Barry Johnson

TRADE ALL ABOUT IT

Headline news is just potentially tradeable data delivered in a spectacularly inconvenient form. You can get it easily enough, but the big question is whether you can use it effectively. Barry Johnson discusses the evolution and practical application of news-analysis techniques to algo trading.

World markets are driven by information. Reacting appropriately to breaking news events can give traders a significant edge over the rest of the market. So, over the last few years, the incorporation of news into automated and algorithmic trading systems has attracted more and more attention.

Technology is also steadily transforming the news itself, both in terms of how information is supplied and how quickly, as well as from where it may be sourced. The major news wires, press services and conventional media now face competition from social media sites, by which I mean everything from blogs and message boards, to wikis and podcasts. Social media can offer glimpses of information well before it reaches the mainstream media. Hence, companies, such as Collective Intellect and Relegence, now provide products which glean useful information from a wide range of both conventional and social media sites.

The relentless increase in both the speed and capacity of computers also offers a massive potential for news-based analysis, as does the increasing

digitisation of news to computer-readable formats. Artificial intelligence and natural language processing techniques may be used to perform linguistic analysis and classification of the information. Pricing models can then adjust their fair values for the given asset, based on this data. However, we can also use news to make short-term forecasts of how market conditions might change, and so try to improve the performance of our execution/trading.

The good, the bad, and the ugly

The advantages of successful news gathering are easy to see. For instance, let's take an example from the pharmaceutical sector. Avastin is Genentech's trade name for Bevacizumab, an anti-angiogenic drug which has been approved for use against colorectal cancer since 2004. On 14th March 2005, the National Cancer Institute (NCI) posted the results of Phase III trials using Bevacizumab combined with chemotherapy for patients with advanced lung cancer. And then, four hours later, Genentech released this news to the press. The immediate result was a 25% hike in the company's stock price. Anyone who had made the connection between Bevacizumab and Genentech within that four-hour window could have had a significant market lead.

But things can also go the other way. Probably the best known example of this is the United Airlines story. Early on 8th September 2008, UA's stock price plummeted from $12.30 to $3 a share. The cause was the mistaken re-release of some old news: the company's 2002 bankruptcy filing. This was a six-year old story that had slipped back out onto the internet and news networks without a date-stamp. More than $1 billion disappeared in a little over ten minutes. And yes, when the story was corrected, the upturn was also steep. It doesn't really matter how the mistake happened, human or technical errors mean that every now and then incorrect news will get broadcast.

Accidents aren't the only challenge. Fraud is not just limited to 'penny stocks'. Those internet 'pump and dump' scams can also affect large cap stocks. For instance, David Leinweber and Ananth Madhavan (2001) point to the example of Lucent Technologies. On 22nd March 2000, a fake earnings warning was posted on a bulletin board, together with a barrage of replies and comments to lend it credibility. The firm's market cap dropped $7 billion. The scam was successful because at the start of the year internet bulletin boards and chat rooms were full of negative rumours about the company, so it seemed plausible.

Clearly, credibility and accuracy are vital in news. Automated trading systems make or lose money based on whether their interpretation gets them on the right side of the market. However, news-driven automated trading doesn't have to be the only answer. Common across all these examples is a sudden shift in market prices and volumes, irrespective of whether the news is correct or not. Though we may not be able reliably to predict exactly what information the various news sources will bring to us, we can still estimate how the market might react to that news.

Reacting to news

A large number of empirical studies have analysed the impact of news for most of the major markets. The impact of macroeconomic news on asset prices varies somewhat for different asset classes. For example, news of an interest rate drop will help increase the price of existing bonds, whilst it weakens the currency. For corporations, the picture is slightly more complex, since they may benefit from lower interest rates but may also suffer from the weaker currency.

Corporate news is obviously much more closely linked to specific assets. That said, the information may also have a considerable effect on other firms in the same industry/sector, or the supply chain. For stocks, news such as earnings announcements can trigger significant short-term price changes, which can last for anything from an hour to days or even weeks. Over-reactions can also lead to reversals over the following days. A nice example of the intraday effect is given in a study carried out for RavenPack by Peter Hafez (2009). Chart 1 shows the abnormal returns for a set of sample stocks (taken from the DJIA and EuroStoxx 50) following bad news, as determined by their sentiment classifiers.

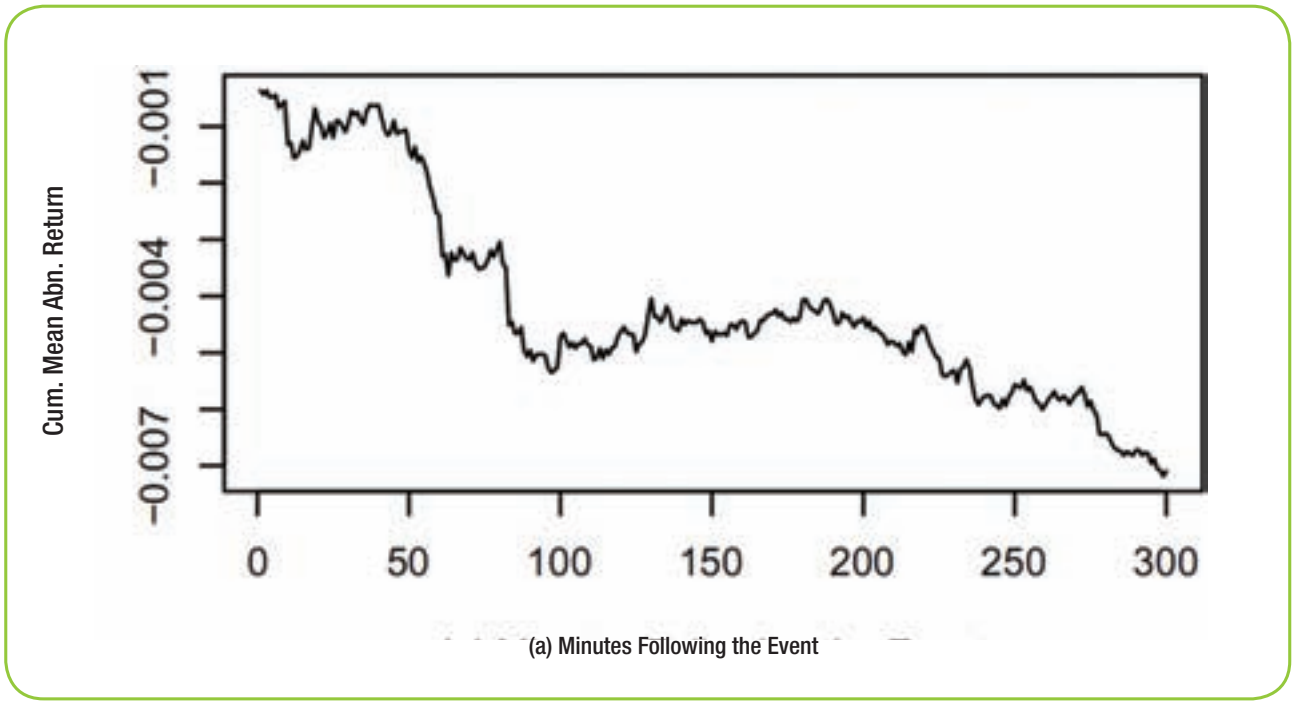

Chart 1: Abnormal returns following a negative sentiment event

In terms of trading volume, and volatility, the reaction is generally a significant short-term increase for all asset classes. For instance, Chart 2 shows the distinctive spikes in volatility for US Treasuries in response to announcements at 8:30am and 10:00am (as a standard deviation of log price changes), taken from a study by Michael Fleming and Eli Remolona (1997).

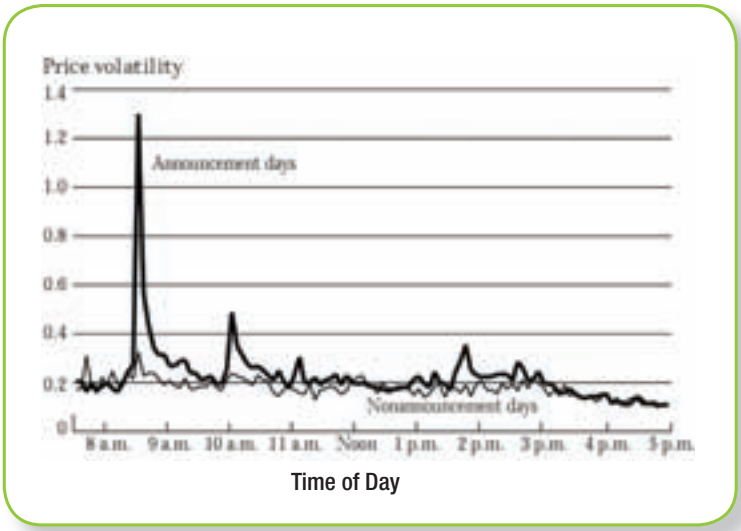

Chart 2: Intraday volatility on announcement days for 5 year U.S Treasury notes

Liquidity is less clear-cut. In many cases, liquidity often increases after announcements. However, there can also be marked decreases in liquidity. Often this is prior to the announcement, and may sometimes be attributed to 'the calm before the storm' for announcements which are widely expected.

Factors which affect market reactions

Obviously, the impact of news depends on the information that it brings to the market. News provides data that either confirms or confounds the market consensus. So the degree of surprise indicates how the market will react. In order to gauge the degree of surprise, we need a common source of expectations, such as consensus forecasts. The difference may then be normalised, helping to highlight the most significant information. Some news obviously lends itself to this approach more than others. Macroeconomic corporate announcements each deliver new numbers to the market. For instance, Chart 3 shows the monthly changes in Non-Farm Payroll data together with the market expectations from Money Market Services International surveys between January 2000 and June 2002, taken from a study by Linda Goldberg and Deborah Leonard (2003).

To quantify the size of any surprises they normalised the difference between actual and expected figures by dividing by the standard deviation of Non-Farm Payrolls relative to expectations for the period. This confirms that events such as the ones at 7/00 and 2/01 are statistically significant.

The timeliness of news also dictates whether it is new information or simply a confirmation of existing data. For instance, some macroeconomic reports (such as Employment) arrive sooner and so often reveal newer ▶

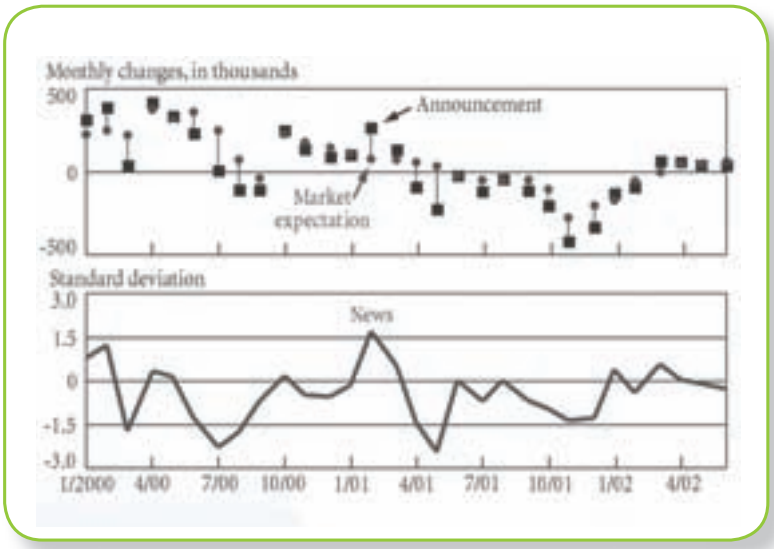

Source: Goldberg and Leonard (2003) Reproduced with permission from Federal Reserve Bank of New York

Chart 3: U.S. Non-Farm payrolls: Announcements, expectations and news

information than others (such as the GDP). Hence, these tend to have more impact on the market. With corporate news, firms tend to delay issuing bad news, so the delay can even act as an indicator.
Market conditions and the business cycle are another important consideration, since these reflect investor sentiment. In a declining market, even good news can fail to buck the trend.

Incorporating news into trading algorithms

So how do we incorporate news into trading algorithms, without creating something from out of '2001: A Space Odyssey'? Algorithms are used to responding to market data and analytics, whether this is the price, spread, volume or volatility. Hence, indicators provide an easy way for trading algorithms to take news into account.

News flow is itself an important trading signal, following the old adage "There is no smoke without fire". So the sheer volume of news items can be just as much an indicator as the actual information they convey. Clearly, we cannot reliably use the number of headlines to predict future prices. But a sudden rush of headlines does suggest that volatility may increase, since uncertainty breeds volatility. Thus, when news levels rise significantly above the historical average we have an immediate indicator that something unusual is happening. For example, Relegence calculates a news heat index, based on the number of stories for a specific company, converting this to a scale ranging from -4 to +4. Risk and cost-based algorithms can then adjust their trading strategies accordingly.

Sentiment analysis is probably the most common means of converting news into a suitable indicator, as outlined in Automated Trader's 2008 feature on the subject ('All the News that's Fit to Trade', page 40, Q1 2008, see www.automatedtrader.net). This provides a straightforward means of interpreting whether news is favourable or adverse. Sentiment models can be based on natural language processing, use feedback from analysts for 'expert consensus' or historical market responses. Numeric sentiments may be assigned for each story, together with a confidence level; this can even span multiple firms. Commercial news analytics solutions, such as RavenPack's News Scores, have been available for some time.

Unlike news, sentiment is a continually changing factor, much like the market price. Hence, the sentiment can also be used in its own right. For assets with frequent enough news events, real-time sentiment analysis is possible. That said, the complex relationship between sentiment and price can make it difficult to reliably take advantage of this information, particularly for individual companies. However, sentiment scores may also be aggregated across sectors or indices to create broader indicators. Sentiment indices have been found to be significantly more reliable than those for single assets.

News-adaptive algorithms

We have become used to trading algorithms, routinely making automatic choices about the location, price and size of the orders they are tasked ▶

with executing. These decisions are often based on a mix of historical and real-time market data, and the responses are driven by rules, which are specified by traders and subjected to rigorous testing. They may also use additional models, to help with making their decisions. Cost models help at the macro level, whilst at the micro level limit order models can offer execution probability for different order types or limit prices.

Increasingly, though, algorithms are becoming even more adaptive, responding dynamically to changing market conditions, such as liquidity. For instance, adaptive implementation shortfall based algorithms take account of short-term price trends and adjust their trading patterns accordingly. Market models may also be used, giving short-term forecasts for how market conditions may change.

So reacting to news may just be viewed as another real-time adaptation. Hence, as Chart 4 shows, news indicators may be incorporated into such market models, helping to make algorithms "news aware".

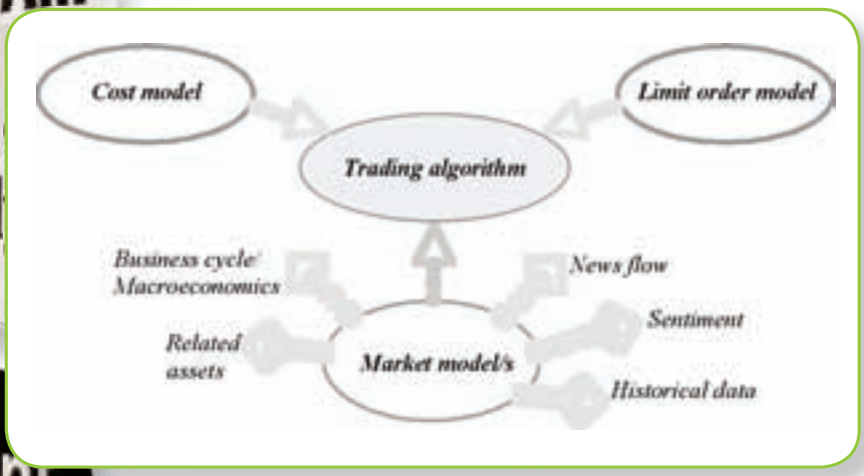

Chart 4: Incorporating news in short-term models of market conditions

For execution systems, simply being able to tell that the market is about to shift can be invaluable. Short-term predictions of market conditions may enable our trading strategies to take advantage of the market response, rather than simply reacting to the rapidly changing conditions. Adding a news handling capability may well improve their performance when news events do occur. For example, VWAP algorithms could modify their target volume profiles to account for the news release, whilst cost-based algorithms might alter their short-term volatility estimates.
In terms of order placement, news-adaptive algorithms could trade a sell order less aggressively when good news arrives, whilst for a buy order they might be more aggressive. Provided the market reaction behaves as

expected, the news adaptation should give slightly better results than more reactive price-based approaches.

Similarly, temporary adjustments could also be made to the limit order models used by cost-based algorithms or execution tactics to determine the optimal approach for order placement. Thus, news just becomes another factor, like price and liquidity, upon which algorithms and execution tactics base their order placement decisions.

Purely news-driven algorithms

Market prices and trades vary continuously throughout the day, hence algorithms such as percent of volume or price inline can constantly track this data and adjust their trading patterns accordingly. In comparison, the arrival of news is a much more discrete event. Each day there might only be a handful of news stories that might affect an order, and often there may be none. Even for a worst case scenario there is likely to only be a few pieces of significant news arriving in a day. There may well be 30 or 40 different stories, but many of these will simply be the same information from different sources, possible revisions, or daily summaries outlining the effect the news has had on the market.

So purely news-driven algorithms are a very different beast to the commonly available trading algorithms that seek to minimise cost or market impact. They are completely opportunistic, in a sense behaving more like a sophisticated conditional order, only being activated if good (or bad) news appears. Essentially, they act much like a stop (or an entry) order, based on information rather than price. In fact, a stop order (or trailing stop) might be used just as well, given a reasonably accurate estimate for the potential impact of the news. No doubt, over time, much more sophisticated news-driven algorithms will develop.

A second pair of eyes always helps

Algorithms evolved as a tool to help busy traders. Incorporating news handling expands their capacity to deal with more complex conditions. That said, handling news means interpretation, so there is always scope for error. So it is important to track the algorithm's performance, ensuring there are no issues. Built-in flagging would let algorithms alert the trader/investor when unexpected conditions arise.

Admittedly, price limits offer some protection, but based on the news event the investor/trader may prefer to cancel the order; so they need to be given this option when appropriate. For instance, let's say we're buying when sudden good news arrives - given that the price looks likely to trend away for the rest of the day it seems reasonable to trade more aggressively. However, it is important to remember that there can be an overreaction to news; so the price may well reverse tomorrow. Obviously, there is no guarantee of this, so the investor/trader must decide how much they want to complete the execution today.

There's still some way to go in the incorporation of news into electronic trading strategies. But it's important to remember that it doesn't just have to be about generating ideas for investment or trading. News can also play a significant role in enhancing execution.

This article was based on material from the new book 'Algorithmic Trading and DMA: An introduction to direct access trading strategies' by Barry Johnson (4Myeloma Press; February 2010; $49.99; 978-0956399205; Paperback) which is available from amazon.com. There is also a dedicated website (www.algo-dma.com) which offers a preview of the book.

References:

Fleming and Remolona (1997) Fleming M.J. and Remolona E., 'What Moves the Bond Market?', Federal Reserve Bank of New York Economic Policy Review, vol:3 (4)

Goldberg and Leonard (2003) Goldberg L. and Leonard D., 'What Moves Sovereign Bond Markets? The Effects of Economic News on US and German Yields', Federal Reserve Bank of NewYork Current Issues in Economics and Finance, vol: 9 (9)

Hafez (2009) Hafez P.A., 'Impact of News Sentiment on Abnormal Stock Returns', RavenPack International S.L. (www.ravenpack.com)

Leinweber and Madhavan (2001) Madhavan A. and Leinweber D.J., 'Three Hundred Years of Stock Market Manipulations', Journal of Investing, Summer 2001