Stress Testing ALM models: A Coherent Approach

Alexandre Canavezes, Phd, Quant Analytics
Mario Schlener, NCMA

This chapter Covers:

- INTRODUCTION
- MAIN TEXT
    - The Frequentist Approach
    - The Subjective Approach
- CASE STUDY
- SUMMARY AND FURTHER STEPS

INTRODUCTION

Typing "Stress Tests" in Google produces the staggering amount of nearly 60 million hits. This result shows the popularity of the topic. Indeed, in light of recent extreme events, both the industry and the regulators have keenly felt the need to complement traditional percentile-based risk management tools (i.e. such as Value at Risk or economic capital) with stress tests and scenario analyses.

Especially model innovation and sophistication in the financial industry and knowledge of cross risk interrelationships (i.e. market and credit risk correlation) developed ALM from the usual practice of restricting itself to the control of interest rate (e.g. duration and convexity) and liquidity risks arising from positions on the balance sheet (i.e. banking book) to a comprehensive and integrated framework.

Following the logic of Dermine[1] , Asset & Liability Management 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 to discuss, in an integrated way, fund transfer pricing, deposit pricing (i.e. 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 risks, the marginal risk contribution, and also the allocation of economic capital.

Learning from past misbehaviors of all market participants (especially the over reliance on quantitative measures with "statistical entropy" and 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. cover tail/ extreme events) and

[1] Jean Dermine, "ALM in Banking", INSEAD, 2003

decision making tools which cover the amount of available capital, leverage adjustment costs and the duration mismatches of assets and liabilities.

The main 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 (i.e. using Monte Carlo) and typically they do not cover tail risks - the so called extreme events. The inclusion of tail risks, for individual risk factors as well as joint outcomes of several factors, through a comprehensive stress testing process needs to be applied for defining price risk and/or value risk of asset and liabilities.

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, 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 in order to express the likelihood of the various stress scenarios. A specific probability can be given to a stress test in usually 2 ways:

- Non objective basis or judgmental basis, e.g. by an economist/ expert (i.e. provides context-sensitive and conditional stress scenarios)
- Objective basis using historical data (i.e. requires long period of history in order to observe stressed situations[2])

The ALM manager has two alternatives to treat the mismatch between assets and liabilities: He/She will either try to hedge that risk, or they can treat it as an asset class on its own right, i.e., the ALM manager can invest in it. Short term horizons are popular because they match the liquidity horizons that investors are willing to accept. It both cases, it is fundamental to have a correct understanding and measurement of the probability of extreme events.

Stress testing as a risk management tool has been in existence already over a decade but was not really applied by the industry as an enhancement of the daily decision making process. The reasons for this reluctance are well explained by Aragones, Blanco and Dowd[3]:

"…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

[2] Historical data is available to estimate probability e.g.

- Credit spreads: Baa and Aaa spreads back to 1920s
- Default frequency by rating: From Rating agencies back to 1920s
- Equities: S&P 500 back to 1920s
- Interest rates: treasury bond yields back to 1920s
- Crude oil prices: back to 1946
- FX: History may not be so meaningful as currencies change roles

[3] J. Aragones, C. Blanco, K. Dowd, 2001, "Incorporating Stress Tests into Market Risk Modelling"

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 …"

One main goal in this article 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.

MAIN TEXT

A financial institution will traditionally perform stress tests by stressing certain variables such as interest rates, default rates, etc, with the view to analyze the impact of such movements on its balance sheet. No assessment of the likelihood of such scenario is attempted. This presents an obvious problem.

This state of affairs is clearly unsatisfactory. In order to price risk, we need the probabilities of outcomes. This includes both the probabilities of frequent events and the probabilities of extreme events.

There are two ways in which to approach this problem:

1. We can make use of extreme value theory (EVT), in order to fit an appropriate joint probability distribution of exceedances to the historical distribution of extreme events. (Frequentist approach) – Let the data speak for itself.

2. 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 (Subjective approach) – 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. In order 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 is the frequentist approach. One tries to fit a joint probability distribution function to what data is available. The difficulty here arises from the different *shapes* of the distribution implied by the data for common or rare events. Whilst 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 why a normal distribution is 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%"[4]. 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%.

Quoting Greenspan (1996):

…"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[5]. The problem comes down to finding the probability distribution that best fits the available data.

A whole section of statistics is devoted to this task: Extreme Value Theory (EVT). 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."[6]

[4] Riccardo Rebonato, "Coherent Stress Testing"

[5] For the purposes of this article we define "Extreme Event" as a market movement larger than three standard deviations away from the mean

[6] Guus Balkema, Paul Embrechts, "High Risk Scenarios and Extremes. A geometric approach"

Embrechts () et al. propose a geometric theory for EVT that we find enlightening. 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)[7]:

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

The shape parameter ξ determines how "fat" the tail is, i.e. how much more frequent extreme events are when compared to the normal distribution case. 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 $G0$ is the standard exponential distribution function $G0(v) = 1 - e^{-v}, v \geq 0$

As an example let us consider the S&P 500 history. The S&P500[8] values and relative movements are shown in Figure 1 and Figure 2 respectively.

Figure 1: S&P 500

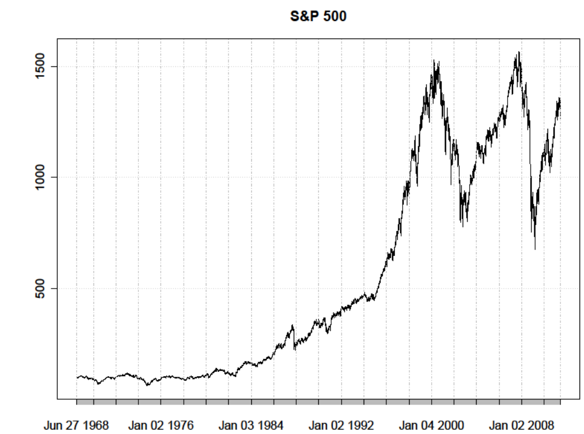

Figure 2: S&P 500

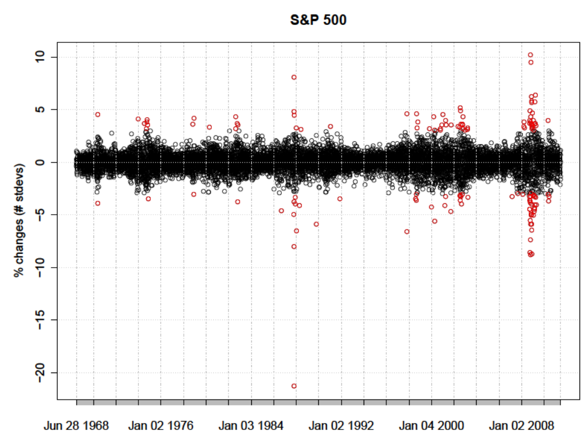

A quick look at daily movements reveals a lot of relatively small movements (shown in black), but also a significant number of very large movements (shown in red). This suggests the existence of a normal core (black points) and a fat-tail (red points).

Indeed, we can fit a GPD to the daily log-differences with a varying number of

[7] Guus Balkema, Paul Embrechts, "High Risk Scenarios and Extremes. A geometric approach"
[8] Source: Bloomberg

exceedances[9] and analyze how the shape of the distribution varies as a function of the threshold[10] that determines the number of exceedances.

Figure 3 shows how the shape of the tail $\mathbf{x}_i$ varies with this threshold.

Figure 3: S&P 500

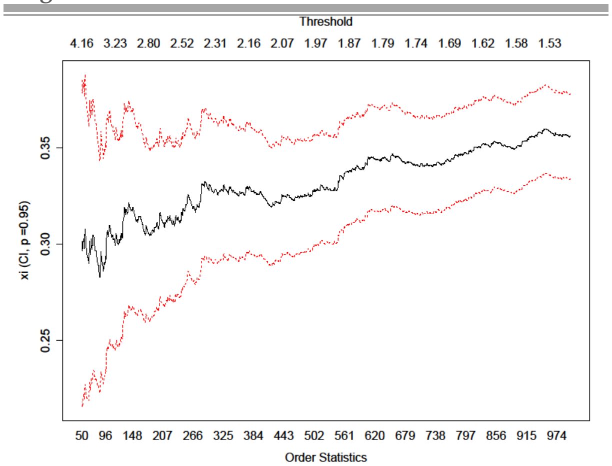

By varying the threshold that determines the number of movements larger than the threshold itself (the order statistics), we obtain a different value for the shape parameter. The furthest we are from the core, i.e., the larger the threshold and the smaller the number of data points to which the GPD is to be fit, the better the fit to the fat tail becomes until, eventually, the number of data points becomes far too small to draw any meaningful conclusions. 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 lower than 50, the associated error in the calculation increases and it is no longer possible to draw any significant conclusion.

In Figure 4 below, 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 in red. What was described by the normal distribution by a once-every-200-years' event is indeed a once-a-year event!

Figure 4: S&P 500

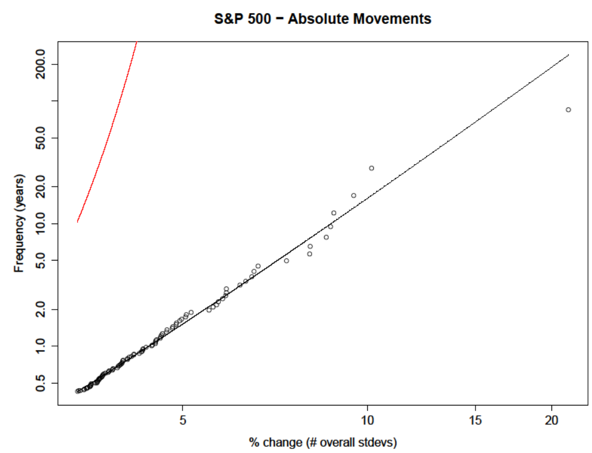

The corresponding multivariate theory is described in great detail by Embrechts and Balkema[11]. It is not trivial to expand to more than one dimension. However one can in general define a metric in the multidimensional space such that size and direction of the movements become well defined.

To illustrate this let us consider the following 2-dimensional example: A portfolio composed of

---

[9] By exceedances we mean the number of extremes to wich the GPD is to be fit

[10] The threshold is the number of standard deviations above which data points are considered to purpose of fitting a GPD

[11] Guus Balkema and Paul Embrechts, "High Risk Scenarios and Extremes, A geometric approach"

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 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 determining 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 3.

Figure 5: Hypothetical 2D Asset Space

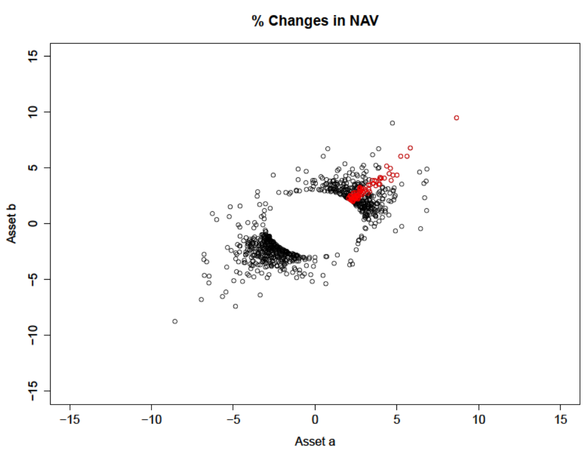

A quick look at Figure 5 suggests that tails can be measured for given directions, such as the tail shown in red (corresponding to an angle of 45°). In our hypothetical data set, extreme movements in the NAV of asset a are positively correlated with extreme movements in the NAV of asset b

This correlation can be viewed in Figure 6

Figure 6: Hypothetical 2D Asset Space

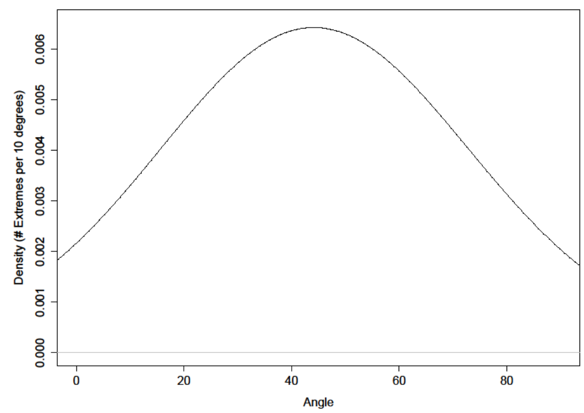

Figure 6 shows how the number of joint extreme events (e.g. r larger than 3) per 10° angle aperture changes with θ and reaches a maximum at θ = 45°, indicating a positive correlation between extreme movements in the NAVs of our hypothetical assets a and b.

Now we can, for example, define the region comprised between the angles θ_1 and θ_2, and look at the distribution of extremes within that area.

Fitting a GPD to the data points comprised 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 amount of data points is large enough , parametrize 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 2D case) representing the relative returns for the two assets. If we were to price, say, a simple 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 compuited 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 market movements's space**
- **When there is enough data, the shape of the tail can be parametrized as a function of an angle defining the direction**

The Subjective Approach

If instead of typing just "Stress Tests" on google, we search for the exact sequence "Stress Tests: A Coherent Approach", the number of hits is dramatically reduced to only four, all referring to the same work by Ricardo Rebonato.

In his book, Rebonato explains how one can draw conclusions about the joint probability of extreme events by making use of causality networks. We shall explain this concept in some detail in this section.

In the previous section, where the frequentist approach was outlined, we have placed all 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 that fact that it is cognitively much easier and *natural*.

To illustrate what we mean, consider the following example.

Suppose that the variable we are interested on 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 whether the church was damaged by the earthquake or not or whether there had been a fire for example. One could take a purely associative approach and build all the relevant probability tables. To do this, the risk manager would need some numbers such as the stand alone probabilities (the marginals), which would be relatively easy to calculate, and some singly conditioned probabilities (i.e. the probabilities of one eveny, conditional on another). This singly conditioned probabilities could be, in turn, either easy and natural, such as asking the question "What is the probability that the church is damaged, given that an earthquake has occurred?", or they could be difficult and awkward such as asking the question: "What is the probability that an earthquake has occured, given that the church is damaged?" The first question is of a *causal* nature and that is why we find it cognitively easier to answer (the number would be close to one), whereas the second question is of a *diagnostic* nature, and because there are many possible causes to the same effect, the answer to the second question is hard to guess.

There is another reason why causal models are more powerful than associative models: the fact that small changes to 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, as is also the case 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 becomes then possible to derive the full joint distribution from a small number of marginals, singly conditioned probabilities and (at most) doubly conditioned probabilities. The way this is done is by applying Bayes' theorem across the causality net and using the concept of conditional independence.

A simple example will help to understand how this is done:

Consider the events **A**, **B**, **C** and **D** defined thus:

A: Earthquake

B: Fire

C: Tsunami

D: Church on the hill is damaged

And the very simple model of our world:

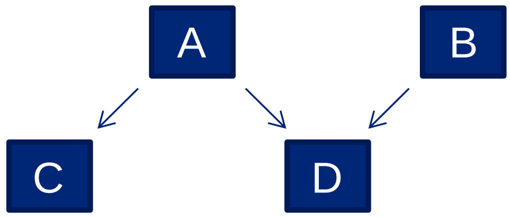

In this model, A causes C and D; B causes D. A and B are assumed to be independent (they are the roots of the causality net).

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 the possible combinations of the 4 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 (4 marginals plus 3 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 is a tsunami, an earthquake has occurred, there was a fire, but the church was not damaged. We will define this as P(A,B,C,~D):

P(A,B,C,~D) = P(C,A,~D,B)

= P(C|A,~D,B)*P(A,~D,B)[12]

= P(C|A)*P(A,~D,B)[13]

= P(C|A)*P(A|~D,B)*P(~D,B)[14]

= P(C|A)*P(A|~D)*P(~D,B)[15]

= P(C|A)*[P(~D|A)*P(A)/P(~D)]*P(~D,B)[16]

= P(C|A)*[1-P(D|A)]*P(A)*P(~D,B)/[1-P(D)][17]

= P(C|A)*P(A)*[1-P(D|A)]*P(~D|B)*P(B)/[1-P(D)][18]

= P(C|A)*P(A)*[1-P(D|A)]*[1-P(D|B)]*P(B)/[1-P(D)][19]

However convoluted this calculation might look, the important result is that we were 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

[12] From Bayes' theorem
[13] From conditional independence, C is independent from D and B, conditional on A
[14] From Bayes' theorem
[15] A and B are independent
[16] From Bayes' theorem
[17] From completeness
[18] From Bayes' theorem
[19] From completeness

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 be then used to solve for the joint distribution.[20]

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 woks 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 focus 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 it illustrates the issues surrounding stress testing that we discuss in this chapter.

We analyse the percentage changes of the NAV() over the last 8 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

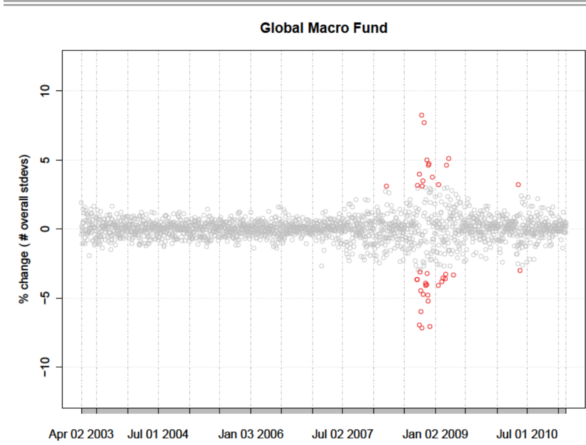

One might be tempted to identify extreme events as the ones corresponding to movements larger than 3 standard deviations (red in the picture). However, we would like to differentiate between two very different types of extreme movement:

[20] Riccardo Rebonato, "Coherent Stress Testing"

1. The type of extreme that is driven by volatility
2. The type of extreme that is a genuine black swan/fat tail event

The first type of extreme appears 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. extremes of type 2), one must re-scale the percentage changes to a moving average measure of the local volatility prior to performing the GPD fit. Figure 8 shows how movements compare to this local definition of extreme (the red lines correspond to the limits +3 standard deviations and -3 standard deviations that define an extreme event) in the case of the Global Macro fund

Figure 8: Global Macro Fund

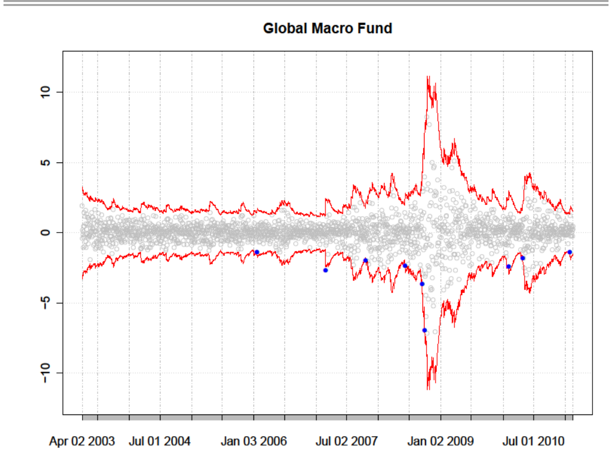

It become now apparent that all the large positive movements in the Global Macro fund occur *after* there is an increase in volatility, making them type 1 extremes. However, some negative extremes occur *before* the increase in volatility making these type 2 extremes (shown in blue). The best fit distribution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

Figure 9: Global Macro Fund

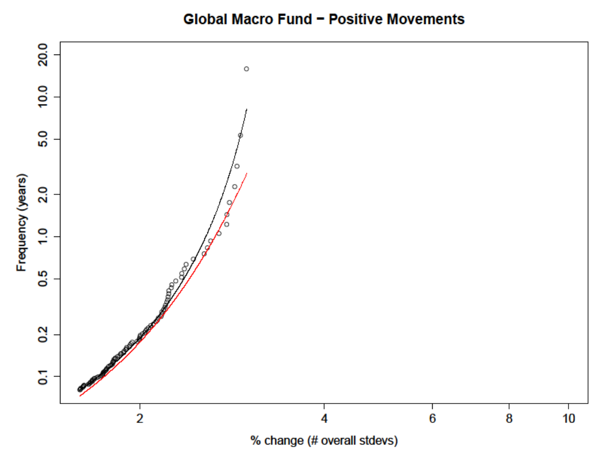

By contrast, the HFR Distressed fund shows genuine fat tails for positive movements, as we can see in Figure 10

Figure 10: Distressed Strategy Fund

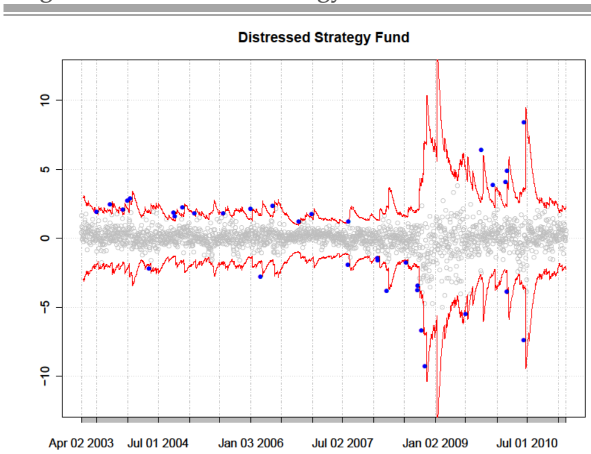

Note that some extreme events occur *before* there is an increase in volatility (shown in blue) – in fact it is themselves who cause the spikes in volatility. This makes these events type 2 extremes.

It is also very 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 fund shows, unlike the Global Macro fund, a very fat tail as shown in Figure 11

Figure 11: Distressed Strategy Fund

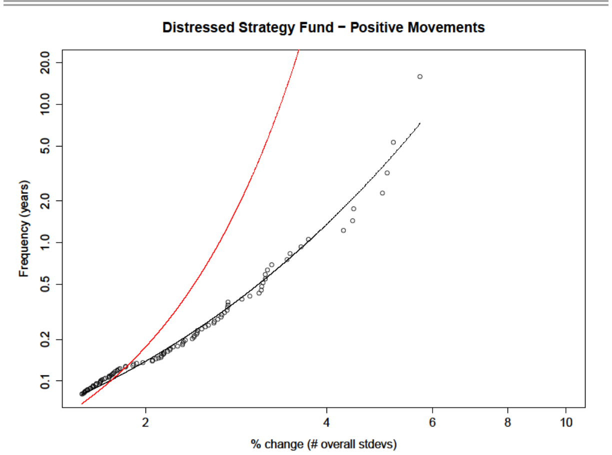

Let's consider an investor with a short position in either of these two funds, concerned with the measurement of his/her risk.

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 NAV (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 *stand-alone* 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. In order 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 "Let the Data Speak" approach), one attempts to fit a probability distribution to extreme events directly, using whatever data is 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 largest the value of ξ, the closest the distribution becomes to the normal distribution.
- When one decides to follow the subjective approach, the addition of extra information in the form of a causality net (or model of the world), allows for the calculation of the full joint distribution of extreme events, starting from a relatively small number of inputs. These inputs are the marginal distributions (the stand alone distributions) and some *natural* conditional probabilities.
- We define two types of extreme event: A Type 1 extreme is the extreme 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 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 step further will be to apply the multivariate EVT techniques outlined in the Frequentist approach section to a space of investment classes, parametrize 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 and combine the frequentist and subjective approaches, in effect testing the robustness of a given model of the world.