advance. Furthermore, the investor sometimes needs to keep his stakes for a certain time or "lock-up" period before he can sell them. Overall, it takes several weeks to

liquidate a hedge fund position. For this reason, the 10-day holding period underestimates risk on some structured products, like options on funds of funds or options on a basket of single hedge funds.

In today's complex markets, no single metric can accurately assess the risk exposure of banks. That doesn't mean, however, that an established measure like VAR isn't useful. What it does mean is that in assessing risk we must also look to additional metrics and data to build on the information that banks' VAR calculations reveal. For these reasons, we've adopted a hybrid model to better assess potential risks in more complex markets. We believe these steps provide a more accurate picture and greater transparency--a goal that more accurate and detailed public disclosure from banks would only further.

Appendix 1: Value At Risk In Detail

General principle

VAR 10-day 99% is the level of losses over a horizon of 10 days that statistically should only be exceeded 1% of the time. Thus, if a bank reports a 10-day 99% VAR of €100 million, it is estimating that 99% of the time a change in the value of its trading portfolio won't result in a decline of more than €100 million over the 10 days. It does not mean that the bank rules out any trading losses beyond €100 million over this horizon.

In practice most banks compute VAR at the one-day horizon and scale up the result to obtain a VAR at the 10-day horizon. Three main types of VAR are used: parametric, historical, and Monte Carlo. For each of them, the bank identifies a set of risk factors (such as the S&P 500 Index and the three-month LIBOR) at play in their trading portfolios. Changes in the value of these risk factors govern the change in mark-to-market of the portfolio. The idea is then to:

- Simulate the joint behavior of the risk factor shocks in different scenarios;
- Compute, for each of the generated scenarios, the new set of risk factors, for example the levels of the S&P 500 Index over the VAR horizon;
- Revalue the positions for each generated set of risk factors; and
- Compute VAR at the 99% confidence level, as the 99th percentile of the generated distribution of the changes in value of the portfolio.

The three different VAR methods

Parametric VAR can be expressed by a closed-form formula by making some simplified assumptions, such as that risk factor returns are multivariate Gaussian, and by omitting convexity effects in valuing the positions. Monte Carlo and historical VAR differ only in the way they model the joint behavior of the risk factor shocks. Historical VAR requires no particular modeling assumptions in the sense that risk factor shocks are the ones that actually occurred in the past--for a given time window, at least equal to one year. Monte Carlo VAR makes some assumptions regarding the univariate distribution and calibrates the parameters on historical data and the dependence among them.

To reduce computational complexity, VAR does not consider some relevant market parameters as risk factors. The model either maps them to existing risk factors or assumes they remain fixed over the VAR horizon.

General risk VAR and specific risk VAR

General risk VAR shocks only general risk factors--like risk-free interest rates and stock market indices--leaving specific risk factors unchanged. Specific risk VAR shocks only specific risk factors--like corporate credit spreads or

idiosyncratic single-stock returns--leaving general risk factors unchanged. Most banks compute only a global VAR combining general and specific risk factor shocks.

The backtesting exercise

Backtesting consists in comparing the VAR estimate with the P&L for a given day. For an internal model at a 99% confidence level to be accurate, we expect the daily VAR to be exceeded not more than between two or three times a year. In practice, VAR computed for a given day D predicts an amount of losses at unchanged positions--those at D-1 end of day--and cannot capture profits or losses generated by trading at D--including commissions and fees collected on client-driven trades. In our opinion, the robustness of the VAR model is better assessed by comparing VAR for D with the "hypothetical P&L"--the trading book P&L the bank would have recorded by simply holding positions of D-1 end of day until D at the end of day--and not the actual P&L.

The regulatory capital charge

The regulatory capital charge for general risk is computed as a multiple of the general risk for 10-day 99% VAR. The national regulator chooses the VAR scaling factor, but the minimum is 3. Depending on the results of the backtesting, an extra factor, between 0 and 1, is added to the scaling factor. The regulatory capital charge also includes a charge for specific risk, which banks may compute separately with a specific risk VAR (or by the standard approach). When a bank computes neither general risk VAR nor specific risk VAR but "global VAR" instead, the regulatory capital charge is given as:

scaling factor x global VAR 10-day, 99% + global VAR 10-day, 99% for the portfolios containing specific risk

The second term is global VAR for the trading portfolios containing some specific risk--this excludes portfolios such as index arbitrage portfolios.

Appendix 2: Constant Proportion Portfolio Insurance (CPPI)

The net asset value (NAV) of "gap risk" products (like CPPI) is allocated between a risky asset--for example, a mutual fund or a fund of funds--and a risk-free asset. In addition to having a return depending on the performance of the risky asset from inception to maturity, the investor benefits from a capital guarantee from the bank. Reallocations take place at some specified dates, typically every week or every month. At each reallocation date, the percentage of the NAV allocated to the risky asset is proportional to the gap between the total NAV and the floor. The floor is the minimum amount that, if totally invested at this date in the nonrisky asset, would allow the investor to fully recover its initial investment. More precisely, the percentage invested in the risky asset is equal to:

coeff x (NAV of the CPPI – floor)/NAV of the CPPI (where coeff is a coefficient decided between the bank and the investor)

Moreover, if at any reallocation date the NAV of the product is lower than the floor, the bank sells the risky asset and invests the proceeds in the nonrisky asset, leaving it with a loss. For example, if at a given reallocation date the risk-free rate is 5%, the residual maturity of the product is one year and the current NAV is €110, the percentage invested in the risky asset is:

coeff x (€110 – €95)/€110

For a coefficient of 5, the percentage invested in the risky asset is equal to 68%. Hence, €75 is going to be invested

in the risky asset and the remaining €35 will be invested in the nonrisky asset. If, at the next reallocation date, let's say in one month, the risky asset loses 20% of its value, the new NAV of the product becomes:

0.8 x €75 + €35 = €95

ignoring any interest earned on the nonrisky asset. This leaves just enough assets for the bank to invest the totality in the nonrisky asset and wait until maturity. If, however, the risky asset loses, say 30% of its value, the new NAV becomes:

0.7 x €75 + €35 = €87.50

The €87.50 is entirely invested in the nonrisky asset, leading to €92 at maturity (with a 5% risk-free rate) and an €8 loss for the bank. (Given the capital guarantee, the bank has to pay back €100 to its client.)

This example shows that the allocation rule insulates the bank from any market losses as long as the risky asset doesn't drop in value by more than 1/coefficient between two reallocation dates. However, the higher the coefficient, the riskier the product is for the bank and, given the capital guarantee provided to the client, losses can be large.

Typically, the probability that the risky asset drops in value by more than 1/coefficient between two reallocation dates is smaller than 1%. Otherwise, the bank would have chosen a smaller coefficient.