

Risk Systems That Read[®] Redux

Christopher Kantos

Northfield Information Services

Asia Seminar Series October/November 2017

Two Simple Truths

- It is hard to forecast, especially about the future
 - Niels Bohr (not Yogi Berra)
- It's really hard to predict the future from the past if you ignore the present
 - Me

“How to Spot the Next Really Big Thing”

- The above phrase was the title of a research paper I published in the *Journal of Indexes* in 1999. It related to how identify new factors impacting financial markets.
- In terms of our main business line of creating risk models to predict the covariance of securities and portfolios, the use of text analytics to improve near-term risk forecasts is “the next really big thing”.
 - Text analytics enhanced risk models are now in daily production at Northfield. We will be releasing the enhanced models to users within a few days.
- We believe this is the biggest advance in risk assessment for asset managers since the 1970s.

Physics Meets Finance

- The analytical basis of the new approach is borrowed from “relativity” in physics. Differences in financial asset volatility over time are treated as a function of how quickly information is flowing to investors.
 - *When there is **intense** arrival of news, time itself is presumed to be passing quickly. When there is a little news of interest to investors about a particular company, sector or country, time itself is treated as passing slowly.*
 - We are treating time as endogenous. Rather than consider time as defined by minutes, hours or days, we define the passage of time in terms of news intensity, just as traders often think of time in terms of trading volume.
 - The intensity of news flow is a function of both the **frequency and the importance** of news articles.

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 typical conditions of the past 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

Some Motivation

- The Global Financial Crisis of 2007-2009 demonstrated that the risk systems of financial institutions are frequently grossly inadequate.
 - The \$1.4 billion fine recently paid by Standard and Poors is yet another illustration of that inadequacy.
 - What is missing from nearly all financial models is a recognition of how the present is different from the past, and therefore, how the near future is also likely to be different from the past.
 - By defining “news” explicitly as the information set that informs us of how the present is different from the past, we can improve our estimates of near time horizon risk levels.
 - *Investors go to great effort to analyze financial news. It is bizarre that the estimation of security and portfolio risk should somehow ignore the information that investors give the greatest attention*

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

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 stocks on which options are not traded.
 - 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 related industry factor.
 - **Very intuitive results when markets reopened after September 11th, 2001. Estimated volatility of an average airline stock was up 65% between September 10 and September 17.**

Publication in *Quantitative Finance* (2009)

- 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

Internal 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).

News and the “Invariance Theorem”

- 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

A Pinch Hit by Louis Scott

- All of the previous research suggests a *multiplicative* relationship between news flow and security volatility. We also 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 highly statistically significant degree.
 - **T stats ranged from seven to nine** in a Heston autoregressive (1993) style model structure wherein tomorrow's intra-day volatility is a linear function of the past one, five and twenty-two day volatility levels.

News Sentiment as a Factor?

- We reject this formulation of using news flow (or some metric thereof such as “sentiment”) as one of the factors in the risk model for a number of reasons.
 - The first and most important is that it misses the profound basis of the entire process. We agree with Kyle, et. al. that one can alternatively *describe changes in asset volatility as time itself speeding up and slowing down.*
 - When there is a lot of news time is passing quickly, so volatility seems high when measured in clock time. When there is very little information coming to investors, time is passing slowly when measured in clock time.
 - The concept of endogenous “information time” was previously explored in Haug (2004).

Functional Form

$$S_i(t) = S(m) * (1 + K_{it})$$

$S_i(t)$ is the security level **annualized** volatility estimate for day T

$S_i(m)$ is the previous month end **annualized** volatility estimate

$$K_{it} = f(RNI, BNI, L, H)$$

RNI = recent news intensity

BNI = baseline news intensity

L = a measure of liquidity of stock I

H = the time horizon of the risk forecast

The Secret Sauce

- We define the speed of the passage of time by the intensity of news flow. It is a weighted count of news stories in a given period.
- The input metrics used to define the “weight” of a story are provided in real time by a “supervised machine learning” process provided by Alexandria Investment Technology based on text flowing through the Dow Jones news system. *Typically, we analyze over 5000 articles per day.*
- The weight assigned to a news story in the count depends on
 - Sentiment (good news, neutral, bad news)
 - Relevance (an article may mention several firms or topics)
 - Confidence (confidence interval on the sentiment score)
 - Classification of the subject matter (a story about a merger or accounting fraud will be more important than a routine dividend) into one of about 30 categories.

Recognition of Time Decay

- The process 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. *Different rates of time decay across stocks are based on liquidity levels estimated from the existing Northfield transaction cost model.*
 - 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.

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. See Shah (2015 on SSRN) for details.

A Positive Definition of Stock Specific Risk

- One of the key benefits of the new technique is that it allows for a positive definition of stock specific (idiosyncratic risk).
 - In most risk models, the idiosyncratic risk of a stock is the standard deviation of the residual returns over time. **We are defining this very important item by what our model does not explain**, rather than by what our model does explain.
 - This is a crucial weakness for quantitative active managers, as standard alpha scaling techniques (Grinold, 1994) take stock specific risk as an input.
 - It should be obvious that stocks that in the midst of controversy (e.g. the Volkswagen emissions scandal) represent a key portion of the opportunity set for active managers.

Other Applications

- News conditioned risk assessments will be immediately useful in a number of applications
 - 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 estimates for horizons as far out as one year.**

Conclusions

- After a great deal of research and 18 years of experience in a similar framework, we have moved ahead commercially with daily conditioning risk models based on our proprietary measures of news flow.
 - We are confident that this is a large 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 the very obvious and elemental fact that news matters to investors.