GRAS Newsletter 11

May 2009

Summary: Like any discipline, risk management needs a vocabulary that is both rigorous and workable. The terms used in risk management are plagued by the seeming familiarity in every day parlance. Credit risk and to a greater degree market risk, have found a way around this dilemma by employing strict mathematical and statistical language. For ORM and ERM, no such luxury exists and they are stuck with every day terminology. The least we should do is to clarify as clearly as possible what the risk in ORM and ERM means. To do that, it will be useful to distinguish Risk from Uncertainty, which are often confused. This newsletter explains why this is important and how the concept of uncertainty may help ORM and ERM.

Dear reader,

Judging from the reaction of professional risk managers to the banking turmoil, it would seem that the proper route forward for risk management is, by and large, 'more of the same'. Hence, the answers of risk management appear to be around setting much tighter credit lines, maintain stricter monitoring of defaults, set far more aggressive margin calls, carrying out way more stress testing, perform more scenario analysis and round up ever more data to feed, calibrate and validate VaR models.

While these enhanced risk management programmes all serve to feed the need for more risks management, the missing ingredient in all this activity is a fundamental reconsideration of the risk management function. A concept that is used in the area of policy analysis may provide a language for that fundamental reconsideration. This concept is that of providing a taxonomy for uncertainty[1] to allow for a more complete appreciation of the risk universe: its methods, its models, its inputs, its outputs, its dependencies and other factors that should give us pause for thought when we are putting our trust in the risk methodologies. Before compounding the current risk approaches into an ERM or IRM framework, it will be worthwhile to have better grip on what level of certainty we can place in our risk management methods. For that, we need a language and a taxonomy of uncertainty.

Factors in Uncertainty

There is something called the uncertainty paradox[2] in many risk approaches. The paradox is that uncertainty is briefly acknowledged at the outset after which all risk statements are couched in terms of certainty. Be it credit scores, value at risk numbers, risk self assessments or loss data, the outcome is taken without much attention for the question whether what we measure is a good reflection of what we are interested in. Uncertainty surrounding the outcomes is only one aspect though. Of far greater significance is what happens at the model building stage, the relevance of the assumptions going into the model; the relations between elements that the model describes and the relations that it ignores; the selection of elements itself as well as our ability to correctly observe these relations and elements. All these combined determine the usefulness of the model. Ascertaining the uncertainty of

---

[1] The notion that risk and uncertainty should be treated seperately was First proposed by F.H. Knight (1921), in "Risk, uncertainty and profit", Houghton Mifflin, Boston, USA.

[2] The term uncertainty paradox was introduced by Van Asselt and Vos in (2006). "The precautionary principle and the uncertainty paradox." Journal of Risk Research, 9(4), 313-336.

the outcome (such as in a confidence interval) only becomes important once we are confident about the level of uncertainty in the model itself

In the words of the statistician George Box: "All models are wrong, but some models are useful". To find out how useful our risk models are, a taxonomy for uncertainty, such as outlined below, will be a good starting point. The one outlined here was modified from a taxonomy first proposed for policy analysis[3]. It is treated in more detail in a forthcoming white paper.

A Taxonomy for uncertainty

Before we start with the taxonomy, it will be useful to give a formal definition of uncertainty. In our analysis, uncertainty is defined as "a measure for the degree of validity of the model choices". We express this measure as a distance from the state of full determinism

Three steps are involved in determining uncertainty: (1) Selecting which *objects* to include in our uncertainty analysis, (2) ascertaining the *level* of uncertainty for each object and finally (3) ascertaining the *variability* of that level of uncertainty.

1) Selecting the *objects* of uncertainty

No model is an island, and in risk management, the models we use make a host assumptions and operate in a setting that is highly complex and variable. The individual assumptions can all be legitimate *objects* in our analysis, as well as any context relevant information. It is important that these assumptions and this contextual information is captured at the atomic level, or as near as we can get to. Any compound, such as treating "Regulation" as a catch-all for internal regulation, national regulation, banking regulation (SEC, FED, OCC, etc), accounting regulation, insurance regulation etc etc will only blur the picture. Next to these elements, the inputs into the model, the static data, the model itself as well as the outcomes are all legitimate *objects* of our study. In a typical risk model, there may be anywhere between 200 and 1000 objects to consider.

2) Ascertaining the *level* of uncertainty for each object

The basic question we seek to address here is "what do we know about the uncertainty regarding each object?" The level is a continuum punctuated by some markers.

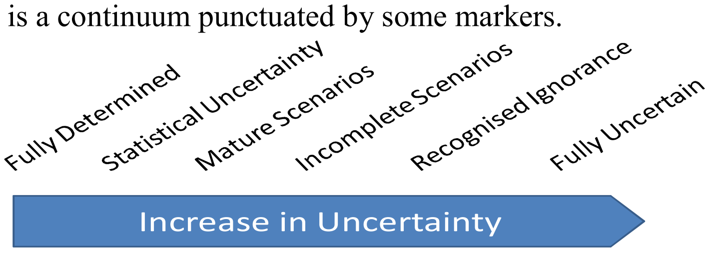

It is worthwhile to pause at this continuum. In risk management, 'Statistical Uncertainty', of the kind where we understand the behaviour of our object and are able to express it in pure statistical terms is

[3] *Defining Uncertainty. A Conceptual Basis for Uncertainty Management in Model-Based Decision Support,* W.E. Waker, P. Harremoës, J. Rotmans, J.P. Van Der Sluijs, M.B.A. Van Asselt, P. Janssens and M.P. Krayer von Krauss, Integrated Assessment 2003, Vol. 4, No. 1, pp. 5–17

often taken for granted. An example of statistical uncertainty would be flipping a coin. The outcome uncertain, not at all fully determined, but easily modelled with mere 'Statistical Uncertainty'. The way we account for trader errors, however, may be through 'Incomplete Scenarios' at best and probably is better described as 'Recognised Ignorance' if not 'Fully Uncertain'.

There are four aspects of each object that are measured on a scale of 0 – 20[4],: (A) How well do we understand the theory underlying our model ? (B) How well do we understand the empirical working of the model ? (C) How well do we understand the structure of the model ? (D) How well do we understand the variables in the model ?

The scale we use for these four aspects is roughly drawn up as follows:

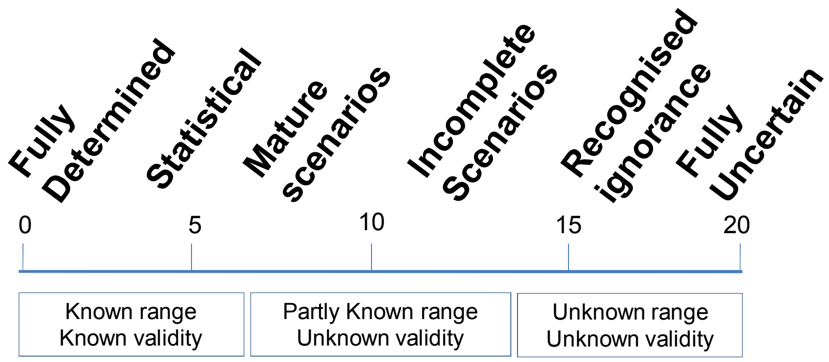

The fifth question to ascertain the *level* of uncertainty refers to the sensitivity of object to changes in the environment. Here, we use a scale of 0-20 again, with the following rough outline:

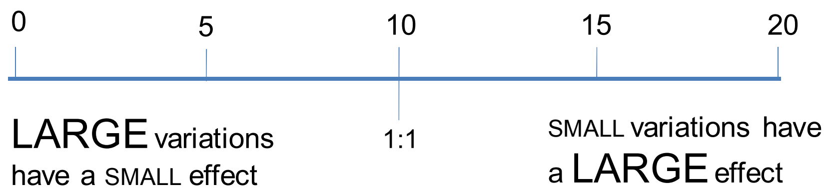

These five aspects give us a range of 0-100, which puts the objects a certain distance from the state of 'Full Determinism'. The higher the level value, the more uncertainty the object represents. Since there may be many objects in a model, much can be gained from examining various combinations of objects. When comparing models as a whole, some form of normalisation should be adopted.

3) Ascertaining the *variability* of that level of uncertainty

The last step is that of determining whether the uncertainty is reducible or not. For some objects, the level of uncertainty is almost a given. Further research or additional data will hardly change the level of uncertainty for a coin toss. The level remains at the statistical level. We may discover a doctored

[4] Adapted from: Martin Krayer von Krauss, "Uncertainty in Policy Relevant Sciences", Ph.D. Thesis, November 2005, Environment & Resources DTU, Technical University of Denmark.

coin, so it would be too rash to call it purely statistical, but for any coin, the level of uncertainty will not change much. For other objects though, such as a client's credit score, the amount of effort we put into determining it will make a difference and we may move the level of uncertainty from recognised ignorance to incomplete scenario's. The range we use here is roughly as follows:

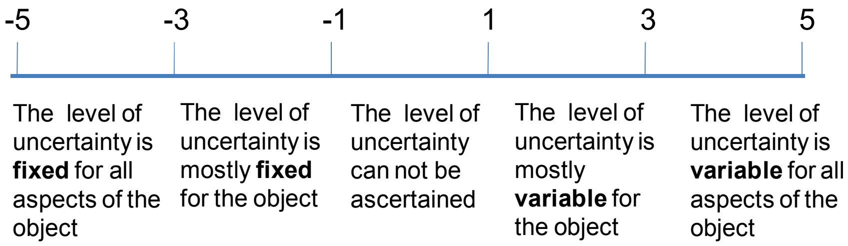

The diagram

The two measures, level and variability of the level can be pictured in a two dimensional diagram as follows. This diagram shows a simplified analysis made for a problem in liquidity risk management.

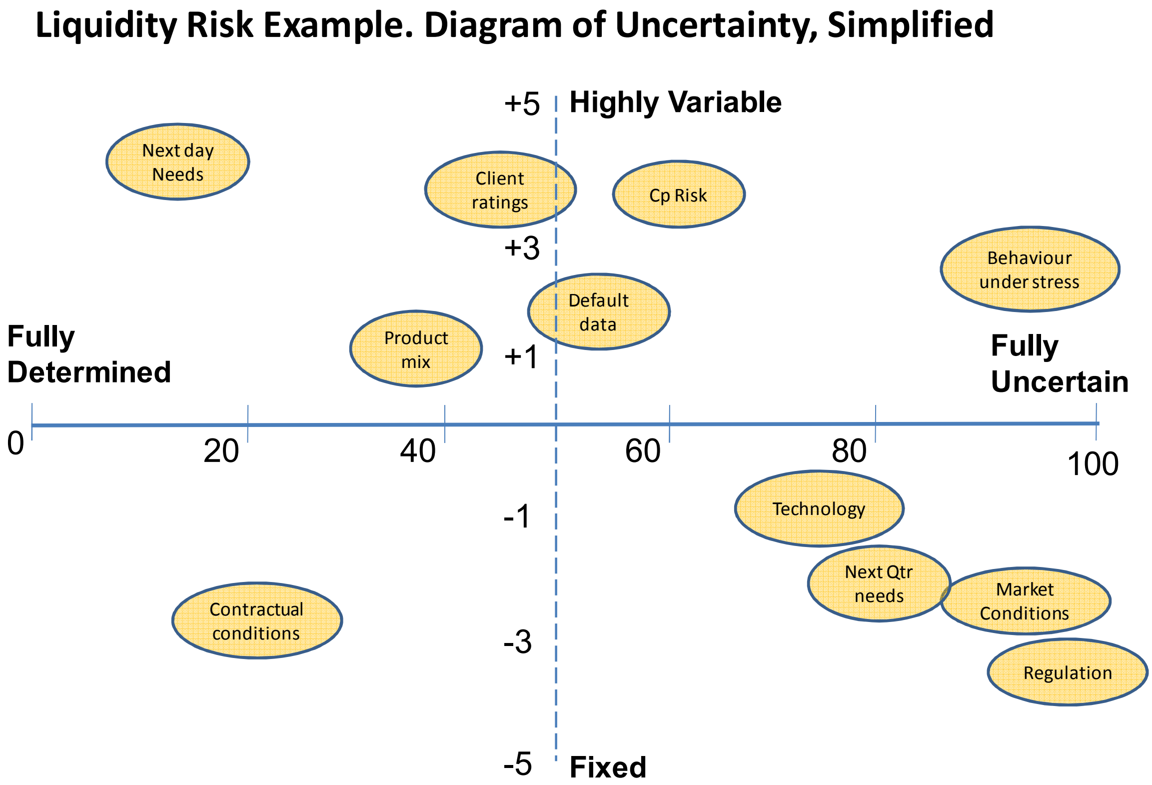

In this example, we have limited ourselves to 11 objects in this overview for illustration purposes. We note already from this small sample, however, that the majority of the objects lean more towards 'Fully Uncertain' than to 'Fully Determined', both of which are abstractions. They are nearly split between 'Variable' and 'Fixed', which gives a good indication of where to run stress tests, where to spend further effort in fine tuning models through sensitivity tests, where to focus further validation efforts, where to develop additional scenarios and, finally, what aspects to highlight in the all important report to senior management (and regulators).

The use of the taxonomy

The taxonomy for uncertainty is primarily an aid to compare and contrast model approaches. In that sense, it could be useful for banks that wish to make an informed decision about their AMA. Should it be scenario based, (external) loss data focused and what mix is better at capturing the risk profile. Secondly, the taxonomy can be used delve into the model's constituent parts to fine tune existing approaches by accounting for the uncertainty introduced for each element of the model.