City Research Online

City, University of London Institutional Repository

Citation: Tsanakas, A. (2012). Modelling: The elephant in the room. The Actuary, 2012(Septem),

This is the submitted version of the paper.

This version of the publication may differ from the final published version.

Permanent repository link: https://openaccess.city.ac.uk/id/eprint/16877/

Link to published version:

Copyright: City Research Online aims to make research outputs of City, University of London available to a wider audience. Copyright and Moral Rights remain with the author(s) and/or copyright holders. URLs from City Research Online may be freely distributed and linked to.

Reuse: Copies of full items can be used for personal research or study, educational, or not-for-profit purposes without prior permission or charge. Provided that the authors, title and full bibliographic details are credited, a hyperlink and/or URL is given for the original metadata page and the content is not changed in any way.

The elephant in the room: model error and solvency regulation[1]

Andreas Tsanakas
Cass Business School, City University London.
a.tsanakas.1@city.ac.uk

July 2012

> "In applying mathematics to subjects such as physics or statistics we make tentative assumptions about the real world which we know are false but which we believe may be useful nonetheless. […] Since all models are wrong the scientist cannot obtain a "correct" one by excessive elaboration. […] Since all models are wrong the scientist must be alert to what is importantly wrong." [2]

In an insurance world preoccupied with Solvency II, internal models have increased in both complexity and business significance. Nonetheless, actuaries are painfully aware of models' limitations in representing the economic world, not least because of their reliance on often arbitrary assumptions. It is fairly commonplace to pronounce that 'all models are wrong but some are useful'. But we do not always stop to think about the implications of such a statement. In what sense may an internal model be wrong? If a model is wrong, how can it be useful? Significantly, what incentives does regulation produce for model development and use?

Solvency capital requirements involve calculating the probability of extreme events, as well as the probability that a confluence of such or less extreme events produces a high financial loss. The focus on rare events makes statistical estimates intrinsically unreliable, as they are obtained from analysing limited – sometimes non-existent – data. Moreover, the complexity of models, mirroring that of insurance enterprises, exacerbates potential errors, by increasing the sensitivity of model outputs on assumptions that cannot be supported by empirical evidence. For example, it is well understood that for highly "granular" internal models, changes in correlation parameters that are easily dominated by statistical error lead to swingeing movements in estimated portfolio VaR.

---

[1] A slightly edited version of this article has appeared in *The Actuary*, September 2012.

[2] Box (1976). Science and Statistics. *Journal of the American Statistical Association* 71(356), 791-799.

Still, such problems do not seem to deter us from quantitatively modelling risks. When *technically valid* estimates are hard to come by, we are quite happy to make do with estimates that are *socially valid*, that is, they are “shared by others, are stable, and are believed in with confidence”[3]. Assumptions are set so as to make sense: to colleagues, regulators, consultants, peers. We may ask: ‘How much capital should be reasonably allocated to operational risk?’ The expert judgement used to answer such a question does not involve a mental calculus of probabilities; instead it considers the social expectations of stakeholders. If ‘12% of the total capital’ happens to be the answer, it is only reasonable because we all agree that it is.

When such issues are acknowledged by insurance practitioners, it can be with an air of resignation. But despairing at models’ lack of technical validity is to misunderstand their function. To start with, whether models are *importantly* wrong depends on the application. While estimating accurately a 1-in-200 years loss is illusory, there are other questions that models may help answer satisfactorily, such as the probabilities of less extreme scenarios or the relative impacts of exposure changes on the total risk profile. More generally, *the usefulness of a model is not reducible to the accuracy of its outputs*. In any scientific discipline, models are tools to conceptualise the reality surrounding us; a model is a metaphor, not a mirror-image of the modelled thing itself[4]. Models can be presented with different inputs and their outputs can be subsequently studied; it is such interrogations of a model that help us make sense of the aspect of reality that is being modelled[5].

In particular, internal models can be used to educate management in aspects of risk, by illustrating concepts, analysing scenario impacts, studying sensitivities, and demonstrating the sheer range of possible outcomes. We learn from the practice of modelling itself, not from summaries of model outputs. Moreover, models are essential tools for communicating risk across organisations and informing commercial transactions. For example, model output is often used to demonstrate the value of a reinsurance product to a potential buyer. As the subprime credit crisis has demonstrated, models can be also be used to convince investors in complex products that they are not taking on much risk. But the resulting criticism of models often misses the point: the problem was not that models were wrong, but that enough people were willing to believe otherwise.

---

[3] March (1994), *A primer on decision-making: how decisions happen*, New York: Free Press.

[4] For models as representations and much more, see Edwards and Hoosain (2012), *The philosophy of modelling*, http://www.sias.org.uk/diary/view_meeting?id=SIASMeetingJune2012

[5] Morgan (2001), ‘Models, stories, and the economic world,’ *Journal of Economic Methodology* 8 (3), 361–384.

So what is the role of regulation in all this? Most obviously, regulation and rating agency requirements have done much to move the focus of risk modelling to quantities that cannot be reliably quantified, such as extreme percentiles. (Though before blaming regulators for the world's ills, it is worth considering that regulatory requirements incarnate societal needs for security and certainty). Regulators are naturally aware of the substantial potential for model error. Such awareness must, at least in part, be behind the increased emphasis that Solvency II places on model validation and documentation, and on embedding models into decision making. This is sensible but not unproblematic. First, the focus on extreme events of low probability makes internal model output not only potentially inaccurate, but also hard to validate. After several years, it is possible to judge whether a long-tail liability portfolio was under-priced, but we may never know whether the portfolio had been capitalised consistently with the regulatory standard. Second, embedding the internal model into decision-making processes is seen as evidence of management's confidence in the model. It would be quite wrong to let such confidence count as evidence of technical validity.

A different sort of problem arises from regulation establishing a causal link between a company's available assets (an economically driven figure) and internal model outputs (a statistical construct). Under Solvency II, internal model approval is often perceived to confer an economic advantage, by lowering the capital requirement in comparison to the standard formula. Consequently, the substantial investment in internal models may reflect the perceived cost of non-conformity, rather than management's own desire to be educated in the statistical aspects of risk. Indeed, if there is a conflict between the pursuits of understanding and of advantage, it is not hard to see which perspective will prevail. If insurance firms perceive that openness about the limitations of their modelling puts model approval at risk, they may try to conceal such limitations and put on a show of confidence. But this has a corrosive effect, as "confidence in the model" rather than "learning from modelling" becomes a key story within the organisation.

As long as regulatory approval of the internal model is business critical, modellers are required, along with other professionals, to make approval happen. Experienced modellers are a scarce resource, paid to deliver confidence, not doubt. Some may not risk undermining their role and status within the organisation by being fully open with management about the uncertainties that their work is riddled with. A deeply embedded risk culture is needed to avoid such perverse incentives. It seems that Solvency II not only requires, but also necessitates strong corporate governance.

The regulator's role is no less challenging. Suppose a firm decides to be completely candid about the potential for model error and demonstrates to the regulator a great sensitivity of capital requirements to unverifiable statistical assumptions. It is not clear how this information should be dealt with. While honesty will surely be appreciated, once the potential inaccuracy of model output is on the record, it cannot be ignored. It seems that to allow for the possibility of even a single internal model being approved, regulators need to be quite tough in the overall supervisory review, but somewhat tactful in explicit enquiries about the accuracy of model outputs.

So, what should we do? Breaking the nexus between regulatory capital requirements and statistical risk modelling is not a realistic choice. There is certainly not much appetite in the insurance market for a return to rules-based, one-size-fits-all regulation. Problems such as the ones described in this article are the price we pay for the benefits of principles-based regulation. The emphasis that Solvency II places on model validation, hard as it may be, should be welcomed. But a thorough validation process forms primarily evidence on the quality of a firm's reasoning around risk; assurance over the accuracy of model outputs can only be of secondary importance. In that context, the model approval process is a platform for having informative conversations about risk. Pretending that model outputs at the 1-in-200 years level can be meaningful may be the premise of such conversations. Even if the pretence is somehow useful, we should question whether it is necessary.