Surveillance Deep Dive:

A TIME TO BE BOLD

FORECAST 2023

smarsh | 1LoD

- Regulatory concerns still impede innovation in the alert triage process
- Tweaking calibration loses valuable data from the surveillance processes
- Machine learning (ML) should be used to give priority to handling alerts and closing out low-risk alerts
- Conservatism within banks hinders improvements in surveillance efficiency and effectiveness
- Model risk management is hard to justify for surveillance models
- A cautious approach to trader profiling hinders progress at most banks, apart from the industry leaders

Non-financial risk managers say they are often torn between meeting regulatory requirements and detecting, mitigating, or preventing genuine risks. And they may be forced to choose between innovation (in both technology and methodological sophistication) and the status quo, or business as usual. 1LoD's Surveillance Leaders' Network in June delved into these opposing requirements, asking whether existing surveillance mechanisms are up to scratch, why model risk governance has any place in non-financial risk management, and why so many banks refuse to go the whole way with trader profiling.

Is surveillance being held back by outdated regulatory requirements or by banks' innate conservatism? Take the suspicious transaction/order reports (STOR) process.

Surveillance leaders look at how well their processes generate STORs (or their equivalents, if they are outside the UK). Most alerts do not result in STORs, so by that measure, the process simply generates noise. "But we have the process because that's what the regulators want", one head of surveillance explained. It also generates noise in terms of broader risk mitigation because the alerts do not generally result in the discovery of material risks worthy of further investigation.

Banks have typically tried to solve this problem by adjusting their calibration parameters to best balance the trade-off between noise generation and missing true positives. This reduces the number of alerts, but all alerts generated are then sent for human review.

The alternative is to worry less about calibration, accept that most alerts will be false positives, and use technology to rank the alerts according to risk: low-risk alerts can be handled using processes which are driven by robotic process automation (RPA) or machine learning (ML) for analysis and probably closure, while the higher risk alerts are escalated to humans for review.

"It's about casting a wide net and ensuring that when analysts do look at them, they're focusing their efforts on those that are worthwhile and having an audit trail that explains the exact process that was followed for each alert," one attendee said.

"If you trim your thresholds and your procedures to make sure you are receiving a low number of alerts, you run the risk of you missing outliers or missing cumulative historical activity that can amount to significant malfeasance. It is better to keep the thresholds broad but use machine learning or robotic process automation to make sure that when you do look at alerts, they are high risk and high impact. Low-quality alerts are closed out by the system, having been analysed according to business rules that the organisation is happy using, but are always available in the system along with an exact audit trail for how and why those low risk ranking alerts have been treated in that way."

The problem is that most banks operate surveillance (and other non-financial risk) processes on an assumption that regulators will not accept imperfect machine decision-making over even more imperfect human reviews.

They don't put it quite like that, but as one explained, "I'd be quite uncomfortable telling a regulator I generated the alert, and it was not reviewed. I'd be more comfortable saying we set up this parameter and we have a defence for why we've done that. Generating an alert and then saying, 'no human actually looked at it and it was closed automatically', I think would be a much more difficult explanation [to make to a regulator]."

Given that the calibration of alerts is periodic — and so drifts out of sync with market conditions quickly and is subject to the same kinds of potential error as the types of explainable algorithm that would close out an alert — this sounds more akin to a Luddite argument than a logical objection to the use of technology.

Attendees also took issue with regulators in other areas. "Take lexicons. That is definitely not the most efficient way of doing things. But we can't stop doing that because the regulators don't understand the alternatives. So, until they change their mindset, we're stuck doing what we're doing. It's very depressing," said one participant.

But the slow regulatory response to new technology is not the only reason for the lack of progress. Regulators do not accept that they are to blame, and some participants at the Leaders' Network argued that the banks' innate conservatism is the more significant problem.

The model risk rabbit hole

First, banks often seem to go well beyond regulatory prescription, as one attendee pointed out: "There is an issue with banks not distinguishing properly between primary legislation and guidance around taxonomies and scenarios — things like the ESMA (European Securities and Markets Authority) guidelines. There's massive conservatism now about people daring to depart from what are quite outdated guidelines."

Perhaps the best example of how bank bureaucracy can derail improvements in surveillance is the imposition of model risk management/governance (MRM) protocols on surveillance models. More than half of those present said that their rules-based surveillance scenarios and calibrations are now subject to MRM, while an even larger proportion said that they still strongly believed that surveillance algorithms are not models as defined by regulation.

"I think we have models. The question is, are they the models that the regulators set out to control? No, they're not. They're not financial. They're not used to make financial decisions. They are fundamentally different in operation and materiality to the models that control credit risk, liquidity risk, curve and position risks, or algorithmic trading models that actually take positions and execute things in the market," said one attendee. "But we have lost that battle so my advice would be just to be very, very nice to your MRM guys."

Frustration over just how little the MRM teams understand about non-financial risk was also evident. "I've never come across a more blinkered set of individuals in my life than the quants we employ in MRM," said one attendee. "I've had shouting matches on the phone in the office with these people because I'm just so astounded by what they're saying.

For example, 'if you can't demonstrate definitively that the model gets rid of all the risk then why not just have humans do it? Prove to me that just having humans doing this wouldn't be more effective than your current process.'

And they don't understand non-financial risk nor that this model — say a rules-based part of SMARTs — is one part of a huge surveillance infrastructure that includes audio, messaging, e-comms, collaboration tools, all of which are noisy pieces of evidence in a forensic process, not a trading-type process."

Other banks complained that once surveillance models are included in MRM, they are subject to multiple reviews whenever calibrations or scenarios are changed or updated, with MRM teams assuming that those changes are as significant as changes to a trading system.

The attention on this area has created a feedback problem in terms of regulators. One participant revealed that in talks with the National Futures Association, the regulator was, "aggressively looking for anything that could constitute a model,… [and] policies have been expanded to give them greater flexibility to bring things [like rules-based surveillance alert algorithms] into scope."

In the UK, the Prudential Regulation Authority too has signalled an increased interest in model risk, but in the context of material risk to financial institutions and the system. It is not clear why surveillance models would be deemed that significant.

Asked whether model risk management teams added value to the surveillance process, most attendees said no. "I think there is definitely no value-add from it. It just thinks it provides a degree of comfort up the hierarchy of the organisation point of view of reassuring that someone has independently judged the tool."

The only defence against MRM creep seems to be partial waivers or dispensations for low-risk situations within surveillance. However, this raises the question of why surveillance models could not be viewed as low risk and entirely exempted. As several bankers explained, "that battle has now been lost.

Too hesitant on trader profiling

The other main area where banks are too conservative — thus holding back developments which could transform the efficiency and effectiveness of surveillance — is trader profiling.

This term means different things to different banks. For some, it means using a focused analysis of traders' mandates, what they trade, who they trade with, in what sizes, P&L movements and other data closely related to trading such as analysis of communications patterns, to build a 'normal' avatar for each trader. This avatar or digital twin of each trader is then used as the baseline against which to measure the behaviour of the human trader.

Depending on how the divergences are used, this could be viewed as a new class of behavioural surveillance, generating its own standalone alerts which need to be reviewed. Or the outlier data can be used as an additional input into a traditional TS engine. Or it can be used as additional context in the escalation and investigation process. As data on divergent behaviours builds up, and the significance of those behaviours is investigated, it would be possible to use it to assign risk scores or employee risk ratings to traders. But this is not the initial objective of these kinds of programmes.

The other way that banks think about trader profiling involves taking data from broader conduct programmes, HR data, compliance training data, and more, as well as trade-related data. Programmes run like this are much more likely to be used to explicitly assign risk scores to individuals — scores that persist beyond, say, just the one-time gathering of behavioural context into an investigation initially triggered by traditional TS.

More than a third of the participants said that they used some form of trader profiling. Most used narrow trade-related versions, and even then, only for particular risk types. For example, communications network analysis was used in some banks' insider dealing surveillance. But the majority were uncomfortable with explicit employee risk rankings, whether narrow or broad, for reasons that seem to boil down to internal conservatism and assumptions about regulatory blowback.

These reservations largely concerned ethical and HR issues. Several attendees felt that the compliance functions and the business would not agree over whether to allow the explicit ranking of traders according to risk, and that it would be difficult to justify such ranking in the absence of specific regulatory demands. Others pointed to specific issues in Europe regarding data protection and workers' councils — although crime prevention trumps these concerns in the major jurisdictions. And others worried that as soon as banks put a system of any kind in place, it becomes a matter for regulatory scrutiny and criticism.

"The regulator is going to want to see everything that you are doing in terms of that process, and they will want to see that it is airtight," said one attendee. "Particularly if you close out an alert around a particular trader based on that risk score then, because you are now targeting the trader and not the trade, you're potentially running regulatory risk."

This fear is based on the argument that using trader profiling as an input into the TS alerting process creates an additional step in the model that then falls under both regulatory and internal scrutiny. Banks which take this approach also argue that explicit trader profiling and risk scoring don't add significantly to the surveillance process anyway. In their view, because you can bring all the trader-related context into the escalation process, there is little to be gained from creating an entirely new surveillance type with its attendant issues.

Splitting the pack

These divergences of opinion reflect a growing split between the largest and most sophisticated institutions and the rest. Several of the largest organisations have well-advanced trader profiling and risk-scoring systems in place. They see these as part of a sophisticated risk-based approach to surveillance and focus on the practicalities of operating the processes and related governance.

As one participant explained, "We do both [narrow profiling against an individual's 'normal' and risk scoring based on a broader set of data than just trading-related data]. We see it as part of supervision – surveillance can be seen as a key part of supervision – and it is an important part of how we demonstrate that we are supervising staff properly."

A core use of this kind of profiling is to prioritise alerts for human review. The process generates alerts which are used as data points in the TS alerting process; it generates employee risk scores which can suggest that TS alerts involving high-risk individuals should be investigated first or more thoroughly.

As for the regulators, these banks say the regulators are keen to move in this direction and support their efforts. If this is the case, then other banks may find themselves with no choice but to adopt these methods. After all, this could become the new definition of best practice.

This information was taken from the Surveillance Leaders' Network event in London, 15 June 2023.