# Regime Switching Compared To Conditional Factor Risk Models

**Dan diBartolomeo**
Northfield Webinar
August 2020

# Starting At the Start

- The recent volatility in global equity markets driven by the COVID-19 pandemic has reminded investors that both perceived and expected risk levels change from time to time, as events unfold.

- Such circumstances beg the question as to how to best structure equity factor risk models to correctly reflect these changes.

- The most basic conception of a model is *unconditional*, wherein we assume that the risk between now and some future horizon date are comparable to the risks experienced during a chosen historical sample period.

- Often, the observations in the sample period are weighted to give more influence to more recent events.

# Autoregressive Models

- There are two methods that have been broadly used to improve the ability of equity factor risk models to reflect new conditions and market information.

- The first is the presumption that volatilities of economic time series are autoregressive as described in prominent papers by Engle (*Econometrica*, 1982) and Bollerslev (*Journal of Econometrics*, 1986).
  - Engle, Ng, and Rothschild (*Journal of Econometrics*, 1990) have proposed the structure of an autoregressive factor model.
  - In this approach, we assume that abnormal volatility occurs in clusters so that large magnitude events are more likely than usual to be followed by large magnitude events, and that quiet periods are more likely to be followed by quiet periods.

- In effect, our expectations for volatility in the near future are *conditioned on the extent to which the recent past is different* from long term norms.

# Conditioning on News

- The second approach which Northfield has pioneered since 1997 is a *fully conditional* factor risk model.

- In such models, expectations of future volatility are conditional on contemporaneously observable information which is explicitly forward looking. Northfield uses three kinds of conditioning information:
  - The most widely used is the content of financial news text, which is evaluated in our *Risk Systems That Read*® process. Each day the content of about ten thousand news articles is evaluated, and adjustments are made to the expected risk of both individual securities and model risk factors (e.g. countries, sectors). Expected risks increase with both the volume and perceived importance of news coverage.

# Comparing the Present to the Past

- The second type of information used to condition risk estimates are currently observable market variables that can be statistically compared to the past.
  - Such observable variables might include the level of the CBOE VIX index (*expectations* of equity volatility) and credit related yield spreads on bonds (*expectations* of future financial distress. Discussion in https://www.northinfo.com/Documents/356.pdf.
- The third type of conditioning information is the option-implied volatility from individual stock options.
  - The model assumes that option implied volatility values are not correct, but time series changes in implied volatility are proportionally correct in describing investor expectations. See diBartolomeo and Warrick (2005). The extension of the process to text analytics is provided in diBartolomeo, G. Mitra and L. Mitra (2009). The full mathematical representation of the conditioning process is provided in Shah (2015).

# Regime-Switching Models

- A third process that has been widely applied in the area of dynamic asset allocation is a "regime switching" model, wherein we believe that at each moment in time, global markets exist in one of two or more identifiable different states.
- Each state has its own probability distribution for returns and volatility (a separate model). Adherence to such concepts has led to the popular idea of "risk on, risk off" wherein investors switch their portfolios between what they believe to be appropriate allocations for the various distinct states.
- The use of these concepts has been largely motivated by a series of four closely related research publications.
  - The first is Chow, Jacquier, Kritzman and Lowry (*Financial Analyst Journal*, 1999) which put forward the idea of using the Mahalanobis distance (a statistical measure to detect multivariate outliers) as way to separate historical observations into quiet and volatile sample periods.
  - A different covariance matrix among various assets is then estimated for the two states.

# More on Regime Switching

- This regime based process is extended to other asset types in Kritzman, Lowry and Van Royen (*Journal of Derivatives*, 2001).

- A further elaboration on use of the Mahalanobis distance to detect "financial turbulence" was provided in Kritzman and Li (*Financial Analyst Journal*, 2010).

- A final paper in the group, Kritzman, Paige, and Turkington (*Financial Analyst Journal*, 2012) introduces the formal concept of "regime switching" using the Hidden Markov Method (HMM) to detect when a change in regime had taken place.

- It should be noted that all of these papers address risk regimes at the asset class or currency level and do not attempt to address how to apply these methods to individual stocks or fixed income securities.

# Limitations on Regime Switching

- We will first discuss likely practical limitations on the implementation of regime switching equity factor models.

- The first issue is that in order to detect regimes and model regimes, a very large set of historical observations is required.
  - The Markov switching process described in Kritzman, Paige and Turkington (2012) uses data back to 1947 to model only a handful of financial times series (equities, inflation, economic growth, etc.)

- While equity factor risk models have existed in their modern form since the early 1980s, a current global model would typically *cover more than forty thousand* individual equities.

- Even if dimensionality of the problem is reduced to a persistent set of less than one hundred common factors, it is unclear that the historical record of factor model data is sufficient.

# Persistence of Individual Equities

- A second problem is that equities themselves are not persistent. New publicly traded companies are constantly created through initial public offerings and existing companies frequently disappear either through merger, acquisition or bankruptcy.

- diBartolomeo (*Journal of Investing*, 2010) used a Merton-style contingent claims model to show that share prices of US traded equities implied that investors expected firms to have a half-life on the order of twenty years for the sample period from 1992 through 2010. This is equivalent to an extinction rate of about 3.5% per annum from failures.

- We estimate a total creation/extinction rate (inclusive of mergers, acquisitions, IPOs, and de-listings) *of about 1% per month.*

# Persistence of Equity Markets

- An even broader problem exists when emerging markets are included in the global set. While many equity markets in many developed countries go back at least to the late 1800s, many emerging markets have had formal equity trading only for a limited number of years.

- The creation of new "frontier" markets as time passes is also problematic as the inclusion of new markets for which no history exists eliminates the potential to apply regime switching to the related factors and the covariance of those eliminated factors with other factors.

- Regime switching models are based on backward looking data in which survival to the current date is implicit.
  - Rare, extinction events such as China (1949), Russia (1917), Zimbabwe (2008) are not included in past data by definition.

# Idiosyncratic Risks are Excluded

- The final limitation of the potential application of regime-switching equity factor models is that a key benefit of factor models cannot be improved by regime-switching.
- Factor models provide an explicit separation of common factor risks from the idiosyncratic risk of specific firms or securities. Since these risks *are mathematically defined to be unrelated to the market conditions*, there is no plausible expectation of improvement in their estimation as a result of changes in the state of the market.
- Firm level idiosyncratic risks are often the predominant portion of benchmark-relative volatility (tracking error) and are also a key input to security level alpha estimation as in Grinold (*Journal of Portfolio Management*, 1994).
- Empirically observed changes in the central tendency of idiosyncratic risk levels arise as a mathematical artifact that most factor risk models estimate idiosyncratic risk from in-sample residuals, so temporary declines in the in-sample explanatory power of the factor set incidentally appear as wholesale changes in idiosyncratic risk levels.

# Empirical Analysis of Regime-Switching

- We will begin our empirical investigation with a simple exercise. We will evaluate the potential to apply a regime-switching model to the most obvious relevant single return times series, the S&P 500.

- Using daily return data from 2004 to date, we apply a Hidden Markov Model (HMM). The concept of an HMM model is that by observing a particular variable X (the S&P 500) we are able to make statistical inferences about the state of an otherwise unobservable variable Y (the "hidden" state of the world).

- Both two and three regime models were created.

# HMM Switching Applied to S&P 500

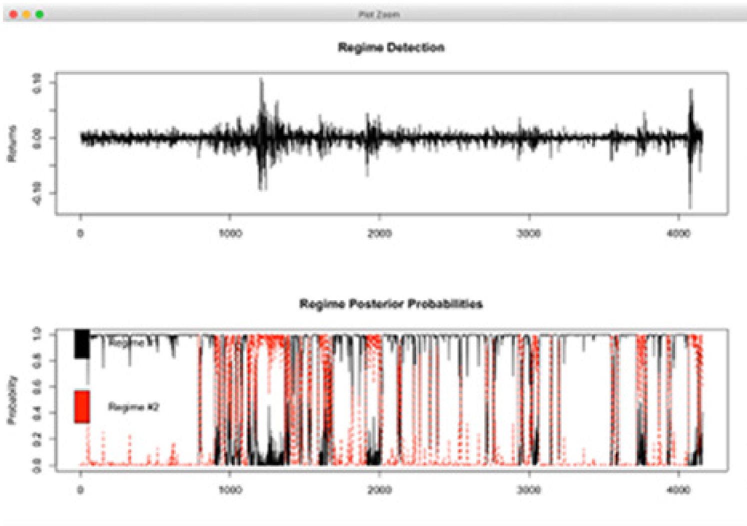


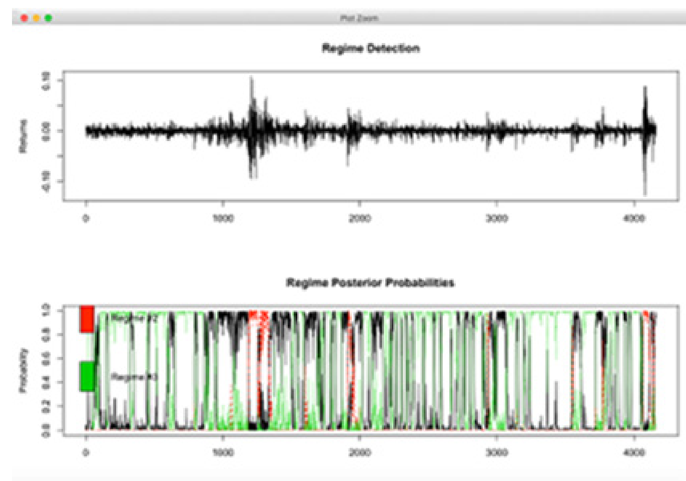

# HMM Results Discussion

- The two regime model (red/black) is in the upper graphic, while the three regime model (black/red/green) is presented below. The X axis is counting of trading days since the start.
- As can be seen, while there are many purported changes in regime (changes in color), unsurprisingly the two really large events that stand out are the periods of the Global Financial Crisis and the current COVID-19 pandemic.
- For a regime switching model to be actually used we must create expectations of whether the markets will switch regimes at each moment in time (i.e. a transition probability matrix).
- This data indicates that the probability of staying in the current state at a particular moment in time is about 90%, which is almost identical to the result in Kritzman, Paige and Turkington (2012). Obviously, as the number of defined states is increased from two to three (or possibly more), the potential for error in the transition probabilities increases.

# Recombining Multiple Regimes Into One

- It should be noted that if we believe that the transition probabilities can be estimated reliably, *any multiple state model can be collapsed back to a single model* as a "mixture of normal distributions" (algebra below).
- The resultant distribution will typically exhibit significant higher moments (skew and kurtosis). The most popular method of converting a four moment distribution to the economically equivalent two moment distribution is the method of Cornish and Fisher (1938).
- Northfield application software incorporates higher moments appropriately for both risk reporting and optimization as described in [https://www.northinfo.com/Documents/901.pdf](https://www.northinfo.com/Documents/901.pdf).

$$
\begin{aligned}
\mathrm{E}[(X-\mu)^2] &= \sigma^2 \\
&= \mathrm{E}[X^2] - \mu^2 && \text{(standard variance reformulation)} \\
&= \Big(\sum_{i=1}^{n} w_i(\mathrm{E}[X_i^2])\Big) - \mu^2 \\
&= \Big(\sum_{i=1}^{n} w_i(\sigma_i^2 + \mu_i^2)\Big) - \mu^2 && (\text{from } \sigma_i^2 = \mathrm{E}[X_i^2] - \mu_i^2, \text{therefore } \mathrm{E}[X_i^2] = \sigma_i^2 + \mu_i^2) \\
&= \sum_{i=1}^{n} w_i(\sigma_i^2 + \mu_i^2 - \mu^2).
\end{aligned}
$$

# Experimenting With Mahalanobis Distance

- Our next empirical exploration was to construct a two state version of the Northfield Global Equity Model (version 3) using the Mahalanobis distance to separate the observations in the style of Chow, Jacquier, Kritzman, and Lowry (1999).

- The data set used was from January 1990 through May 2020 with more than 360 observations.
  - We applied the Mahalanobis measure to the history of the factor returns to arrive at two approximately equal length subsamples.

- From each subsample a new factor covariance matrix was created. This gave us four different factor covariance matrices to compare as of May 31, 2020.
  - These are the "quiet" subsample, the "turbulent" subsample, the regular Northfield Global Model, and the Northfield "near horizon" Global model.

# Experimental Data Details

- For all four representations, the subject portfolio was the 1000 largest capitalization stocks in our global universe and the benchmark was the same 1000 stocks weighted by capitalization.
- Factor exposures of individual securities were taken from the May 31, 2020 Global model.
- The “market portfolio” is the global equity index we internally define for the purposes of the risk model.
  - It is comparable to a FTSE All World or MSCI ACWI.
- The experimental design implies an investor at May 31, 2020 trying to decide what portfolio volatility and tracking error will be going forward.
  - The results of risk analyses are summarized in the next table.

# Two Regimes, Unconditional, and Conditional

| | | Forecast Volatility Annual % | | | |
|---|---|---|---|---|---|
| | | | | | |
| | | Port | Bench | TE | Market |
| | | | | | |
| | Quiet | 10.97 | 10.86 | 2.65 | 11.05 |
| | Turbulent | 18.33 | 17.68 | 3.09 | 17.19 |
| | Regular | 15.57 | 15.67 | 2.63 | 15.08 |
| | Near Horizon | 21.23 | 20.85 | 2.19 | 21.31 |

# Comparative Results Discussion

- All results should be intuitive.
  - The expected volatility associated with quiet periods is materially less than the expected risk associated with turbulent periods.
  - Our regular long term forecast was between the levels forecast using the two subsamples.

- Our conditional near horizon forecast was materially higher than any of the others reflecting the influence of forward looking measures such as VIX, credit spreads and equity option volatility on the expected risk of the end of May 2020.
  - A very extensive discussion of how conditional models have functioned during the pandemic period is presented in our June 20 newsletter, https://www.northinfo.com/Documents/946.pdf

# The Trouble with Tracking Error

- A key differentiator between the "turbulent" and the current near horizon forecast (a pandemic would be presumed turbulent) is that the turbulent model forecast that benchmark relative tracking error would be higher than normal, while the conditioning of the near horizon process predicts lower than typical tracking error as correlation across stocks increases in a crisis period.

- As correlation is a bounded function (absolute value <=1) it is a mathematical artifact that we should expect higher than normal correlation across assets as volatility increases.

- This effect is widely observed during market crashes when all assets seem to move together.
  - A behavioral explanation of this phenomenon has been advanced in Falbo and Grassi (*Discrete Dynamics of Nature and Society*, 2015).

# An Unfair Test?

- One might argue that we're being a bit unfair to the regime-switching concept by splitting the historical sample of more than thirty years into two equally sized subsamples.
- It seems likely that the market conditions might be "normal" most of the time and have occasional high volatility conditions a minority of the time.
- To test this hypothesis, we construct a simple test. We take the monthly returns to our "global market portfolio index" over the 365 month sample and rank the observations by their absolute value.
- We can then calculate the mean for various subsets of the sample (e.g. most volatile 36 months).
  - The mean absolute deviation (from zero) and the usual measure of volatility (standard deviation) are very closely related as mean monthly returns are small as compared to the volatility.

# Sorted Sub-Samples of MAD

Mean Absolute Deviation for Sorted Subsamples

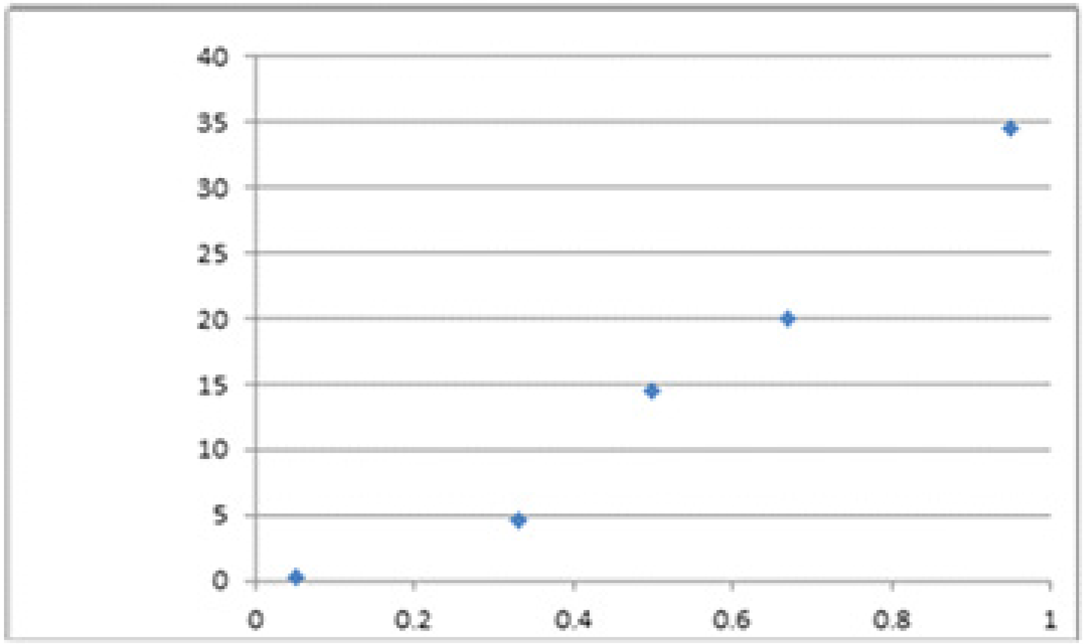

# Results of MAD on Sorted Sub-Samples

- Using mean absolute deviation, in the graph below the center dot is the *annualized* return volatility for the entire sample of over 360 months at 14.43%.
- The leftmost dot represents 36 months with the lowest magnitude returns at .31%
- The second from the left is the lower half (182 months) based on magnitude of returns at 4.69%.
- The rightmost dot is the 36 months with the highest magnitude return events at 34.48%.
- The second from the right is the upper half (183 months) at 19.96%.

# An Inference on MAD Sorted Sub-Samples

- As can be seen, one can put a rather straight line in to connect the dots, except for the lowest which is bounded to be greater than zero.

- Since the line is straight, this suggests that the distribution of global return volatility is not broken into two regimes representing distinct low and high volatility periods.
  - If there were regimes, the line should curve in the shape of an arctangent function.

- We also tested the time series of absolute market returns using a "sequential probability ratio test" for changes in regime known as CUSUM.
  - While CUSUM tests suggest that the time series could be broken into regimes the number of such breaks is not higher than might arise randomly.
  - For more information on CUSUM methods, see the Northfield newsletter from February 2005, https://www.northinfo.com/Documents/72.pdf.

# Would Daily Data Help?

- We did not evaluate subsamples smaller than 36 months as such a small sample would be grossly inadequate to estimate the high number of factors necessary to estimate global equity risk at the individual security level.

- One might be tempted to increase the number of observations by resorting to daily returns
  - However, there are a number of biases that distort risk estimation due to the asynchronous nature of "daily" returns across global time zones.
  - This issue has been well recognized for decades and there is an extensive literature such as Shanken (*Journal of Finance*, 1987), Wood and McInish (*Review of Economics and Statistics*, 1985), Perry (*Journal of Financial and Quantitative Analysis*, 1985) and Burns, Engle, and Mezrich (*Journal of Derivatives*, 1998).
  - Daily returns also typically exhibit fat tails which complicates risk assessment. See diBartolomeo (*Professional Investor*, 2007).

# MAD Autoregressive

- Another piece of evidence which is not supportive of the regime-switching concept was a simple test for auto-regressive properties in the original time series of the absolute value of global market returns.

- Over the sample of 365 observations, a statistically significant first order autocorrelation coefficient of .18 (T > 3) is present for the first lag. This indicates that the if the absolute value of the market return was higher (lower) than the long term average in the prior period, we would expect a higher (lower) than normal absolute value of the market return for the next period.

- The autocorrelation coefficients for the second and earlier lags are all approximately zero. This is result is supportive of an ARCH-GARCH type process at the monthly periodicity but not of a regime-switching process where regimes were likely to persist for lengthy periods.

# Conditional Models: 2000 Words?

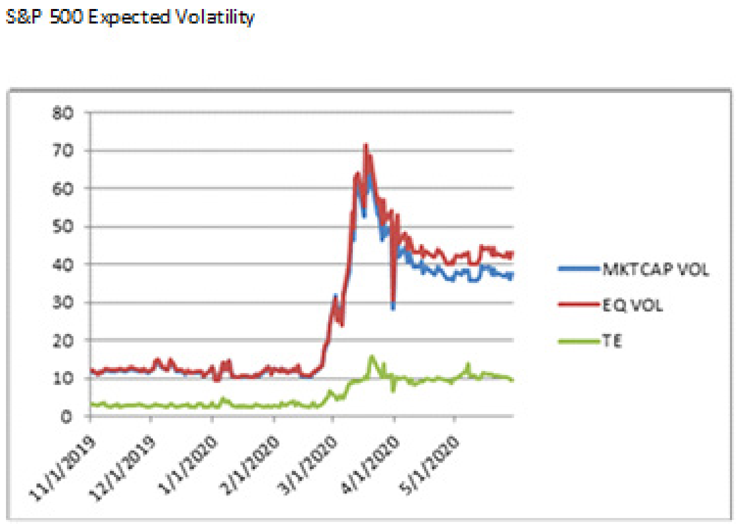


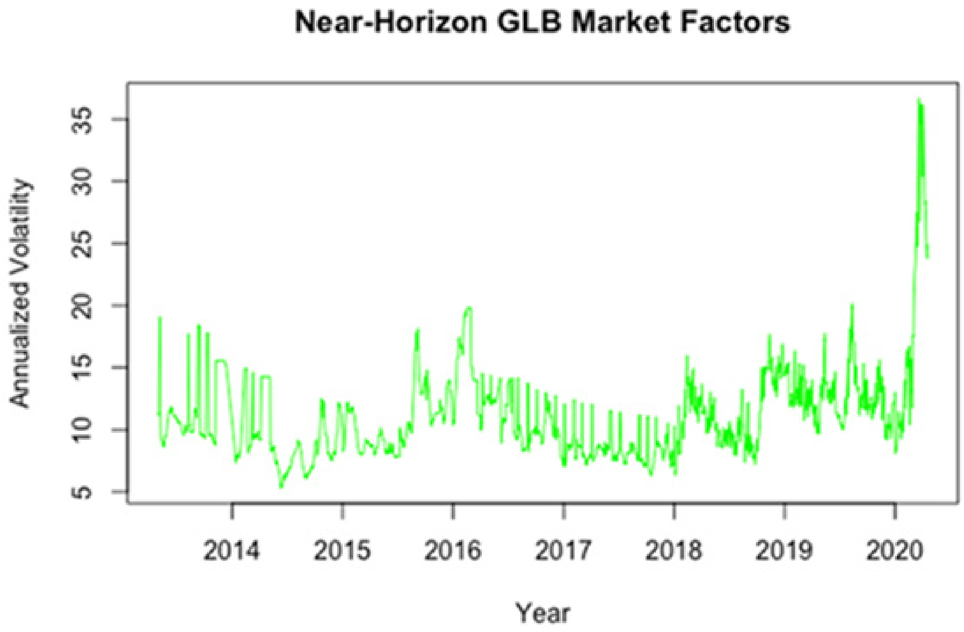


https://www.northinfo.com/Documents/946.pdf

# Conclusions

- We have considered the possibility of a regime-switching model being a viable alternative to conditional equity factor models.
- At first look, there are substantial structural impediments to the tractability of such models being applied to portfolios of individual equity securities.
- Three different types of empirical tests on three different data sets have been carried out to determine if there was substantial evidence that a regime-switching equity factor model would be a reasonable alternative to our existing conditioning processes.
  - In all three tests, support regime-switching was lacking.
- Finally, it should not be overlooked that regime-switching models require that a transition probability matrix can be effectively implemented.
  - However, if there is evidence of stable parameters for the transition probabilities, multiple models can be merged into a single model by the use of "mixture of normal" and the Cornish-Fisher (1938) methods.