Detecting Regime Changes in Financial Markets using Hidden Markov Models and Directional Changes

Devansh Joshi[1]

[1]CSAM Undergraduate
IIIT Delhi

Abstract: The purpose of this study is to construct a multivariate input based Hidden Markov model based on directional changes to detect regime changes in financial markets. For this study, a Hidden Markov Model with multivariate inputs was used. Directional changes were used on historical S&P 500 index returns, additionally, Chicago Board Options Exchange's CBOE Volatility Index along with commonly used index performance indicators were used as inputs to the Hidden Markov Model. The Hidden Markov Model with Directional Change indicators is able to effectively classify regimes on the basis of their statistical properties viz. Mean and standard deviation. The motivation of this study is to model and detect regime changes in US financial markets over the 22-year period from 2000 to 2022. Hidden Markov Models have historically been used by modelling index returns using time series analysis and realised volatility, this paper uses directional changes along with other inputs as observed states to the Hidden Markov Model like the Chicago Board Options Exchange's CBOE Volatility index, 22 and 66 day returns of the S&P 500 index in addition to the 22-day volatility of the S&P 500 index.

Keywords: HMM: Hidden Markov Model, DC: Directional Change, OS: Overshoot, TMV: Total Price Movement, R: Returns

Introduction

Financial markets can undergo capricious fluctuations in the wake of macroeconomic and political changes. Moreover, black swan events like COVID-19 and the Russian -Ukraine war induce extreme pressure on the cycle of supply and demand, thereby causing extreme market fluctuations. These fluctuations can be easy to notice in terms of the share prices of individual stocks, in terms of modelling these, however, a more potent variable to look at is the change in the statistical properties of stock returns -namely mean, variance and correlation. Regime switching models can be used to detect such sudden changes. [1]

This paper outlines a Directional Change (DC) based Hidden Markov Model originally proposed by [Tsang et al.] to fit a regime switching model on the S&P 500 index. The S&P 500 index has been modelled over the time period January 2000- July 2022 as this covers a wide gamut of market conditions and regimes. The advantage of Directional Change over Traditional time series-based models is that DC indicators can sample values at irregular intervals of time vis-à-vis time series based Hidden Markov Models that measure realised volatility over fixed intervals of time. [2][4][5][6]

The proposed DC based Hidden Markov Model modelled using multivariate inputs [defined in the section titled Data and Methodology] returns the optimal number of hidden states i.e. regimes to be 4. The model divides the 4 regimes with regime 0 having a mean and standard deviation of 3 -.000134 and 0.013091, regime 1 having a mean and standard deviation of 0.000131 and 0.027900, regime 2 having a mean and standard deviation of 0.000339 and 0.006519 and regime 3 having a mean and standard deviation of 0.000668 and 0.007111 respectively.

Literature Review

1. Directional Changes: Directional changes are used to summarise price movements over time. Unlike time series analysis that uses realised volatility over seconds, minutes, hours, days etc. DC indicators are sampled at irregular time intervals. This helps better model random transactions in financial markets that follow no regular rate like x transactions per time period. The figure below illustrates a DC event:

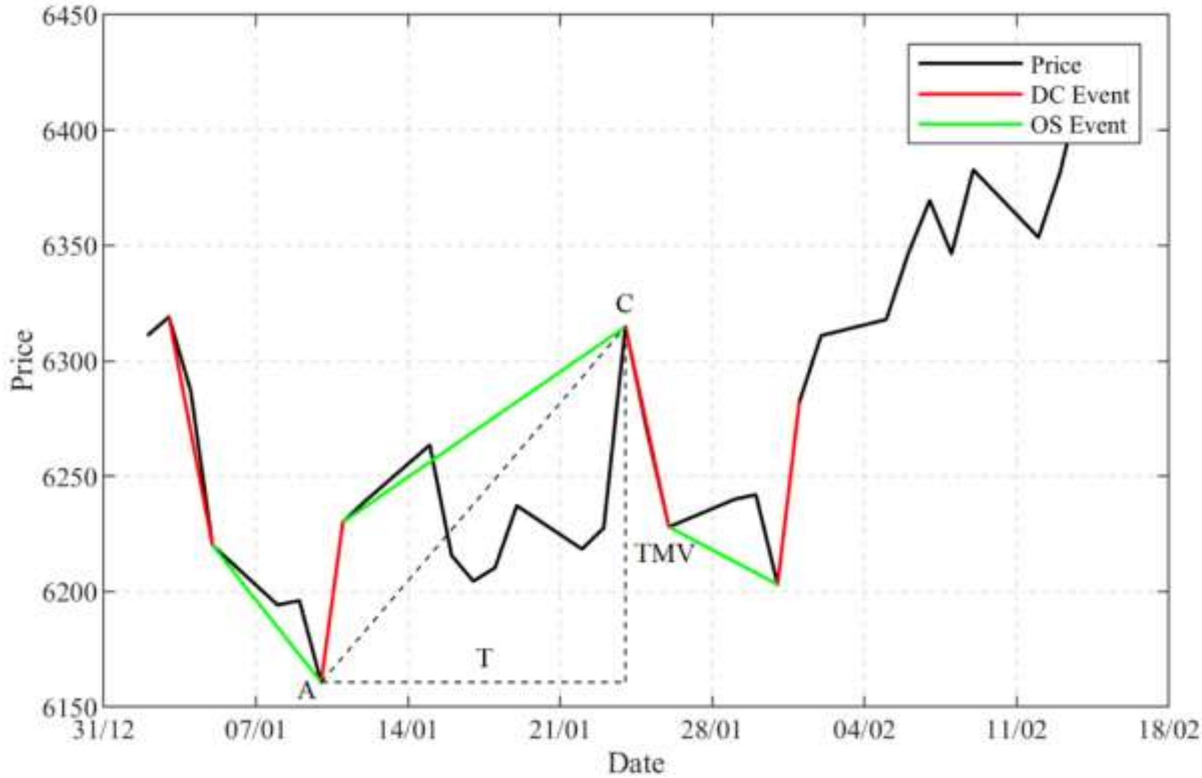

A representation of a Trend with DC and OS events Source: [2][4][5][6] [Fig 1]

Consider an asset price is on a downtrend, at point A in the figure above the price of the asset on the subsequent day increases by more than a predetermined threshold value say Q. This is a Directional Change event. The threshold determines how sensitive the model is to market fluctuations- a lower threshold would capture market fluctuations at a higher frequency and vice-versa. The choice of the threshold depends on the frequency of data available, in case of high frequency data like per second data a lower threshold might be more suited, however for relatively low frequency data like daily data, a higher threshold might be more suitable. Moreover, point A is now considered to be the extreme point (low). A DC event is followed by an OS or Overshoot Event where the asset continues its uptrend till point C before falling by more than the threshold Q. Subsequently, point C is considered to be the extreme point(high). A trend comprises a DC event and an OS event. Similarly, a trend can be an uptrend followed by a downtrend. [Tsang et al.][2][4][5][6] also defines the concept of TMV or Total Price Movement and Returns R (These returns will be later used as the observable set of states in the Hidden Markov Model which will be described in the next section.)

TMV and R have been described as follows [Tsang et al.][2][4][5][6]:

1. TMV=(|Ps−Pe|)/(Ps*Q)
2. R= (|TMV|×Q)/T

Where:
- Ps=Price at the start of the trend
- Pe=Price at the end of the trend
- Q=Threshold
- T= Time to complete a trend

2. Hidden Markov Model: Hidden Markov Models have a wide range of applications like sequence matching, signal processing, speech recognition and computational finance. Hidden Markov Models can be used effectively to detect regime changes in financial markets and in subsequent asset and portfolio allocation.[8] used a Hidden Markov Model based investment strategy using pricing trends. [7] used hidden Markov models to construct a regime switching factor investing model. A Hidden Markov Model is a Markov process and comprises a set of hidden states and visible observable states. The hidden states are influenced by the observable states and the objective of the model is to infer the hidden states from the set of observed states . A Markov chain or Markov process is a stochastic model describing a sequence of possible events in which the probability of each event depends only on the state attained in the previous event and not on any of the states attained before the previous event.

Figure 2 shows a representative diagram of a hidden Markov model:

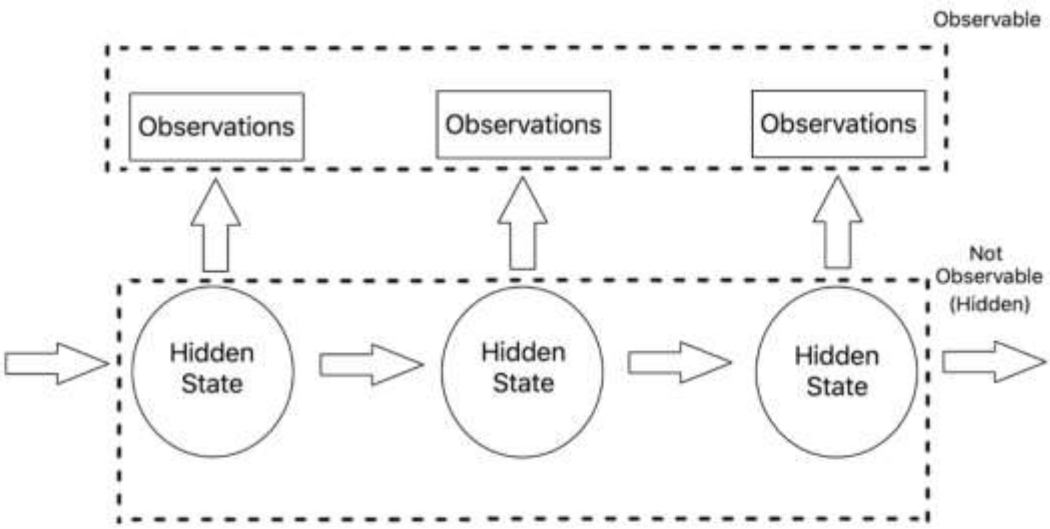

Source:[7][Fig 2]

A Hidden Markov Model or HMM comprises of the following components:[9]

A. A set of H of n hidden states h_i.

B. A transition probability matrix T, where t_{ij} corresponds to the probability of a transition from state i to state j.

C. A sequence O of observable states o_i

D. An emission probability matrix E representing the probability of observable state o_i being generated from the hidden state h_i.

E. An initial probability distribution π where each π_i represents the probability that the Markov chain will start in state h_i.

In addition to the Markov property which states that the probability of being in the current state only depends on the previous state i.e., $P(h_i|h_1...h_i-1) = P(q_i|q_i-1)$, a Hidden Markov Model also satisfies $P(o_i|h_1 ...h_i,...,h_T ,o_1,...,o_i,...,o_T) = P(o_i|h_i)$ i.e. the probability of being in observed state o_i only depends on the state h_i that produced it and not on any other states.[9]

Rabiner et al. proposed that HMM's should be utilised to address these 3 problems:[10]

A. Likelihood
B. Decoding
C. Learning

For the purpose of regime detection, the Hidden Markov Model will be used to Decode the hidden set of states i.e. the regimes using multivariate inputs i.e. observed states. 7 The Viterbi algorithm, which is a dynamic programming problem, is used for the purpose of decoding the most probable set of hidden states from the observed states.[9] For learning the parameters (transition probability matrix and emission probabilities) of the Hidden Markov Model, the Baum Welch algorithm is used which is a special case of the Expectation Maximisation algorithm.[9][11][12]

Data and Methodology

1. Data: All the data used for the purpose of this research was obtained from Yahoo Finance using python's dataReader. This includes historical S&P 500 index data and Chicago Board Options Exchange's CBOE Volatility Index which measures the expected volatility of the Stock Market based on S&P 500 index options. Furthermore, commonly used index performance indicators like 22 day returns,66 day returns and volatility of the S&P 500 index have also been used. The time period of the data is from January 2000 till July 2022. To account for gaps and discrepancies across different time series datasets, data values have been carried forward within a dataset. 2. Methodology: For implementing the Hidden Markov Model, python's hmmlearn library was used. Directional Changes were applied on historical data for the S&P 500 index. Since a trend i.e. direction change event followed by an overshoot event takes place over different time periods with a single value for R, to create daily data for R, the same value of R for a trend was taken for the entire duration of the days in the trend. The same was done for the TMV values and time taken for the completion of a trend value. The R values obtained were transformed using log transformation to model the data better and change the distribution to gaussian distribution. 8 For the purpose of this research, the GaussianHMM from the hmmlearn python library has been used. GaussianHMM uses the expectation maximisation algorithm to learn the parameters of the Hidden Markov Model like the transition probability matrix, emissions probability matrix and the initial distribution probabilities. The expectation maximisation algorithm is an iterative algorithm and might get stuck at a local optima, for this purpose, whilst fitting, the random state of the GaussianHMM needs to be changed continuously to ensure that the algorithm converges for different values of the random state. Furthermore, another consideration

with the Hidden Markov Model is that it is an unsupervised learning algorithm, as a result, the number of hidden states of the model need to be predetermined. To find the optimal number of regimes, the model is fit on the same dataset for different values for the number of hidden states. The score of each model is evaluated using model.score and appended in an array then the model with the best score is used for predictions. **Results and Discussion**

Fig 3 shows the PDF of log transformed values of R, the transformation is done to introduce symmetry so that the R values can be modelled better.

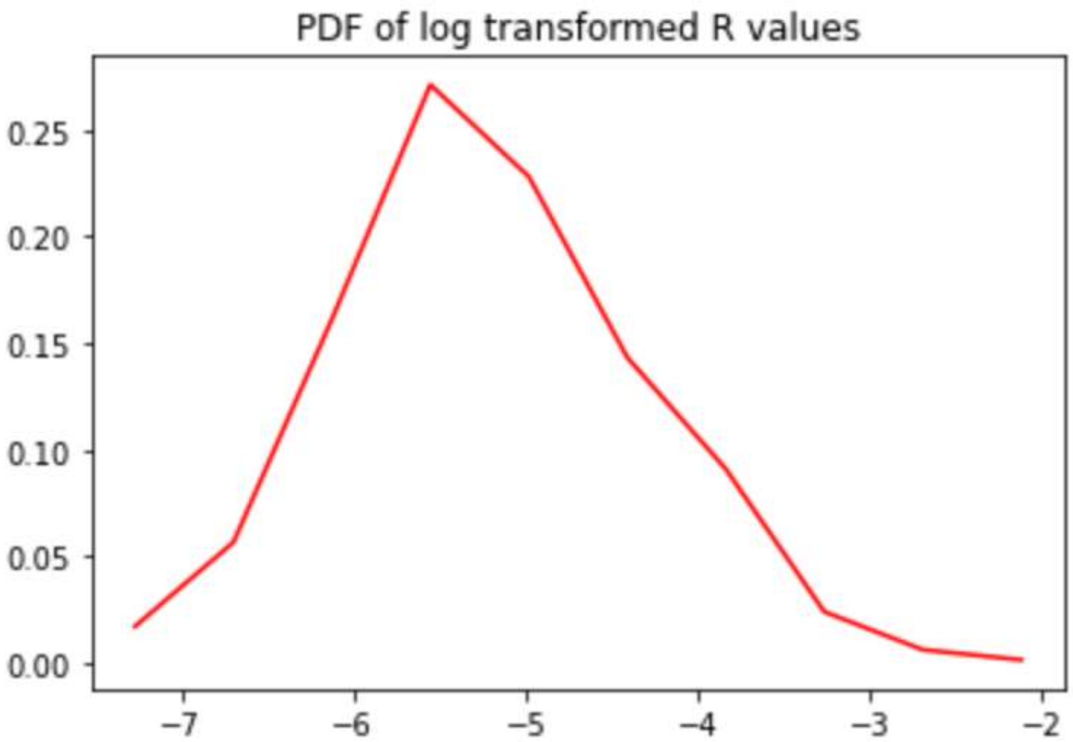

PDF of log transformed R values [Fig 3]

Fig 4 shows the values of R plotted over time, the massive spikes in 2008 and 2020 indicate the ramifications of the 2008 financial crisis and COVID-19 on the stock market and its volatility.

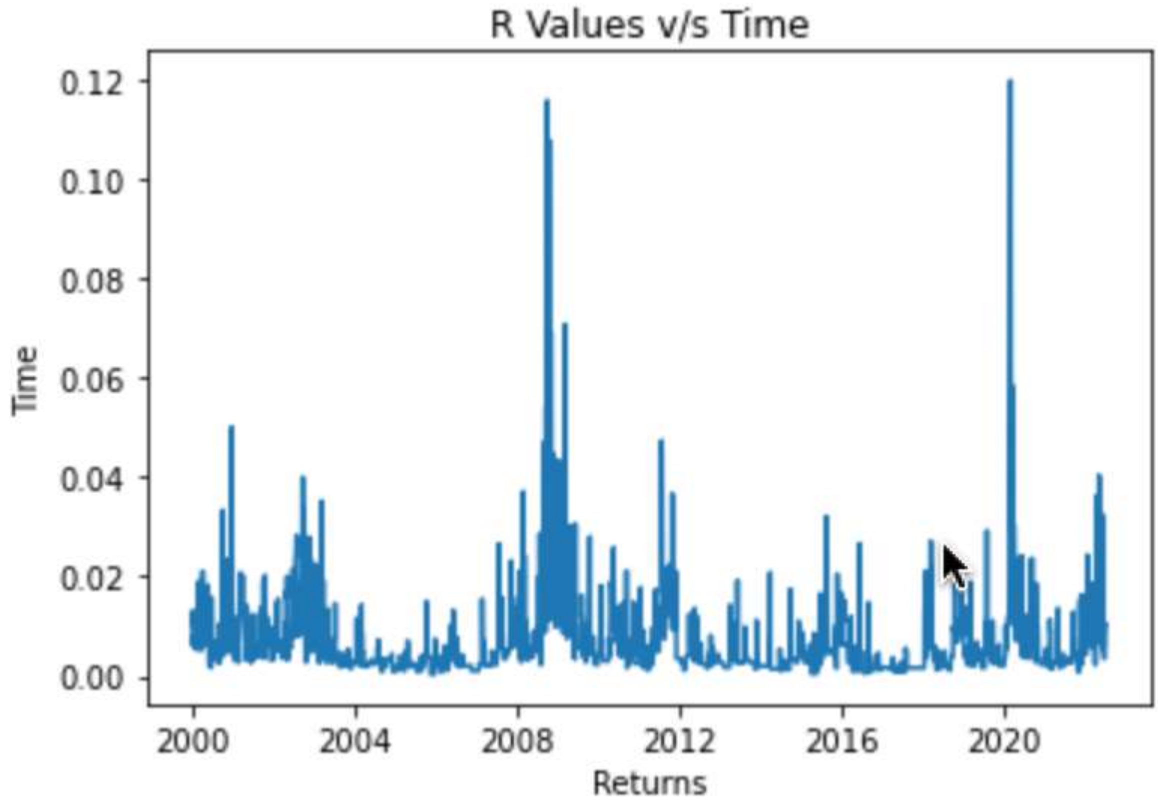

R values v/s time [Fig 4]

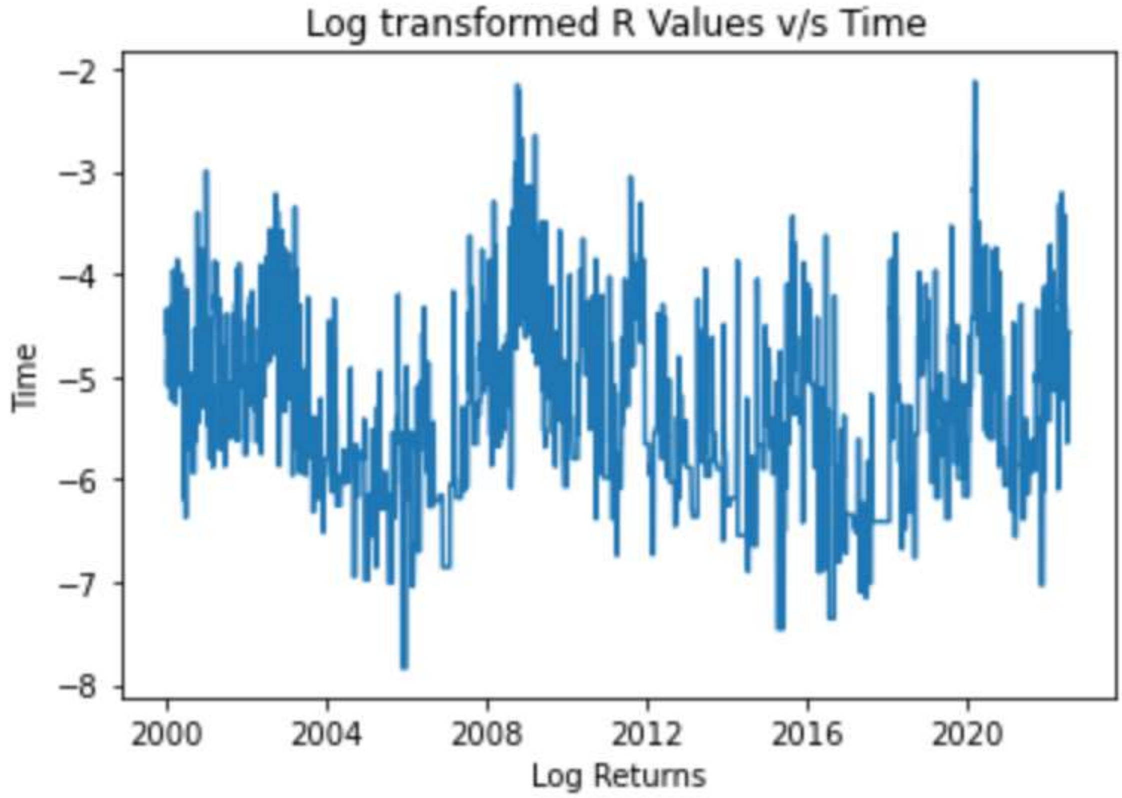

Log transformed R values v/s time [Fig 5]

For the number of hidden states equal to 4, the model gives the best score, thus the optimal number of regimes is 4.

Number of regimes	Score
1	-17330.43753719413
2	-3330.3621706324507
3	1330.6377142413733
4	3290.299785451294

[Table 1]

[Tsang et al.][2][4][5][6] visualised market regimes in an indicator space by plotting the normalised average value of Time and normalised average value of TMV. Fig 6 below visualises the 4 market regimes based on their normalised average value of Time(y-axis) and normalised average value of TMV(x-axis).Fig 6 clearly shows a distinct separation among the 4 different regimes.

Normalisation has been done as per the formula:

x(i)=x(i)-min/(max-min) where max and min denote the maximum and minimum values of TMV and time for each of the 4 regimes.

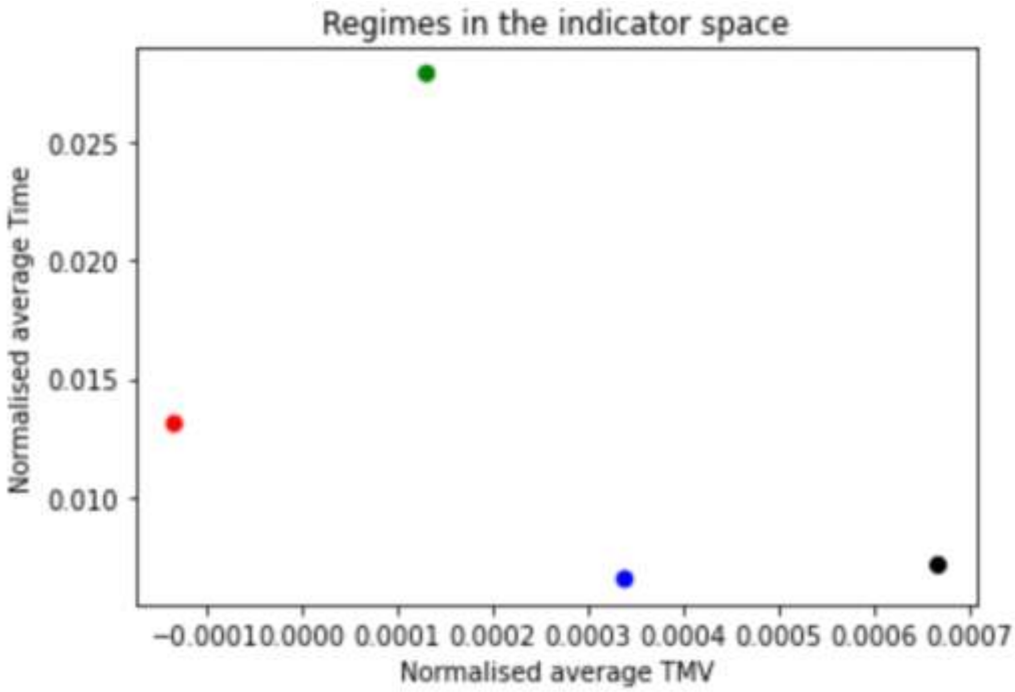

Representation of regimes in the indicator space [Fig 6]

The value of the threshold for the directional change indicator R has been chosen using trial and error at 0.01 for optimal separation of the 4 regimes. Fig 7 below plots the mean(x-axis) vs std deviation(y-axis) values of the daily returns for all the 4 regimes.

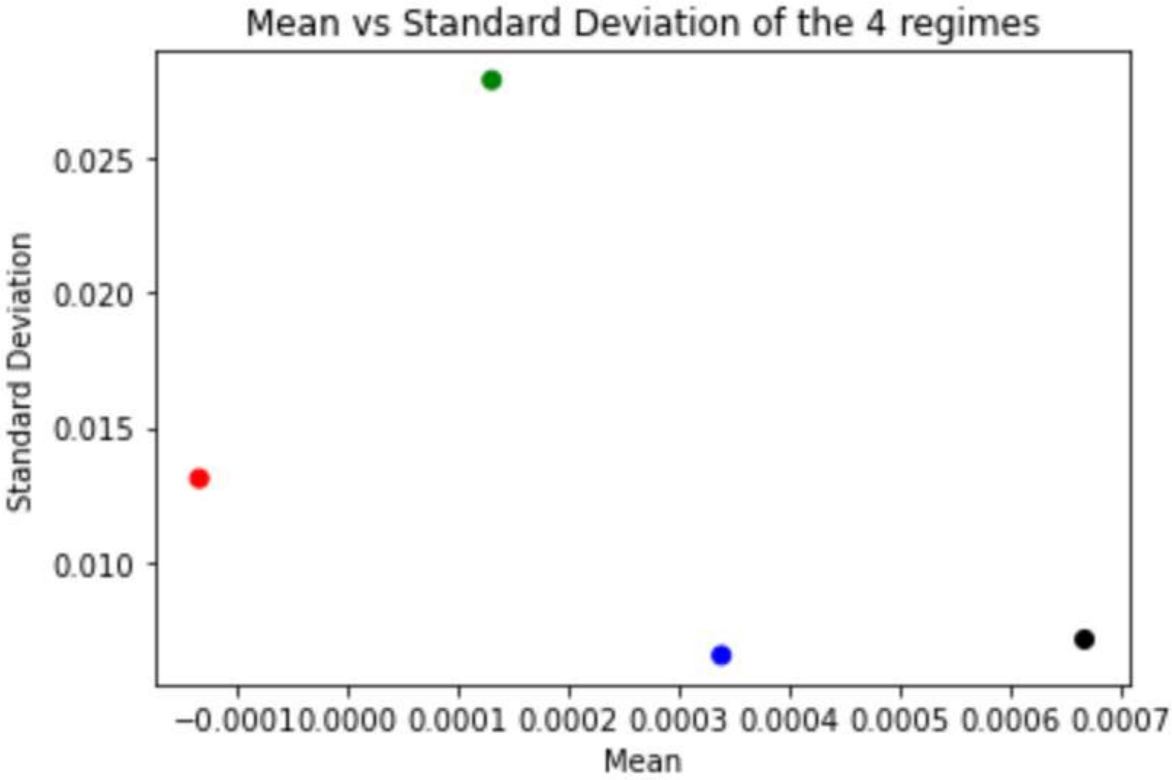

Mean vs Standard Deviation of the 4 regimes [Fig 7]

Regimes when the mean daily return is high and the standard deviation is low are better for investing whereas those with low mean returns and high standard deviations are not suited to investing. The table below shows the values for mean and standard of the daily returns for all 4 regimes. Regime 0 is the worst as it has negative mean daily returns and the second highest standard deviation whereas Regime 3 is the best for investing in equity as it has the highest mean returns and second least standard deviation.

Regime	Mean	Standard Deviation
0	-0.000134	0.013091
1	0.000131	0.027900
2	0.000339	0.006519
3	0.000668	0.007111

[Table 2]

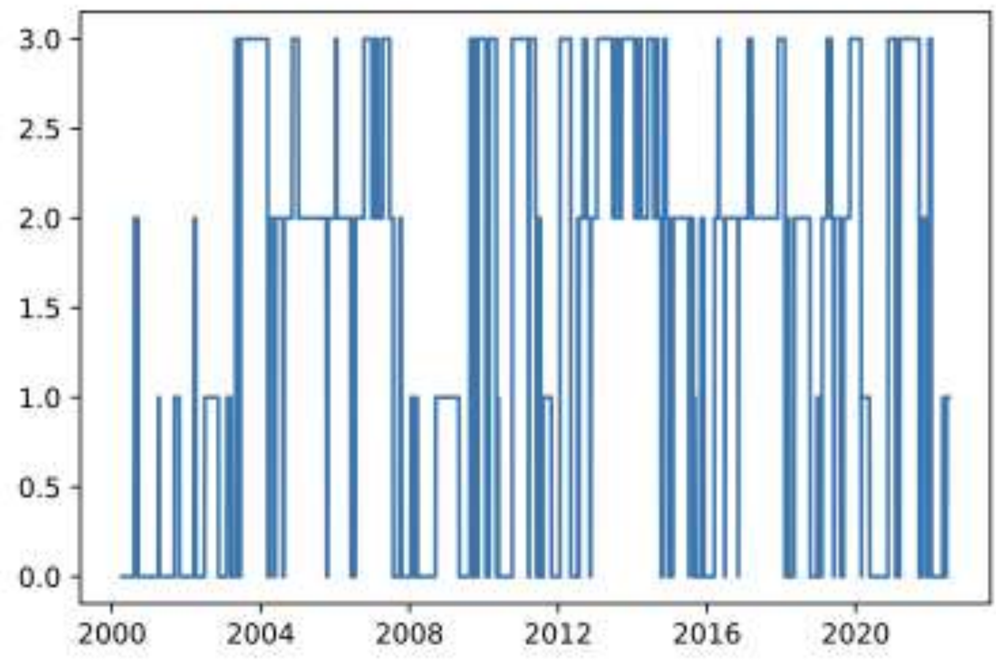

Regime prediction for historical data [Fig 8]

Fig 8 depicts the predicted regime for dates spanning from January 2000 till July 2022 based on the Viterbi algorithm which predicts the most likely set of hidden states i.e regimes using the maximum a posteriori probability estimates of the regimes.

Conclusion

As shown by the results, a Hidden Markov Model using directional changes can determine regimes during different macroeconomic and political conditions. The model is able to differentiate regimes suitably based on mean and standard deviation of the daily returns of the close price of the S&P 500 index in the 4 different regimes. Moreover, the model is able to differentiate between the 4 regimes when plotted on the indicator space that depicts normalised average TMV values versus the normalised average time values.

Additionally, there are further developments that can be made to the model- high frequency macroeconomic data, housing data, growth statistics, employment statistics etc. can also be incorporated as observable states as a part of the multivariate input to the Hidden Markov Model, however, due to limited computing power and lack of open source data, these could not be incorporated in this study.

References

1) Ang, Andrew and Timmermann, Allan, Regime Changes and Financial Markets (June 20, 2011). Netspar Discussion Paper No. 06/2011-068, Available at SSRN: https://ssrn.com/abstract=1919497or http://dx.doi.org/10.2139/ssrn.1919497

2) Chen, J.; Tsang, E.P.K. Classification of Normal and Abnormal Regimes in Financial Markets. *Algorithms* **2018**, *11*, 202. https://doi.org/10.3390/a11120202

3) Hamilton, James D. "A New Approach to the Economic Analysis of Nonstationary Time Series and the Business Cycle." *Econometrica*57, no. 2 (1989): 357–84. https://doi.org/10.2307/1912559.

4) Aloud, Monira & Tsang, Edward & Olsen, Richard & Dupuis, Alexandre. (2011). A Directional-Change Event Approach for Studying Financial Time Series. Economics. 6. 10.2139/ssrn.1973471.
5) https://www.bracil.net/teaching/CFE/Notes/DirectionalChanges-slides.pdf
6) Tsang, E.P.K.; Tao, R.; Serguieva, A.; Ma, S. Profiling high-frequency equity price movements in directional changes. Quant. Finance. **2017**, 17, 217–225.
7) Wang, M.; Lin, Y.-H.; Mikhelson, I. Regime-Switching Factor Investing with Hidden Markov Models. *J. Risk Financial Manag.* **2020**, *13*, 311. https://doi.org/10.3390/jrfm13120311
8) Kim, E.-c.; Jeong, H.-w.; Lee, N.-y. Global Asset Allocation Strategy Using a Hidden Markov Model. *J. Risk Financial Manag.***2019**, *12*, 168. https://doi.org/10.3390/jrfm12040168
9) Jurafsky, Daniel & Martin, James. (2008). Speech and Language Processing: An Introduction to Natural Language Processing, Computational Linguistics, and Speech Recognition.
10) Rabiner, Lawrence R. "A tutorial on hidden Markov models and selected applications in speech recognition." *Proceedings of the IEEE* 77.2 (1989): 257-286.
11) Baum, Leonard E.. "An Inequality and Associated Maximization Technique in Statistical Estimation for Probabilistic Functions of Markov Processes." Paper presented at the meeting of the Inequalities III: Proceedings of the Third Symposium on Inequalities, University of California, Los Angeles, 1972.
12) Dempster, A. P., N. M. Laird, and D. B. Rubin. "Maximum Likelihood from Incomplete Data via the EM Algorithm." *Journal of the Royal Statistical Society. Series B (Methodological)* 39, no. 1 (1977): 1–38. http://www.jstor.org/stable/2984875.