

A Coherent Aggregation Framework for Stress Testing and Scenario Analysis

Jan Kwiatkowski & Riccardo Rebonato^{*,+}

RBS

OCIAM - Oxford University

⁺Tanaka Business School, Imperial College

March 22, 2010

Abstract

We present a methodology to aggregate in a coherent manner conditional stress losses in a trading or banking book. The approach bypasses the specification of unconditional probabilities of the individual stress events, and ensures via a Linear Programming approach that the (subjective or frequentist) conditional probabilities chosen by the risk manager are internally consistent. The admissibility requirement greatly reduces the degree of arbitrariness in the conditional probability matrix if this is assigned subjectively.

The approach can be used to address the requirements of the regulators on the Instantaneous Risk Charge.

1 Introduction and Motivation

This paper¹ outlines a simple idea to aggregate in a coherent fashion the potential stress losses arising from a given set of trading-book positions.

In the wake of recent events, both the industry and the regulators have keenly felt the need to complement traditional frequentist, percentile-based risk management tools – such as Value at Risk or Economic Capital – with subjective stress tests and scenario analyses. One of the main unresolved questions is how to aggregate the results of these scenarios and stress tests. This is of particular relevance if a capital charge is to be levied against the combined stress losses, as the recently introduced Incremental Risk Charge (IRC) requires. See BIS (2008a), BIS (2008b). In the wake of recent market events the urgency of this capital charge has become evident, and the financial industry has been invited to submit proposals to address the problem. This note is a contribution in this industry effort.

¹We gratefully acknowledge the help of Dr Sukhdeep Mahal with coding.

Clearly, the main problem arises from the fact that simply adding the hypothetical losses from all the scenarios would result in a number that is of little meaning, and would certainly be prohibitively and unrealistically large. On the other hand, the regulators are correctly and understandably opposed to the endorsement of rules of thumb such as ‘square root of sum of squares’. The onus has been placed on the financial industry to come up with a realistic way to combine and aggregate the stand-alone stress losses in a coherent and justifiable manner.

The magnitudes of the stand-alone losses associated with the individual stress events can either be estimated using traditional frequentist approaches, or it could be based on expert judgement and/or subjective probabilities. Given the regulators’ examples of ‘events’ (eg, the depegging of a currency that has never depegged before), we strongly prefer a subjective-probability, non-frequentist framework. The risk manager, in our preferred approach, would start from an understanding of the portfolio under his watch, and determine severe but plausible moves of risk factors that would put under stress the largest positions in the portfolio. (See below for a stylized example.) However, we stress that the approach we propose works just as well whichever meaning (frequentist or subjective) is given to the stand-alone (marginal) probability of the individual stress events.

If the risk manager also believed that the data at his disposal were of sufficient quality and relevance to allow the estimation of the joint co-dependence of the stand-alone stress events (perhaps using a copula approach), then the procedure presented in this paper would be of no use. Our proposals are of interest in those situations where the available data are deemed to be either insufficient, or not sufficiently relevant to the problem at hand (eg, ‘backward-’ rather than ‘forward-looking’) or where the risk manager wants to stress the ‘objectively-determined’ co-dependences. In this latter sense the present paper can be seen as a generalization of the work by Rebonato and Jaeckel (1999)². It is a generalization because we deal with the much richer concept of causation rather than correlation (see, eg, Williamson (2005)).

The method we present requires the risk manager to provide conditional probabilities of rare events. While this is clearly not an easy task, it can be simpler than specifying the stand-alone probabilities of the individual rare events: trivially, while it may be extremely difficult to estimate the probability of a person being involved in a car accident today, or of the same person being taken to hospital, it is relatively easy to give a plausible probability of the person being taken to hospital, conditional on a car accident having occurred. In practice, very often the risk manager will have at his disposal heuristics or imperfect models of the workings of sections of the financial markets in certain market conditions (eg, stylized facts such as: ‘conditional on an equity market crash, volatilities are likely to increase and credit spreads to widen’). These partial models and rules of thumb fall well short of a coherent model of the market in

²We refer to the works quoted in Rebonato and Jaeckel (1999) for a discussion of ways to alter on the basis of subjective inputs a statistically-determined correlation matrix.

its entirety. However, as former Chairman Greenspan recently pointed out in a related context, disregarding these heuristics because of their incompleteness would be a great mistake: “Policymakers often have to act [...] even though [they] may not fully understand the full range of possible outcomes, let alone each possible outcome’s likelihood. As a result, [...] policymakers have needed to reach to broader, though less mathematically precise, hypotheses about how the world works...” (quoted in Frydman and Goldberg, 2007).

Finally, we point out that the aggregate loss associated with these potential stress events can be directly linked to the event-risk charge (see, eg, BIS (2008a), (2008b)) that regulators have recently discussed. The question addressed in this paper is therefore how to aggregate these losses and how to take into account possible offsetting gains.

2 Links with Existing Literature

Moskowitz and Sarin (1983) present a linear programming approach that is similar to what presented here. In particular, they assume that marginal or conditional probabilities have been exactly assigned (possibly from frequentist estimation). They then use a Linear Programming technique to determine bounds on the joint probabilities that are compatible with the assigned marginal and conditional probabilities. They assume that the conditional and marginal probabilities are consistent, and suggest in the first part of the or part how to use joint probabilities to ‘cure’ possible inconsistencies. In this paper we do not deal with joint probabilities, and we look instead at the internal consistency of an exogenously assigned (and possibly inconsistent) conditional probability matrix.

An anonymous referee has brought to our attention the work by De Kuyuver and Moskowitz (1984). There are similarities with our work, in that the authors use interactive goal programming (as opposed to Linear Programming) for assessing scenario probabilities in long-range forecasting and decisional analysis. Their approach, however, requires both marginal and ordinal or interval first-order conditional probabilities, while our approach only requires conditional probabilities. See also the references therein for related approaches.

Gilio (1995) and the references quoted therein look at the situation where qualitative judgements or interval-based probabilities are assigned. The general strategy is to check the consistency of conditional probability bounds of the form $P(A|B_i) \leq K_i$, $i = 1, 2, \dots, n$. The numerical technique employed by Gilio (1995) requires an iterative manual procedure, and does not make use of Linear Programming. In our approach the conditional probabilities (as opposed to bounds) are assumed to be known by the risk manager, and the upper and lower bounds that we use are purely a numerical technique to ‘relax’ the coherence constraints and allow for the search of a coherent solution in the neighbourhood of the candidate set of conditional probabilities.

We are not aware of direct applications of the techniques in these papers to financial stress testing. Also, we have not found in the literature the equivalent

of the ‘triplet constraints’ that we present in Section 7. Simple as these checks are, we have found that they greatly improve the quality of the output: if the proposed conditional probability matrix is not sufficiently close to a coherent one, the output can bear little resemblance to what the risk manager intended to provide, and can become very unstable – small changes in the incoherent input conditional probabilities can produce very different coherent conditional probabilities. This is certainly true with our method, and we suspect that this will be the case for the similar approaches referred to above.

3 Description of the Methodology

Suppose that we have n positions, $\pi_1, \pi_2, \dots, \pi_n$, which have been identified as large by the desk risk managers. Let E_1, E_2, \dots, E_n , $i = 1, 2, \dots, n$, denote the significant events that would give rise to large profits *or* losses arising from our positions $\{\pi_i\}$ at a given point in time. Let P_i and L_i , $i = 1, 2, \dots, n$, be the profits and losses, respectively, associated with position π_i if event E_i occurs.³ We follow the convention that $P \geq 0$, $L \leq 0$. So, to fix ideas, position 1, π_1 , could be a large long equity position in the S&P. There are two events associated with position 1, E_1 and E_2 : E_1 could be a 1987-like equity market crash; E_2 could be a large market rally. The rise and fall need not be of the same magnitude. Associated with event E_1 there is a profit, P_1 and a loss L_1 . Both profit P_1 and loss L_1 are a function of the event E_1 (the equity market crash) and of the position π_1 : $P_1 = P_1(E_1, \pi_1)$, $L_1 = L_1(E_1, \pi_1)$. Similarly, there are a profit, P_2 and a loss L_2 associated with the equity rally.

It must therefore be stressed that these events do not encompass only the large losses but also the large possible gains arising from the most significant positions $\{\pi_i\}$. Of course, for a given position, the expected profit or loss need not be just the same number with the opposite sign. This can be either because the moves in the underlying are not symmetric (equity markets ‘crash’ in different ways than they spike, volatilities do not fall in the same way as they suddenly rise, etc); or because the underlying position may be not linear (eg, a short-gamma position). The notation used in the following assumes that, say, a yield curve steepening and a yield curve flattening are two distinct events. If our portfolio contains, say, a steeper position, then the loss for the steeper scenario would be zero, and the gain for flattening scenario would be zero. In short, given a large position, the associated events will, in general, be the largest plausible moves that give rise *to profit or losses*.

As mentioned in the introductory section, we assume that the risk manager

³Ideally, these should be the profits or losses that our portfolio could incur *over and above the hypothetical profits or losses that our VaR engine, run with the longest time series and volatility rescaling, produces*. As we represent events as Boolean random variables, the loss can either be uniquely specified by the event (eg, if the event is a default with zero recovery); or can assume a range of possible values (eg, if the event is of the type ‘move in the yield curve by 200 basis points or more’). In this latter case, the loss could be the smallest loss incurred if the event happens (as in the case of VaR). If the risk manager felt confident enough, she could try to assign a stand-alone expected conditional loss given that the event has happened.

rally, but we do not consider this in this example.) Let event 3 be a flattening of the US\$ yield curve and event 4 a steepening of the US\$ yield curve. Let's assume that we are long the S&P and we have a yield curve steepener on. Given an equity market crash we have a conditional probability of a curve steepening of the US\$ curve, $p_{4|1}$, of 80%, and a conditional probability, $p_{3|1}$, of a flattening of 20% (the two probabilities only happen to add up to 1 by chance – ie, we happen to have given a zero probability to the event ‘the yield curve neither steepens nor flattens’..)

Since we have a long equity position and a yield curve steepener we then have

- a gain, P_4 , and a loss, L_4 , of, say, \$100m and \$0m, respectively, associated with event 4 (the steepening of the curve) – L_4 is 0 because there is no loss from our position (a steepener) if event 4 happens, ie, if the curve indeed steepens;
- a gain, P_3 , and a loss, L_3 , of, say, \$0m and \$100m, respectively, associated with event 3 (the flattening of the curve) – it is P_3 that is now 0 because there is no gain from our position (a steepener) if event 3 happens, ie, if the curve flattens;
- a stand-alone loss, L_1 , of, say, \$200m on the S&P position if the equity market crash occurs.

The conditional losses associated with event 1 (equity market crash) are given by:

$$y_1 = p_{3|1}(L_3) + p_{3|1}(P_3) + p_{4|1}(L_4) + p_{4|1}(P_4) = \\ p_{3|1}(L_3) + p_{3|1}(0) + p_{4|1}(0) + p_{4|1}(P_4)$$

The stress loss 1 associated with the equity market crash is then given by:

$$SL_1 = L_1 + (P_3 + L_3)p_{3|1} + (P_4 + L_4)p_{4|1} = \\ L_1 + (0 + L_3)p_{3|1} + (P_4 + 0)p_{4|1} = \\ -\$200m + \$100m * 0.8 - \$100m * 0.2 = -\$140m$$

It must be stressed again that the profits and losses L_2 , P_2 and L_3 and P_3 do not depend on event 1 (the equity market crash), and are unconditional (marginal) quantities. The link with the equity market crash only comes via the conditional probabilities, $p_{2|1}$ and $p_{3|1}$.

As our intuition suggests, in the example above the steepening position mitigates the stand-alone loss in the naked equity position, reducing it from \$200m to \$140m.

5 Some Observations

The following obvious points are worth noting:

- In general

$$\sum_{j=1,n} p_{i|j} \neq 1 \quad (3)$$

- In general

$$p_{i|j} \neq p_{j|i} \quad (4)$$

ie, the matrix \mathbf{p} is not symmetric: the probability of, say, equity volatility increasing given a market crash is not the same as the probability of a market crash given a an increase in equity volatility.

- In the worked-out example above, a (linear) position in a steepener gave a profit and zero loss if the curve did steepen, (and vice versa for the flattener.) This need not be case: if we had a short gamma position on the level of a yield curve, for instance, we could record a loss both on yield curve moving up and down.

6 The Link To the Event Risk Charge

On the basis of this interpretation it is reasonable to equate the Event Risk Charge, ERC, to

$$\text{ERC} = \max\{SL_i\} \quad (5)$$

This measure has some nice properties.

- If we add a new scenario, its marginal contribution on ERC is not purely additive, reflecting the intuition that all the stress losses will not occur at the same time.
- If we add a new scenario, say, the $n+1$ th scenario, its marginal contribution on ERC can be positive even if the loss L_{n+1} is not the largest, provided that the conditional losses for some other event ($i \neq n+1$) add up to the largest SL .
- If we add a new scenario, its marginal contribution on ERC can be negative, provided that the marginal conditional profits for events $i \neq n+1$ reduce the largest SL .
- As one position becomes larger and larger, it will tend to dominate the ERC number, both via its stand-alone loss, and via the conditional profits and losses. In other words, the measure ERC knows both about concentration and diversification.

7 Internal Consistency of the Conditional Probabilities

In the presentation above the conditional probabilities are supposed to be exogenously provided by the risk managers. However, as these conditional probabilities ultimately reflect dependence among events (and actually can be linked to causal relationships in a way that the correlation coefficient cannot), it is not surprising that they cannot be arbitrarily assigned.

When correlation matrices are exogenously assigned, the consistency requirements are that the real symmetric matrix should be positive (semi)-definite. We derive similar conditions for the matrix of conditional probabilities. We note that an arbitrarily-chosen set of conditional probabilities are extremely unlikely to be consistent, the more so when some conditional probabilities are ‘large’ and some ‘small’ – as will typically be the case with the conditional probability matrix of interest in our application.

We stress that the constraints are actually of great help in guiding the intuition of the risk manager when, as we advocate, the conditional probabilities are subjective. As we shall see, the space of admissible solutions can be remarkably narrow. This implies that, given one broad type of assumed joint co-dependences, the possible solutions cannot differ too much and remain admissible. The degree of uncertainty and arbitrariness in assigning the conditional probability matrix is therefore greatly reduced.

7.1 Preliminary Analysis and Notation

We present before a Linear Programming approach to find an admissible solution ‘close’, in some sense to be discussed, to a subjective input conditional probability matrix. It is useful to apply some ‘pre-processing’ of this input matrix in order to ensure that the risk manager’s intuition is fully reflected in the matrix that is input into the more opaque Linear Programming algorithm. This ‘pre-processing’ takes the form of a couple of preliminary tests that should be run to eliminate obvious inconsistencies. Highlighting the inconsistencies at an early stage forces the risk manager to re-analyze his input matrix.

We have events x_i for $i = 1, 2, \dots, n$. In this section and for the rest of the paper, in order to lighten notation we will use (i) to represent the (unknown) unconditional probability of x_i and $[i|j]$ to represent the conditional probability x_i of given x_j . The following identity (Bayes’ theorem) is used extensively in the following:

$$\frac{(i)}{(j)} = \frac{[i|j]}{[j|i]} \quad (6)$$

Obviously, in specifying $[i|j]$ and $[j|i]$ the risk manager is already making a statement about the relative likelihood of (i) and (j) – a useful first ‘sanity check’ can be carried out here.

Our first constraint (**Constraint Triplets**) refers to triplets, and requires

that

$$\begin{aligned}
[i|j] &= [j|i] \frac{(i)}{(j)} = [j|i] \frac{(i)}{(k)} \frac{(k)}{(j)} = [j|i] \frac{[i|k]}{[k|i]} \frac{[k|j]}{[j|k]} \\
\implies 0 &\leq [j|i] \frac{[i|k]}{[k|i]} \frac{[k|j]}{[j|k]} \leq 1 \text{ for } i \neq j \neq k
\end{aligned} \tag{7}$$

where the first line follows from Bayes' theorem, and the second from the obvious limits that apply for any probability. We have 6 combinations of $[i|j]$ for $i \neq j$, **Constraint Triplets** shows that, given the conditional probabilities of 5 combinations, the 6th ($[i|j]$) is uniquely determined (and, of course, must be smaller than or equal to 1 and non-negative). A simple routine can easily be written that returns all the inconsistent triplets, and allows the risk manager to make changes accordingly.

There is a further constraint on triplets that can be used to identify other obvious inconsistencies (**Constraint Triplet Limits**).

Consider, as before, 3 events, i , j , and k . We use the notation \bar{i} , \bar{j} , and \bar{k} to represent the non-occurrence of these, and the notation $i \cap j$, for example, to denote the probability of the joint occurrence of i and j .

We then have:

$$\begin{aligned}
(i) &= [i \cap k] + \{i \cap \bar{k} \cap j\} + \{i \cap \bar{k} \cap \bar{j}\} \leq \\
&[i \cap k] + \{i \cap \bar{k} \cap j\} + \{i \cap \bar{j}\}
\end{aligned} \tag{8}$$

We can also write:

$$\{i \cap \bar{j}\} = (i) - \{i \cap j\} \tag{9}$$

and therefore by substitution we get

$$(i) \leq \{i \cap k\} + \{i \cap \bar{k} \cap j\} + (i) - \{i \cap j\} \tag{10}$$

i.e.

$$\{i \cap j\} \leq \{i \cap k\} + \{i \cap \bar{k} \cap j\} \leq \{i \cap k\} + \{\bar{k} \cap j\} \tag{11}$$

Writing

$$\{\bar{k} \cap j\} = (j) - \{k \cap j\} \tag{12}$$

we get:

$$\{i \cap j\} \leq \{i \cap k\} + (j) - \{k \cap j\} = \{i \cap k\} + (j) - \{k|j\} (j) \tag{13}$$

Now, dividing by (i) we get:

$$[j|i] \leq [k|i] + (1 - [k|j]) \frac{(j)}{(i)} = [k|i] + (1 - [k|j]) \frac{(j|i)}{(i|j)} \tag{14}$$

Re-arranging, we finally get:

$$[j|i] \left\{ 1 - \frac{1 - [k|j]}{[i|j]} \right\} \leq [k|i] \tag{15}$$

which in the symmetric case ($[j|i] = [i|j]$) reduces to:

$$[j|i] \leq 1 - [k|j] + [k|i] \quad (16)$$

For example, suppose in the symmetric case we have we have $[k|j] = [j|k] = 0.9$, so that events j and k are closely related. Suppose also that $[k|i] = [i|k] = 0.1$, so there is relatively little chance that events i and k occur simultaneously. Then it is obvious that, given the close relationship between j and k , there must be relatively small chance that i and j simultaneously occur; in fact from above we have:

$$[j|i] \leq 0.2 \quad (17)$$

with the same constraint on $[i|j]$.

We note in passing that setting $[i|j] = [j|i]$ ensures that **Constraint Triplets** is automatically satisfied for any triplet, but this is in general not a good place to start – the probability of a general equity market crash given a widening in the credit spread of firm X can be very different from the probability of widening in the credit spread of firm X given a general equity market crash.

If two scenario are mutually incompatible, there are severe constraints on the other conditional probabilities. We make use of the identity

$$[i \cup j|k] = [i|k] + [j|k] - [i \cap j|k] \quad (18)$$

(where \cup and \cap indicate union and intersection, respectively). If $[i|j] = [j|i] = 0$ for all other scenarios k it must hold that

$$[i \cap j|k] = 0 \quad (19)$$

Therefore (**Constraint Zeros with Incompatibility**)

$$[i|k] + [j|k] = [i \cup j|k] \leq 1 \quad (20)$$

A simple routine can be written to identify violations of this inequality for any zero conditional probability.

These ‘sanity checks’ are very simple, but they also are extremely important. In our experience, the systematic, Linear-Programming-based solution we present in the next section only returns a solution close to what the risk manager intended if the starting point is not too distant from a coherent solution in the first place. Also, if the starting point for the systematic search is ill-constructed, the output is likely to be unstable, in the sense that relatively small changes in the non-coherent inputs can give rise to large differences in the coherent output.

7.2 Systematic Solution

The above analyses can be used to remove the more obvious inconsistencies, but there are still likely to be subtler inconsistencies. The general methodology consists in checking the proposed conditional probabilities for consistency, and until consistency is proven we allow the conditional probabilities to be moved

within a widening range of values. At each iteration the algorithm suggests a consistent solution, and at any time this solution can be accepted (even if it is outside the current range).

The formulation is as follows. With each of the scenarios we associate an indicator variable, $I_i, i = 1, 2, \dots, N$, where $I_i = 1$ if scenario i occurs, and $I_i = 0$ otherwise. There are 2^N mutually exclusive and exhaustive events corresponding to the distinct combinations of the set of I_i s. We signify any given combination of I_i s by the vector \mathbf{I} , and the set of all \mathbf{I} s by \mathcal{I} . To each \mathbf{I} we associate a probability $P[\mathbf{I}]$.

In order for any set of proposed conditional probabilities to be consistent there must exist (at least one) set of $P[\mathbf{I}]$ with

$$\sum_{\mathbf{I} \in \mathcal{I}} P[\mathbf{I}] = 1 \quad (21)$$

$$P[\mathbf{I}] \geq 0 \text{ for any } \mathbf{I} \in \mathcal{I} \quad (22)$$

The corresponding conditional probabilities are given by:

$$[i|j] = \frac{\sum_{\mathbf{I}: \{I_i=1, I_j=1\}} P[\mathbf{I}]}{\sum_{\mathbf{I}: \{I_j=1\}} P[\mathbf{I}]} \quad (23)$$

Since any proposed set of conditional probabilities is very unlikely to be internally consistent, we allow the conditional probabilities to lie anywhere between lower and upper limits, $\sigma_{i|j}$ and $\vartheta_{i|j}$, respectively, for $[i|j]$. These upper and lower limits can be specified by the risk manager. Hence we have for each i and j , the two linear constraints:

$$\sigma_{i|j} \sum_{\mathbf{I}: \{I_i=1\}} P[\mathbf{I}] - \sum_{\mathbf{I}: \{I_i=1, I_j=1\}} P[\mathbf{I}] \leq 0 \quad (24)$$

and

$$\vartheta_{i|j} \sum_{\mathbf{I}: \{I_i=1\}} P[\mathbf{I}] - \sum_{\mathbf{I}: \{I_i=1, I_j=1\}} P[\mathbf{I}] \geq 0 \quad (25)$$

We then define non-negative ‘slack’ variables, $s_{i|j}$ and $t_{i|j}$:

$$s_{i|j} = -\sigma_{i|j} \sum_{\mathbf{I}: \{I_i=1\}} P[\mathbf{I}] + \sum_{\mathbf{I}: \{I_i=1, I_j=1\}} P[\mathbf{I}] \quad (26)$$

$$t_{i|j} = -\vartheta_{i|j} \sum_{\mathbf{I}: \{I_i=1\}} P[\mathbf{I}] + \sum_{\mathbf{I}: \{I_i=1, I_j=1\}} P[\mathbf{I}] \quad (27)$$

where

$$s_{i|j} \geq 0 \text{ for any } i, j \quad (28)$$

$$t_{i|j} \geq 0 \text{ for any } i, j \quad (29)$$

We refer to any set of probabilities that obey Equations (21), (22) a ‘feasible’ solution, and any set that in addition obey Equations (24) and (25) a ‘coherent’

solution.⁴ If we can find a coherent solution then the given constraints are consistent, otherwise they are not.

In order to find a coherent solution we use Phase 1 of the Revised Simplex Method (see, eg, Press et al (1996)). This involves defining non-negative artificial variables to give an initial ‘basic’ solution⁵, and minimize the sum of the artificial variables, say, z , subject to the given constraints. If the minimum value of z , z_{min} say, is greater than zero, the constraints do not allow a coherent solution. However, in this case our optimal solution will be, in some sense, as ‘close’ as we can get to a coherent solution. The risk manager can then either accept the resulting ‘close’ solution or widen the limits, $\sigma_{i|j}$ and $\vartheta_{i|j}$, within which the conditional probabilities are allowed to lie.

It is worth noting that any feasible solution gives rise to an infinite number of related solutions. Representing the given solution by $P_{opt}[\mathbf{I}]$, any solution such that:

$$P[\mathbf{0}] = \frac{1 + \alpha}{1 + \alpha P_{opt}[\mathbf{0}]} P_{opt}[\mathbf{0}] \quad (30)$$

and

$$P[\mathbf{I}] = \frac{1}{1 + \alpha P_{opt}[\mathbf{0}]} P_{opt}[\mathbf{I}] \quad \text{for } \mathbf{I} \neq \mathbf{0} \quad (31)$$

is also feasible for $-1 \leq \alpha$. (In the expressions above, $\mathbf{0}$ is the vector \mathbf{I} whose components are all zero, and $P[\mathbf{0}]$ is the associated probability that no scenario occurs.) Scaling all the individual probabilities (except $P[\mathbf{0}]$) by the same proportion clearly does not affect the conditional probabilities. This means that the solution can be associated with unconditional probabilities that are arbitrarily small or large (provided, of course, that none of them is greater than 1).

Note also that, given any proposed solution with $P[\mathbf{I}] \geq 0$ for any $\mathbf{I} \in \mathcal{I}$, we can always find a *feasible* solution (i.e. the constraint $\sum_{\mathbf{I} \in \mathcal{I}} P[\mathbf{I}] = 1$ can always be satisfied) simply by dividing each $P[\mathbf{I}]$ by the sum of the $P[\mathbf{I}]$ s, and this will not affect the conditional probabilities.

8 Implementation

We prototyped the simplex solution in VBA; for 12 scenarios this was taking of the order of 5-10 minutes to run. Coding it in C++ makes the routines runs 10 to 20 times faster.

Preliminary analysis suggests that each additional scenario will increase the run time by a factor of about three; considering that the analysis is likely to be required infrequently (say, on a weekly basis), the computation time should not be a problem with up to, say, 15-20 scenarios. In fact we are more likely to run into memory-space problems first.

⁴To avoid confusion, it should be pointed out that there is no link between our usage of the word ‘coherent’ and the word ‘coherent’ in coherent risk measures.

⁵A ‘basic’ solution is one in which all but M of the variables are set to zero (where M is the number of constraints – excluding the non-negativity constraints). At each iteration each of the basic variables and the objective function (z in our case) are expressed as affine functions of the non-basic variables. See the worked-out example below.

9 A Worked-Out Example

In order to illustrate the procedure in practice we present in this section the results of a fictitious case study.

We have four scenarios, A, B, C, and D, and we suppose that the risk-manager supplied the following conditional probabilities:

$$\begin{bmatrix} & \mathbf{A} & \mathbf{B} & \mathbf{C} & \mathbf{D} \\ \mathbf{A} & 1 & 0.60 & 0.50 & 0.60 \\ \mathbf{B} & 0.80 & 1 & 0.22 & 0.20 \\ \mathbf{C} & 0.50 & 0.20 & 1 & 0.00 \\ \mathbf{D} & 0.40 & 0.40 & 0.00 & 1 \end{bmatrix} \quad (32)$$

where, for example, the figure 0.5 in the first row represents $[C|A]$.

We first observe that $[C|A] + [D|A] = 1.1$. But, from the matrix, C and D are mutually exclusive ($[C|D] = [D|C] = 0$). Therefore we need to reduce either or both of $[C|A]$ and $[D|A]$ so that their sum is ≤ 1 . In fact we would make the sum strictly less than 1 to avoid making the overly strong statement that if A occurs then either C or D must occur.

Suppose that the risk manager therefore sets $[C|A] = 0.4$ and $[D|A] = 0.5$. If we now look at the triplet $\{A, B, D\}$, we should have:

$$[A|D] = [D|A] \frac{[A|B]}{[B|A]} \frac{[B|D]}{[D|B]} = 0.5 \frac{0.80}{0.60} \frac{0.40}{0.20} = 1.333 \quad (33)$$

So the 5 probabilities on the right hand side are clearly inconsistent. To fix the problem, the risk manager could set $[D|B] = 0.3$, for example, which gives $[A|D] = 0.89$. Our revised matrix therefore becomes:

$$\begin{bmatrix} & \mathbf{A} & \mathbf{B} & \mathbf{C} & \mathbf{D} \\ \mathbf{A} & 1 & 0.60 & 0.40 & 0.50 \\ \mathbf{B} & 0.80 & 1 & 0.22 & 0.30 \\ \mathbf{C} & 0.50 & 0.20 & 1 & 0.00 \\ \mathbf{D} & 0.89 & 0.40 & 0.00 & 1 \end{bmatrix} \quad (34)$$

which shows no obvious inconsistencies. Note, however, that, to ‘fix the problem’, the risk manager has had to change substantially his original estimate of $[A|D]$ (from 0.40 to 0.89). The plausibility of this change should be questioned. Was the original suggestion ill-thought-out? Or is the ‘fix’ forcing assumptions on the risk manager she is not comfortable with? We view this as an important and useful part of the process (the ‘auditable decision process’).

We now move to the next phase, to ensure consistency of the plausible and ‘cleansed’ matrix. Suppose that the risk manager allows each of the conditional probabilities⁶ to lie within a range

$$[i|j]_1 (1 - \delta) \leq [i|j]_2 \leq [i|j]_1 + \delta (1 - [i|j]_1) \quad (35)$$

⁶Except, of course, those that are 1 or 0.

where the $[i|j]_1$ s are taken from the revised matrix above (the subscripts 1 and 2 indicate the first or second iteration of the procedure, respectively, and that we are not necessarily dealing with *bona fide*, internally consistent conditional probabilities yet.) Taking $\delta = 0.01$ this gives lower and upper bounds:

$$\begin{bmatrix} & \mathbf{A} & \mathbf{B} & \mathbf{C} & \mathbf{D} \\ \mathbf{A} & 1 & 0.59 & 0.40 & 0.50 \\ \mathbf{B} & 0.79 & 1 & 0.22 & 0.30 \\ \mathbf{C} & 0.50 & 0.20 & 1 & 0.00 \\ \mathbf{D} & 0.88 & 0.40 & 0.00 & 1 \end{bmatrix} \quad (36)$$

and

$$\begin{bmatrix} & \mathbf{A} & \mathbf{B} & \mathbf{C} & \mathbf{D} \\ \mathbf{A} & 1 & 0.60 & 0.41 & 0.51 \\ \mathbf{B} & 0.80 & 1 & 0.23 & 0.31 \\ \mathbf{C} & 0.51 & 0.20 & 1 & 0.00 \\ \mathbf{D} & 0.89 & 0.41 & 0.00 & 1 \end{bmatrix} \quad (37)$$

respectively.

We now define the 16 indicator vectors and their corresponding probabilities:

$$\begin{bmatrix} \text{Indicator} & \mathbf{A} & \mathbf{B} & \mathbf{C} & \mathbf{D} & \text{Prob} \\ \mathbf{I}_0 & 0 & 0 & 0 & 0 & P(0) \\ \mathbf{I}_1 & 0 & 0 & 0 & 1 & P(1) \\ \mathbf{I}_2 & 0 & 0 & 1 & 0 & P(2) \\ \mathbf{I}_3 & 0 & 0 & 1 & 1 & P(3) \\ \mathbf{I}_4 & 0 & 1 & 0 & 0 & P(4) \\ \mathbf{I}_5 & 0 & 1 & 0 & 1 & P(5) \\ \mathbf{I}_6 & 0 & 1 & 1 & 0 & P(6) \\ \mathbf{I}_7 & 0 & 1 & 1 & 1 & P(7) \\ \mathbf{I}_8 & 1 & 0 & 0 & 0 & P(8) \\ \mathbf{I}_9 & 1 & 0 & 0 & 1 & P(9) \\ \mathbf{I}_{10} & 1 & 0 & 1 & 0 & P(10) \\ \mathbf{I}_{11} & 1 & 0 & 1 & 1 & P(11) \\ \mathbf{I}_{12} & 1 & 1 & 0 & 0 & P(12) \\ \mathbf{I}_{13} & 1 & 1 & 0 & 1 & P(13) \\ \mathbf{I}_{14} & 1 & 1 & 1 & 0 & P(14) \\ \mathbf{I}_{15} & 1 & 1 & 1 & 1 & P(15) \end{bmatrix}$$

Recall that indicator variables represent the occurrence of combination of scenarios. Thus, for example, the indicator variable \mathbf{I}_3 represents the outcome that scenarios C and D occur and scenarios A and B and do *not* occur, and $P(3)$ is the corresponding probability (to be solved-for).

We have 12 constraints of the form:

$$\sigma_{i|j} \sum_{\mathbf{I}:\{I_i=1\}} P[\mathbf{I}] - \sum_{\mathbf{I}:\{I_i=1, I_j=1\}} P[\mathbf{I}] \leq 0 \quad (38)$$

and 12 of the form

$$\vartheta_{ij} \sum_{\mathbf{I}:\{I_i=1\}} P[\mathbf{I}] - \sum_{\mathbf{I}:\{I_i=1, I_j=1\}} P[\mathbf{I}] \geq 0 \quad (39)$$

For example, corresponding to $[D|A]$ we have:

$$0.5 \{P(8) + P(9) + P(10) + P(11) + P(12) + P(13) + P(14) + P(15)\} - \{P(9) + P(11) + P(13) + P(15)\} \leq 0 \quad (40)$$

ie,

$$0.5 \{P(8) + P(10) + P(12) + P(14) - P(9) - P(11) - P(13) - P(15)\} + s_{D|A} = 0 \quad (41)$$

where $s_{D|A}$ is the ‘slack’ variable corresponding to this constraint.

Trivially, all the constraints, as formulated, can be satisfied by setting $P(\mathbf{0}) = 1$ (the probability of no scenarios occurring is one), and all other variables to zero. Clearly, however, for this solution the conditional probabilities are not defined and we do not have an economically interesting solution. Therefore we introduce non-negative ‘artificial’ variables, $a_{X|Y}$, of the kind:

$$a_{D|A} = 0.5 \{P(8) + P(10) + P(12) + P(14) - P(9) - P(11) - P(13) - P(15)\} + s_{D|A} \quad (42)$$

and in order to get the algorithm started we initialize all of these to some arbitrarily small value, ε .

Our initial (infeasible) solution is given by setting all the artificial variables to ε , and all others to zero.

We immediately have an expression for the sum of the artificial values in terms of the currently zero variables⁷, by summing the right-hand sides of all the equations exemplified by Equation (42), and we use the revised simplex algorithm to minimize this expression.

This gives the following solution:

$$\left[\begin{array}{l} \text{Indicator} \\ \mathbf{I}_2 \\ \mathbf{I}_1 \\ \mathbf{I}_{14} \\ \mathbf{I}_9 \\ \mathbf{I}_{13} \\ \mathbf{I}_{12} \\ \mathbf{I}_4 \\ \mathbf{I}_{10} \end{array} \quad \begin{array}{cccccc} \mathbf{A} & \mathbf{B} & \mathbf{C} & \mathbf{D} & K \\ 0 & 0 & 1 & 0 & 1.606 \\ 0 & 0 & 0 & 1 & 0.336 \\ 1 & 1 & 1 & 0 & 0.702 \\ 1 & 0 & 0 & 1 & 0.865 \\ 1 & 1 & 0 & 1 & 0.946 \\ 1 & 1 & 0 & 0 & 0.537 \\ 0 & 1 & 0 & 0 & 0.664 \\ 1 & 0 & 1 & 0 & 0.734 \end{array} \right] \quad (43)$$

⁷This is what is required for the Revised Simplex Algorithm: at each iteration we have a set of ‘basic’ (in general non-zero) variables, and ‘non-basic’ (zero-valued) variables. The value of each of the basic variables and the value of the ‘objective function’ (the expression to be optimized) are expressed as functions of the non-basic variables only. See, eg, Press et al (1996)

where K is a quantity proportional to the probabilities and all other probabilities, except $P(\mathbf{0})$, have been set to zero.

The corresponding solution in terms of conditional probabilities is:

$$\left[\begin{array}{ccccc} \text{Trial Solution} & \mathbf{A} & \mathbf{B} & \mathbf{C} & \mathbf{D} \\ \mathbf{A} & 1.00 & 0.58 & 0.38 & 0.48 \\ \mathbf{B} & 0.77 & 1.00 & 0.25 & 0.33 \\ \mathbf{C} & 0.47 & 0.23 & 1.00 & 0.00 \\ \mathbf{D} & 0.84 & 0.44 & 0.00 & 1.00 \end{array} \right] \quad (44)$$

Comparing this with Tables (36) and (37) we see that the largest violation of the constraints is for $[A|D]$, where the above solution lies 0.04 below its lower limit. We could accept this solution, make *ad hoc* changes to the limits, or increase δ and re-optimize.

10 Conclusions

We have presented a methodology for aggregating in a coherent manner the ‘stress’ losses associated with a number of pre-chosen scenarios. The risk manager is required to provide a non-symmetric matrix of conditional probabilities conjoining the various scenarios. A procedure is then suggested to assess whether the suggested matrix is internally consistent. If not, we show how to find a ‘closest’ admissible solution. We provide a worked-out example to describe the procedure in detail.

The procedure can be used to address the regulators’ requirements for the Incremental Risk Charge.

References

- Bank for International Settlements - BIS, 2008a, *Proposed Revision to the Basel II Market Risk Framework*, July 2008
- Bank for International Settlements - BIS, 2008b, *Guidelines for Computing Capital for Incremental Risk in the Trading Book*, 22 July 2008
- de Kluyver, Moskowitz, H, (1984), *Management Science*, Vol 30, No 3, March
- Frydman, R, and Goldberg, (2007), *Imperfect Knowledge Economics*, PUP, Princeton, New Jersey
- Gilio, A, (1995), *Probabilistic Consistency of Conditional Probability Bounds*, Lecture Notes In Computer Science, Vol 945, 200-209
- Moskowitz, H, Sarin R K, (1983), *Improving the Consistency of Conditional Probability Assessment for Forecasting and Decision Making*, *Management Science*, Vol 29, No 6, (June 1983), 735-749
- Press, W H, Teukolski, S A, Vetterling, W T, Flannery, B P, (1996), *Numerical Recipes*, Cambridge University Press, Cambridge, UK
- Rebonato, R, Jaeckel, P, (1999) ‘The Most General Methodology to Create a Valid Correlation Matrix for Risk Management and Option Pricing Purposes’, *Journal of Risk*, Vol 2, Number 2, Winter 1999/2000, 17-28.
- Williamson, J, (2005), *Bayesian Nets and Causality*, OUP, Oxford, UK