

Risk Systems That Read

Dan diBartolomeo

Webinar

May 2014

Big News and You Heard It Here First!

- After a series of research projects going back to 1997, Northfield is about to commercially introduce risk models augmented with quantified news flows for investors.
- News conditioned models will be available to Northfield clients starting in July (sorry but it won't be free).
- *We believe this will be the biggest step forward in risk modeling for asset management since the creation of the multi-factor risk model in the 1970s.*

Outline of Today's Presentation

- Concepts of conditional risk models
- A formal definition of News
- Potential applications
- Review of prior literature and research
 - Seven separate studies
- The Northfield process
 - “Ingredients”
 - Description of the functional form
 - Two aspects of time decay
 - Controlling “noise”
 - Separating news impact on factors and specific risk
- Conclusions

Concepts of Conditional Risk Models

- Almost all available risk models are “unconditional”.
 - They are based on a sample of past history that deemed relevant, possibly giving more weight to recent observations, or assuming a simple trend in volatility (e.g. GARCH)
 - Sample periods range from 60 days to more than 20 years.
 - Once the sample period is determined, the *heroic assumption is made that the future will be like the past.*
- This process omits everything we know about the present, and how the present is different from the past average conditions of the sample period.
 - Using the information about the present to adjust the risk estimates has been standard in some Northfield models since 1997, and in all models since 2009

A Definition of News

- For our purposes, “News” is the set of information coming to investors that tell us how the present is different from the past.
 - This definition implies that routine information affirming the “status quo” is not news irrespective of how it is delivered. Investors respond differently to “announcements” (time of information release anticipated) than to “news” where both the content and timing are a surprise
 - Only a minority of large asset price moves are a direct response to investors responding to news. There are a lot of “information-less” trades (see Livnat, et. al. 2013). We need to be selective.
- It should be very intuitive that risk assessments should also respond to news

Applications

- News conditioned risk assessments will be immediately useful in a number of applications
 - Alpha estimation, portfolio risk and optimization for high turnover portfolios, especially hedge funds and proprietary trading
 - Algorithmic trading including “high frequency” and optimal execution algorithms, especially those like the Northfield “algo” which are framed as multi-period optimization problems
 - Hedge fund investors include “fund of funds” and asset owners
 - Making short term “compliance” risk assessments for mutual funds and asset managers (e.g. UCITS regulations)
 - Daily updating of credit and counterparty risks
 - Our research indicates that conditioning with news can account for **10-15% of changes in portfolio risk for horizons as far out as one year.**

Review of Previous Literature and Research

- diBartolomeo and Warrick (2005)
 - Uses percentage changes in option implied volatility to adjust volatility estimate of individual stocks *daily since 1997*
 - Uses a regression method to separate adjustments to security level volatility into adjustments factor volatility and security specific risks. Allows for partial adjustment of risk for non-optional stocks.
 - If there is a big jump in the implied volatility of one stock it is assumed to security specific but if most stocks in an industry see a jump in implied volatility most of the change applied to the industry factor.
- Very intuitive results when markets reopened after September 11th, 2001.

Review of Previous Literature and Research

- diBartolomeo, G. Mitra, L. Mitra (2009)
 - Followed the analytical structure of diBartolomeo and Warrick (2005)
 - Replaced option implied volatility with measures of quantified news flow and sentiment
 - Empirical tests on high liquidity stocks (Dow Jones 30 and EuroStoxx 50 names) for next day intra-day volatility
 - Findings were that *news driven metrics were more efficient predictors of changes in volatility than metrics based on changes in option implied volatility*
 - We believe that news metrics work better than implied volatility because option markets have trading costs so the changes in implied volatility are muted compared changes in the beliefs of investors

Review of Previous Literature and Research

- We had two separate teams of MIT graduate students conduct their own research projects over two years
 - Changes in one day security volatility was highly statistically significantly associated with changes in news flows at the individual stock level
 - More than a dozen functional forms of the relationship were tested on a large samples of hundreds of stocks over several hundred trading days
 - Intuitively, the predictive power of news flow changes decayed rapidly for the most liquid stocks and more slowly for less liquid names.
 - The impact of news also decayed more quickly for firms with more public recognition (e.g. Apple or Google).

Review of Previous Literature and Research

- Kyle, Obizhaeva, Sinha and Tuzun (2012)
 - Shows that a theoretically predicted relationship between the frequency of news articles on companies, and the volatility of their stocks was fit almost perfectly by the empirical data over hundreds of companies and many years.
 - Clever construct suggesting that a function of stock volatility and trading volume across stocks is constant when the rate of time passage is defined in “numbers of articles” which they call “business time”
 - Prescribed functional form is a power function that also includes a “expected changes in trading volume” component
 - Related papers show how this structure can be used to predict bid/asked spreads and more generally trading costs

Review of Previous Literature and Research

- Northfield internal research (2013-2014)
 - All of the previous research suggests a *multiplicative* relationship between news flow and security volatility
 - Tested an *additive* functional form (H1: tomorrow's volatility goes up when a threshold value of news flow is hit today)
 - This is an easy structure to fit in an existing factor model as a dummy variable. The factor exposure is 1 if there was “enough” news today and zero otherwise.
 - Main data set had 1.7 million data points (stocks * days)
 - We tested multiple providers of news flow analytics. They all worked to a statistically significant degree. T stats ranged from 7 to 9 in a Heston (1993) estimation structure with sample periods as short as 22 trading days

Where We Stand Today

- Northfield internal effort 2015
 - Independently replicated the 2013-2014 research to confirm differences between news analytics providers
 - Additional tests to compare effectiveness of multiplicative versus additive functional forms
 - Wrote article for publication (forthcoming) on how our use of news analytics for individual stock volatility can be extended for evaluating credit risk in corporate bonds, government bonds and counterparty risk
 - Finalizing analytical details and production processes for July introduction of daily updates of most Northfield models and delivery platforms.

Ingredients to the Conditioning Process

- We are using Dow-Jones news feeds as summarized by Alexandria Investment Research
 - Coverage of many thousands of stocks and several thousand “topics” such as sectors and countries. Quantitative summarization of each article takes no more than 30 milliseconds. Alexandria “reads” both English and Japanese.
 - Each article summary provides numerous metrics including the “sentiment of the article” (good news/bad news), “relevance” “confidence”, “novelty” (has similar news been previously reported recently) and “event type” (a merger is probably more important than a routine dividend)
 - Northfield then compiles “today to date” aggregations of the article count on each subject. Alexandria metrics are used to weight the importance of individual articles in the overall count.

Functional Form Part 1

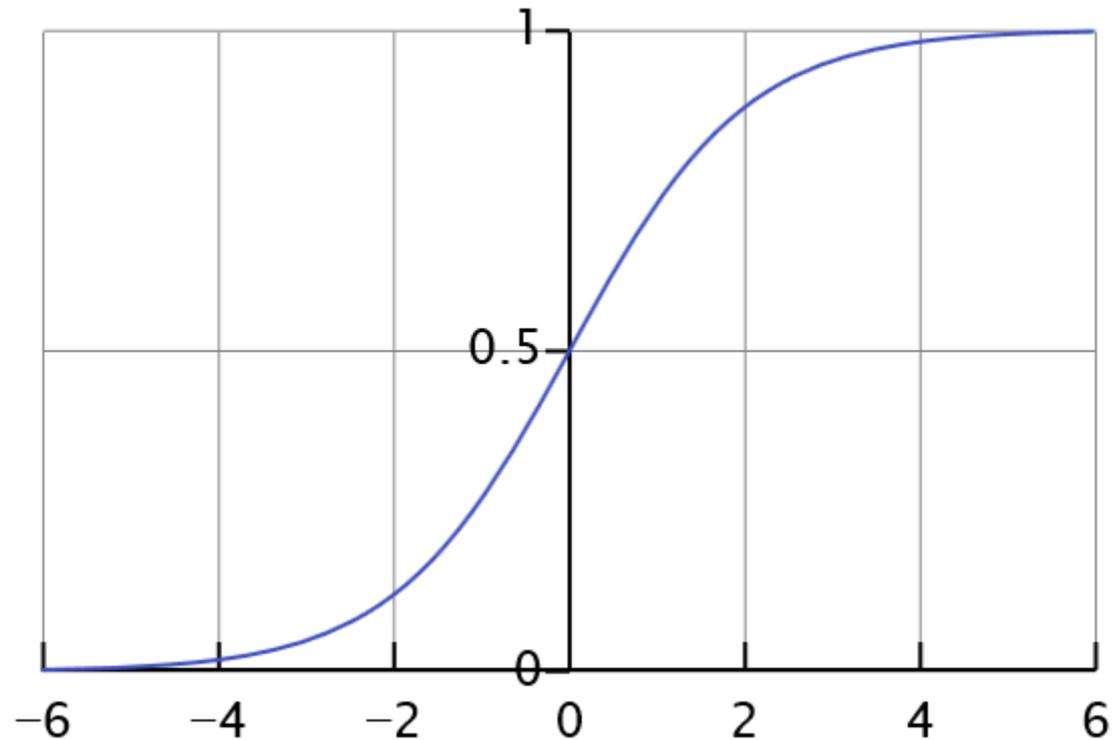
- I'm not going to provide the exact functional form of our process just yet. We're usually very transparent but,
 - We decided not to when we saw announcement that one of our competitors is starting to condition their models on security level implied volatility (as we did in 1997). *Better late than never.*
- What I will say is:
 - The functional form is multiplicative so our conditional forward estimate of risk is our unconditional forward estimate of risk times some scalar (with default value one).
 - The scalar is derived from the average amount of news flow in recent days (as weighted by the Alexandria metrics) in comparison to a long term average of news flows on the same company. Scalars can be above or below one.

Functional Form Part 2

- The conditioning captures multiple aspects of time decay
 - How long ago did the news take place? For example, there may have been a spike in news volume three days ago which will still be important, but less important than if the spike in news volume occurred today.
 - How fast will investors notice the events? For high volume, US liquid names the impact of news events will decay a lot faster for an obscure firm with no analyst coverage.
 - Separately we consider the impact of time decay based on the forward risk horizon. For example, if we are trying to forecast intra-day volatility for tomorrow, an increase in news volume will have more impact than if we are trying to forecast average daily volatility over the next ten trading days.

Functional Form Part 3

- We want to ensure that the conditioning on news does not cause the risk forecasts to “wobble around” based on insignificant changes (i.e. minor random noise).
- To address this we include a logistic function to smooth out small changes.



Separation of Factor and Specific Risk

- diBartolomeo and Warrick (2005) shows how adjustments to the security volatilities can be “fed back” into the model to adjust factor variances and volatility estimates for stocks on which no options are traded
 - See equations 7 through 9
 - Same process for news flow in diBartolomeo, Mitra, Mitra (2009)
 - If the factors are orthogonal this process can be reliably estimated with a simple OLS regression
 - If the model factors are not orthogonal you either use a non-linear optimization process (used in Northfield models since 2009) or generate an orthogonal transform of the factors, estimate using OLS, and then translate the factors back to the original basis.

Conclusions

- After a heck of a lot of research and 18 years of experience in a similar framework, we are moving ahead commercially with daily conditioning of Northfield risk models based on our proprietary measures of news flow.
 - We are confident that this is a huge improvement in near horizon risk estimation in a variety of portfolio and trading applications.
 - There is little doubt that the strategies of most investors involve some form of response to financial news as it comes forward.

It is implausible that estimation of security and portfolio risk should somehow ignore this very obvious and elemental fact.