Post-crisis financial risk management: Some suggestions

Received (in revised form): 17th November, 2009

Riccardo Rebonato
is Head of Front-Office Risk Management and Client Analytics for GBM, RBS, Visiting Lecturer at Oxford University (mathematical finance) and Adjunct Professor at Imperial College (Tanaka Business School). He sits on the Board of Directors of ISDA and on the Board of Trustees for GARP. He holds doctorates in nuclear engineering and in solid state physics and was a research fellow in physics at Corpus Christi College, Oxford, UK.

Royal Bank of Scotland, 135 Bishopsgate, London EC2M 3UR, UK
Tel: +44 (0)20 7085 5070; Fax: +44 (0)20 7085 6894; E-mail: riccardo.rebonato@rbs.com

Abstract This paper argues the case that, when it comes to the management of financial risk, it is essential to provide interpretative models of reality (data do not 'speak by themselves'); that these interpretative models are not unique; that too little attention has been devoted to explaining rather than describing; and that the existence of a number of competing views of the world can give rise to co-ordination among traders. The implications of this for the distribution of returns are highlighted. The limits of purely statistical descriptions of risk in terms of marginal or joint distributions of risk factors are also discussed.

Keywords: *risk management, co-ordination, herding, bubbles, return distributions*

WHAT WENT WRONG?

People learn from their mistakes — or so the saying goes. If this is indeed true, then risk managers must be among the wisest folks on the planet. But what is the source of this wisdom, and what has been learned? More concretely: had any of the various novel ideas that are now being touted as the new best practice in risk management been established back in, say, late 2005, would things have turned out any better? If not, how can we try to avoid repeating the same mistakes?

This paper sets out to answer these questions by calling for a rethink in the way in which financial risk is controlled. To articulate the analysis and proposals, the paper first puts into context the current thinking about financial risk management; it then highlights some of the shortcomings of this approach; finally, it offers suggestions to help regulators and the industry rethink the risk management practices in which they are engaged.

HOW DID WE GET HERE?

Great upheavals demand a radical revision of the way in which people make sense of reality. Finance and economics are no exceptions. Indeed, the events of the late 2000s are bringing to the fore a critical review of many aspects of the neoclassical economics orthodoxy — its reliance, for instance, on the ability of rational investors and private sector firms to self-regulate economic activities. Vigorous criticism of the orthodox way of thinking is now coming not from conspiracy theorists or

mavericks 'out on the fringes', but from mainstream economists and central bankers: 'The modern risk management paradigm held sway for decades. The whole intellectual edifice, however, collapsed in the summer of last year.'[1]

Change is in the air. For this reason it is particularly important to put the current thinking about quantitative risk management in its intellectual context.

The current situation was not arrived at by accident or 'from scratch', ie by thinking in the abstract, 'how should financial risk be managed?' The logical underpinning of quantitative risk management as it has been practised in the run-up to the crisis is intimately enmeshed with the prevailing (neoclassical) conceptual framework of financial economics. Take, for instance, one of the cornerstones of the neoclassical finance edifice — Markowitz's portfolio diversification. It shares important elements with the risk management statistics, such as value-at-risk and economic capital, that underpin the estimation of regulatory capital: Rational Man, equipped with a perfect computing machine, is supposed to be able to estimate all the statistical properties of return distributions (down to the highest percentiles in the case of economic capital), and to use this information to make financial decisions under uncertainty. Unfortunately, both the neoclassical finance theory and contemporary risk management 'assume that we have a knowledge of the future of a kind quite different from that which we actually possess'.[2]

This way of looking at financial risk has become the financial orthodoxy. But if the global economy is at a juncture when the very foundations of the neoclassical finance edifice are being questioned — which would seem to be the case — then this is also a good time to revisit the way in which financial risk is controlled.

Alan Greenspan, a strong proponent of the Rational-Investor school, gives another insight into what went wrong:

> 'The extraordinary risk-management discipline that developed out of the writings of the University of Chicago's Harry Markowitz in the 1950s produced insights that won several Nobel prizes in economics. It was widely embraced not only by academia but also by a large majority of finance professionals and global regulators. But in August 2007, the risk management structure cracked. All the sophisticated mathematics . . . rested on one central premise: that the enlightened self-interest of [market agents] would lead them to maintain a sufficient buffer against insolvency.'[3]

It is silly to claim that the current situation occurred because the assumptions that underpin this way of looking at financial problems are 'wrong'. Model assumptions are always wrong, and, to use McKenzie's metaphor,[4] a model is neither a mirror nor a camera. It may well be that, when it comes to neoclassical finance, practitioners have been beguiled by the elegance of the construction and begun taking the scaffolding (ie the assumptions) too seriously and literally. But, in the end, it is far too simplistic to claim that at the root of all evil were 'unrealistic assumptions' or 'blind faith in normal distribution'. It is not reasonable to expect a model with 'perfect assumptions' (whatever that means). Risk managers would do well to appreciate

that, in finance, models interact with reality in a more complex and feedback-prone manner than in other disciplines. It is to this aspect that the paper now turns.

WHAT IS DIFFERENT IN THE PROPOSED APPROACH?

Recent events suggest that financial risk may be better managed if the following aspects of model–reality interaction are considered:

- Data (eg return distributions) do not speak by themselves, but make sense only in the context of models of reality. Model-free statistical analysis of past data is of little use in reaching effective risk-management choices for events of systemic relevance. Without a plausible structural model, data analysis is blind.
- It is vital to recognise the plurality of plausible interpretative models of reality, ranging all the way from fully-fledged 'structural' models, to the heuristics and rules of thumb used by traders.
- The simultaneous coexistence of these interpretative models can give rise to coordination among market participants and to positive feedback mechanisms.
- When it comes to risk management, the search for a unique 'true' model (ie a unique 'correct' mapping from new information to price changes) may be misguided, futile, and at times even harmful.

It is worth elaborating on these points. Within the economic sciences, the current orthodox view assumes a very inflexible relationship among data, models and how economic agents combine the two. There is but one true model of financial reality, and this one true model is known by all the agents in the economy. In Sargent's words, there is a 'communism of models': 'all agents inside the model, econometricians and God all share the same model'.[5] The existence of different opinions among rational individuals about, say, the future prospects of a business project, makes common sense, yet is only awkwardly accommodated by mainstream financial theory.

This position can be, and has been, criticised,[6] but at least it has a respectable, if slightly otherworldly, logical coherence. In what follows, this paper argues that *for risk management purposes*, this view is unhelpful. Indeed, risk managers must actively entertain the possibility that different agents may have different interpretative models of reality. 'Model risk' understood in this sense is one of the most important, and least recognised, risk factors.

There is a second unfortunate feature of current risk management thinking. In common with finance theory, the statistical distribution of returns is central to quantitative risk management — if anything it has assumed a status of almost religious idolatry. However, the debatable (but at least logically coherent) link in current financial thinking between these returns distributions and decisions via expected utility maximisation has been either forgotten or diluted beyond recognition. In risk management, the accepted (implicit) assumption has been that, once obtained, impossibly accurate statistical distributions of returns would effectively 'speak by themselves' and self-evidently point to the correct risk management choices. Simplistic, *ad hoc*, and cognitively unappealing rules of thumb have been offered as crutches to assist the decision-making process. High percentiles — the higher the

better — have been presented as the be-all-and-end-all of risk management (Rebonato offers a non-technical discussion of these points elsewhere[7]).

The lack of logical coherence is regrettable. However, this is not to suggest that the way out of the current risk management impasse is the integration of the statistical data analysis with a single correct model of reality (be it the neoclassical one or any of the many alternatives that are being, and have been, proposed). As said previously, what is more productive *for risk management purposes* is awareness of the different plausible interpretative models of reality alluded to above, and how different agents, by embracing these different plausible models, may react to new information.

This is the case because the feature strongly affects market dynamics. Market agents are exposed to an information flow far greater than any human being can possibly digest in the time required to execute the next trade. Traders therefore need heuristics, rules of thumb and simplified models that can help them make sense of this flood of data. Most of the time, novel pieces of information will be interpreted in the light of the current market consensus. Occasionally, however, particularly ambiguous or 'disturbing' new data will prompt a 'Gestalt switch' from one interpretative model to another.

Given this plurality of models, the greatest danger for any one trader faced with a new piece of information is not so much to be 'wrong' (whatever that may mean), but to be out of step with the way his fellow traders interpret reality. Keynes, of course, wrote about this back in the 1930s, when he likened investment decisions to a competition run by a contemporary newspaper. In those politically-unaware days, a parade of pretty faces was shown in an evening paper, and a prize was offered to the readers who could pick not the most attractive face, but the one *chosen by the majority of readers as the most attractive*. A similar beauty parade of models is a crucial feature of investment and trading decisions.

The view in current finance thinking is radically different. When the possibility of different market views is entertained at all, it is mainly ascribed to different degrees of information among the rational agents. If many views are present among similarly informed traders, all but one (the 'true' one) are assumed to be 'irrational noise' that will wash out in the end. And, finally, when coordination of irrational views is allowed, the pseudo-arbitrageur enters as a *deus ex machina* to bring prices back to the rational fundamentals.[8] But in these days of balance sheet and capital constraints, of skittish risk managers and of worried regulators, the life of a pseudo-arbitrageur caught on the wrong side of the consensus interpretation of a market event may not be brutish and nasty, but it is certainly short. If it is true that the 'market can remain irrational far longer than [traders] can remain solvent',[9] then coordination behaviour becomes extremely important because it can produce positive feedback mechanisms, herding behaviour, cascades and, at times, even bubbles — the stuff that risk managers' nightmares are made of. Needless to say, these coordination phenomena are greatly exacerbated by the presence of leverage in the financial system. Coordination plus leverage can therefore produce cumulative price moves that are well in excess of what

would apply if every agent acted independently using their own 'model'. When this happens, the distribution of returns can easily become populated by the 'ten standard deviation' events at whose occurrence newspapers never cease to marvel.

What is the relevance of all of this from the perspective of the management of financial risk? If this author's analysis is correct, it is essential for the risk manager not to be wedded to any one view of the world, but to be able to see the same financial reality through the eyes of different market players. To go back to Keynes' beauty parade, the task of the risk manager is to gauge which of today's 'pretty faces' is in the running to be chosen by most readers tomorrow — or, to abandon the metaphor, to be aware of today's plausible interpretative models of reality, so as to be prepared when the one that prevails tomorrow asserts itself as the 'winner', causing the accompanying feedback-reinforced price moves. This approach is radically at odds with the prevailing view, which places great emphasis on *the* correct interpretative model of reality.

RISK MANAGEMENT AS THE NEW CHARTISM?

This view of market dynamics and of the role of the risk manager puts the study of the statistical properties of return distributions into a different light. Much too much has been written about whether the normal distribution provides an adequate description of financial returns (for many applications it clearly does not); and about which alternative statistical distributions may be up to the task. Leave aside the fact that, given a finite amount of relevant data, choosing among these alternatives is a statistically very difficult task (even the innocuous log-normal distribution can easily masquerade as a Levy distribution for a sample of the length typical in many financial applications). Even if it were possible to ascertain beyond reasonable statistical doubt that, say, a Levy distribution with a characteristic exponent of exactly 1.75 best described the data, *in itself* this would be of little practical use. Yes, with this piece of information one would admittedly know the probability of events (price changes) of all magnitudes. And, yes, the probabilities for extreme events would certainly be much higher than had the more 'tame' distribution been used. But, even if attainable, such an *unconditional* estimate (an estimate, that is, unrelated to *present* market conditions) is of very little use for practical risk management. Systematically buying insurance against extreme adverse events is far too expensive to be contemplated as an everyday strategy: as any trader knows well, the ongoing purchase of insurance against tail events will (more than) negate the advantage that one may expect to reap from entering into risky trades in the first place.

What is needed, again, are imperfect but plausible models of reality that indicate when it may be worth using the 'silver bullets' and spending money on protection, or when it is simply better to reduce the size of exposure. What a risk manager should provide are conditional assessments that today, given the market conditions, certain extreme events are more likely to occur than they normally would be. This is only possible for the risk manager who has at their disposal an imperfect but plausible structural model that links market information to outcomes.

These alternative models of reality can take different forms, and inform courses of action at different levels. Consider a few examples, starting from a macro perspective. The prevailing neoclassical model of reality is one of general market equilibrium based on rational expectations. Being less wedded to this view, and considering Keynesian or Minsky views as plausible alternatives, may have induced central bankers to take more robust action against the formation of speculative bubbles. At a lower level, questioning the efficient-market view according to which 'market implied' quantities always contain the best estimate of fundamentals could have raised some doubts in 2006 as to whether the financial sky was as blue as spreads and volatilities suggested. The 'decoupling story', the 'money glut' or 'savings glut' theories and the 'balance-sheet recession' view by Koo (2009)[10] all provide concrete examples of different ways of interpreting reality that can suggest different ways to manage the associated risks.

Of course, the very existence of competing models of reality indicates that no one will ever be able to predict with certainty when and where the next storm will hit. But risk managers should be aware of when it is more likely to do so, either because one plausible model is flashing red, or because several competing models all show a raised level of concern. A better understanding not of the statistical properties of the return distribution, but of the plausible market dynamics, can give a different urgency to the suggestion from a risk manager that *now* may be the time to spend that protection money or to cut the 'profitable' position.

This leads to the second point. Statistical analysis of past data, of course, has value. It can give indications about the likely losses that may be incurred once a week, a month, a year, or perhaps even once every few years. But when it comes to the real 'black swans', the events that can create havoc of systemic dimensions, this type of analysis is of very little use. In a way, blind reliance on a fitted loss distribution is the ultimate reduced-form model, sharing, in spades, all the limitations of such models. In general, reduced-form models have been abstracted from the detailed mechanisms that generate the phenomenon at hand, and rely on observed regularities to fit to the data the model's 'free parameters'. They are loved by engineers, who have to deal with challenging but well-defined and 'within-the-paradigm' problems. Unfortunately, when one uses reduced-form models (just because they bypass a specification of how the data are generated in the first place), one has no way of telling when yesterday's parameterisation ceases to be of use. The fitting of statistical distributions (no matter how sophisticated) to past data shares all the advantages and the unavoidable shortcomings of reduced-form models: absent an understanding of how the exceptional returns in the beautifully-fitted fat tail have been generated, there is no way of telling if yesterday's parameterisation is still valid today. And, when it matters, it invariably is not. Believing that the history of price changes 'speaks by itself' shares the same intellectual foundations and credibility as Chartism, and may be just as useful. Once again, imperfect, but plausible, structural models are infinitely more powerful in making financial decisions because they give a possible explanation of why people should be worried *today*.

To be of practical relevance, any suggestion regarding how to control financial risk must lend itself to being embedded in the organisational structure of a financial institution, of which risk management is only one part, and often far from the most important one. How do these suggestions fit with this reality? Admittedly, introducing competing models and the need for expert judgment in the risk analysis take 'objectivity' away from the process (even if this objectivity was never there in the first place). Formula-based approaches, for all their limitations, do have the positive features of allowing standardisation, reproducibility and comparability across firms. Even as they recognise the limitations of traditional quantitative measures of risk, neither private firms nor regulators can afford to do away with them. It is the balance between the quantitative and the subjective, however, that could be profitably altered. The current Basel II regulatory framework (with its two 'pillars' — one formula-based, and one qualitative) readily allows for this change in weights. This, of course, will require greater engagement, by regulators and by senior managers, with the more subjective aspects of risk management. This is certainly a very demanding process. But not everything becomes more difficult when one shifts the weight away from statistical measures of risk towards 'expert judgment'. Many senior decision makers of financial institutions have neither the technical ability nor the inclination to challenge the mathematical assumptions of a statistical model; but they are perfectly capable and, given their likely professional background, often naturally inclined to question the subjective reasoning process that leads to a well thought-out stress scenario. When it comes to influencing decisions and prompting action, the power of a 'story' should never be underestimated. And a 'plausible model of reality' is exactly that — a 'story' that connects a variety of visible and readily understandable inputs to more or less extreme outcomes. These inputs should be much richer than the statistical loss distributions or their percentiles, and, as an anonymous referee suggested, could include 'non-traditional risk signals such as valuations, flows, liquidity, spreads, etc'. The cognitive resonance of this approach may make a big difference when it comes to informing behaviour and, perhaps, deciding when to use the silver bullets of protection referred to above.

The regulators have become aware of the need to think outside the narrow statistical, model-based approach that they had until recently so strongly emphasised, and are engaging with the financial industry in a far-reaching rethink of their stress-testing and scenario-analysis framework.[11] Having had the privilege of personal involvement in this process, the author can confidently say that the regulators are genuinely open to new ways of looking at risk — including more subjective ones. As their remit requires, they will not dictate how individual firms will carry out these new tasks, but they will opine on the quality and acceptability of the outcomes. Here, one would suggest that, in their assessments of industry efforts, they may want to consider how the 'stressful outcomes' produced by the banks have been reached: have they been obtained by looking at financial reality via the narrow prism of a single 'best model'? Or have they taken into account the plurality of plausible models of reality and

the coordination among market players that this can generate? Just as importantly, they will want to consider how in practice these analyses fit in the organisational structure of the regulated institutions: who gets to see these results, what kind of challenge and questioning they generate, and what actions are prompted.

These suggestions are, of course, no panacea for the problems of risk management, and do not even begin to address some of the more intractable risk management problems in financial institutions, such as the difficulty in aligning the interests of the agents who manage the bank, of the shareholders who own it and of society at large. But perhaps the approach suggested here could make the financial community more conscious of the 'unknown unknowns'; more accepting of the fact that reality can be looked at from different, similarly plausible, competing perspectives; and, as a consequence, make the sector more cautious and somewhat humbler. This, in itself, would be no mean feat.

References

1 Greenspan, G. (2008) 'Congressional testimony', 23rd October.
2 Keynes, J. M. (1938) Letter of 4th July to R.F. Harrod, reprinted in 'The Collected Writings of John Maynard Keynes', Vol. 14, No. 173, edited by D. Moggridge, Macmillan, London, pp. 270–272.
3 Greenspan, A. (2009) 'We needed a better cushion against risk', *Financial Times*, 27th March, p. 11.
4 McKenzie, D. (2004) 'An Engine, Not a Camera', MIT Press, Cambridge, MA.
5 Interview with Thomas Sargent in Frydman R, Dyson (2005) 'Imperfect Knowledge Economics', PUP, Princeton, NJ.
6 Shefrin, H. (2005) 'A Behavioural Approach to Asset Pricing', Elsevier Academic Press, Burlington, MA and London, UK.
7 Rebonato, R. (2007) 'The Plight of the Fortune Tellers', Princeton University Press, Princeton, NJ.
8 Shleifer, A. (2000) 'Inefficient Markets — An Introduction to Behavioural Finance', Clarendon Lectures in Economics, Oxford University Press, Oxford.
9 Keynes, J. M. (1937) 'General Theory of Employment', *Quarterly Journal of Economics, Interest and Money*, quoted in 'Keynes the Spectator', available at: http://www.maynardkeynes.org/keynes-thespectator.html (accessed 26th January, 2010).
10 Koo, R. C. (2009) 'The Holy Grail of Macroeconomics — Lessons from Japan's Great Recession', Second Edition, revised and updated, John Wiley & Sons (Asia), Singapore.
11 Bank for International Settlements (2009) 'Principles for Sound Stress Testing Practices and Supervision', consultative paper.