IPE INVESTMENT & PENSIONS EUROPE

Quant surveyors

9 hours ago

As the recent AXA Rosenberg episode revealed, systematic quantitative investment promises the ideal in transparency and integrated risk-management – as long as managers' corporate governance is equally systematic. Martin Steward reports

What do we mean when we talk about 'risk management'? Generally, we mean constraints. Things that prevent us from doing what we want to do freely – because we recognise that what we want to do can go wrong.

Cars get us from A to B quickly. But that is risky, which is why cars are fitted with seatbelts and subject to speed-limit laws. But given that we have to honour these laws, and that cars are mechanical devices, wouldn't it be better if all cars had built-in mechanisms to slow themselves down if we try to break those laws, just as they have built-in seatbelts?

This is one of the arguments for computer-driven quantitative investment management. A discretionary stockpicker should be focused on stockpicking. It is the job of risk management to define his constraints – on concentration, exposures, volatility, potential capital loss, and so on. But those constraints are quantitative – which means that a computer-driven stockpicker can programme them in right from the start, like the car with the automatic braking mechanism. For 'quants', risk management isn't a constraint on the strategy – it is part of the strategy.

"For 20 years I've heard people tell me 'I'm just more comfortable seeing the whites of the eyes of my discretionary trader'," says Martin Lueck, director of research at trend-following managed futures specialist Aspect Capital. "But why? The beauty of a quantitative strategy is that it's explicable, repeatable, systematic – and therefore controllable." Matthew Sargaison, chief risk officer at another trend-follower, Man AHL, agrees: "It's clearly difficult to build a model that will at some point in the future veer off according to a researcher's 'hunch'. We are able to build risk management into the models themselves – while also having independent validation."

David Schofield, president of INTECH International, the quants specialist in Janus Capital Group, notes that his firm does not have anyone with "the specific title of chief risk officer": "That's partly because risk management is right at the heart of our investment process," he says.

We see the logical end of this thinking in certain high-frequency trading or benchmark-replication strategies that automate everything from stock selection to order execution. But automatic speed controls are one thing. Build in automatic steering as well and, sooner or later, drivers will start falling asleep at the wheel.

That is why Sargaison's point about independent validation is so important. Schofield explains that INTECH's Princeton-based research team, which builds its models, is complemented by a Florida-based team dedicated to analysis – making sure they work properly. And here is Louis Iglesias, chief compliance officer with PanAgora Asset Management: "We are able to code the constraints in the investment management agreement into our investment and compliance processes, but also, independently, into a compliance process which will check orders as they are being sent to brokers for execution then check the portfolio holdings after execution. But any breaches result in full escalation to senior management and also to investors."

These are not 'quants' issues; they are corporate governance issues – but those who argue that they do not apply more especially to quants are wrong. They do, precisely because quants put systematic risk management at the heart of their strategies.
"A formulaic approach can create the perception of safety when things are not safe," notes Antony Barker, managing director of investment advisory at JLT Benefit Solutions. "If everything was completely objective there'd be no need for governance structures at all: implicit in the importance of governance structures is the fact that no-one has an objectively perfect process."

Laboratory
There is no better illustration of this than the unfortunate events at AXA Rosenberg during 2009. After all, few know more about quantative investment risk management than the firm's founder, Barr Rosenberg, who set up Barra to develop quant tools in the 1970s, started managing money in quant strategies in the mid-80s and took that business to AXA Investment Managers (AXA IM) in 1999.

AXA Rosenberg was effectively a self-contained laboratory dedicated to sophisticated quant strategies, and it was during development of techniques to model rapidly emerging, transitory risk in June 2009 that a coding error was discovered in the part of its optimiser designed to control common-factor risk.

The bad code had been optimising clients' portfolios for almost two years, but AXA Rosenberg management was not informed until it had been corrected in mid-November 2009 – and, according to the Securities and Exchange Commission's administrative proceeding file, even that was when a Barr Rosenberg Research Center (BRRC) employee "felt compelled" to inform the group's CEO in defiance of a "senior official's directive not to disclose the error".

The SEC's file describes clear governance failings. Within BRRC "an informal but undisclosed 'micro group'" restricted full access to the model, even though no one in this group had compliance-related responsibilities. Those responsibilities lay with AXA Rosenberg's investment committee, which the SEC found "rarely convened". Unsurprisingly, when it did meet on 24 September to approve the code correction, the investment committee members were not informed – and appeared not to know – that this was a correction. As far as they were concerned, this was a routine update to the model.

Nine months later, the firm's Global CEO Stéphane Prunet and its vice-chairmen were having to write to clients informing them that Barr Rosenberg and director of research Tom Mead were resigning, having failed to "act in accordance with the firm's escalation policy and code of ethics". Global CIO Agustin Sevilla was moved to "a senior research role", also for breaching escalation policy. At the end of 2010, after shepherding AXA Rosenberg through the crisis quite admirably, Prunet stepped down to make way for a new CEO, Jeremy Baskin, hired from Northern Trust.

It had been an expensive episode: assets under management plunged (the Ohio School Retirement System and Florida Retirement System Pension Plan were among investors that lost confidence) and the firm agreed to pay $217m (€150m) in compensation and a $25m (€17m) civil penalty.

Some quants are surprised that it took so long for BRRC to discover this error. Michael Markov, the CEO of Markov Processes International (MPI) well-known for having pulled apart Bernie Madoff's return stream, blogged about running daily data from one of the firm's US large-cap funds through very simple tests for beta and tracking error in May 2010: he found tracking error "unusually high" through 2009 and beta well below 1.0 during the March-June market rally, and also substantial deviation from peers. Analysis against S&P500 sectors suggested significant overexposure to healthcare and cash and underexposure to financials. "Tests using very simple metrics show that this error could have been detected months before it was noticed by the industry," he argues.

That is probably unfair: even if one accepts that Markov had selected a reasonable peer group for the AXA Rosenberg funds, divergence from them only became egregious after March 2009 (as one might expect of a portfolio tilted systematically to low-risk stocks during a 'dash for trash'); and returns-based attribution is not always reliable during periods of exceptionally high volatility (a spokesperson for AXA Rosenberg funds has told IPE that its funds had no meaningful cash allocation, for example, which should raise questions about Markov's results). On the other hand, the SEC file notes that tests on the new risk model before roll-out in 2007 "did not detect the error". For all the difficulties of constructing ex-ante measures of risk for highly-volatile scenarios, one might have expected out-of-sample testing that included a 'dash-for-trash' scenario to have picked up anomalies.

This underlines the importance of independence – a second pair of eyes – which goes beyond the issue of coding errors and right to the heart of the robustness of quant models per se.

"The tools that so many firms use for oversight are based on the same tools and techniques that they have used to build their models," notes Markov. "Why would you go to the same doctor for a second opinion? Effectively you are getting a risk report from the portfolio manager sent up to the board, so if there is an error embedded in that from the start, this communication is flawed, the board ends up with a distorted picture and agents outside the company that rely on this reporting also get a distorted picture. You need independent techniques, tools and data. And of course it helps to have an independent body within the company to use them."

AXA Rosenberg was highly successful, agrees another risk software vendor, but it had one big flaw: "It was doing absolutely everything in-house. If you use risk models and optimisers that you've built yourself, you might well be the last person to find out when an issue comes up."

Any bug in an optimiser used by a thousand clients is likely to be picked up long before the same bug is discovered in an optimiser with just one user. Of course, there is a case for maintaining the 'edge' of proprietary technology, but it does not obviate the case for a second opinion.

"At the very least, quants should use a third-party product as a benchmark, so that if something is wrong it will be evident to them through comparison with the more widely-used software," suggests Sebastian Ceria, CEO at Axioma. "Quant strategies have lost assets over the past few years, and of course the common explanation for that has been poor performance. But I actually think it's more about risk management. Looking at 2007-09 you can see quant portfolios performing badly, all right, but you also see tracking errors of 8% on strategies that were supposedly completely risk-controlled. This is what hurt the most."

Software vendors will talk their game, but that should not blind us to the fact that it is common sense to deploy "protocols that may lie outside of the framework of the alpha model", as Andrew Lo and Mark Mueller put it in their useful 2010 paper on 'Physics Envy'. They also make the important point that popular criticisms of statistical models have given "too much weight to quants and too little to senior management" as causes of recent difficulties.

When quants and senior managers become too specialised in their roles, they argue, a potentially dangerous risk-management gap is created: "This gap usually emerges gradually and without notice, perhaps through the departure of quantitatively sophisticated senior management, the rapid promotion of business managers without such training, or the rising influence of inexperienced quants who happened to have enjoyed a period of profitable trading."

Or, one might add, with the acquisition of a specialist quants group by a non-specialist. (See panel, Integrated circuit). It is easy to see how (going back to the built-in risk-management idea again) the acquirer's senior management might feel unqualified or not required to exercise oversight. It is notable that AXA Rosenberg's chairman, Dominique Carrel-Billiard, made much of Baskin's deep quant experience when announcing his appointment. The entire research function now reports direct to Baskin, backed up by a new chief risk officer.

Academic
This brings the firm into line with industry standards on independent risk management (indeed, its engagement of an external compliance consultant goes well beyond them). But 'standards' is perhaps the wrong word here. The fact is that there are different ideas about how to balance the ideals of independence and collegiate spirit and, again, this seems peculiar to the rather 'academic' quants context.

Pierre-Yves Moix, chief risk officer in Man Investments' multi-manager business, is in no doubt about the ideal structure. "When we look at quants we put a lot of emphasis on the research process," he says. "Is there a split between the people doing the research, the people calibrating and backtesting the model and the people implementing the model? That's very important: human nature means the research team will be essentially optimistic about the quality of the ideas and the model."

But if that is the ideal, other structures are presented as virtues, not compromises. Thirty of PanAgora's 70-plus employees are investment professionals involved in research, says Iglesias. "They hire people like themselves, so we really have a holistic approach. We don't have a research department – we are researchers." COO Michael Turpin contrasts firms with "very academic but not particularly hands-on" research teams feeding less well-compensated portfolio construction groups who do what the black box says and carry no responsibility for performance. "It's that kind of structure in which responsibility gets lost," he argues. "Here, the same group is responsible for research, performance and interacting with clients."

Jean-Charles Bertrand, CIO for quantitative strategies at HSBC Global Asset Management, makes a similar point. At weekly investment meetings portfolio managers' first responsibility is to understand and assess their models' output signals and, if necessary, exercise discretion in implementation.

"Those portfolio managers have strong quants backgrounds and most have been researchers in the past," he observes. "In practice, that means the fund management team monitors the models that researchers have built and generates all of the backtesting." But, he adds, that role is made more efficient by close collaboration between researchers and fund managers at design stage – precisely to ensure models are based on sound investment hypotheses, rather than blind
data mining.

That seems a far cry from the separation of roles described by Moix – one group to build the models, another to test them to breaking point. But we find a similar culture among the modern managed futures fraternity.

"There are different ways of [ensuring independence], depending on the size of the business," says Sargaison at Man AHL. "With more than 82 professional research staff sporting very strong quant backgrounds, AHL is large enough not to have a situation where a single model is being developed by a small research team and then being tested to destruction by a small, independent risk team."

Just like on a large-scale physical science research project, different groups work on different elements of the overall model, and research papers are treated as they might be in the academic world: "They are presented to other quants who have nothing to do with that original idea and whose interests may be in a completely different asset class or part of the model," Sargaison explains. "Those peer reviews then go to a committee which includes myself and AHL's chief scientist. That peer review process helps us to assess the original hypothesis and the work done on it." The quants industry gets tarred with the 'black box' label, he concedes: "But obviously within AHL the box is completely transparent."

Compare the practice at Aspect Capital. "There is an obvious tension between the need for scientific rigour and commercial sanity," notes board director John Wareham. But he identifies 'scientific rigour' not with separation of roles but with peer review among fellow model-builders, and 'commercial sanity' with the need to protect intellectual property. He describes a continuing management review of research, a "periodic, mid-cycle peer review of all the research being done", and extensive pre-implementation testing of individual pieces of research by risk management, portfolio risk and other teams.

"The consequence of this process is that quite a lot of people in the firm know quite a lot about the code, but we feel that's necessary if we are to maintain a collaborative effort to evolve our models," he says. "We feel that's much better than having each individual working on one tiny sub-component of the code such that they couldn't leave with a tremendous amount of proprietary information."

Which is not to say that there is no risk management function separate from the research function. There clearly is, and it reports direct to Aspect's CEO rather than to the CIO or head of research. "We also have a separate and dedicated team of quantitative developers who re-code all the models," Wareham notes. Likewise, Sargaison concedes that when things do occasionally break down – not usually coding errors, but perhaps particular data streams that turn out not be as clean as researchers anticipated – they are flagged up to Man Group-level compliance and risk assurance committees for analysis independent of AHL.

But it appears difficult to establish watertight protocols about what matters should be escalated to which level. Within quantitative strategies at HSBC GAM, potential errors are immediately escalated to the top management of HSBC France and HSBC Global Asset Management and reported to the operational risk and internal control committees. But asked if he could spell out the protocols, Bertrand said no: "We had internal discussions when the AXA Rosenberg issue came to light, and our conclusion was that, given the intuitive models we use and the strong collaboration and full access between research and fund management, there was no need to add
any
new procedure."

Again, the virtue – not the risk – of the lack of separation of roles. Like Moix, Antony Barker, managing director of investment advisory at JLT Benefit Solutions, thinks such a separation is important. "But realistically, in the small organisations where the more exotic quants are performed, this independent function is somebody's job for Friday afternoon, once they've finished their day job, and even in the biggest financial houses there is not one team of quants creating and another team of quants checking," he says.

This is picked up by Lo and Mueller in 'Physics Envy'. Can it be right that risk management is ultimately the responsibility of profits-focused CEOs and CIOs, "where risk is viewed primarily as a constraint", they ask; or that chief risk officers, like CEOs and CIOs, get compensation linked to profits? Shouldn't CROs report direct to the board and receive compensation dependant on corporate stability rather than profits?
No-one wants to go down the route of differentiated compensation. Everyone cites the need to have employees 'pulling in the same direction' – in other words, to lift the constraints. Governance structures at AHL ensure clear independence, as Sargaison puts it, but no one is incentivised to knock down colleagues' research. "Criticism is based on real perceptions of the quality of the work presented and people are hired with the requisite expertise, judgement and indeed character to analyse other people's work appropriately."

This takes us back to the point made by Turpin at PanAgora: the important thing is that employees making risk management decisions are compensated to a level commensurate with that responsibility, and feel they can escalate awkward questions. And in the cold, commercial world that means employees engaged in making money for the firm day-to-day.

Systematic
In the end, then, front-office pay for a middle-office reporting function does seem to be the genuine logic, in quants, to having the engineers monitor their own creations. That leaves investors reliant on robust escalation procedures being applied as advertised. After all, as the SEC file makes clear, AXA Rosenberg "misrepresented to clients that internal controls processes and procedures covered BRRC when, in fact, certain of these controls [….] were not implemented." Turpin observes how difficult it is to assess the internal workings of an investment organisation if you are not an insider. "It becomes relational," he says. "Does the commentary and description you get of processes relate well to the personalities and culture that you pick up on?"

But perhaps investors that mandate quant specialists to manage their money should recognise other, particular responsibilities. Just as a willingness of quants to subject their models to tests by third-party risk-management tools can help achieve a more robust research environment, so investors should be ready to deploy their own tools, to be the third pair of eyes.

Lo and Mueller suggest that one lesson of the financial crisis is that "we need more financial-engineering expertise and training at the most senior levels of management at banks, broker-dealers, insurance companies, mutual funds, pension funds, and regulatory agencies, not less". As Moix puts it, you can get good information from most quant managers, but as a manager selector "we do find it valuable to have quantitative analysts talking with quantitative managers." Markov, who showed how easy it could be to pick up anomalies in AXA Rosenberg's performance, adds: "The lesson is not so much about governance within quants firms as it is about the controls at pension funds and their consultants – who all missed these obvious red flags. Are they over-reliant on the managers' own tools and analysis?"

As he observed in his 2010 blogs, daily data is crucial if an investor wants to carry out truly effective performance monitoring. But that is not an unreasonable demand, and once they have that data it is arguably a lot more useful for analysis in the quant world than it would be for analysing qualitative stockpickers. Remember Lueck's words – "explicable, repeatable, systematic"? If the numbers start failing to add up then either something has gone wrong or the underlying hypothesis was flawed in the first place. With quants, there should be no place to hide.

This happens already. As PanAgora's chief compliance officer Iglesias observes, the most sophisticated institutional investors have their own very large compliance teams looking at the very same benchmarks and constraints as PanAgora: "That is another check and balance in place outside PanAgora itself." Industry analysts and consultants often demand very large amounts of data for their own analysis, concurs Schofield at INTECH: "To that extent there have always been outside eyes looking at the way the model behaves." Indeed, one of the more striking findings of the SEC file on AXA Rosenberg is that, long before Markov's after-the-event analysis, there were clients "voicing substantial concerns" about anomalous performance even before the coding error was discovered by BRRC.

Evidently, quants are not 'black boxes'. They are inherently transparent, and they grow more transparent the more they are interrogated by independent analysts using a diversity of tools and techniques. Systematic investment strategies do present particular risks; but the appropriate response is to uphold equally systematic corporate governance, and for investors to apply equally systematic monitoring of their mandates.

Author: Martin Steward