RiskMetrics Group™ Research Monthly

Exposed!

Christopher C. Finger
chris.finger@riskmetrics.com

June 2006

In Wikipedia, the online encyclopedia, there is a concept called a disambiguation page. Such a page brings together articles associated with a single topic, when the topic itself may have a variety of different meanings. A search on "exposure" brings the user to such a page: a collection of links to articles in biology, photography, law, and so on. Notably, the Wikipedia disambiguation page for exposure does not (at least as of this writing) contain a section on finance or risk.

Exposure is perhaps the most frustratingly ambiguous term that we ever encounter. It is frustrating first simply because of its ambiguity, but second also because of the lack of appreciation for its ambiguity. When I am asked how we price an option, I am never chastised for asking "equity or interest rate", "American or European." My request for clarification is expected and appreciated.

On the other hand, the meaning of exposure is simple and obvious in enough contexts that it is easy to slip into the assumption that it is simple and obvious always. Unfortunately, not all of the simple and obvious answers are the same. Exposure is usually not an end in itself, but rather a shortcut risk measure; and it is not always a shortcut to the same thing. To choose from the multiple obvious (or not so obvious) answers, we must make explicit what we are trying to get out of our exposure measure.

Moreover, as we will see, what is a natural and useful concept for a single position may be awkward or even useless for a collection of positions. In these cases, we are faced with the choice of generalizing the concept to again be useful, accepting the flaws of the exposure measure for collections of positions, or abandoning the exposure measure altogether.

Our aim, then, is a disambiguation page for financial or risk exposure, and perhaps to elicit a bit of sympathy as we try to sort out a problem that is much more difficult than it initially appears.

Three notions

At the most general level, we can think of three distinct notions of exposure:

- Exposure as sensitivity,
- Exposure for leverage measurement, and
- Exposure to counterparty credit risk.

As we proceed to discuss these three notions in detail, we should always bear in mind that exposure measures are meant to be shortcuts. They will never tell us the whole story, since we will always be discarding some details. But we should be conscious of which notion of exposure we mean, since the details we discard may be different.

Exposure as sensitivity

Exposure as a sensitivity measure is the answer to seemingly simple questions as "How much am I invested in equities?" or to the telecommunications sector, or AA-rated bonds, or hedge funds. Phrased this way, the exposure measure is meant to provide relative information: if my exposure to equities is twice as high as that to bonds, then a gain in the equity market is twice as beneficial to me as a gain in bonds. It may also be utilized along with basic information about the riskiness of the asset classes to give us some sense of the risk of the overall holdings.

As phrased above, it is implicit that our exposure to a category is equal to the sum of the exposures of all of the positions in that category. This is true because it is obvious that a given investment is either an equity, a bond, or a hedge fund, rather than a mix of these. Moreover, it is obvious that the exposure of any position in the category should be simply the current market value of that position. Or is it?

Consider short positions. Should the exposure on a short position be defined as a positive or negative number? Or stated differently, should the exposure for a category be equal to the sum of the signed sizes of the positions in the category, or to the sum of the absolute sizes of the positions? This second question gives rise to definitions of net (sum of signed positions) and gross (sum of absolute positions) exposures. But this just begs the question of whether gross or net exposure is the appropriate answer to the questions at the beginning of this section.

If what we are asking is our sensitivity to a market, then the net exposure in the category representing the market is a reasonable rough approximation. Assuming that the individual positions are tightly correlated and have similar volatilities, then the net exposure does give us a rough sense of how much we stand to gain or lose given a particular move in the overall market. For broad categories—for instance, general global equities—positions have varied volatilities and on average weak correlations; it is much harder to argue that the risks on long and short positions cancel, and our interest in gross exposure grows. In general, we may think of gross and net exposure as rough bounds on our risk profile in a particular category, with net exposure being most relevant for categories that are most tightly defined.

Positions with embedded funding present a more subtle problem. We have said that the exposure of an individual position is obviously the position's market value, with the one complication of the exposure's sign. This approach works for positions in actual assets, but what of futures positions? A futures contract that is settled daily will have a market value of zero, and yet an investor in the contract can clearly make or lose money. The complication is that the contract packages a long position in the future's deliverable with a funding (short) position in cash. Viewed this way, the exposure of the futures contract obviously should be the equivalent position in the underlying deliverable.

Derivatives add another level of complexity, in that not only do they package funding with a position in an asset, but the position in the asset changes dynamically. Products with multiple risk factors—including derivatives, but also equities denominated in a foreign currency and credit-risky bonds—raise the inevitable question of "exposure to what?": foreign exchange or equity markets, underlying or implied volatility, rates or spreads. Because the notion of exposure appears deceptively simple, we of-

ten do not consider the "to what" question, and assume whatever comes to mind first without appreciating the multiplicity of possible answers. And how can we begin to think about aggregating once our exposure measures represent such a hodgepodge of underlying markets?

This leads us finally to the notion that there really is no single exposure measure to characterize an individual position's risk. On the same set of positions, we may be interested in our exposure to currency, equity, or spread movements; the appropriate definition of exposure depends on our question. Our question here is likely phrased using the word "sensitivity". Somehow, with sensitivity, the "to what?" part of the definition is usually made explicit; as a consequence, we may meaningfully aggregate sensitivities. The downside is that we have move away from the notion of a single exposure measure for each position; further, we wind up dependent on pricing models from which to compute sensitivities, rather than on just a market value as in the simpler case.

So our first pass at disambiguation has made a bit of a mess. We started with a seemingly simple question, and obvious answers, and managed to end with a reliance on pricing models and Greek letters. The lesson is that exposure as a generic risk proxy is just that, generic. For simple asset allocation information, we should do simple things, but we should not try to extract more precision from an exposure measure without accepting greater complexity.

Exposure for leverage[1]

The second notion of exposure would seem to be easier to handle, in that it is attached to a more specific question: what is the leverage of a portfolio? Leverage, intuitively, is defined as the ratio of the value of assets that a portfolio controls to the portfolio's capital. If we have a quantity 100 to invest (our capital), and we invest the entire amount in straight equities, then our leverage is one. If we borrow an additional 100, and invest our initial capital plus the borrowed funds in straight equities, then we control 200 of assets, and our leverage ratio is two.

In practice, the denominator in the leverage ratio—the capital—is usually clear. The numerator is less clear, and is typically not called the controlled assets[2] but, alas, exposure. So here, exposure is defined implicitly: it is whatever is needed in the numerator to produce a meaningful leverage ratio. Of course, this bit of semantics only delays our real work: it is now incumbent on us to define leverage.

Leverage, like exposure in the previous section, is a shortcut to characterize the risk of a portfolio, used in particular when the portfolio employs a significant amount of borrowing, whether explicit as in our example above or implicit as embedded in a derivative contract. As in the previous section, the "obvious" definitions get cloudy as our situation becomes complex. It is crucial, then, that we are explicit as to what we ask of the leverage measure.

The first thing we ask of the leverage measure is a simple relative indication of risk. The leverage ratio tells us how large the potential gains or losses on a portfolio are, relative to those for the analogous un-

[1] The discussions and references in Barry Schacter's web log, Bel Ranto, were extremely helpful in formulating this section.

[2] This terminology is taken from Breuer (2000). This author is not sure if it is standard, but feels that it should be.

leveraged portfolio. In our example, the leveraged portfolio is twice as risky than the first, provided the two invest in the exact same set of assets. Greater leverage does not necessarily imply greater risk; the measure in this sense must be used as a relative indicator.

It is this notion of leverage (and exposure) that the recent UCITS III Directive[3] addresses. The directive, among other things, allows investment funds to utilize derivatives, provided they maintain appropriate controls. One of these controls is expressed:

> ...ensure that the global exposure relating to financial derivative instruments may not exceed 100% of the UCITS net asset value (NAV), and hence that the UCITS overall risk exposure may not exceed 200% of the NAV ...

This suggests that the exposure notion the regulators are tackling is leverage, a point they state explicitly later, while at the same time asking for help:

> ...the methods for assessing the leverage of a UCITS need further refinement, in particular with respect to the maximum VaR/stress-test value corresponding to a total exposure of 200% of a UCITS NAV.

Importantly, following the UCITS Directive is not a formulaic exercise in compliance, but an opportunity to contemplate how to create simple limits that achieve the goal of controlling the risk contributed by derivative instruments.

Defining leverage

Breuer (2000) sets out to measure the institution-level leverage of banks. As such, his chief obstacles are off-balance-sheet items, such as forward contracts and derivatives. Breuer develops a paradigm whereby leverage is defined by decomposing positions, making implicit borrowing explicit.

Consider a forward contract where we commit to pay a fixed price (X) for some (non-income producing) security at a future date. The value of this position today (f) is the difference between the price of the security today (S), and the fixed price discounted at the risk-free rate of interest (Xe^{-rt}). We may replicate the position by borrowing Xe^{-rt}, augmenting this amount by f, and purchasing the security in the spot market for S. Seen this way, the contract is no different from our simple example above; leverage is the ratio of the controlled assets (S) to the total value of the position, f. Of course, f could well be negative, a complication we will come to later.

The next case Breuer considers is a simple option. Here, we may replicate the option, as under the Black-Scholes framework, by a short cash (that is, borrowing) position and a long position (the delta-equivalent) in the underlying security; the option premium is equal to the difference in value of these. The controlled assets are equal to the delta-equivalent position, and the leverage is the ratio of this to the value of the option. Of course, the option decomposition, and thus the leverage ratio, changes

[3] UCITS (Undertaking for Collective Investments in Transferable Securities) III was implemented in 2004 for European Economic Area states. It is designed to allow providers of retail investment products a wider scope of activities, including the use of derivatives. In return, fund providers are required to establish a risk management process, and to work within certain risk and leverage limits.

dynamically with the market. But for the forward contract, the ratio S/f changes as well. So even for the simplest of positions, leverage is a point-in-time indicator, not a fixed quantity.

Moreover, it is wrong to state that any type of position inherently contributes leverage. A futures position may be utilized for leverage—if used alone to speculate on an asset without a significant cash outlay—or it may be used as simply an efficient alternative to buying the underlying asset—if it is matched with an equivalent cash position. To understand a portfolio's leverage, and by implication its exposure, it is necessary not only to know its risky assets, but also its cash positions.

Breuer does not consider long-short equity trades, but we can still apply his logic. Suppose, as above, we have 100 to invest. We invest the 100 in stock A. We wish to invest a further 100 in A, but instead of borrowing, we finance this further investment by entering a short position of 100 on stock B. The denominator in our leverage ratio is obviously 100, our initial stake. But what of the numerator—the controlled assets, or exposure? Here, we can argue for three different choices.

Recalling our discussion in the previous section, we could argue that if stocks A and B are relatively uncorrelated, then the best characterization of our exposure is the sum of our absolute positions (300), making our leverage three. On the other hand, if the two stocks are strongly correlated, then we could argue that much of our risk is offset, and our exposure is best characterized by our net position (100); our leverage in this case is one. Finally, as a middle ground, we could argue that the short position, particularly if stock B carries very little risk, is simply for financing, and that our position is no different from one where we borrowed the additional 100. In this case, our controlled assets are the 200 invested in stock A, and our leverage is two.

To attempt to resolve the argument, we return to the application of the leverage measure itself. We stated that leverage can be used as a relative measure of risk, but in this case, relative to what? When we leverage by simply borrowing, there is a natural analogous unleveraged portfolio. Here, there is no such notion, and none of the three possible values of the leverage ratio really tells us anything. What we really want to know about is risk.

So we have arrived already at a case where our principles are insufficient to provide us an answer. We respond by looking for more principles.

A second use for a leverage measure is to assess the risk of ruin. A fund employing an unleveraged strategy can only fail if the value of its assets falls to zero. A fund with a leverage ratio of two fails if its assets experience a loss greater than 50%, in which case it will have insufficient funds to repay its loan, and its capital will be gone. As unlikely as this is, it is certainly more likely than a 100% loss.

In our example, the fund will fail if the value of our long position on stock A is at some point insufficient to cover our short position on stock B, for example if stock A falls by 25% and stock B appreciates by 50%. We can argue that this is more likely than stock A losing all of its value, and therefore that it is not appropriate to characterize the long-short position as unleveraged. That eliminates one of our possible values, but is it more appropriate to characterize this portfolio with a leverage ratio of two or three? Is the risk of ruin comparable to that of a position where we borrow 100 in cash and invest a total of 200 in stock A, or one where we borrow 200

and invest 300 in stock B?

The problem is that leverage is a relative measure. In situations where it is traditionally applied, there is a natural unleveraged position; we thus at least know what a leverage ratio of one means, and can define riskier positions relative to this. But in the event that there is no natural unleveraged position to use as a starting point, there is little hope of defining leverage, or consequently exposure, in any meaningful way.

From leverage to risk

The situation appears somewhat hopeless, and yet we are still faced with interpreting the directives quoted previously. We can start with a simple case: consider a fund that at a certain point is invested in an unleveraged index portfolio, and then borrows cash to invest in exactly the same portfolio. Since the new investment is entirely funded by borrowing, the fund's NAV is unchanged. Here, the exposure is obviously the total investment in the index portfolio, and the leverage is the ratio of the exposure to the fund's NAV. The meaning of the directive is clear: the exposure must not be more than twice as large as the NAV, or equivalently, the leverage ratio must not be larger than two. The effect of this limit on leverage, at least in this simple case, will be to limit the risk of the actual fund to no more than twice the risk of the original unleveraged portfolio.

Exposure and leverage, as we have seen, are difficult concepts to generalize, but risk is easier. Thus, our general interpretation of the directive is that the risk of a fund, including its derivatives and other leverage-producing positions, should be no greater than twice the risk of the fund's unleveraged counterpart, what we might call its risk benchmark. Of course, measuring risk is a decidedly more complex task than summing position values, but on a complex portfolio, we have no alternative.

We are left then with the choice of the risk benchmark. In practice, though it is not imperative, we rely typically on the fund's performance benchmark as the risk benchmark as well. For this purpose, we are not measuring tracking error against the risk benchmark, but rather just using the benchmark to provide an indication of an acceptable level of (unleveraged) risk. Effectively, the risk limit here, in addition to controlling leverage, constrains a fund manager from moving into significantly more risky positions (even if they are unleveraged) than his benchmark contains. On the other hand, it does not restrict a manager from investing in an altogether different portfolio from the benchmark; a measure of tracking error would be necessary for this.

In the end, with leverage (and by extension, this notion of exposure) we are still talking about a shortcut risk measure. We should not be surprised, then, that the measure has flaws, and situations will arise when we need a more rigorous approach.

Exposure to counterparty credit

We arrive now at our final, and most concrete, notion of exposure: the characterization of how much credit we extend as part of a swap or derivative contract. With certainty, we can value a contract today; if that value is positive to us, we are exposed to our counterparty, in that it would cost us something to replace the contract should the counterparty default. In the future, whether and how much we are

exposed to the counterparty depends on the evolution of our contract; while the notion of exposure is the same, the actual measurement is a matter of estimation and forecast rather than of certainty. At best, we can forecast the distribution of future exposures.

Still, because many institutional policies regarding credit exposure (for instance, limits and economic capital measurement) are defined for straightforward lending positions, it is necessary to summarize the exposure distribution with one number, sometimes referred to as a loan equivalent. The loan equivalent gives us a shortcut with which we may treat a derivatives position, for purposes of limits or capital, exactly as a simple loan.

Here then, we have a concrete question to define exposure: how much would we lose on (or equivalently, how much would it cost to replace) a set of positions if a particular derivatives counterparty defaulted at a specific time? Crucially, this definition extends naturally from an individual contract to a collection of contracts with an individual counterparty. This is not to say that counterparty credit exposure is an easy problem to solve, but rather only that it is an easy problem to define (in contrast to other cases we have seen). Indeed, there are numerous issues outstanding in this area: how to perform long-term forecasts of the underlying risk factors; how to account for collateral, netting, and other credit risk mitigation techniques; and how to recognize correlations between counterparty credit quality and the value of contracts with the counterparty.

Beyond these well-specified problems, there is the challenge of recognizing when one simple loan equivalent measure is sufficient. Is it always true that a derivative, whose future exposure is uncertain but whose estimated loan equivalent exposure is 100, should contribute equally to the utilization of a credit limit as a straight loan of 100? Thus, when it is necessary to go beyond a shortcut risk characterization? In this respect, we face a similar problem as with our other exposure notions.

Conclusion

So what have we learned? First, the word "obvious" is a dangerous one. While it is certainly nice to have obvious answers in simple cases, this can hide the subtleties that make exposure, and other notions, difficult to define in generality.

Second, we must bear in mind that exposure in all its forms here is a means but not an end. Thus what is important is that we arrive at the right answer to whatever our real question is, be that sensitivity, leverage, or credit. If a shortcut in the form of an exposure measure is useful, then by all means, we should use it, but we should not stubbornly pursue an exposure definition in the face of all logic if our end goal is really to limit risk.

Further reading

- Breuer, P. (2000). Measuring off-balance-sheet leverage. Working Paper, International Monetary Fund.
- Finger, C. (1997). Toward a better estimation of wrong-way credit exposure. *RiskMetrics Journal*, **1**(1): 25–40.