Stress-Testing in Asset and Liability Management: A Coherent Approach

by Alex Canavezes[a] and Mario Schlener[b]

[a] Quant Analytics, London, UK
[b] NCMA, Europe

This Chapter Covers

- How traditional stress tests are performed and why they are meaningless.
- How to assign a probability number to a given stress event.
- Exposition of the frequentist methodology.
- Exposition of the subjective methodology.
- Application of the frequentist methodology to a case study in asset management.

Introduction

In light of recent extreme events, such as the collapse of Lehman Brothers in 2008, both the financial services industry and its regulators have keenly felt the need to complement traditional percentile-based risk management tools (such as value-at-risk (VaR) or economic capital) with stress tests and scenario analyses.

Following the logic of Dermine (2003), asset and liability management (ALM) can be interpreted as the main management tool for controlling value creation and risks in a financial institution. Additionally, ALM should be the main management tool for discussions, in an integrated way, of fund transfer pricing, deposit pricing (for fixed and undefined maturities), loan pricing, the evaluation of credit risk provisions, the measurement of interest rate risk for fixed and undefined maturities, the diversification of risk, the marginal risk contribution, and also the allocation of economic capital.

Learning from the past misbehaviors of all market participants (especially the overreliance on quantitative measures with "statistical entropy," on diversification, and the assumption that capital is always available), risk management evolved from being just used as a risk-minimization, insurance, or diversification tool to an optimization tool for managing the risk–return profile. This implies that financial institutions have to develop forward-looking models (i.e. to cover tail/extreme events) and decision-making tools that cover the amount of available capital, leverage adjustment costs, and the duration mismatches of assets and liabilities.

The fundamental basis for every ALM model is to define future scenarios of risk parameters and value assets and liabilities. One of the main challenges in that process is to come up with scenarios. These scenarios are usually based on historical observations or forward-looking simulations (using Monte Carlo) and typically do not cover tail risks—the so-called extreme events.

It is clear that stress tests are much needed in order to complement the usual VaR measures as a foundation for risk-adjusted decision-making. However, the traditional stress-testing approaches used by market participants and/or requested by the regulators suffer from a fundamental problem: there is no attempt to assign probabilities to the scenarios considered.

A framework is needed to express the likelihood of the various stress scenarios. A specific probability can be given to a stress test in, usually, two ways:

- on a nonobjective or judgmental basis—for example, by an economist/ expert who provides context-sensitive and conditional stress scenarios;
- on an objective basis using historical data—i.e. one requiring a long period of history in order to observe stressed situations.[1]

Stress-testing as a risk management tool has been in existence for more than a decade but was not really applied by the financial services industry as an enhancement of the daily decision-making process. The reasons for this reluctance are well explained by Aragones, Blanco, and Dowd (2001) (quoted by Rebonato, 2010):

"…the results of [traditional] stress tests are difficult to interpret because they give us no idea of the probabilities of the events concerned, and in the absence of such information we often don't know what to do with them. …As Berkowitz [1999] nicely puts it, this absence of probabilities puts 'stress testing in a statistical purgatory. We have some loss numbers, but who is to say whether we should be concerned about them?' …[we are left with] two sets of separate risk estimates—probabilistic estimates (e.g. such as VaR), and the loss estimates produced by stress tests—and no way of combining them. How can we combine a probabilistic risk estimate with an estimate that such-and-such a loss will occur if such-and-such happens? The answer, of course, is that we can't. We therefore have to work with these estimates more or less independently of each other, and the best we can do is use one set of estimates to check for prospective losses that the other might have underrated or missed…"

The main goal in this chapter is to explore ways in which a probability number can be assigned to stress tests in order to make sense of them and be able to integrate them within ALM in a meaningful manner.

Asset–Liability Management: Stress testing

A financial institution will traditionally perform stress tests by stressing certain variables such as interest rates, default rates, etc., with a view to analyzing the impact of such movements on its balance sheet. No assessment of the likelihood of such scenarios is attempted. This state of affairs is clearly unsatisfactory. In order to price risk, we need the probabilities of outcomes. This includes both the probabilities of recurring events and the probabilities of extreme events.

There are two ways in which to approach this problem.

- We can make use of extreme value theory to fit an appropriate joint probability distribution of exceedances to the historical distribution of extreme events. We could call this the "frequentist" approach, where the data are left to speak for themselves.

- We can postulate a model of the world in which the causal links between extreme events are determined, leading to a more intuitive determination of the joint probability of extreme events. This is known as the "subjective" approach and makes use of Bayesian theory.

Important Results

Stress tests are performed by financial institutions as part of their asset–liability management to assess the impact of large movements of underlying economic variables on their balance sheet.

- Stress tests by themselves are meaningless. To make sense of stress tests, the probabilities associated with large movements need to be determined.
- There are two ways in which one can attempt to determine the joint distribution of extreme events: the *frequentist* approach and the *subjective* approach.

The Frequentist Approach

The first approach we will consider is the frequentist approach. Here one tries to fit a joint probability distribution function to the data that are available. The difficulty here arises from the different *shapes* of the distribution implied by the data for common or rare events. Although for usual events (relatively small movements in the underlying variables) the normal distribution can be a good fit, this is not the case for rare events.

To see why, let us remind ourselves of the central limit theorem (CLT). CLT states that the sum of a very large number of independent variables, each with finite variance, is normally distributed. Relatively small movements in the underlying variables tend to happen under normal market conditions, when the underlying risk factors are largely independent of each other. So it comes as no surprise that a normal distribution should be a good fit under normal market conditions. However, during times of market turbulence, "correlations among asset classes become more polarized, tending towards +100% or –100%" (Rebonato, 2010). This means that in times of market turbulence the CLT does not apply, and indeed we observe that the normal distribution is a very bad fit. The obvious example that comes to mind is the recent credit crisis. After the collapse of Lehman Brothers, it became very difficult for an investor to diversify his or her market position efficiently, because the correlations between the various asset classes converged to 100%.

To quote Greenspan (1995): "From the point of view of the risk manager, inappropriate use of the normal distribution can lead to an understatement of risk, which must be balanced against the significant advantage of simplification. …Improving the characterization of the distribution of extreme values is of paramount concern."

A sophisticated ALM model must be able to include the right probability distribution for extreme events.[2] The problem comes down to finding the probability distribution that best fits the available data. A whole section of statistics, known as extreme value theory (EVT), is devoted to this task. Standard EVT techniques can be efficiently applied when the dimensionality is relatively low. Dimension reduction techniques can be employed, but even after a successful reduction "an effective dimensionality between five and ten, say, still poses considerable problems for the application of standard EVT techniques. By the nature of the problem extreme observations are rare. The curse of dimensionality very quickly further complicates the issue." (Balkema and Embrechts, 2007).

Balkema and Embrechts (2007) propose an enlightening geometric theory for EVT. From a mathematical perspective, a geometric theory is appealing as the theory applying to the *objects* (vectors representing portfolio positions) will be invariant under coordinate transformations.

In the univariate case, i.e. for one variable only, the condition that extreme scenarios can be described by a probability distribution leads to a one-parameter family of fat-tail shapes, the generalized Pareto distribution (GPD) (Balkema and Embrechts, 2007):

$$G\mathrm{xi}(\upsilon)=1-(1+\xi\upsilon)^{-\frac{1}{\xi}}, \xi \geq 0$$

The shape parameter ξ determines how "fat" the tail is—i.e. how much more frequent extreme events are than in the normal distribution. A large value of ξ means a distribution close to the normal distribution, whereas a small value of ξ means a very fat tail. By continuity, G_0 is the standard exponential distribution function $G0(\upsilon)=1-e^{-\upsilon}, \upsilon \geq 0$.

As an example, let us consider the history of the S&P 500 Index. The values and relative movements of the S&P 500 over the period 1968–2008 are shown in Figure 1 and Figure 2, respectively. A quick look at the daily movements in Figure 2 reveals a lot of relatively small movements, but also a significant number of very large movements. This suggests the existence of a normal "core" and a fat tail.

Figure 1. Values of S&P 500 Index, June 1968–January 2008. (*Source*: Bloomberg)

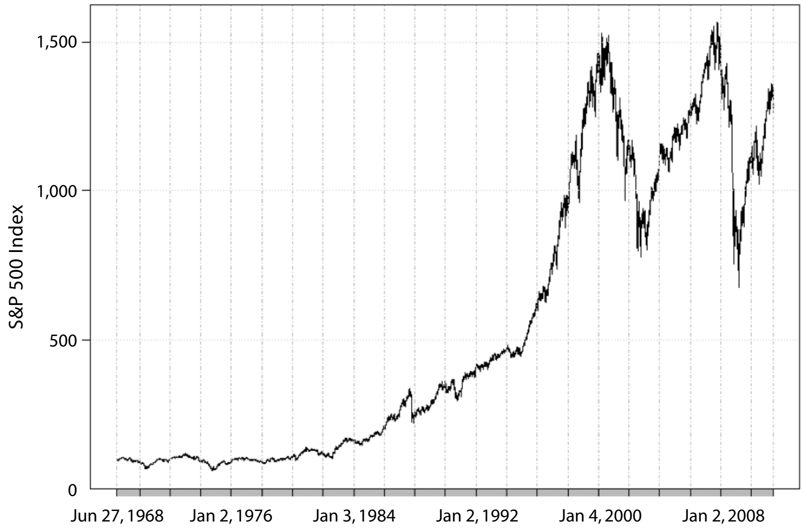

Figure 2. Relative movements of S&P 500 Index, June 1968–January 2008. (*Source*: Bloomberg)

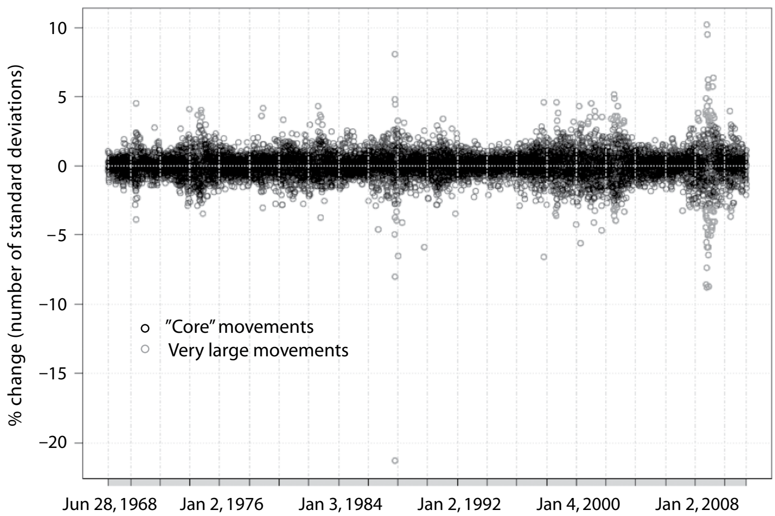

Indeed, we can fit a GPD to the daily log-differences with a varying number of exceedances[3] and analyze how the shape of the distribution varies as a function of the threshold[4] that determines the number of exceedances. Figure 3 shows how the shape of the tail x_i varies with this threshold.

Figure 3. The shape parameter xi as a function of exceedances for relative movements of S&P 500 Index, June 1968–January 2008.

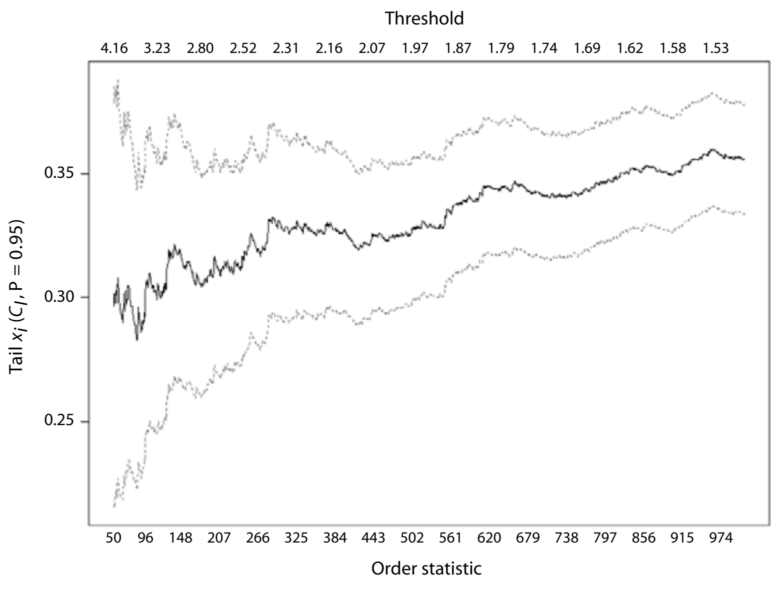

By varying the threshold that determines the number of movements larger than the threshold itself (the order statistics), we obtain a different value for the shape parameter. The further we are from the core—i.e. the larger the threshold and the smaller the number of data points to which the GPD is to be fitted—the better the fit to the fat tail becomes until, eventually, the number of data points becomes far too small to draw any meaningful conclusion. In the case of the S&P 500 we can see that the value of the shape parameter that best fits the tail lies somewhere 0.25 and 0.35. For a number of exceedances less than 50, the associated error in the calculation increases and it is no longer possible to draw any significant conclusion.

In Figure 4 we can see how nicely our GPD distribution function fits the extreme events in the S&P 500 data. For comparison, we show the normal distribution that fits the core. What was described by the normal distribution as a once-in-200-year event is seen now to be a once-a-year event!

Figure 4. The GPD fit to extremes in absolute movements of S&P 500 Index, June 1968–January 2008.

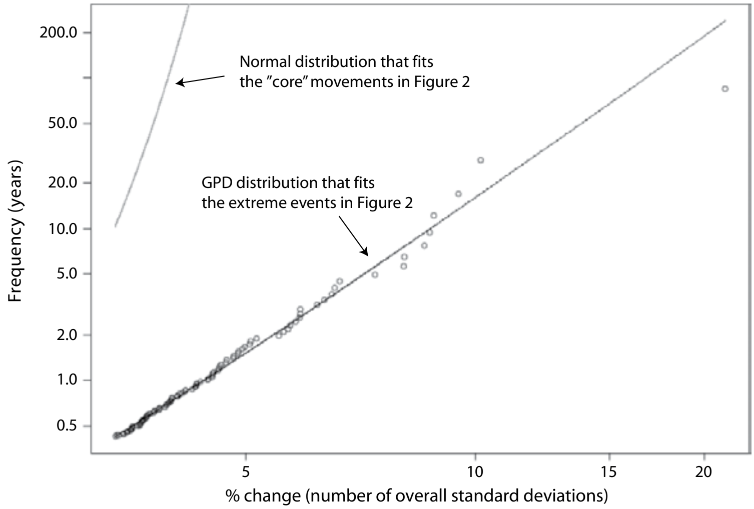

The corresponding multivariate theory is described in great detail by Balkema and Embrechts (2007). It is not trivial to expand to more than one dimension. However, one can in general define a metric in the multidimensional space such that the size and direction of the movements become well defined.

To illustrate this, let us consider the following two-dimensional example: a portfolio composed of only two hypothetical assets a and b. We shall define Δa and Δb as the number of standard deviations away from the mean for movements in the net asset value (NAV) of asset a and asset b, respectively. We then define the two-dimensional distance $r = \sqrt{\Delta a^2 + \Delta b^2}$ and the angle θ that determines the direction of the joint movement as $\theta = \tan^{-1}(\frac{\Delta b}{\Delta a})$. Figure 5 shows the result of plotting the points for which the joint relative movements are larger than three.

Figure 5. Hypothetical two-dimensional asset space showing percentage changes in NAV for assets *a* and *b*

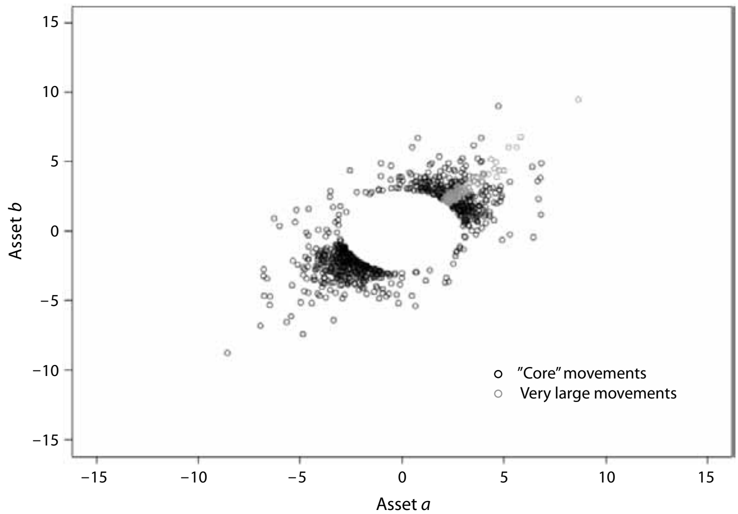

A quick look at Figure 5 suggests that tails can be measured for given directions, such as the tail shown in gray corresponding to an angle of 45°. In our hypothetical data set, extreme movements in the NAV of asset a are positively correlated with extreme movements in the NAV of asset b. This correlation can be viewed in Figure 6, which shows how the number of joint extreme events (e.g. r larger than 3) per 10° angle aperture changes with θ and reaches a maximum at $\theta = 45°$, indicating a positive correlation between extreme movements in the NAVs of our hypothetical assets a and b.

Figure 6. Number of extremes as a function of the "angle" defined by the movements in the hypothetical two-dimensional asset space.

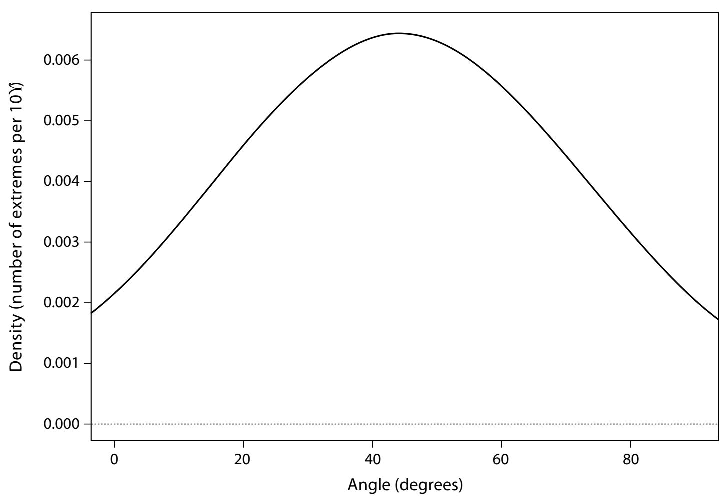

Now we can, for example, define the region between the angles θ_1 and θ_2 and look at the distribution of extremes within that area. By fitting a GPD to the data points located between θ_1 and θ_2 we could, in principle, extract the value of the shape parameter. Fitting a GPD to data points defined by another angle area would, in general, yield a different value of the shape parameter. This shape parameter could then be indexed with the angle for a given angle aperture. The issue here resides with the quality of the data available. To be able to index the shape with the angle requires a certain number of extreme events to be sampled per angle aperture. We can, if the number of data points is large enough, parameterize the shape of the tail distribution of returns with the angle defined by the two asset returns.

This is indeed a very interesting result. The direction is itself determined by a pair of numbers (in the two-dimensional case) representing the relative returns for the two assets. If we were to price, say, a simple out-of-the-money (OTM) hybrid determined by this pair of numbers, the shape parameter of its corresponding tail distribution would be the crucial quantity to calculate. The price would have a one-to-one relationship with the shape parameter ξ.

Now, imagine that prices for OTM hybrids are computed for each direction using this technique. It is not difficult to see how one can take advantage of an arbitrage situation: there will be an arbitrage situation whenever there are significant differences between the prices computed using the GPD fit and the market prices.

Important Results

- The central limit theorem (CLT) is appropriate during normal market conditions. CLT implies a normal (Gaussian) distribution of market movements.
- CLT is not appropriate for extreme market movements. The *tails* of the distributions of market movements are *fat*, i.e. extreme events are more probable than otherwise predicted by CLT.
- A generalized Pareto distribution (GPD) can be used to fit the tails of the distribution of market movements. Its shape parameter ξ determines how fat the tail is.
- A generalization to a multivariate theory is not trivial. However, one can, in general, define *directions* in the space of the market movements.
- When there are enough data, the shape of the tail can be parameterized as a function of an angle defining the direction.

The Subjective Approach

In his book, Riccardo Rebonato (2010) explains how one can draw conclusions about the joint probability of extreme events by making use of causality networks. We will explore this concept in some detail in this section.

In the previous section, where the frequentist approach was outlined, we placed all the emphasis on the level of *association* between variables. The subjective approach, on the contrary, places the emphasis on the *causal links* between variables. The main advantage of this approach is the fact that it is cognitively much easier and more natural.

To illustrate what we mean, consider the following example. Suppose that the variable we are interested in is whether a particular church in Lisbon is damaged or not. We know that in 1755 an earthquake and tsunami destroyed vast areas of the city. The other variables in this example could be, say, whether the church was damaged by the earthquake or not or whether there was a fire. One could take a purely associative approach and build all the relevant probability tables. To do this, we need some numbers such as the standalone probabilities (the marginals), which are relatively easy to calculate, and some singly conditioned probabilities (the probabilities of one event, conditional on another). These singly conditioned probabilities could be, in turn, either simple and natural, such as the probability that the church was damaged given that an earthquake had occurred, or they could be difficult and awkward, such as the probability that an earthquake had occurred given that the church was damaged. The first formulation is of a *causal* nature, which is why we find it cognitively easier to arrive at an answer (the probability would be close to one), whereas the second is *diagnostic* in nature, and because there are many possible causes for the same effect, the answer in the second case is hard to guess.

There is another reason why causal models are more powerful than associative models: the fact that small changes in the causal structure of the model can give rise to large changes in the joint probabilities. It is easy to encode changes in the causal links between variables. However, from a purely associative point of view, they may be very difficult to explain.

Let us outline the goals of the subjective approach: the final goal is—just as in the frequentist approach—to gain access to the joint distribution of extreme events. However, instead of attempting to fit a generalized Pareto distribution to the data directly, as we do in the frequentist approach, we inject more information about how we expect the world to behave. It then becomes possible to derive the full joint distribution from a small number of marginals, singly conditioned probabilities, and (at most) doubly conditioned probabilities. This is achieved by applying Bayes' theorem across the causality net and using the concept of conditional independence.

A simple example will help to explain how this is done. Consider the events A, B, C, and D, which are defined thus:

- A = Earthquake;
- B = Fire;
- C = Tsunami;
- D = Church on the hill is damaged.

And the very simple model of our world:

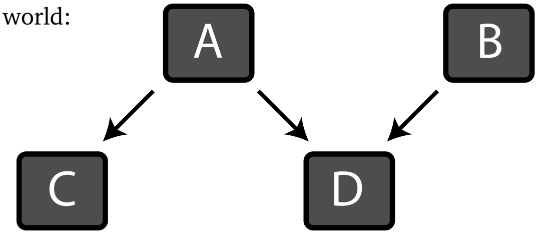

In this model, A causes C and D; and B causes D. A and B, the earthquake and the fire, are assumed to be independent (they are the roots of the causality net). Note that all the information affecting C originates from A. Hence, given A, C, and D are independent, C and D are said to be *independent*, *conditional* on A. This characteristic of Bayesian nets is crucial to the evaluation of the full joint distribution.

In this very simple model, the joint probability distribution is defined by $2^4 - 1 = 15$ numbers (all possible combinations of the four Boolean variables minus 1 from the condition that the total cumulative probability must equal unity). Utilizing the information provided by the causality net and making use of Bayes' theorem allows us to derive all 15 numbers from only $4 + 3 = 7$ numbers (four marginals plus three singly conditioned probabilities). In general, and as long as we keep the causality net simple, we are reducing a $2^n - 1$ problem to a $2n - 1$ problem. For a large n, this is a massive simplification.

Let us calculate the probability of one joint event in our mini-model to see how this is done in practice. Starting from the marginals P(A), P(B), P(C), and P(D), and the conditional probabilities P(C|A) and P(D|A), let us calculate the probability that there was a tsunami, that an earthquake has occurred, that there was a fire, but that the church is not damaged. We will define this as P(A, B, C, ~D):

$P(A, B, C, \sim D) = P(C, A, \sim D, B)$

$= P(C|A, \sim D, B) \times P(A, \sim D, B)$[5]

$= P(C|A) \times P(A, \sim D, B)$[6]

$= P(C|A) \times P(A|\sim D, B) \times P(\sim D, B)$[7]

$= P(C|A) \times P(A|\sim D) \times P(\sim D, B)$[8]

$= P(C|A) \times [P(\sim D|A) \times P(A)/P(\sim D)] \times P(\sim D, B)$[9]

$= P(C|A) \times [1 - P(D|A)] \times P(A) \times P(\sim D, B)/[1 - P(D)]$[10]

$= P(C|A) \times P(A) \times [1 - P(D|A)] \times P(\sim D|B) \times P(B)/[1 - P(D)]$[11]

$= P(C|A) \times P(A) \times [1 - P(D|A)] \times [1 - P(D|B)] \times P(B)/[1 - P(D)]$[12]

However convoluted this calculation might look, the important result is that we are able to obtain the joint probability only from the marginals and the singly conditioned probabilities, making use of Bayes' theorem and our specific model of the world, the causality net.

Note, however, that it is sometimes impossible to obtain a meaningful value for the joint probability (a number between zero and one), given a specific set of inputs. This imposes bounds on the initial marginals and singly conditioned probabilities defining the subset of feasible inputs.

A fully automated system can be built, given a particular causality net and set of feasible inputs. The topological structure of the causality net must be characterized in a way that can be understood by a computer algorithm. Linear programming can then be used for the joint distribution (Rebonato, 2010).

Important Results

- Some conditional probabilities seem more *natural* than others. This is explained by the causal links between variables.
- If we postulate that we understand the way the world works through a causality net, we add more information to the *natural* probabilities that are easy to compute.
- Using Bayes' theorem, and in the case when the inputs constitute a *feasible* solution, one can, in general, recover the full joint distribution of extreme events.

However, the choice of model remains subjective.

Case Study

Our case study is an EVT analysis of a global macro fund and a distressed fund. Our analysis focuses on the distribution of extreme events for both funds and their classification, and concludes with a comparison with other "traditional" risk measurement techniques, such as VaR. Although this could be considered pure asset management (rather than asset–liability management), we think that it illustrates the issues surrounding stress-testing that are discussed in this chapter.

We analyze the percentage changes in the net asset values over the last eight years. The results for the HFR Global Macro fund are shown in Figure 7, where the percentage changes are quoted as the number of total standard deviations for the period.

Figure 7. Global macro fund: changes in net asset value 2003–10

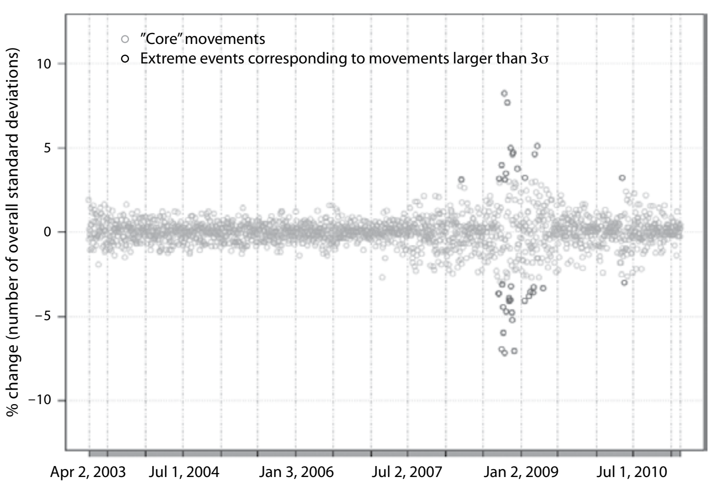

One might be tempted to identify extreme events as those corresponding to movements larger than three standard deviations (see Figure 7). However, we would like to differentiate between two very different types of extreme movement:

- type 1: the type of extreme that is driven by volatility;
- type 2: the type of extreme that is a genuine "black swan" (fat-tail) event.

The first type of extreme appears to be an extreme simply because the volatility has increased, whereas the second type is a genuine extreme because the volatility has not increased, yet the event still occurred. In order to correctly identify the genuine extremes (i.e. those of type 2), one must rescale the percentage changes to a moving average measure of the local volatility prior to performing the GPD fit. Figure 8 shows how the movements compare to this local definition of extreme in the case of the global macro fund. Here the dark lines correspond to the limits plus three standard deviations and minus three standard deviations that define an extreme event.

Figure 8. Global macro fund

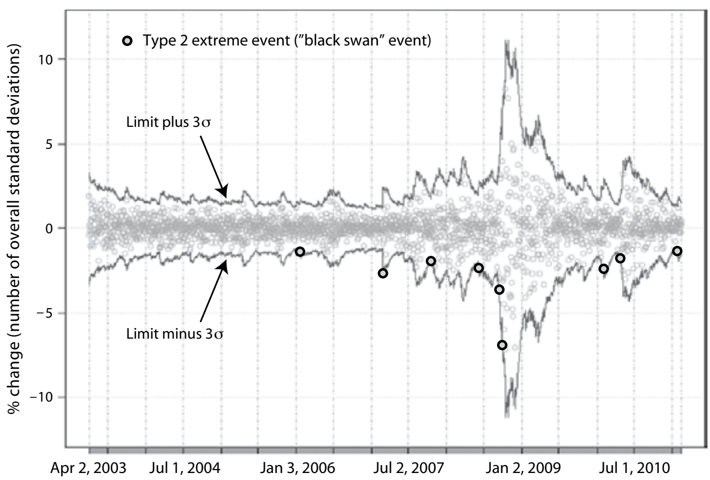

It now becomes apparent that all the large positive movements in the global macro fund occur *after* there is an increase in volatility, making them type 1 extremes. However, some negative extremes occur *before* the increase in volatility, making these type 2 extremes. The best-fit distribution for positive movements is indeed the normal distribution, as shown in Figure 9.

Figure 9. Global macro fund: best-fit distribution for positive movements

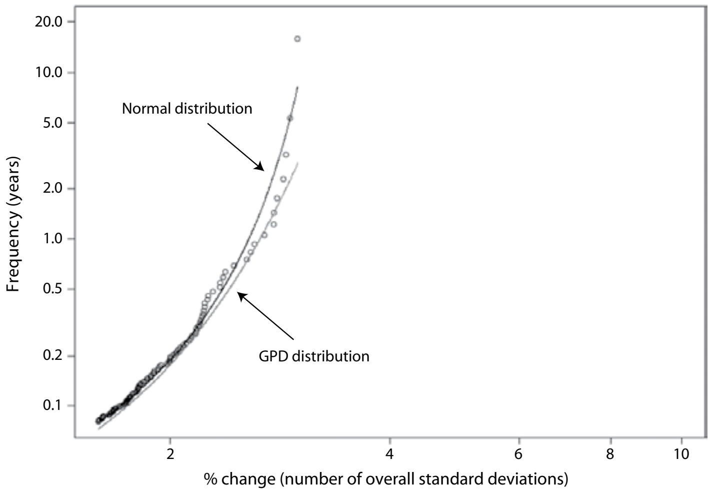

By contrast, the HFR Distressed-strategy fund shows genuine fat tails for positive movements, as we can see in Figure 10. Note that some extreme events occur *before* there is an increase in volatility—in fact it is they themselves that cause the spikes in volatility. This makes these events type 2 extremes.

Figure 10. Distressed-strategy fund: Changes in net asset value 2003–10

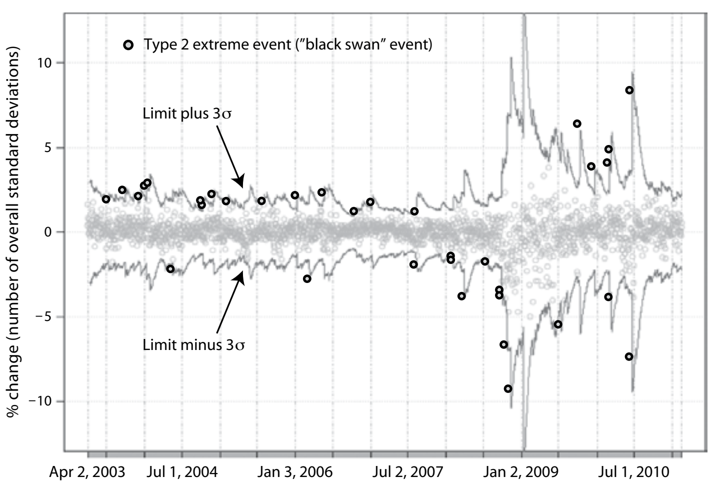

It is also interesting to note that there are large positive movements during normal times when the volatility of the market is relatively small and stable. This indicates that the distressed strategy is working for this fund.

A GPD fit for positive movements in the HFR Distressed-strategy fund shows, unlike the HFR Global Macro fund, a very fat tail, as shown in Figure 11.

Figure 11. Distressed-strategy fund: best-fit distribution for positive movements

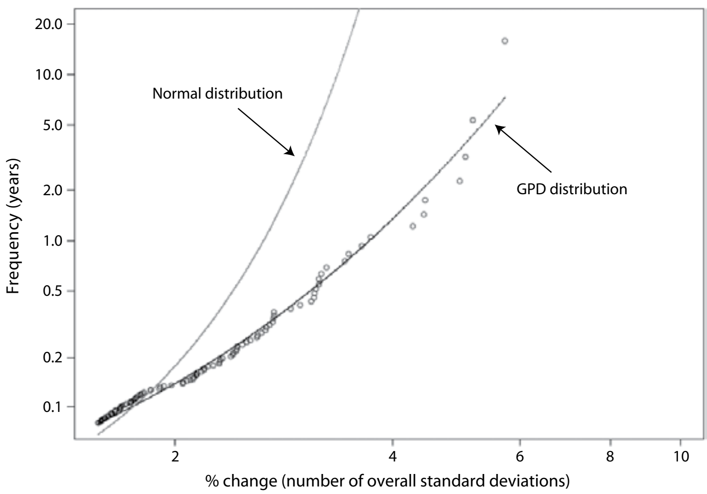

Let us consider an investor with a short position in either of these funds who is concerned with the measurement of his or her risk. In the case of the global macro fund, the *usual* way of calculating VaR—i.e. measuring the standard deviation and assuming a normal distribution—would produce a realistic assessment of risk because the normal distribution is a good fit for positive increments in the net asset value (negative movements in the investor's position). On the other hand, the same calculation for the distressed fund would produce a highly unrealistic assessment of risk as it would fail to capture the tail.

Summary and Further Steps

- "Traditional" stress-testing is done on a *standalone* basis. It is then not possible to combine probabilistic estimates of risk (such as VaR) with the loss estimates produced by the stress tests. This situation renders the (traditional) stress tests meaningless. To make sense of stress tests, the probabilities associated with extreme events need to be determined.
- There are two ways in which one can attempt to determine the joint distribution of extreme events: the *frequentist* approach and the *subjective* approach.
- In the frequentist approach (aka the "let the data speak" approach), one attempts to fit a probability distribution to extreme events directly, using whatever data are available.
- In the subjective approach, a model of the world is postulated in which the causal links between the various variables are established.

- Normal distributions are appropriate during normal market conditions, when underlying variables are largely independent of each other. During periods of market turbulence, however, correlations become more polarized, the central limit theorem no longer applies, and the normal distribution is no longer a good fit.
- The generalized Pareto distribution (GPD) is the appropriate fit to extreme events. Its shape parameter ξ determines the fatness of the tail. The smaller the value of ξ, the fatter the tail. The larger the value of ξ, the closer the distribution becomes to the normal distribution.
- If one follows the subjective approach, the addition of extra information in the form of a causality net (or model of the world) allows calculation of the full joint distribution of extreme events starting from a relatively small number of inputs. These inputs are the marginal distributions (the standalone distributions) and some natural conditional probabilities.
- We define two types of extreme event. A type 1 extreme is one that is driven by volatility, and, as such, it happens *after* there is an increase of volatility. A type 2 extreme is the genuine black swan that happens *before* there is an increase in the volatility—one could say that the volatility is driven by the type 2 extreme.
- Traditional calculations of VaR assume normal distributions.
- Whereas for assets that are prone to type 1 extremes the traditional calculation of VaR might produce a realistic assessment of risk, for assets prone to type 2 extremes the same calculation is wholly unrealistic, as it fails to capture the tail of the distribution.
- One further step that can be taken is to apply the multivariate EVT techniques outlined in the section on the frequentist approach to a space of investment classes, parameterize the tail shape parameter as a function of direction, and explore arbitrage opportunities between different directions.
- Another interesting line of research would be to try to combine the frequentist and subjective approaches, in effect testing the robustness of a given model of the world.

More Info

Books:

Balkema, Guus, and Paul Embrechts. *High Risk Scenarios and Extremes: A Geometric Approach*. Zürich: European Mathematical Society, 2007.

Bouchaud, Jean-Philippe and Marc Potters. *Theory of Financial Risk and Derivative Pricing*. Cambridge: Cambridge University Press, 2009.

Rebonato, Riccardo. *Coherent Stress Testing: A Bayesian Approach to the Analysis of Financial Stress*. Chichester, UK: Wiley, 2010.

Articles:

Aragones, J., C. Blanco, and K. Dowd. "Incorporating stress tests into market risk modelling." *Derivatives Quarterly* 7:3 (2001): 44–49.

Berkowitz, Jeremy. "A coherent framework for stress-testing." *Journal of Risk* 2:2 (1999): 5–15.

Dermine, Jean. "ALM in banking." Finance and banking working paper, INSEAD, Fontainebleau, July 17, 2003. Online at: tinyurl.com/76tk8gm

Greenspan, Alan. Presentation to Joint Central Bank Research Conference, Washington, DC, 1995.

Notes

1. Historical data are available to estimate probability—for example:
   - credit spreads: Baa and Aaa spreads back to the 1920s;
   - default frequency by rating: from rating agencies back to the 1920s;
   - equities: S&P 500 Index back to the 1920s;
   - interest rates: Treasury bond yields back to the 1920s;
   - crude oil prices: back to 1946;
   - foreign exchange: historic data may not be so meaningful because currencies change roles.
2. For the purposes of this chapter we define an "extreme event" as a market movement larger than three standard deviations away from the mean.
3. By exceedances we mean the number of extremes to which the GPD is to be fitted.
4. The threshold is the number of standard deviations above which data points are considered for the purpose of fitting a GPD.
5. From Bayes' theorem.
6. From conditional independence, C is independent of D and B, conditional on A.
7. From Bayes' theorem.
8. A and B are independent.
9. From Bayes' theorem.
10. From completeness.
11. From Bayes' theorem.
12. From completeness.