RiskMetrics Group™ Research Monthly

What's the worst that could happen?

Christopher C. Finger
chris.finger@riskmetrics.com

March 2005

Introduction

In a recent speech describing how he evaluates funds in which his firm is considering an investment, Emmanuel Derman spoke about the various models he has at his disposal, and then stated that he liked to ask the traders and portfolio managers what market scenario they thought could cause them significant losses. Putting ourselves in the position of the portfolio manager, we imagine ourselves, panic stricken, muttering an answer as vacuous as "I worry that my long positions will go down and my short positions will go up," and losing our bid for a mandate. The problem is that we can often talk in great detail about how a certain trade or bet can go wrong, without knowing what is likely to cause the portfolio as a whole to underperform or depreciate significantly. This is possibly sufficient for the portfolio manager himself, but a catalog of trades is far too much information for an investor or risk manager with responsibility for many such portfolios.

A natural reaction to this is the somewhat overused battlecry of stress testing. Unfortunately, most stress testing frameworks solve only the easy part of the problem: they tell us how much we lose if a particular scenario occurs, but do not offer guidance in choosing a scenario that is at once relevant to our portfolio and realistic (or at least plausible) in current market conditions. In practice, the scenarios chosen are in the end more or less arbitrary. Consequently, it is often difficult to even know whether the relationships implicit in these arbitrary scenarios represent a departure from or a continuation of relationships that have prevailed in the marketplace.

In this note, we discuss a technique, Maximum Loss, that combines portfolio and market information to derive realistic and relevant scenarios. Maximum Loss is not a new technique, but previous discussions of it have emphasized it as a risk measure, rather than as a method to choose scenarios. We also discuss a more practical problem. Often we are faced with filling in the specifics for a general scenario: how to decide, for instance, the specific depreciations to apply in a "Currencies Depreciate" scenario. Derman is asking us to do the opposite: distill a specific market scenario for popular consumption.

Maximum Loss and the Loss Scenario

The Maximum Loss technique was introduced in a number of articles[1] in the late 1990s. The articles introduce Maximum Loss as a risk measure that is similar to Value-at-Risk, but with a few more desirable properties. To define Maximum Loss, we first define a set of risk factor scenarios, referred to as the Trust Region; we then find the worst portfolio

[1] See Studer (1999) and Studer and Luthi (1996).

loss over the scenarios in the Trust Region, and call this the Maximum Loss. In theory, the Trust Region is completely arbitrary, but in practice, we typically define this in terms of a particular confidence level.

As a simple example, suppose we have one normally distributed risk factor. For a confidence level of 95%, it is natural to consider a Trust Region consisting of all risk factor return scenarios that lie within 1.64 standard deviations of the mean risk factor return. Of course, there are other regions comprising 95% probability, but it is most natural to consider the one symmetric, contiguous region with this property. Over this Trust Region, we search for the worst case portfolio loss. If our portfolio is linear, this loss will occur at one extreme of the Trust Region, while if the portfolio depends on the risk factor in a more complex way, the maximum may be achieved in the interior of the region. As with Value-at-Risk, Maximum Loss can be described generically as the worst case loss that a portfolio might suffer within some confidence interval; further, Maximum Loss and Value-at-Risk are identical for linear portfolios that are a function of a single risk factor. Thus, the intuition for the Maximum Loss risk measure is very similar to that for Value-at-Risk, though the specifics of the measures differ.

At the time it was introduced, the attraction of Maximum Loss was that it is similar to Value-at-Risk, perhaps the most common risk measure in practice, but satisfies (as Value-at-Risk does not) the properties of a coherent risk measure.[2] Discussions of the actual scenarios identified have been limited to cases with a small number risk factors. For our purposes, it is not the coherency of Maximum Loss, or even the statistic itself, that is the main attraction, but rather the particular market scenario at which the Maximum Loss is realized.

To be a bit more specific, denote by $v(\omega)$ the portfolio profit and loss for a specific realization ω of returns on the n portfolio risk factors. Further, assume that the set of risk factor returns follows a n-dimensional multivariate Gaussian distribution with covariance matrix Σ. The natural extension of the one-dimensional trust region discussed earlier is the set

$$A = \{\omega : \omega'\Sigma^{-1}\omega \leq c\},$$

where c is the quantile of the Chi-squared distribution with n degrees of freedom corresponding to our chosen confidence level. In one dimension, the set A is simply an interval centered at zero. In two dimensions, for independent and equally volatile risk factors, A is a circle centered at the origin; for strongly correlated factors with equal volatilities, A is an ellipse oriented along the 45 degree line from the origin.

The portfolio Maximum Loss is defined as the solution to the optimization problem

$$\min_{\omega \in A} v(\omega).$$

Of particular interest to us will be the scenario $\omega^\star$ at which the minimum value of v occurs. We will refer to $\omega^\star$ as the Loss Scenario. Note that $\omega^\star$ has precisely the characteristics we discussed initially: it is a worst case market scenario that is realistic given current market conditions (that is, volatilities and correlations) and relevant to our particular portfolio.

Now suppose further that the portfolio profit and loss is a linear function of the risk factor returns, that is,

$$v(\omega) = \delta'\omega,$$

[2] For further discussion of coherent risk measures, see Artzner et al (1999).

where δ is the vector of portfolio sensitivities to the risk factors. In this case,[3] there is a closed-form solution to the optimization problem and our Loss Scenario is

$$\omega^{\star} = -\sqrt{\frac{c}{\delta'\Sigma\delta}}\Sigma\delta.$$

Note that this is simply the vector $-\Sigma\delta$ multiplied by a scaling factor that depends on our (arbitrary) choice of confidence level. Thus, the Loss Scenario for a linear portfolio consists of the same relative moves in the risk factors, regardless of the confidence level chosen.

For a linear portfolio, the Loss Scenario is related to the portfolio's Implied Return.[4] In a standard optimization framework, we consider as optimal the set of portfolio weights δ that maximizes the quantity

$$\delta'\mu - \frac{\lambda}{2}\delta'\Sigma\delta,$$

where μ is the vector of expected returns on the risk factors, and λ is a risk aversion parameter. Thus, the optimal portfolio is chosen to maximize the portfolio expected return, with a penalty for portfolio variance. Expected returns being difficult to estimate, we can consider the optimization problem from a different angle, and ask which expected returns μ would need to prevail in order that the current portfolio be optimal. We refer to this vector μ as the portfolio's Implied Return. In practice, we may examine a portfolio's Implied Return to ascertain whether the portfolio has been constructed consistently with a manager's actual stated views on the market. In the unconstrained case we are discussing here, the Implied Return is the vector $\Sigma\delta$ scaled by a constant that depends on the parameter λ.

Comparing results, we see that the Loss Scenario of a portfolio is simply the downside of the portfolio Implied Return. This is quite intuitive. The Implied Return of a portfolio may be interpreted as the market scenario against which the portfolio has been best positioned to profit. With a linear portfolio and symmetric risk factor distributions, it follows that the downside of this scenario is that against which the portfolio is most vulnerable, corresponding to the Loss Scenario.

Of course, the definition of Maximum Loss and the Loss Scenario is not specific to linear portfolios, and the literature is well developed on calculation techniques for Maximum Loss on quadratic profit and loss functions. Further, with quadratic portfolios, the Loss Scenario at an arbitrary confidence level is no longer a simple scaling of one basic scenario; the relative size of Loss Scenario risk factor moves changes as the confidence level does. Nonetheless, the linear case discussed above provides useful intuition about the Loss Scenario. In the examples that follow, we will consider diverse fixed income portfolios. Though these portfolios do exhibit some convexity, the Loss Scenario results are nearly identical for the linear and quadratic cases. We present the results for the linear case only.

Global bond portfolio

To illustrate the Maximum Loss framework, we consider a portfolio of government bonds in twenty-one countries, issued in eleven currencies. We will consider both the absolute risk on our portfolio of fifty bonds and the relative risk of this portfolio versus a benchmark consisting of over six hundred bonds.

[3] See Studer (1996) for details.

[4] See Sharpe (1974).

Figure 1: Relative positions by currency

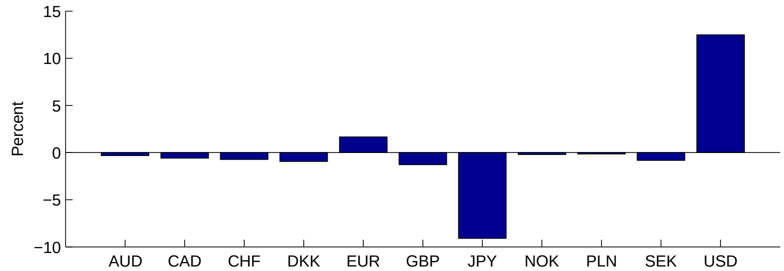

We consider a risk horizon of one month and a confidence interval of 95%, and report risk from the point of view of a USD investor. We examine both the absolute risk of portfolio loss and the relative risk of portfolio underperformance measured against the benchmark. The portfolio's relative currency positions versus the benchmark are presented in Figure 1. Note that the portfolio differs from the benchmark in duration as well, though our focus in this example will be on currencies.

We perform our analysis as of February 1, 2005, and utilize historical daily returns over the previous year to produce our covariance forecast. Combining the individual vertices on the government interest rate curves and the foreign exchange rates, we are exposed to well over 200 distinct risk factors. It is therefore impractical to report the Loss Scenario as a simple vector; we summarize the Loss Scenario instead.

For our portfolio, the primary risk driver is foreign exchange. We present the Loss Scenario for the foreign exchange rates in Figure 2. Not surprisingly, when we consider absolute risk, the Loss Scenario is a scenario where all currencies fall relative to the USD. By contrast, for relative risk, the Loss Scenario consists of depreciations in only half the currencies. Furthermore, the sizes of risk factor moves in the relative case are much smaller than in the absolute case. Relative to the benchmark, we have opposite positions in strongly correlated currencies (eg long in EUR, short in GBP). The event that these currencies each move against us is relatively unlikely; it follows that the absolute size of moves in such a scenario is smaller than the size of moves in the USD depreciation scenario to which we are exposed in the absolute. Interestingly, the relative case Loss Scenario contains a depreciation in JPY, even though we are net short this currency relative to the benchmark. This appears to be a case where the balance between realism and relevance is tipped toward the former.

It is important to note here that the Loss Scenario provides us with a more relevant, and realistic, scenario than the standard practice of moving each currency by an arbitrary number of standard deviations. Applied blindly, this simple approach fails to address our relative positions. On the other hand, if we stress the currencies opposite our relative positions, the simple standard deviation stress produces unrealistic scenarios, since it is quite unlikely that strongly related currencies would move significantly in opposite directions.

With respect to interest rates, our positions are straightforward, and thus our effective exposure is to overall rate shifts, and not to more complex curve moves. It is sensible, then, to summarize the interest rate component of the Loss Scenario in terms of the average shift in each government curve. Our large relative position in USD bonds is reflected in a large Loss Scenario shift of 122bp in the USD curve. The shifts in the remaining curves are mostly in the 50-60bp range, with the GBP shift a bit smaller at 25bp and the PTE and AUD curves actually falling by roughly 20bp.

The impact of the combined foreign exchange and interest rate Loss Scenario is an underperformance of 121bp relative to the benchmark. (By contrast, the Maximum Loss of the portfolio in absolute terms is 15%.) We present the contributions to this underperformance by currency in Figure 3. Note that the underperformance is driven by our positions in USD, EUR and JPY bonds, while our positions in the remaining currencies actually appreciate in the Loss Scenario. Though the largest underperformance comes from the USD bonds, our position in EUR bonds looks inherently riskier: our position in EUR is almost ten times smaller than that in USD, and yet the EUR bonds contribute more than half as much to the underperformance in the Loss Scenario as do the USD bonds.

To this point, we have relied on some knowledge of our portfolio to filter out the most salient parts of the Loss Scenario. In addition to this, we would like an approach that can be automated, and yet still give some intuition for the Loss Scenario. We do this by comparing the Loss Scenario to actual historical returns. Though we have used daily returns to estimate our covariance matrix, our risk horizon is one month, and so it is most appropriate to make our comparison with historical monthly returns. We compute the vector distance between the Loss Scenario and each historical return scenario.

Three months stand out as closest to the Loss Scenario: November 2002, February 2004 and July 2004. For these three months, at least six of the ten foreign exchange rates moved in the same direction as in the Loss Scenario. Interestingly, November 2002 was not part of the dataset we used to estimate our covariance matrix, and yet still appears to be close to the Loss Scenario. Note that these three months do not necessarily represent the worst historical portfolio losses, but rather those months with returns that both produced losses and that are relatively likely to recur.

Before concluding, we should revisit our emphasis that the Loss Scenario be realistic. Throughout, our definition of this notion has been that the scenario be consistent with our latest estimate of the covariance matrix. In a stress testing framework in particular, it would be foolish to build in a reliance on a single market forecast. The framework here admits its own historical stress tests, however, in that we may recalculate the Loss Scenario using a forecast from a different period. To illustrate this, we replace our covariance matrix with one derived from daily returns from 2002. We present the resulting foreign exchange Loss Scenario in Figure 4.

Conclusion

To a manager of a large portfolio or set of portfolios, it is a particular challenge to describe succinctly the types of market events that could produce significant losses. Though the Maximum Loss framework is not new to risk management, its attraction in the past has been as a coherent counterpart to Value-at-Risk. We have utilized the framework to another end: as a way to generate stress scenarios that are at once realistic based on a market forecast and relevant to a particular portfolio.

At this point, we are only half finished, and it remains to characterize the specific Loss Scenario in terms general enough to communicate to other managers or investors. Much of this task must be based on a knowledge of the portfolio, and a sense of which risk factors are most critical. Still, some more objective approaches exist, such as determining which historical scenarios are in a sense closest to the Loss Scenario. To respond to Emmanuel Derman's query, we may offer that we worry most about a repeat of July 2004, which at very least is an improvement over our initial response.

Further reading

- Artzner, P., Delbaen, F., Eber, J.-M., and Heath, D. (1999). Coherent measures of risk, *Mathematical Finance*, **9**(3): 203–228.
- Sharpe, W. (1974). Imputing expected returns from portfolio composition, *Journal of Financial and Quantitative Analysis*, June, 463–472.
- Studer, G. (1999). Market risk computation for nonlinear portfolios, *Journal of Risk*, **1**(4): 33–53.
- Studer, G. and Luthi, H.-J. (1996). Quadratic maximum loss for risk measurement of portfolios, RiskLab Technical Report.

Figure 2: Loss Scenario for foreign exchange rates

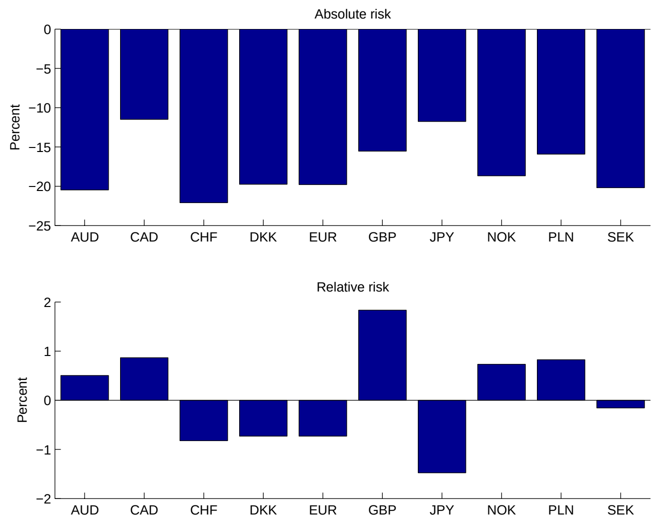

Figure 3: Relative losses by currency based on the Loss Scenario

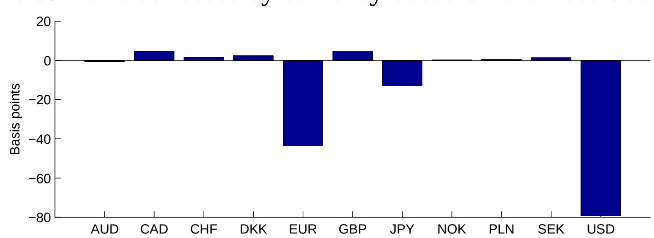

Figure 4: Loss Scenario for foreign exchange rates. Covariance matrix based on 2002 data