

All the *unusual* suspects

Once upon a time, town dwellers were familiar with the nightly callings of the night-watchman – “*Three o’clock and all’s well!*” It may have woken you up but at least it represented a comforting feeling that there was someone keeping an eye out for suspicious and unusual goings on in the middle of the night. However, with the passage of time and growing urbanisation, the individual night-watchman was not up to keeping watch over sprawling conurbations. It’s all too big and too busy. It’s too hard to know who should be about and who is up to something unusual.



So for the modern day urban dweller the equivalents to the night watchman are automated alarm systems. However, automation can bring problems of its’ own. Some may be familiar with this situation - lying in bed at night, trying to get to sleep, but in the distance is the background noise of alarms on cars or buildings ringing on and on. No-one reacts to them anymore because they are always ringing. Everyone assumes they’ve been set off by the wind, a stray animal or that it’s just an oversensitive system. They assume that they are false alarms – just noise.

What we know

Monitoring and surveillance of financial transactions and trading has been growing inexorably for the last decade, especially since the advent of the Market Abuse Directive in Europe in 2003.

As trading has been automated and volumes have increased dramatically over the period, manual oversight and monitoring of what is going on hasn’t been able to keep up. Just like the night-watchman, it has required automated help. In fact, most regulators and supervisors now insist that any significant market participant installs some form of automated surveillance system to monitor their own trading. In turn, this has led to the rise of several, different types of trade monitoring and surveillance systems. The common factor between them is that they all tend to be based on the concept of rules. These were very often drawn from the foundations of specific exchange rules and added to as new legislation was brought in to prohibit and restrict certain activities. Rules might cover anything from minimum and maximum tick sizes, order to trade ratios, large trades, pricing away from the market best bid and offer prices, sequences of orders on an order book and so on. Essentially, all these rules are based on some quite specific measures of things we know and can easily observe in the trade data itself. You could call them the ‘known knowns’.

Calibrating for what we know

But whilst rules can be specific about factors such as time, volume, value and so on, they tend to be non-discriminatory, that is they don't tend to differentiate between individual traders, customers or clients – they are non-specific and general in their approach to factors such as who the participants actually are.

Often rules will include some parameters and calibrations to control when they trigger and raise an alarm. However, despite this, one of the biggest practical problems for the surveillance analyst today, is that their monitoring systems may produce too many false alarms that waste time and effort being investigated before being dismissed. Anecdotally, there have been examples where some rule types are known to produce hundreds of such false positives and end up being ignored or even switched off.

One of the causes of this type of problem is that the calibrations can be set to the wrong absolute levels or, even if they have dynamic parameters, can be over sensitive. And it is not always a quick and easy thing to fix with existing systems. Another cause of false alarms may be that the system is unable to distinguish between what is acceptable for one party but not for another. Many types of alerts can be triggered by certain parties, not as a result of any wrong-doing on their part, but purely as an accidental result of their usual activities – prop trading desks, HFT engines, market makers are all frequently the cause of numbers of spurious alerts. Rather like our urban dwellers' problem of burglar and car alarms ringing endlessly in the night, they can create 'noise', large numbers of false positives which distract the monitoring analysts from more productive work, focusing on what is truly unusual and suspicious.

Unknown, unknowns

If the first generation of surveillance systems sought simply to automate monitoring, attention is now turning towards how surveillance solutions can be improved and made more efficient and effective. An increasing amount of interest has been shown in the potential to go beyond the current fixed rules-based systems - the 'known knowns'.

One thread of the discussion touches upon advanced theories for self-learning algorithms. Although certainly an interesting theoretical topic, today we are still some way away from being able to deploy practical monitoring solutions based on this.

Recent regulatory initiatives have focussed on extending the scope of monitoring beyond trade data to voice communications, especially using mobile phones. It should be said however, that the practical relevance and success of this type of monitoring is widely questioned in many quarters.

Other thoughts have turned to using more specialist monitoring systems previously deployed in the intelligence sector and more recently as fraud detection systems in retail finance and insurance settings with some success. Typically, such systems can sift through very large volumes of data from many different sources looking for previously unknown connections – to paraphrase US Secretary of Defense Donald Rumsfeld's well-known phrase – the 'unknown unknowns'.

The approach is certainly different from a pure rules-based approach and certainly a tempting thought for the curious – the equivalent to an enormous join the dots puzzle, perhaps. However, it does depend on having as large a dataset as possible, showing as much of the market activity and background of all the participants. In reality, this may be a possibility for regulators and supervisors, who may have sufficient oversight of the goings on of whole markets and even through collaboration, across borders, but less feasible for the market participant who only has sight of their own activities. And it also depends on having the resources to follow up a potentially large number of ‘interesting’ possible links which may or may not be directly relevant.

As illustrated below, successful Surveillance and monitoring is certainly about being able to ‘join the dots’ between perpetrators, their strategies and behaviours and the symptoms that can be observed in the available data.



However, trying to ‘boil an ocean’ of data looking for any possible links or connections can be costly and time consuming, without necessarily addressing the more mundane but relevant questions of false positives and monitoring efficiency. Perhaps the answer to how improve current monitoring lies elsewhere?

Why, how, what, where, when and who?

Avoiding the pitfalls of being seduced simply by clever shiny new technology, it is worth taking a step back and considering some fundamentals about surveillance. Probably the 6 best and most obvious questions in trying to examine anything are why, how, why, what, where, when and who.

In terms of market surveillance, the why question is generally the easiest one to answer – the reason that it’s undertaken at all is for profit or avoidance of loss.

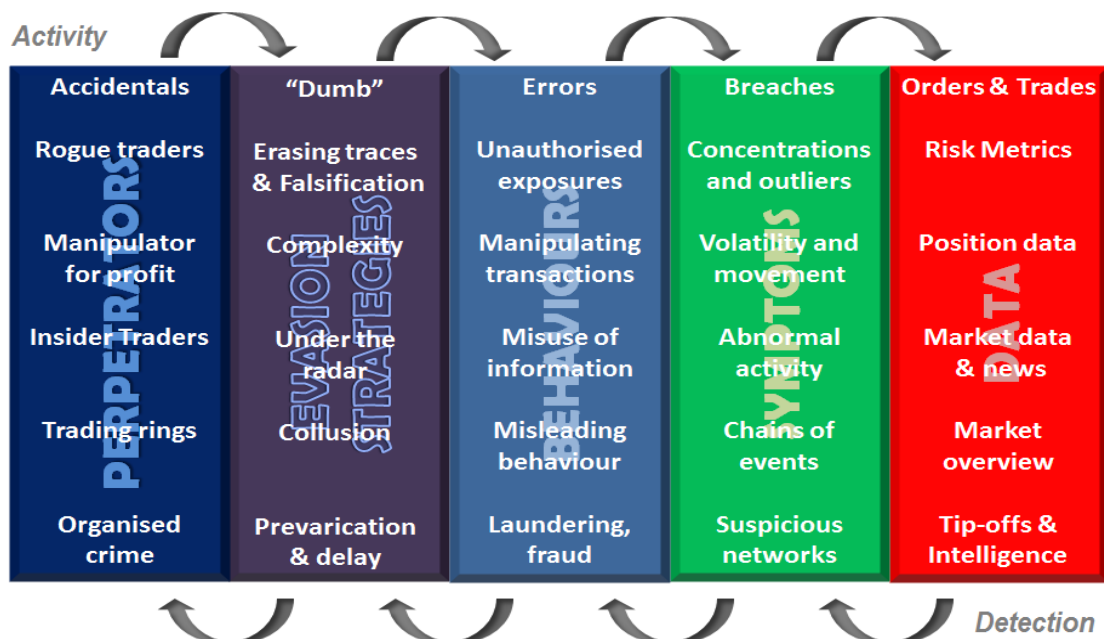
The How question is the modus operandi, the behaviours and strategies being used by the potential perpetrator. The more we understand this, the better we become at structuring our surveillance systems to spot the symptoms that can alert us to suspicious events.

What, where and when are the basic elements of trading data – what instrument, quantity and price, market, executed time and date. These are the basic and essential data elements for any trading or monitoring system.

Which leaves us with the question of who? Modern surveillance systems now tend to allow for more accurate calibration that can be specifically tuned to the actual nature of trading and the markets that are particular to a given firm. They may even allow for some discrimination between parties, clients v internal trading desks for example. But they don't allow for unique calibration for each and every party – that would become unnecessarily complex. Furthermore, market volatility and the dynamic nature of businesses mean that things are constantly moving and changing. Even a system that was well calibrated when it was installed can become out of date. So this only goes part way to addressing the issue of the rules being non-specific and general about the parties involved. What maybe normal and acceptable trading behaviour from one trading party may be symptomatic of abnormal, unusual and suspicious in another. For example, a sudden change in trading patterns and volumes or a move into trading different instruments can be a very useful indicator that the event requires further investigation. Better still, there may be market announcements or significant internal events that can be linked to this indicator.

Taking a leaf out of the Risk Management book, it could be very useful to have an understanding of what 'normal' trading behaviour is for a specific party, which can then be used to inform and improve the surveillance rules that are applied to each.

Limiting factors on the effectiveness of a surveillance system include the cost and practicalities of what of accessing the necessary data. Boiling an ocean of data can be prohibitively expensive and impractical. Conversely, too little data or data of poor quality can severely compromise the effectiveness of monitoring. But there is an optimum and practical balance to be found. The illustration below shows some of the data types that can be relevant in surveillance models.





Order and trade data are already available in most surveillance systems. Fundamental reference data about instruments and parties is also a given. Position and risk data exists in most firms and could be brought into the mix, if desired. This data, built up over time, can provide a good basis to determine a historic view of the trading activities of a party. So we can begin to move from the non-specific to party specific and build and maintain a dynamic picture of what these parties 'normally' do and so determine whether the behaviours being looked at are really usual, normal trading behaviours, similar to peer groups or unusual and abnormal. In other words, we can begin to usefully include the 'who' factor to reduce false positives and create much more targeted, relevant and specific alerts. This capability is generally known as 'Profiling' and is of increasing interest to many large market participants for its potential application. It seems to offer a more practical and logical next step in the development of Surveillance and Monitoring solutions, as well as for further Business Intelligence uses.