


The Shock of the News

Machines can read the news now, but can they be trusted to act on it? There's been a lot of talk about machine-readable news since our Q1 2008 feature, but does that go any further than adding bells, whistles and meta-tags to the text streaming across the bottom of the screen? Spurred on by recent events, William Essex has returned to the search for a genuinely machine-usable news solution.*

There's a lot of news about these days, and much of it has the capacity to impact trading outcomes. The frustrating part, though, is that the line from event to impact is neither predictable nor necessarily direct. There are green-amber-red statistics on US mortgage arrears/defaults, and the tale has been told of how the rating agencies heralded the downturn with sweeping downgrades just as soon as they saw the lights change. But would you, today, programme a causal link between those mortgage numbers and, let's say, the banking sector in Iceland?

The global financial sector is not going to make this self-same mistake next time. But for the rest of us, the more useful lesson (apart from "Nobody knows anything", page 84, and "Thomas Jefferson was right", page 18) is that all news is potentially dangerous. It seems obvious to say that news is a factor that might usefully be taken into account in any algorithmic or automated trading strategy. But the further you stretch your definition of news, the less obvious it becomes.

The machine-readable news business has tended to define news as directly relevant data that can be digitised into numerical form. This is not a criticism. If you're trading an asset class, you need the data relevant to that asset class. Stephen Mitchell, CEO of Weather Insight, says: "At the end of the day, trading is just an information arbitrage. You get the information faster than the next guy, you win. Second place doesn't pay." True enough. But there is a trend, accelerated by recent events, towards a "layered" approach to news. ►



What is news?

News might best be defined as “inputs capable of having an impact”. But that isn’t half as useful as drawing up a pyramid of layers in which predictable news is the bottom and the shocks are at the top. If you trade on scheduled, predictable data, you’re down there at the bottom. If your algorithm grunts and rolls over every time the words “Breaking News” scroll across your screen, you’re off the top.

Optimum news sensitivity is the as-yet-unachievable situation where a trading system picks up and (re)acts appropriately to any “input capable of having an impact”, but ignores everything else. Such a situation may be unachievable, but it is useful as a target/benchmark.

The foundation layer is the immediate trading data. Then there’s the tagged news about the company, asset class, market, that can also be readily digitised. Then there’s the economic news, then, then, then.

Then finally there’s the weird stuff. Example. You’ve set your machine up to read all the financial news about an apparently friendly merger, and react appropriately. The numbers



Alan Slomowitz

look good, the share prices are behaving predictably, the whole thing is beginning to look really very boring indeed. But you’ve also tagged the social news. So your machine picks up, let’s say, the gossip-column story about the two CEOs’ spouses punching each other at a charity dinner (and, reacting appropriately, it makes an excited squealing noise). You start watching the two CEOs’ body language. They’re smiling, but after that, can they ever work together on this deal?

At the top layer, up in the blue sky, is the marginally relevant detail (we’re looking at a pyramid shape here) that might prompt you to up the weighting of any negative news. Everybody has the basic stuff and most of the intermediate stuff, but not everybody knows that those two CEOs go home every night to a tirade of, “Drop this deal or I’ll divorce you.” Back in the Q1 2008 feature, Armando Gonzalez, President and CEO of RavenPack International, was quoted as saying that



some models “lacked a sense of awareness of the world beyond fundamentals and technical analysis”. It’s time to take a step back and look at that world.

There are two catch-up plays with machine-readable news. The first is that people are getting those basic layers in an actionable form with low latency. If news itself is dangerous, so is not getting the news as fast as the other guys. The second catch-up is: news consumers are moving “up the pyramid”. This is not a progression from relevant to irrelevant, but from consumption of the news that everybody knows, to a

more user-defined, thus customised, flexible, perhaps idiosyncratic, perhaps creative mixing-in of details that might make the difference between moving along with the market, and being first to call the change.

It's that playing field again

So how's it done? Start in the time-honoured way by cutting latency. Clint Rhea, COO of Need To Know News (NTKN), says: "Clients connect to our servers at our colocation facilities in major financial centers, and receive messages via a proprietary protocol that is compact and easy to parse. To get the lowest latency, many clients put their hardware in the same facility and cross-connect to our systems." All of which will no doubt sound familiar to any data user. This is levelling the playing field.

Next, switch on. Rhea says: "NTKN specializes in providing market-moving economic data from government press rooms and private organizations. We provide this data as a numerical feed to clients." Numerical feeds are intrinsic to trading, of course, and they're easily tagged. Similarly, Alan Slomowitz, Director, Algorithmic and Trading Products, Dow Jones Content Technology Solutions, says: "All the news we produce is in a feed that is machine-readable in terms of its tagging. The company symbol is tagged, as is the sector, the industry, et cetera. People can do additional analysis of the actual stories and headlines." So this is where we get to talk about tagging. Rhea adds: "Our reporters also write news stories on market developments throughout the day, and we distribute this via a desktop client and as a textual news feed."



News isn't like a butterfly

The "butterfly effect" is both central and peripheral to the development of machine-readable news. It is the chaos-theory notion that a small initial variance may lead to significant and unpredictable changes in outcome. A home-owner in Michigan, for example, defaults on his mortgage. The entire world banking system collapses. Or, in the conventional exposition, a butterfly moves its wing on one side of the world, and the cumulative atmospheric change is sufficient to alter the course of a tornado on the other.

But note this. The butterfly does not cause the tornado. There's already a tornado building up in the atmosphere. If that farmer, months before his default, had come home boasting that he'd got himself a huge mortgage despite having no income and no assets, and if all his friends had gone and got themselves mortgages, and if that micro-mortgage boom had got itself into the local paper – no, you're right. We wouldn't have picked it up. And even if we had, we couldn't have used it to predict the wider outcome.

For the wing-movement that redirects the tornado, there has to be a butterfly. For the default, there has to be a farmer. But the butterfly effect is still peripheral in that there is no direct cause and effect: you couldn't have found the farmer and predicted the collapse of Lehman Brothers. The "message" of the butterfly effect is: if you think you've found a causal event, don't prejudge its effect. There may be another butterfly out there.

Admittedly, Rhea continues: "Automated traders are more interested in the numerical feed," but there's a trading edge here, if we can find it, and tagging looks like a way forward. Being first to find an effective methodology for "fuzzy tagging" (or "flexible tagging" or even "programming an opinion into an algorithm", to take a few phrases from the blue-sky end of the market) will feel a lot like being first with the news. Tagging is a largely qualitative activity. Even factual news (earnings figures, etc.) requires a qualitative

The Shock of the News

judgement as to its likely impact. The question is how to render such concepts as “good news”, “bad news”, “disappointing results”, “downturn”, “total global meltdown”, “encouraging figures” in a form that is not only machine-readable but also, ideally, machine-usable. And idiot-proof. If you tag “disappointing”, be ready for “After last year’s disappointing results, this year saw a dramatic turnaround.” Then you get to expend brainpower on linking those tags to some kind of system for determining likely impact.

In Q1 2008, we discussed sentiment analysis at length, and highlighted the other main driver of machine-readable news evolution: simple archiving. There is enough old news in storage, along with an extensive lexicon of the words CEOs and analysts typically use, to enable “good”, “bad” and “aargh!” to be assigned a usable value. Discussing this from his own company’s perspective, Don Williams, Managing Director and Senior Vice President of Sales at RavenPack, says: “RavenPack Analytics combine four main categories of information: facts, sentiment, aboutness and genre. RavenPack systems identify the facts and sentiment of news stories with multiple proprietary algorithms. New analytics are backtested over many years of sample data.” There’s years of the stuff back there.

But in machine-readable news as in risk rating and management, as in so many other areas of activity, past performance is not necessarily a guide to the future. At a simplistic level, this is just to say that tomorrow’s mortgage default won’t necessarily lead to the day after tomorrow’s global financial meltdown. Fine. The news-specific problem is this. News isn’t news unless it’s an exception. As Slomowitz says: “If you’ve got a



scoop and others don’t have it, there’s a lot of alpha there.” But it has to be a scoop. If you’re not geared to detect, let alone use, the truly left-field stuff, you’re not as alert for a trading edge as you might be.

Equally, it’s not the numbers that come in on time, as expected, that make the news worth machine-reading. It’s the exceptions. Therefore, the value of a machine that reads the news, putting it crudely, is that it knows what to do when something unexpected happens. This can be problematic. Discussing the inherent unpredictability of news, Frederic Ponzio, MD of NET2S, says: “Two things are missing. The first is getting the news in a structured format. News is unstructured data. The second problem is, how much can you trust?” Nine times out of ten, as Ponzio says, news is correct and verified. But ... see also the box ‘What if the news is wrong?’.

News is dangerous, remember. So the approach to machine-usage of news is likely to be defensive. What does that mean in practice?

Making the news

Larry Rafsky, CEO of Acquire Media, makes a thought-provoking distinction between “anticipated news” and “unanticipated news”. The anticipated variety comes on schedule. You may not know what it is, whether it’s going to be good or bad, but you do know that it’s due out next Thursday at 4pm. Unanticipated news is the other kind. Rafsky says: “All of a sudden, the week before a scheduled announcement by a company, there’s a story on that company. There’s volume. So you programme your black box to unwind your position. You just don’t care



whether it's good news or bad news. You just don't want the volatility." You might have an order out that hasn't been filled yet; you kill the order because you don't have to know the detail of the unanticipated news, to know that the ground has shifted under your trading decision.

It is an attractive risk-reduction strategy to avoid the unexpected, and as Rafsky says, there is significant interest in solutions that do little more than flag up unscheduled news announcements. This is machine-readable "news that there is news", and in these uncertain times, it's a no-brainer to input the guess that unexpected news is likely to be negative. Defensive indeed. But Rafsky develops a further distinction between "hard" and "soft" news. Rafsky says: "Hard news is reported. Soft news is manufactured. 90 per cent of the business news in the US is manufactured not reported. The companies don't want to manufacture it, but the regulatory authorities force them to do it."

Mandatory disclosure is a form of news manufacture. Occasionally, a story might get out before the press release has been approved, but typically, any item of business news, even unanticipated news, will have been through a manufacturing process before release. This is interesting in several ways. We might speculate, for example, on whether a company's news-manufacturing process is detectable in itself. Have board members suddenly disappeared from view, for example?

Has the CEO stopped returning your calls? If there are no results due, but they're not talking in the way that they usually talk, maybe this counts as an indicator in its own right.



Larry Rafsky

What if the news is wrong?

Mark Palmer, CEO of Streambase, made the news himself with a 10th September 2008 blog post on the "unfortunate algorithmic trading run" on United Airlines of the preceding Monday. The run was caused when an inaccurate news story was fed into trading systems. Palmer says: "It's not just algos, but the algos tend to get hit first because they react to news." Firms have to demonstrate that they provide best price when they trade, Palmer says, so news organisations should have to demonstrate basic news standards.

And the most basic is: the date on which the story was first released. The inaccuracy in the United story was not that it was wrong, but that it was six years out of date. That's the one detail the algos missed.

The drop in United's share price was from \$12.50 to \$3. By the end of Tuesday, it was back up to \$10.60.

More importantly, the key features of a news-manufacturing process are, first, that you can plug into it, and secondly, that both manufacture and release will be structured. This applies to anticipated and unanticipated news alike. Manufactured news will also typically have a life cycle: from the first intimation that it's coming, to the ripple in the market as it passes. Mandatory disclosure also tends to be formatted: over time, you could develop fields for much of the core data that any given source is likely to release. A central bank will probably be announcing a bail-out package (or an interest-rate decision), just as a commercial entity wants to talk, it will tend to be saying something about a new product or its sales figures.

Now let's turn this on its head. Technology enables the management of complexity. In a magazine concerned with automated and algorithmic trading, that's another no-brainer. But the actual trading, although it's not exactly simple, tends to be – what shall we call it? – focused. It is perfectly feasible, for ▶

Continued on page 51

Continued from page 47

example, to build an entire trading strategy around a given subset of news such as corporate actions, just as it is possible to tag a sufficient number of potential news sources (companies due to announce results; any federal agency collating price-sensitive data) to ensure a supply of trading opportunities.

What isn't possible (yet) is to set up a system that will take into account the impact of everything that happens. News is by definition "chaotic", in the strict sense of that term whereby reliable forward-looking statements cannot be made (see also the box 'News is not a butterfly'). Any trading activity has to move away from chaos. This suggests that, while the ongoing sophistication of news is welcome, the more effective approach to using it may still be either to remain defensive (as in: if there's unscheduled news, get out), or to limit the news universe to a manageable subset (corporate actions).

But there is a chaotic field in which forward-looking statements can be, and are, used as the basis for trading decisions: the weather. Discussing Weather Insight's methodology, Mitchell says: "Just to cover all the different weather forecasts that come in during a day, I have to build a spreadsheet that has over 100,000 cells. What I have to do is provide the name of the government forecast model, the forecast date of that model, then the run-time of that model. Instead of having one cell update lots of times, I have to have a whole bunch of cells update once in a day." Not surprisingly, Mitchell emphasises standardisation and the need to track the accuracy of forecasts.

Brian Haag, Sales Manager at RTS Realtime Systems, says: "I think weather is popular because it has such immediate and tradable impact in areas such as energy markets. If people buy a machine-readable weather feed they can plug it in and see it has an immediate effect. That's why there is real uptake on that."

Indeed. But there may be more to this. Mitchell says: "Since enough people follow the government's weather forecast, it will move traded markets almost instantly. More times than not, traders would be on the correct side of the trade. The swings from one model run to the next model run of a particular government's weather forecast is fairly high, and these types of directional trading opportunities tend to present themselves (statistically) about every seven business days. This trading style has nothing to do with accuracy; rather, the market's general weather perception at any given moment in time, becomes reality."

What about the extremes?

How would the near-collapse of western capitalism be handled in news analytics, and would its negative weighting distort news analysis in future?

Don Williams, Managing Director and Senior Vice President of Sales at RavenPack, says: "One technique is to gauge the impact of a current news item based on the impact similar items have had in the past. If the news was about a strike, changes in management or other corporate events it is possible to understand, based on historical news and market history, the likely impact of a new event. This of course is only possible with access to a reliable, deep archive of historical data with lightning-quick query and response.

"Other techniques include the ability to assess the 'novelty' of a news item. This means analyzing the news and subsequent versions of a story put out by a single publisher or as republished by other media outlets. You can imagine that the first instance of a story would receive a high score based on its historical impact and uniqueness, while subsequent versions of the story would receive a weaker 'novelty' rating that would dilute its impact on an aggregated news sentiment score.

"If we play this out in recent events, the first instance of the news about the failure of Lehman Brothers would receive a very high market response and novelty score. Subsequent news rehashing the same story would receive weaker novelty scores which could dilute the effect of recycled news on a trading algorithm. However the algorithm may also take into account the sheer volume of news behind a particular event.

Past performance may not be an indicator of the future, but even in weather forecasting, perception can become reality.

Would you have predicted that, as a conclusion?

