Stealing others' strategies: lessons on model risk

David Jamieson Bolder[1]

Abstract

Mathematical and statistical models are used extensively in decision-making across the financial industry. Complex, high-dimensional decisions are generally improved through the addition of mathematical structure and rigour. Models, as a simplified representation of an intricate reality, are nonetheless not without their own risks. Weak or unrealistic assumptions, poor calibration of parameters, erroneous implementation and improper interpretation of results can lead to poor decisions, financial losses and, in extreme cases, systemic shocks. Negative financial outcomes resulting from such events are referred to as model risk. Model risk is essentially a quality control problem. The finance industry is not alone; other areas of endeavour are confronted with similar challenges. Quality control issues arise in scientific discourse, the production of computer software and the construction of comprehensive reference material. Logical, empirical or theoretical errors in these areas are a constant danger. Such a possibility has led to the consequent development of strategies to manage those risks. Risk management professionals also draw numerous lessons from these other disciplines. Examples are peer review, actively seeking multiplicity of perspective, openness regarding technical choices, and independent verification of findings. This perspective provides not only useful motivation for the burgeoning field of model risk management, but similar measures might represent useful additions to the model risk manager's toolkit.

"Truth is the daughter of debate, not of sympathy"
– Gaston Bachelard

[1] Head of Capital and Portfolio Credit Risk, Risk Management Unit, Nordic Investment Bank.

Introduction

Mathematical and statistical models are used extensively in decision-making across the financial industry. They make an appearance in pricing and hedging financial instruments, the computation and attribution of market and credit risk, the estimation of economic capital requirements, the determination of strategic portfolios, and are increasingly employed in the assessment of operational risk. Complex, high-dimensional financial decisions are generally improved through the addition of mathematical structure and rigour.

Models, as a simplified representation of an intricate reality, are nonetheless not without their own risks. Weak or unrealistic assumptions, poor calibration of parameters, erroneous implementation, and improper interpretation of results can lead to poor decisions, damage to an institution's reputation, financial losses and, in extreme cases, to systemic shocks. Negative financial and reputational outcomes resulting from such events are referred to as model risk.

Model risk has, as a concept, been present in the world of finance and economics for the last 20 to 30 years. A famous critique by Lucas (1976) appears to have been something of a starting point, followed by Derman (1996) and, a few years later, by Rebonato (2003). Despite financial crises precipitated at least in part by financial modelling failures – such as the 1987 stock-market crash or the Long-Term Capital Management (LTCM) situation – only recently, as chronicled by Brown et al (2015), in the aftermath of the 2008 Great Financial Crisis, has model risk captured the collective attention of the financial community. OCC (2011) is an example of one of the first regulatory forays into this realm. It introduces the notion of effective challenge as well as some basic governance structures to mitigate model risk. Other comprehensive works, such as those by Morini (2011) and Danielsson et al (2014), have followed in recent years.

While OCC (2011) established – or, at least, initiated – the rules of the game for regulated entities, model risk remains a challenge for non-regulated entities or for model-based analysis performed outside of prescribed regulatory oversight. Many unregulated public institutions such as international organisations, central banks and sovereign wealth funds are still struggling with an appropriate response to this dimension of risk. Even within the regulatory realm, model risk activities appear to follow more the letter than the spirit of the law. Despite a broad-based understanding of the risk posed by poor use and the implementation of financial models, a general framework for its consideration eludes us.

This short paper argues that model risk is, in fact, not a new problem faced solely by financial-market participants. It is closely related to the notion of quality control in scientific discourse, software development, and compilation of reference materials. These disciplines are, as will be argued in the following discussion, fraught with issues closely aligned with the notion of model risk facing the financial community. Logical, empirical or theoretical errors are a constant danger. More importantly, strategies have been developed to minimise this danger and, upon closer inspection, there appear to be important commonalities among these disciplines.

Something can, therefore, be learned from examining the historic response of the academic, software developer, and encyclopaedia writer to these risks. This perspective offers not only ideas for the mitigation of these risks, but also a conceptual framework for their management. This work walks through the pertinent elements in the development of these strategies. It uses these ideas to construct lessons for modern risk managers, which include the criticality of openness regarding assumptions, techniques, and parameter determination, the necessity of peer review, the essential role of documentation, the role of a healthy degree of scepticism, and the general need for an inclusive, open, and honest discussion of ideas. In a proprietary financial setting, attaining the requisite level of openness regarding modelling decisions can be challenging. Similar constraints, however, do not apply to public institutions. This may suggest slightly different policy responses for these different types of organisations, but the underlying challenges and mitigation techniques remain the same.

A key driver

When one thinks about risk management, it is natural to focus on market, credit, or operational risk. We have a reasonable understanding of the underlying factors driving these types of risk. Market risk is, by and large, well described as the potential for financial loss arising from changes in systematic variables such as interest, exchange rates and general levels of credit spreads. Credit risk stems principally from idiosyncratic or specific factors influencing the ability, or willingness, of credit counterparts to meet their obligations.[2] Operational risk, in contrast, is somewhat trickier to handle. Hemrit and Arab (2013) argue that operational risk "arises from lack of awareness and deficiencies of skill in detecting threats related to [operational] risk." While certainly helpful, it lacks the concreteness of market and credit risks.

Model risk finds itself in a similar conundrum. Identification of its underlying drivers is not obvious. It is tempting to identify poor, or potentially dangerous, modelling choices such a linearity, time invariance or Gaussianity. This viewpoint certainly has some merit, but there are cases where such assumptions are entirely justifiable. A model needs to be fit for its purpose and since various models serve varying objectives, clear prescriptive, technical guidelines for model construction have limited value. Restricting our scope to these real, but purely technical, dimensions of model risk feels a bit like confounding the symptoms and sources of an illness.

An alternative logical path provides some insight into this problem. Let us pose the question: what are the optimal conditions for model risk? In other words, when are we most worried about model risk? While a definitive answer will be hard to find, we can nonetheless formulate a reasonable response. Model risk is, potentially at least, greatest when the details of a given model are unclear, unknown, unchallenged, or undocumented. By their very construction, models are approximations. We know a priori that they will often be wrong. It is thus foolish, and dangerous, to talk of a perfect model; such a thing does not exist. The argument is that model risk is most acute when model users have a limited understanding of key model choices and how they might lead the model astray.

2 This is, of course, not perfect. Market risk has idiosyncratic elements, while systematic risk also arises in credit risk.

Model risk, like mushrooms, grows best in the dark — Figure 1

Photo by: Aleksey Gnilenkov / Flickr.

Model risk can thus be considered – only partly facetiously – like a mushroom: ie it grows best in the dark! What is the implication of this conclusion? Models constructed by individuals or small teams, with little external oversight or review, are consequently the most potentially dangerous. A key underlying driver of model risk is limited model knowledge, exposure and understanding. This is broadly consistent with a powerful statement from Morini (2011): "model misunderstanding is the core of model risk." To abuse the English language somewhat, we might refer to this as "black-boxedness".

A bit of caution is required. Open, well documented, thoroughly discussed models are, of course, not free of risk. Every player in a modelling process can have a thorough and nuanced appreciation of the key assumptions, but the implementation code could possess serious errors, or the parameters might be improperly estimated. It seems reasonable to conclude, however, that the model owners are better equipped to identify such shortcomings when they have a thorough understanding of the model. Black-boxedness is thus not an exhaustive descriptor of the sources of model risk, but rather an important one.

The point is that models have less risk when they are well understood. This is not a particularly novel statement and it certainly lies implicitly at the bottom of much, if not all, of the extant work on model risk. It is rarely, however, if ever, explicitly identified as a key driver of model risk. Such a statement has value, because as in the market and credit risk settings, identification of the drivers of risk is the first step in the development of strategies for its mitigation.

The main idea

Having identified a key driver of model risk, the next natural step is to seek mitigation strategies.[3] Our model risk driver provides, in this regard, a benefit beyond the mere conceptualisation of model risk. It allows us to identify endeavours, or disciplines, that face a similar situation. To be clear, we are thinking of complex, multidimensional tasks – with clearly defined inputs and outputs – typically performed by individuals or small groups. The principal objective of these tasks is to enhance our understanding of some phenomena and take decisions about it. While perhaps not the formal definition of a model, this description gets to the heart of the matter.

What other endeavours follow this pattern? An obvious choice is scientific discourse. Science is principally about developing testable theories to contribute to our understanding of the world around us. On a dramatically less general level, software development also falls into this category. Effective software is certainly complex and multidimensional; moreover, it is often employed to solve problems or take decisions. A final example relates to the construction of comprehensive reference material; an encyclopaedia is a good example. These are three disciplines, of increasing specificity, that have strong qualitative similarities to the construction of financial models.

In these three undertakings – as in model construction – logical, empirical or theoretical errors are a constant danger. A variety of strategies have been employed in the aforementioned activities to mitigate such errors. Some of these solutions are long-standing, whereas others are more recent and innovative. They nonetheless share some important commonalities and, ultimately, model risk management can be inspired from examining their responses. Some of the advice is similar to that found in the current model risk literature, whereas some is, hopefully, somewhat original. In all cases, it appears to provide something of a foundation for discussion of the model risk problem.

In the following sections, the plan is to examine each of these three areas separately and attempt to identify the common aspects of their strategic response to the possibility of important and costly errors. The final section will pull these ideas together and offer something of a loose framework for model risk management. The paper concludes, in the spirit of the underlying thesis, with a careful examination and identification of potential shortcomings of this approach.

Scientific discourse

"Take nobody's word for it (Nullius in verba)."
– Motto of the Royal Society

Let us begin with some brief historical perspective. In 1453, Constantinople fell to the Ottoman empire. This is commonly accepted, among historians, as the beginning of the Renaissance period. Around the same time, in the 1440s, Johannes Gutenberg developed the first functional printing press. These two fundamental elements set the

[3] One might argue that risk measurement is a superior starting point, but a bit of reflection suggests that our notion of black-boxedness does not immediately lend itself to direct quantification.

stage for the Enlightenment in the 1600s and 1700s, where science and philosophy took centre stage. Leading thinkers – such as Locke, Leibniz, Descartes, Bacon, and Newton – solidified the scientific method and the notion of empiricism. Thus, in the lead up to the scientific revolution, a modern way of thinking was emerging. New, controversial, and world changing ideas were in circulation. All of these facts are well known.[4] Let us ask an interesting question related to our model risk problem: how did they manage the quality of these ideas?

By all accounts, a free and ready exchange of ideas was essential. Hatch (1999) indicates that, in the first half of the 1600s, such an exchange of ideas occurred through private societies (ie salons) and correspondence networks. The mechanics are fascinating. In a time of potential church censorship, a clever approach was necessary. The response, it appears, was the so-called republic of letters with intelligencers acting as hubs and de facto editors. Letters, mixing mundane discussion of daily life with logical exposition, were employed to communicate, critique, and develop scientific ideas; this was something of an ink-and-paper precursor to the World Wide Web. In the latter part of the 1600s, there was a gradual movement to state-sponsored academies and printed journals. Both of these approaches were a response to a common problem: how to overcome geographical, cultural and political boundaries to communication.

Letter writing had its limits. In the mid 1600s, as chronicled by Connor and Robertson (2004) and others, a number of intellectuals gathered in London to discuss philosophy. Schonland (1958), among others, refer to this as the Invisible Society. This ultimately led to the establishment of the Royal Society and, within a few years, the first scientific journal entitled Philosophical Transactions.[5] This was an organised and formal response to a priority of claim, a dissemination of ideas and the creation of repository for the accumulation of knowledge.

Spier (2002) indicates that it took about a hundred years before the Royal Society started to internally review publications. In 1831, the first peer-review reports were produced. Nevertheless, it was not until after WWII that double-blind peer review became the standard. Peer review is expensive, time-consuming, and places burdens on the academic community. It appears, however, to be a useful strategy against black-boxedness. That scientific ideas require replication and verification is a key tenet of the scientific method. Interestingly, given commercial and career incentives for scientists, there are ongoing attempts to circumvent these constraints. The Economist (2018a, 2018b) chronicles a trend towards predatory journals that seek to undermine these controls. The quality of these unreferenced journals naturally suffers and the risk of dissemination of misleading or erroneous ideas is heightened. This is an ongoing challenge for the scientific community and speaks to the importance of quality control.

Over the centuries, a number of concrete strategies have evolved to further scientific thought and exercise a degree of quality control. Critical, therefore, to the effective exchange of scientific ideas in academia are:

- discussion and formal presentation of one's ideas to qualified and discerning colleagues;

[4] See, for example, Wootton (2015) for a detailed description of the scientific revolution.

[5] It is simultaneously amazing and humbling to look at the list of ground-breaking ideas first published in the pages of this journal.

- publishing one's results in the public domain; and
- subjecting one's analysis to independent peer review.

Scientific thought certainly existed without these concepts, but all of these ideas seek to bring ideas out of the dark. The bright light of constructive criticism helps identify, but not eliminate, weak assumptions and logical inconsistencies.

It is relatively easy to see that all of these elements have links to model development by finance practitioners. The model risk analogue of letter-writing and societies and journals is discussion with colleagues. Publication would take the form of model documentation as well as internal or external publication of results. Peer-review, as already indicated, is quite obviously equivalent to performance of an independent validation exercise; this is, indeed, the idea behind OCC (2011)'s idea of effective challenge. The guiding principle is simple: the greater the audience of competent, discerning eyes, the higher the degree of quality control. This has a Darwinian dimension; those ideas that survive harsh but honest criticism, continue, while those that do not, fail. This is essentially the opposite of black-boxedness.

A potential challenge

"I often compare open source to science. To where science took this whole notion of developing ideas in the open and improving on other peoples' ideas and making it into what science is today and the incredible advances that we have had. And I compare that to witchcraft and alchemy, where openness was something you didn't do."
– Linus Torvalds

Although the academic model is encouraging and enlightening, it is based on the open exchange of concepts and results. Publicly sharing ideas in the financial realm can have costs. Some models have proprietary value and sharing one's development could lead to free-riding and loss of economic rents. It might even, in the limit, reduce incentives to innovate. Public-policy models, if exposed, might change the behaviour of market participants. A classic example is the Bank of Canada's so-called monetary conditions index – Freedman (1995) is an excellent overview. The publication of specific parameters for this measure undermined, in this case, its effectiveness as a monetary policy tool. Transparency in public policy is a subject of ongoing debate. Economists (see eg Jensen (2002)) have established a trade-off between credibility and flexibility. The pendulum has certainly swung towards transparency over the last few decades, but full transparency does not appear to have been viewed as optimal.

There is an interesting counterexample to this argument hiding in plain sight. Innovation and progress need not necessarily require propriety interests. There are examples of the two living in cooperation within the software industry. Lerner and Tirole (2002) indicate that early operating systems and internet protocols were developed collaboratively in academic or research centres (ie Berkeley, MIT, Bell Labs, Xerox PARC). This involved extensive sharing between institutions. In the early 1980s, a desire to enforce intellectual property rights emerged. Richard Stallman and others, in response, launched the GNU Project and the Free Software Foundation. The open-source software license permits code to be used, modified, or shared by other developers. In this manner, many open-source software products can be viewed as collaboratively developed. The guiding idea behind this movement is thus to maintain the collective spirit found in early development circles.

Despite initial scepticism and even derision, the open approach to software development has proven highly successful. Many commercial entities not only employ open software tools – such as Linux, Apache, Perl, Python, or a broad variety of compilers – in their day-to-day activities, but actively encourage staff to participate in the development process. Against all expectation, open-source programming has proven a success. Von Krogh and von Hippel (2006) actively seek to highlight the utility of aspects of this movement for many other fields in terms of governance, organisation, and innovation.

Software development, like model construction, is a technical and complex undertaking, which can often be broken into smaller, modular tasks. The proprietary model, of course, also works reasonably well in this field. Perhaps surprisingly, the open-source approach works as well or, on occasion, even better. Furthermore, it appears to offer the strongest benefits when users are most sophisticated (ie Apache, Linux). This point makes it particularly interesting for complex financial applications. Key elements involve putting source code in the public domain, sharing of knowledge, and exposure of design to critical review. It also appears to foster creativity and innovation. This concrete precedent suggests that exposing, in a cautious and controlled manner, internal models to a larger group can thus potentially enjoy these benefits.

The power of crowdsourcing

"You shouldn't use anything as the sole source for anything."
– Jimmy Wales

Creating an encyclopaedia is also a technical endeavour, which is certainly prone to error. Reference books incorporate a multitude of facts and are relied upon by professionals of all ilks for important background information. Wikipedia, the online encyclopaedia, was launched in 2001. Content creation, in this set-up, is placed under the control of registered volunteer editors.[6] At first blush, this does not seem to be a particularly compelling organisational structure. Individual entries are often constructed by multiple individuals with oversight and a control structure involving the entire editing community. Over the last few decades, however, it has become a preeminent source of information.

Popularity is not, it should be stressed, a universal criterion for success. Although there are questions about its accuracy – and a few horror-stories – there is also evidence that it works quite well. Giles (2005) suggests, for the sciences, that it is roughly as accurate as its commercial competitor, Encyclopaedia Britannica. It is also significantly more reactive to new information, events, and ideas. Brown (2011) further holds that it demonstrates surprising accuracy for political science.

A word of clarification is warranted. This work is certainly not advocating Wikipedia as a research platform. Greenstein and Zhu (2012) raise important questions about its biasedness, which is particularly acute with controversial, subjective, or difficult to verify topics. It is also constructed to place equal weight on expert and non-expert contributions, which seems conceptually difficult to defend.

[6] Interestingly, Richard Stallman was an early proponent of this idea. For more on this point, please refer to Stallman (2000).

These important caveats notwithstanding, it is a successful model. What can we learn to use in model risk management? Once again, openness of ideas, exposure to criticism, and engagement with a broader, discerning community is at play. Wide-ranging access and openness appear to work, in this case, at least as well as small groups of experts. The key lesson, it seems, is the usefulness of multiplicity of perspective. This argues, in the financial setting, for the incorporation of the viewpoints of both technical and non-technical staff into the modelling process.

Beginnings of a framework

We have identified three related, but impressively disparate, complex, technical and error-prone endeavours. Each of these activities possess important qualitative similarities to those involved in the construction of financial models. They cover the spectrum of high-level conceptual to low-level implementation details, which is precisely the scope of complexity found in most financial models. All also face significant quality control issues that can importantly undermine their success. Each of them, even in the face of proprietary interests, have adopted quality control strategies that involve a high degree of openness and exchange of central ideas.

Table 1 highlights a broad range of mitigants aimed at our key driver of model risk: the black box. Each of these elements is inspired from our examination of these conceptually similar disciplines. The scientist's quality control dilemma, in particular, is a rich source of inspiration. Publication, peer review, conferences, and formal degrees all play a role in ensuring a healthy scientific discussion. Model risk managers can draw from this example. Independent model validation is already a key pillar, but more effort can be made to publish methodologies and results and engage with other well trained modellers both inside and outside one's institution.

Organising parallels — Table 1

Parallel	Model risk
Publication	Documentation and white papers
Peer review	Independent model validation
Conferences and societies	Internal and external networks
Courses	Subject-matter training, knowledge transfer
Degrees and diplomas	Professional designations
Openness	Inter-institutional cooperation
	Code sharing and repositories
Crowdsourcing	Multiplicity of perspective

The underlying table attempts to draw a number of practical parallels between extent strategies in scientific discourse, software development and construction of reference material with possible responses to financial model risk.

Knowledge-sharing and transfer are popular concepts that appear, in recent years, to have attained buzzword status in business circles. It is, in the experience of the author, nonetheless quite rare to find model owners formally presenting and defending their work to a broad audience of peers. It would be unthinkable, by contrast, for academic discourse to proceed absent such formal presentation of results. There are certainly costs to both presenters and attendees. Distilling and digesting complicated ideas is hard work, which requires consumption of scarce resources. It is, however, a key mitigant of model risk, which if we deem it serious, makes this a worthwhile undertaking.

Internal presentation is already a useful step forward, but widening the audience to external parties improves the situation further. The challenge is that, quite often, sharing modelling ideas can potentially create business conflict. Managing this proprietary element is, in general, difficult although this aspect is, for public institutions, somewhat attenuated. The software-development and encyclopaedia-construction examples nonetheless suggest that there is the possibility of overcoming this challenge. The open-software movement indicates that inter-institutional cooperation can work. The financial community could, for example, look to cooperate in a limited manner. Sharing of model code through repositories is another step towards mitigating the black-boxedness of one's internal models; this might be most potentially helpful for complex, non-proprietary modelling sub-tasks.

Even faced with tight proprietary constraints, there is a case for enhanced openness within one's organisation. Too often, in the author's experience, modelling choices are viewed as the technical domain of a select few. Rarely are such decisions questioned, not to mention discussed, in senior management circles. Such openness has, of course, costs for both non-technical senior and technical staff. Senior management must be encouraged, or incentivised, to be more demanding in their consumption of modelling outputs and discerning of their inputs. Indeed, more time should be allocated to key inputs, assumptions, and methods and less to final results. Technical staff, for this to succeed, also need to become better communicators. Taking both of these groups out of their respective comfort zones is entirely consistent with the general thesis of this work.

The Wikipedia case appears, along these lines, to stress the power of multiplicity of perspective. This idea has its origins in psychology. Campbell and Fiske (1959) formally introduced the idea of validating results predicated on having demonstrated it along a variety of different dimensions. Although not as formal, the strength of Wikipedia appears to stem, at least in part, from this same fundamental idea. It is a kind of diversification of viewpoint that also argues for bringing non-technical and technical players from a range of areas into the details of the modelling process. Model risk managers, with a finance professional's appreciation for the benefits of diversification, would be well advised to exploit this same idea.

The commonality among these three ideas is openness. It is easy, and common, to make important errors in the performance of technical tasks. Few are better placed to identify these errors – be they methodological, conceptual, or simply operational – than fellow finance practitioners. In some cases, deep technical knowledge is required – to opine, for example, on the reasonableness of a mathematical methodology. In other cases, it is conceptual understanding of the problem that matters. Non-technical portfolio managers can often provide invaluable insight into the veracity of key modelling assumptions. If one does not share one's results, these insights cannot be harvested. The consequence is black-boxedness and a commensurate increase in financial or reputational loss inherent in model risk.

Walking the talk

"You're worse off thinking you have a model and relying on it than [...] simply realizing there isn't one."

– Emanuel Derman

Much of the preceding discussion is fairly commonsensical. The intent is to offer a conceptual structure to what is a difficult and important challenge for the financial community. Identifying a key driver, finding related disciplines, and extracting lessons hopefully has some value in this regard.

As a final step, it is useful to apply the previous appeal for openness and critique to this analysis. One could legitimately argue, for example, that a certain amount of cherry picking, or selective choice of examples, was involved in this argumentation. Openness seems to work in these settings, but one might argue that other areas of technical endeavour operate in the opposite direction. Law, medicine and engineering involve, on some dimensions, tasks that are, qualitatively at least, similar to the construction of financial models. Not being an expert in these fields, it is difficult to comment definitively. These conclusions could thus come with some important conditionality. All, however, are subject to varying degrees of oversight. The obligatory audit of financial statements and the appeal process for judicial decisions would appear – superficially, at least – to resemble the notions of openness, challenge and healthy criticism suggested in the previous analysis.

Another reasonable point of attack relates to one of the key assumptions inherent in this treatise: that black-boxedness is a key driver of model risk. If this is not true, then the logical arguments potentially collapse. An alternative driver is related to an important idea credited to Knight (1921); which is often referred to as Knightian uncertainty. These are the unknown unknowns famously entering the public consciousness and lexicon via former US Secretary of Defence Donald Rumsfeld. Such risk cannot be measured and, as such, cannot be mitigated. An explicit assumption of this analysis is that financial modelling deals principally with known unknowns. In a positive sense, this appears to be true. Most financial modelling involves the specification of value and loss distributions either estimated or calibrated to market conditions under different probability measures. This, by construction, leaves little scope for Knightian uncertainty; since such outcomes are out of the range of our current information sets, we cannot hope to include them in this way. From a normative perspective, however, this is thoroughly debatable. Taleb (2007), among others, criticise current modelling practices harshly for their failure to incorporate these risks. These are valid criticisms and, quite clearly, this challenge is not directly resolved within this framework. At the same time, unless one insists upon taking a very extreme stance on the underlying drivers of model risk, it also does not negate the underlying thesis. On the contrary, it opens the door for sceptical parties, such as Taleb (2007), to vehemently argue for the explicit incorporation of Knightian uncertainty.

Yet another competing driver of model risk is implicit in the work of Derman (1996, 2009, 2012): an over-reliance or overconfidence in modelling results. His work is logically compelling, practically useful and always colourful. Derman (2009) states, for example, that "the greatest danger in financial modelling is the age-old sin of idolatry." By this, he means he views the key driver of model risk as a failure to appreciate the inherent fallibility of a simplified mathematical construct (ie a model) attempting to describe an unknowable reality. Opening one's models to organised external and internal critique, along the previously discussed lines, does not directly address this issue. An argument could even be made that an intense process of ongoing discussion and refinement might even accentuate this risk by lulling model users into a sense of complacency. This would be false. The spirit of the preceding arguments is that a healthy degree of scepticism on the part of the modeller – and his/her actively solicited colleagues in different domains – acts not only to mitigate logical and practical errors, but works to avoid the very hubris Derman (2009) so rightly cautions us against. It is not a one-off exercise to rubber stamp a specific model, but rather an ongoing commitment. Thus, although this paper differs slightly from Derman (2009) in terms of the key underlying driver of model risk, the recommended strategic response appears appropriate in both cases.

Concluding thoughts

"One special advantage of the sceptical attitude of mind is that a man is never vexed to find that after all he has been in the wrong."

– William Osler

Through identifying a key driver of model risk and highlighting a number of qualitatively similar activities facing similar risk drivers, a prescriptive set of strategies has been identified for its management. Openness, sharing of knowledge, and explicitly exposing one's ideas and implementation to criticism form the foundation of the recommended strategic response. Independent validation of one's models is one important step in this direction, but one can, and should, seek to go much further. Formal internal and external presentation of ideas, code sharing and repositories, limited publication of results, inter-institutional cooperation, and subject-matter training are additional tools. Based on the previous logical arguments, these policies or initiatives could usefully be added to the model risk manager's toolkit. Like all elements of risk management, however, the ultimate success will rely upon the hard work and vigilance of all those involved.

References

Brown, A (2011): "Wikipedia as a data source for political scientists: accuracy and completeness of coverage", *Political Science and Politics*, vol 44, no 2, pp 339–43.

Brown, J, B McGourty and T Schuermann (2015): "Model risk and the great financial crisis: the rise of modern model risk management", *Wharton Financial Institutions Center Working Paper*, vol 15, no 1.

Campbell, D and D Fiske (1959): "Convergent and discriminant validation by the multitrait-multimethod matrix", *Psychological Bulletin*, vol 56, no 2, pp 81–105.

Danielsson, J, K James, M Valenzuela and I Zer (2014): "Model risk of risk models", *FEDS Working Paper*, no 2014-34.

Derman, E (1996): *Model risk*, Goldman Sachs quantitative strategies research notes, April.

——— (2009): "Models", *Financial Analysts Journal*, vol 65, no 1, pp 28–33.

——— (2012): *Models behaving badly: why confusing illusion with reality can lead to disaster, on wall street and in life*, Free Press, New York.

Freedman, C (1995): "The role of monetary conditions and the monetary conditions index in the conduct of policy", *Bank of Canada Review*, Autumn.

Giles, J (2005): "Internet encyclopedias go head to head", *Nature*, vol 438, no 15, pp 900–1.

Greenstein, S and F Zhu (2012): "Is Wikipedia biased?", *The American Economic Review*, vol 102, no 3, pp 343–8.

Hatch, R (1999): "Correspondence networks", in W Appelbaum (eds) *The Scientific Revolution*, Garland Publications.

Hemrit, W and M Arab (2013): "The major sources of operational risk and the potential benefits of its management", *Journal of Operational Risk*, vol 7, no 4, winter, pp 71-92.

Jensen, H (2002): "Optimal degrees of transparency in monetary policymaking", *The Scandinavian Journal of Economics*, vol 104, no 3, pp 399–422.

Knight, F (1921): *Risk, Uncertainty and Profit*, Houghton Mifflin Co, Boston.

Lerner, J and J Tirole (2002): "Some simple economics of open source", *The Journal of Industrial Economics*, vol 50, no 2, pp 197–234.

Lucas, R (1976): "Econometric policy evaluation: a critique", in K Brunner and A Meltzer (eds) *The phillips curve and labor markets*, pp 19–46, Carnegie-Rochester Conference Series on Public Policy.

Morini, M (2011): *Understanding and managing model risk: a practical guide for quants, traders, and validators*, Wiley.

OCC (2011): *Supervisory guidance on model risk management*, technical report, Office of the Comptroller of the Currency, April.

O'Connor, J and E Robertson (2004): *The royal society*, http://www-history.mcs.st-and.ac.uk/Societies/RS.html, School of Mathematics and Statistics, University of St Andrews, Scotland.

Rebonato, R (2003): "Theory and practice of model risk management", *Quantitative Research Centre (QUARC) Technical Paper*.

Schonland, B (1958): "The invisible college", *Royal Institute of International Affairs 1944-*, vol 35, no 2, pp 141–50.

Spier, R (2002): "The history of the peer-review process", *Trends in Biotechnology*, vol 20, no 8, pp 357–8.

Stallman, R (2000): *The free universal encyclopedia and learning resource*, https://www.gnu.org/encyclopedia/anencyc.txt.

Taleb, N (2007): *The black swan: the impact of the highly improbable*, Random House, New York.

The Economist (2018a): *Some science journals that claim to peer review papers do not do so*, June.

——— (2018b). *What are "predatory" academic journals?* July.

Von Krogh, G and E von Hippel (2006): "The promise of research on open source software", *Management Science*, vol 52, no 7, pp 975–83.

Wootton, D (2015): *The invention of science: a new history of the scientific revolution*, Harper Perennial, New York.