The Evolution of Counterparty Credit Risk Management

David M. Rowe

SunGard Trading and Risk Systems

Pre-settlement counterparty credit exposure came to the fore shortly before *Risk* Magazine was launched in 1987. It gained additional visibility when the Bank for International Settlements (BIS) insisted that such credit exposure be recognised in the first Basel Capital Accord. In the past 15 years there has been much progress in measuring and managing the associated risks, but the author argues that much remains to be done.

IGNORANCE IS BLISS – PRE-1985

It is said that success has many parents but failure is an orphan. Surely the massive expansion and pervasive influence of derivatives during the last two decades of the twentieth century stands as one of the great success stories of that period. True to form, there are many claimants to the title of first swap arranger. By now the truth is probably lost in the mist of time, but the date of the first swap appears to have been sometime in 1981. At that stage, these were truly arranged contracts. I have even heard of deal parties to celebrate each swap like those often held after the closing of large syndicated loans.

The primary form of credit risk associated with trading in those days was settlement risk. Recognition of this had been driven home painfully by the failure of Bankhaus Herstatt during the trading day on June 26, 1974.[1] Even in the early 1980s, daily turnover in the foreign exchange (FX) spot and forward markets was in the hundreds of billions of dollars. Such volumes gave rise to serious

concern about the credit risk of such contracts, from the time payment instructions became irrevocable to the time the agreed exchange amount was received. As the new swap dealers were quick to point out, however, settlement on most of their transactions was on a net basis, which meant there was no settlement risk.

Looking beyond the settlement risk issue, there was only limited understanding of the pre-settlement credit implications of term derivative transactions. Some people thought that since there was no exchange of principal on an interest rate swap there was no credit risk. After all, one cannot claim loss of future interest payments on a loan when a borrower defaults. The fact that most credit officers were thoroughly steeped in the historical cost accounting framework was an additional obstacle to a clear understanding of these mark-to-market instruments. Furthermore, derivative volumes were small and counterparties were virtually all investment grade names. These factors combined to make pre-settlement credit exposure a latent and largely ignored risk aspect of derivative trading in the early 1980s.

INCREASED MATERIALITY AND VISIBILITY

By 1986, swap market volumes had grown significantly and so had current mark-to-market exposures. Gradually the realisation spread that these positive market values represented material credit risky balance sheet assets similar to corporate and industrial bank loans. As such, counterparties to such trades clearly should be subjected to credit review and transactions should only be done within an approved credit exposure limit.[2] But there was one nagging problem. When making a traditional loan it doesn't require complex mathematical analysis to answer the question "how big is the loan?" The answer to that question is clearly stated in the proposed term sheet. While some would point out that the fair value of a loan fluctuates just as the price of a bond does, historical cost accounting treats the principal amount as fixed.

For swaps, there are two problems. First, they are usually done near par. Initial market values only reflect the slightly off-market impact of a dealer's bid/offer spread. This initial value, however, is only a small fraction of the potential future exposure that may materialise as a deal ages and market conditions change. Thus current exposure, especially current exposure on newly executed

deals, is not a realistic measure of the true potential loss from a downgrade or default on the part of the counterparty. Second, market fluctuations lead to constant revisions in the exposure, however measured. Such unstable exposure presented an entirely new dimension of uncertainty that many traditional credit control personnel found difficult to incorporate into their thinking.

THE FIRST BASEL CAPITAL ACCORD

The derivative credit exposure issue received added attention when the Basel Committee on Banking Supervision began to develop its initial rules for minimum regulatory capital requirements in 1986 and 1987. The committee realised that the current market values of a bank's derivative contracts did not capture their full potential credit exposure. As a result, they set out to derive reasonably simple rules to calibrate potential increases in such credit exposure, based on volatility in the underlying market data and the resulting impact on fair values of broad categories of trades. The result was the now well-known add-on approach to measuring potential future exposure.

Several comments are in order regarding the structure and parameters of the initial Basel add-on calculation.

a. The approach was intended to be simple to implement so that even small banks with a few hedge contracts could perform the calculation without difficulty.
b. It was only intended to capture the increase in *aggregate* exposure. It was neither proposed nor intended as a satisfactory approach to measuring individual counterparty exposures.
c. It provided only very limited recognition of the tenor of exposure by offering different parameters for contracts with more than or less than one year to maturity.
d. The same add-ons were applied to both at-the-money and away-from-the-money contracts, despite the amortisation toward zero in the value of many contracts as they approach maturity.
e. It treated each transaction in isolation. While this greatly simplified the mechanics of the calculation, it completely ignored the degree of co-variation in value among multiple deals.

f. Recognising that portfolio effects had been ignored, the parameters of the method were calibrated to reflect an assumed average degree of portfolio diversification.
g. While the potential increase in *aggregate* exposure was probably not unreasonable, the *marginal change* in potential exposure from adding or deleting a single deal was highly unreliable and could even be directionally incorrect.

A negative consequence of Basel's add-on method was to provide apparent regulatory sanction for this approach. While adequate for estimating aggregate potential exposure, it was, and is, quite unsuitable when applied at the individual counterparty level or to evaluating the marginal exposure of a new deal. The essential problem is the implicit assumption of an average degree of diversification. In fact, individual counterparty portfolios exhibit widely differing diversification characteristics. The most obvious and striking example is an offsetting transaction where the market value moves in a directly opposite direction to that of an existing trade in response to changing market conditions. If the two trades are not legally nettable, the second trade has almost zero impact on potential exposure. If they are legally nettable, then the second trade significantly reduces potential exposure. The add-on approach, however, will produce an unrealistic increase in measured exposure in both cases.

Despite these widely recognised shortcomings, most institutions adopted some variation in the add-on method for tracking and setting limits for counterparty credit exposure. This was driven by the comparative simplicity and modest cost of deploying such a system. Given the huge trading losses experienced in the mid-1980s, market risk was viewed as the most serious issue requiring attention and resources. Building a more sophisticated system for controlling trading credit risk simply was not a serious consideration at most institutions.

While adopting the add-on approach, bank credit departments invariably insisted on more conservative parameters internally than those employed in regulatory capital calculations. Partly this was driven by a recognition of the fact that the purposes for the two calculations were different. Primarily, however, it reflected an inherent distrust among traditional credit officers of these fancy

new instruments. By making the numbers larger, it was felt the volumes could be limited and the associated risks constrained. The result tended to be greatly inflated potential exposure estimates that often had only an infinitesimal chance of materialising as actual exposure-at-default in the future.

NETTING COMES INTO ITS OWN

Another trend in the late 1980s and early 1990s was the legal battle to make netting enforceable under the bankruptcy laws in various countries. The International Swaps and Derivatives Association (then known as the International Swap Dealers Association) (ISDA) lead the way in most countries around the world. ISDA formulated standard contract language and waged multiple campaigns to gain legislative recognition of the enforceability of netting under the terms of such contracts when one party declared bankruptcy.

While we think of netting as the accepted norm today, it is still far from universal. Indeed, even in the US netting is of questionable enforceability when dealing with certain types of institutions. Examples are insurance companies and public utilities where bankruptcy claims are adjudicated in some type of state administrative process rather than under the US bankruptcy code or laws governing federally insured financial institutions.

The most frequent circumstance today is for netting to be partially enforceable across relevant sets of transactions. Certain deals are entered into under terms of an enforceable netting agreement. These are nettable against each other but their combined net market value, if negative, cannot be offset against exposure in other nettable pools or against transactions done outside any enforceable netting agreement. This introduces an added layer of complexity to the analysis. It also means that legal contract information must be captured and properly reflected in the calculations if they are to be robust in reflecting the impact of enforceable netting.

ORGANISATIONAL OBSTACLES TO IMPROVEMENT

A number of organisational obstacles stood in the way of better credit exposure measurement systems. As already mentioned, cost was an important issue, especially when market risk was viewed as the really serious hazard. Most credit officers just wanted

to limit an activity they viewed as fraught with risks they didn't understand and that they weren't sure the traders understood either.

Perhaps equally serious, there was little support from the business side for the cost of improved credit exposure estimation systems. On reflection, this is not too surprising. Traders live and die on the basis of three things:

a. accurate pricing;
b. clear and reliable understanding of their open positions (usually expressed in terms of Greek letter sensitivities); and
c. realistic assessments of their potential losses (either in the form of a formal value-at-risk (VAR) estimate or, more likely, in terms of their own intuitive understanding of market volatility relative to their open positions).

These three factors are part of their daily lives, of the very air they breathe. Credit risk, on the other hand, seems like a remote contingency with little relevance. Traders usually regard credit risk oversight as a distasteful nuisance, a series of hurdles to be overcome. The point is to get the deals done. If a counterparty's limit is full, the usual attitude is "just raise the limit and let me get on with my job!"

Unfortunately, unsophisticated credit exposure measurement systems tended to intensify the inherent cultural conflict between traders and credit officers. Traders realised that add-on based exposure estimates, at least at the counterparty level, were inconsistent and inflated. They particularly recognised that marginal exposure implications were unreliable. This became painfully obvious when counterparties did offsetting trades to neutralise future changes in their net positions, and the system indicated that these increased credit exposure! All this reinforced traders' cynicism about credit oversight in general. They insisted that they were being constrained by a system that was inconsistent and arbitrary while credit staff were left to defend the indefensible.

Another obstacle to improved systems was limited analytical talent in the risk management side of the organisation. The big financial opportunities were in trading and it was difficult to attract the best quants and system engineers to the risk management side

of an institution. As a result, while the business had little interest in more sophisticated credit exposure measurement, risk management staff often lacked the analytical vision to formulate a better approach. The prevailing attitude on the credit risk side often tended to be, "make the numbers bigger and that will make us safer." Unfortunately, this tended to undermine the credibility of the exposure estimates even among credit approvers. The result was an unhealthy behavioural feedback loop. In effect, outstanding loan balances and potential trading credit exposures were two different currencies with an uncertain exchange rate between them. Each limit approver was required by circumstances to decide how to translate the one to the other. In such a situation, it was impossible to generate consistent credit decisions across the organisation.

THE IMPETUS FOR CHANGE

With all these obstacles, one might ask how did any improvement occur? In some cases there was a fortunate convergence of technical expertise, management vision and institutional emphasis on deploying best practice risk methods. More often than not, however, the impetus for change came at the point when credit availability began to constrain the ability to do business. As long as volume was modest and counterparties were solid investment grade names this was not a serious problem. Credit limits were generally available to cover even the inflated exposure estimates produced by a conservative add-on approach.

Moving into the 1990s, however, volumes continued to grow while smaller and less creditworthy counterparties entered the market. Gradually the old arguments about the exposures "not being real" and the estimates "being inflated" became less compelling. Credit approvers grew more reluctant to approve increased trading credit limits in the face of already large measured exposures for good names and for new exposure to lower quality names. Eventually situations arose where credit officers would say, in effect: "I don't care how inflated the numbers are, I'm not putting my name on an approval bigger than the current limit, period. Justify more realistic estimates and then maybe we can talk." This resulted in a real business incentive for more realistic exposure measurement and a willingness to spend some money to make it happen.

ESSENTIAL FEATURES FOR CONSISTENT EXPOSURE ESTIMATION

Given the incentives and resources to build an improved exposure estimation system, it was useful to consider what features were essential and what features, though not essential, were desirable in such a system. This discussion generally came down to the following essential features.

a. *A foundation in historical behaviour*. Generally it was agreed that historical market behaviour should be the foundation on which such a system was built. While some argued for a more conservative "worst case" approach, most recognised the dangers of inflated estimates losing their credibility. This inevitably leads to varied and highly subjective approval standards across the organisation and to an ultimate loss of control over the process.
b. *Proper recognition of offsets and diversification.* Failure to treat the dynamic interaction among trades in a counterparty's portfolio was recognised as the most obvious shortcoming of the add-on approach. This was particularly obvious relative to marginal exposure calculations. An add-on based system would show exposure increasing when a new trade was added even when it actually reduced exposure by offsetting pre-existing imbalances in the counterparty's portfolio. Few things do more to undermine the credibility of a credit risk system than getting the sign of marginal exposure wrong!
c. *Incorporation of market data correlations.* Clearly correlations between pairs of rates and/or prices are among the less stable parameters of market dynamics. Nevertheless, most analysts agreed that incorporating reasonable estimates of such correlations was an important aspect of realistic exposure estimation.[3]
d. *Exposure profiles rather than single loan equivalent amounts.* Another major flaw of the add-on approach was that it gave no insight into the timing of potential future exposure. Sometimes the potential exposures by deal would be "stacked" in the order of their maturity, but this ignores the fact that swap exposures peak in the middle of a contract's life, not at the end. Any credit analyst will say that the timing of when exposure occurs is often central to whether the risk is acceptable. Whatever else it did,

a revised trading credit methodology needed to produce insights into the timing of exposure.

e. *Rigorous treatment of netting.* By the 1990s, netting was becoming more widely accepted in law around the world. This made a reliable treatment of netting increasingly important for accurate exposure estimates. Capturing the impact of enforceable netting in the current exposure is easy. The difficult issue is reflecting it in potential future exposure. A minor revision to the Basel capital rules in April of 1995 introduced a very crude way of reflecting the growth in the enforceability of netting. This was based on the portfolio ratio of current net exposure, to the extent netting is deemed enforceable, versus current gross exposure. Again, while arguably appropriate for an aggregate exposure adjustment, this approach is much too crude to be reliable when applied to individual counterparty portfolios. A more rigorous approach was clearly necessary.

The common thread running through these essential features was an attempt to incorporate all significant structural factors into the exposure estimation process. The goal was to produce results that were conceptually consistent across counterparties and throughout the projection period. Specifically, the probability of the projected potential exposure actually materialising on any relevant future date should be as consistent as is practically possible. It should not matter whether a portfolio contains one deal, a small number of similar deals, or a large and varied range of deals with two-way sensitivities and multiple market drivers. Likewise, the statistical properties of the exposure estimates should not change over the course of a projection as some deals mature and the complexity of the portfolio changes.

DESIRABLE FEATURES FOR EFFECTIVE RISK MANAGEMENT

The initial goal for improved exposure estimation was simply to get more realistic results on a periodic basis. Daily or even weekly runs were the normal target frequency. This provided a sound basis for setting and managing credit exposure limits. There was an issue of how to provide availability information to traders in order to hold them accountable for complying with the limits. This usually was done by one of two methods.

- ❑ A parallel set of limits might be set up on the basis of a much simpler measurement scheme. For example, the old add-on calculations might be used.
- ❑ Alternatively, the sophisticated exposure profiles could be imported into the trading limit system. These would be augmented by conservative exposure increases for new deals between updates from batch runs of the more sophisticated exposure estimation system.

In both cases, incremental exposure of new deals would be added to the existing exposure based on fairly simple additive rules and tracked against limits. Such systems would also provide traders with availability reports or what-if capabilities. It was usually the job of a credit limits administration unit to assure that authorised trades between full exposure simulations would only rarely result in excesses when the batch update was run.

Even in the early 1990s, however, a set of more ambitious goals began to emerge. Progress has been made on some of them, but many remain future aspirations even today. Among these longer-term goals were the following.

a. *Real-time updates on a global basis*. The lag between deals being committed and their credit implications being reflected in the limit system has long been recognised as a source of additional risk. It has generally been necessary to settle for one of two compromises in this area. Global real-time exposure capture can be accomplished if the metrics used are quite simple. Such an approach is used for some products by at least a couple of dozen banks today, usually for foreign exchange trading. Alternatively, a sophisticated measurement method can be combined with real-time deal capture at the local trading room level if the exposure calculations are integrated with the front-office dealing system. This is more common in the fixed income derivative area. Combining global real-time deal capture with sophisticated measurement techniques to answer availability queries remains a future aspiration.
b. *What-if capabilities linked to the front-office booking systems*. Ideally the limits system should be closely integrated with the front office trading environment. Traders should not have to do double

entry to determine availability. Often a counterparty will have sufficient availability that a full what-if check is unnecessary. Visual inspection will assure that a proposed deal can be accommodated within the limit. If a full what-if simulation is needed, however, it should be possible for the exposure engine to pick up the necessary details from the front office trading system.

c. *"Trading floor response times" using sophisticated measurement methods.* Anyone who has been there knows that "trading floor response times" usually means almost instantaneous. Especially in the foreign exchange trading arena, response time requirements are extremely short, usually two to three seconds. The computational complexity traditionally required to produce acceptably robust exposure measures has, to date, been impossible to reconcile with these demanding response times on a commercially acceptable basis.
d. *Exposure decomposition and "wrong-way exposure" reporting*. As described above, measuring exposure consistently and reliably is an important goal. It also is important, however, to be able to dissect the portfolio imbalances that give rise to such exposure. This can be done by recording the impact on exposure of controlled deviations in individual market variables from their expected paths. This conveys the type of market events that would drive exposure higher and is an important guide to what types of trades would be risk reducing for a given counterparty. Proper organisation of such results also allows risk managers to determine which counterparties will exhibit rising exposure in response to an unfolding market event such as a big currency devaluation or a commodity price shift.

   Combining such sensitivities with user-specified wrong-way exposures also can be very useful. This allows reporting that highlights those counterparties who are likely to exhibit weakened credit status in response to the very same market events that result in increased exposure.
e. *Optimal professional counterparty selection*. Market makers in derivatives are deal takers when it comes to their end-users. Clients have a fundamental risk to hedge and come to the dealer to take an offsetting position. When professional market makers need to hedge their own open market positions, however, they are deal givers. They will need to pay away the bid or offer

spread to the other professional market maker with whom they place the trade. At that moment, there is almost free credit risk reduction available to the market risk hedger. By choosing one professional counterparty, potential credit exposure may be reduced, whereas choosing another counterparty will result in a potential credit exposure increase. Careful risk source sensitivity analysis and exposition can provide the necessary information to make the appropriate risk reducing choices as these opportunities arises.

f. *Marginal credit cost implications of a new transaction*. This can reasonably be called the Holy Grail of desired functionality. Today exposure is estimated only as a measure to compare to approved limits. As long as a trader is within limit, there is no other impact on incentives or behaviour. In effect, traders are compensated for maximising gross trading profits subject to constraints in the form of market risk and credit risk limits. If marginal exposure could be calculated accurately in a few seconds, it would be possible to present traders with the incremental credit charge associated with a proposed trade. This would mean the spread on the trade would have to cover this charge before the trader realised a net profit. In effect, it would allow market makers to compensate traders on the basis of maximising risk-adjusted profits. In the process, it would bring traders' incentives into much closer alignment with the goals of the institution.

ALTERNATE APPROACHES TO ACHIEVING ESSENTIAL MEASUREMENT OBJECTIVES

There are three potential approaches to achieving the essential objectives needed for consistent exposure estimation.

a. *Monte Carlo simulation with full revaluation*. This involves revaluing all uncompleted terms and conditions for every deal at multiple future dates under several thousand hypothetical sets of market conditions. Needless to say, this is very computationally intensive. Depending on the complexity and volume of the deals in the portfolio, it typically takes several hours to process all the desired exposure profiles. In some cases, the processing time exceeds the available daily window and can only be run on

weekends. While this is analytically the most robust approach, the computing power and processing time required often make it impractical for day-to-day use. This has given rise to a search for alternatives that offer almost the same reliability with considerably less computational burden.

b. *Monte Carlo with grid-based pricing*. The most obvious alternative to full-valuation Monte Carlo is one that simplifies the generation of hypothetical future values of the transactions. An attractive approach is to use grid-based pricing. In this method, a small number of full valuations are performed for every deal at every future simulation date based on controlled deviations from the status quo evolution of market variables. These results are saved and reused in the subsequent Monte Carlo step. Instead of performing full revaluations for each of several thousand hypothetical market scenarios, prices are derived by interpolation off the price grid. The interpolations are based on where the market variables in a given scenario fall relative to the controlled changes used to construct the grid. The grid may be more or less complex, incorporating different numbers of market drivers for a given deal. Likewise the interpolation methods may vary in complexity from linear to quadratic or higher order approximations to the value surface.

   Obviously the more complex the interpolation the smaller the gain in computational efficiency. How complex to make the calculation is a matter of time and budget resources and is also affected by the complexity of the trades involved. In general, however, grid-based pricing can make the update process several times faster than full valuation Monte Carlo. The scenario specific transaction values are not as precise in the grid-based pricing approach as they are using full revaluation. It is important to realise, however, that the characteristics of the hypothetical future market conditions are themselves subject to considerable uncertainty. As a former associate of mine once said, it is easy to refine within the margin of error. Doing so gains little in the way of actual precision (although it may create a false sense of security) and often creates serious operational and maintenance issues.

c. *Advanced analytic approximation*. Within the past two years, advanced analytic methods have been developed that provide

remarkably robust approximations to the results of full-blown Monte Carlo simulations. While not foolproof, these methods usually give results that fall within the range produced by Monte Carlo simulations with small versus large numbers of draws. Certain ill conditioned combinations of a small number of trades can lead to non-trivial errors. As a result, periodic Monte Carlo runs are important to catch such situations in those rare instances where they arise. Nevertheless, these methods are generally quite reliable and can reduce the time to evaluate the exposure of a large counterparty portfolio from minutes to seconds.

ALTERNATE ASPIRATIONS FOR USE OF SIMULATION-BASED EXPOSURE ESTIMATION

Given the advances of the past ten years, there is little excuse for any institution with significant derivative activity to operate with nothing but add-on style credit exposure estimates for internal measurement and control. Saying that, there are three distinct levels of sophistication in the application of simulation-based exposure estimates.

1. *As an overnight supplement to less sophisticated measures*. This is where most institutions are today. Daily trading activity is controlled by fairly unsophisticated measures at least relative to newly booked transactions. Even if sophisticated exposure profiles are calculated periodically and fed back into the trading limit system, deals done between updates are typically assessed in a much simpler fashion. This inevitably leads to overstated exposure and may result in turning away deals that could be done within limit if more sophisticated exposure estimates were available on a timelier basis. Such a system also cannot offer meaningful insight into the marginal impact of a proposed trade. Nevertheless, it is a giant step beyond having nothing but add-on based exposure estimates. Properly configured, such batch systems also can provide valuable portfolio concentration and risk source sensitivity information. A key success factor is making such information available to relevant decision-makers in multiple locations.

2. *Intra-day simulated exposure updates for new trades*. This approach requires capturing new trades as they are made and transferring them to a central exposure evaluation system. Once received centrally, the counterparty exposure profile or profiles affected by a new deal are updated in the background. Depending on the size of the portfolios affected, and the volume of trades being processed in this way, the latency is likely to be from a few minutes to an hour or longer. The advantage of this approach is that portfolio effects are reflected faster than with nothing but batch updates of the simulations. Also, if a counterparty's exposure is close to the limit, it is possible to configure such a system to allow a small number of what-if simulations based on proposed new deals. While not able to support a large volume of what-if inquiries, this limited capability can provide support for booking risk-reducing deals that would otherwise be denied based on limit excesses.
   While this may seem like a small departure from the first configuration, there is a major difference. To achieve intra-day updates requires moving from a batch to an event driven system architecture. Sufficient trade details need to be captured to allow the simulations to run and exposure updates need to be made while new limit inquiries and transactions are taking place. This requires a significant revamping of the entire process when starting from a batch mode approach. Given the rather modest improvement in business functionality, the cost of such a change is often viewed as excessive.
3. *Use of simulation-based exposure for what-if limit checking and pricing implications*. This is the last "desirable feature" described previously. If accurate simulation-based exposure could be performed on a what-if basis in a matter of seconds, it would offer significant advantages. By providing an accurate estimate of the amount and timing of incremental exposure, reflecting all major effects of offsets and diversification, this would be the basis for reasonable estimates of the cost of the associated incremental credit risk. In addition to creating a minimum spread for exposure increasing deals, such a system would encourage aggressive pricing for exposure reducing deals. Traders' performance could be measured on the basis of risk adjusted returns not gross trading profit. Until the development of the advanced analytical approximation methods described above, this aspiration was

commercially impractical. While such a system could have been built technically using Monte Carlo techniques, the cost in hardware and support has heretofore made it commercially impractical. With such methods, such a system becomes far more commercially justifiable.

THE NEXT 15 YEARS

Looking back on the last 15 years, it is clear that there has been a massive improvement in our understanding, measurement and management of trading counterparty credit exposure and credit risk. I fully expect that looking back to today on the 30th anniversary of *Risk*, current methods will look just as outdated as those of 1987 appear to us currently. Nevertheless, I think the broad outlines of that world can be discerned in unexploited capabilities already available.

A general trend I expect to see over the next 15 years is increasing use of risk information for tactical as well as strategic decisions. Today most risk information is used to make high-level resource allocation decisions using some variation of risk adjusted return on capital (RAROC), or to enforce risk limit controls. Most tactical decisions are still made on the basis of limited local information. Loans and other traditional credit facilities are evaluated on the basis of the financial strength of the obligor with only limited attention paid to portfolio implications. Derivative trades are priced based on how the contract fits the current market risk imbalances and, to a lesser extent, the credit worthiness and sophistication of the counterparty. In neither case are facilities available to assess broader portfolio risk implications. A number of capabilities already available, however, support an expectation of considerable progress along these lines in coming years.

a. *Massive expansion in communication capacity*. The current sad financial state of the telecommunications industry is a reflection of significant excess capacity. This is largely a reflection of massive expansion combined with slower than expected penetration of broadband access among consumers in industrial countries. Such penetration eventually will occur, however, driven in significant measure by access to on-demand streaming video. In that setting,

voice, text and numerical data will be minor portions of the total information flow. This will facilitate the much greater communication volume required to provide real-time updates and what-if inquiries to support thousands of daily decisions.

b. *Evolving semantic standards based on Extensible Mark-up Language.* Extensible mark-up language (XML) has rapidly been accepted as the syntax for self-describing messages. To be practical, however, XML must be supplemented with a series of semantic conventions relevant to specific types of information. Financial products mark-up language (FpML) is one such effort in the financial derivatives arena and there are several others in related fields. Development of these semantic standards is a slow and laborious process. The task is magnified by the need to build a critical core of institutional support through consensus. Progress is being made, however, and such protocols can be expected to be an important tool supporting, real-time, event-driven data exchange in coming years.

c. *Browser-based information access capabilities*. An essential feature of an effective risk information system is not just reliable source data and sound analytics, but effective delivery of the resulting information. To realise the potential of risk information to improve tactical decisions requires that such results be delivered to many thousands of staff members. Furthermore, it must be tailored to their particular job needs and responsibilities as well as their personally preferred means of absorbing information. Finally, it must be possible to upgrade the system with added or modified functionality on a continuing basis without excessive coordination problems between the central server and the remote users' desktops. Fortunately, the rise of the Internet offers technology that is ideally suited to address this problem. Browser-based technology still cannot deliver the full richness and performance of client server applications, but the gap is narrowing. In addition, continued improvements in communication and computing speeds may well make the differences perceptually insignificant in a few years. It may be that a few key "power users" will continue to be supported by client-server configurations, but the vast majority of users will be fully and effectively served by web-based information access and display tools.

d. *Robust approximations to Monte Carlo results*. As noted earlier, recently developed techniques offer the prospect for robust approximations to the results of computationally intense Monte Carlo runs for a small fraction of the processing resources. Deployment of such methods would avoid the need for two measurement systems, one robust but slow and the other crude but fast. Such capabilities would also support risk adjusted pricing by supplying traders with the marginal cost of risk *before* a quote is given to the client. This carries the potential for a major overhaul in how traders' performance is evaluated and their compensation is determined.

CONCLUSION

It is said that forecasting is difficult, especially if it's about the future. I am sure, when the reality arrives, that there will be much to criticise in the vision I have articulated for the year 2017. Nevertheless, systems to capture and present the "macro" implications of "micro" business decisions are becoming progressively more feasible and less expensive. Regardless of the details, I feel confident that in 15 years time risk management will be far less separated from the business and be far less of an after-the-fact function than it is today. Rather it will be a pervasive influence on day-to-day decisions at all levels.

1 See http://newrisk.ifci.ch/134710.htm for an interesting summary.

2 In current terminology, this is a limit on the maximum loss given default.

3 A less commonly recognised point was that correlations appropriate for market risk may not be appropriate for credit exposure estimation. Credit exposure should be based on estimates of longer-term trend correlations rather than daily change correlations, since credit exposure must be simulated many months and even years into the future.