Using Fuzzy Logic to More Accurately Test Technical Analysis

Benjamin Reinach

Technical analysis is a field of study in finance that attempts to predict future price movements of securities by analyzing past market statistics such as price and volume. Technical analysts believe that there is useful information contained in past market statistics that, if properly analyzed, can provide significant information about future price movements. The individual tools used to analyze past market statistics are called indicators. There are a wide variety of indicators contained within technical analysis that are believed to provide some predictive abilities. Indicators usually fall into one of four categories: chart patterns, momentum indicators, trend indicators, and volatility indicators. Many technical analysts believe market prices often move in recognizable and repeated patterns. Observing the formation of these specific chart patterns is used as an indicator to help traders predict future price movements. Momentum indicators such as Relative Strength Index or Stochastic Oscillator measure the speed with which prices are moving. Technical analysts believe that these indicators can help signal a trend reversal in prices, as they expect price movements to slow down before changing direction. Some indicators, such as moving averages, are used by traders to help determine price trends. Another popular indicator is Bollinger bands, which measures volatility in prices and is used in a variety of ways. Technical trading strategies usually combine multiple indicators from different indicator classes in an attempt to reduce noise from trading signals.

Technical analysis is an interesting topic to study for many reasons. First, if any indicators or technical models are found to have significant predictive ability they would impact ideal investing and trading strategies. More importantly, technical analysis stands in harsh contradiction with almost all mainstream financial theory, such as the efficient market hypothesis and random walk theory. Random walk theory states that prices move randomly, and no information contained in past price movements can be used to predict future prices. The related efficient market hypothesis states that current prices already reflect all currently publically available information, and no analysis of any kind can provide traders with additional, profitable information. Together they imply that the ideal investing strategy is a passive buy-and-hold approach of an index, as there is no way to consistently "beat the market." If new information changes prices very quickly as the Efficient Market Hypothesis suggests, it wouldn't be possible to earn excess profits by timing price movements. Yet technical analysis attempts to do just that. Rejecting the implications of the Random Walk Theory and the Efficient Market Hypothesis, technical analysts attempt to both measure predictive patterns in past prices, and time future price movements to earn excess profits. Due to its contradiction with widely accepted financial theories, technical analysis is a black sheep in finance. Academics and large financial institutions alike criticize technical analysis, comparing it to astrology. Anecdotal reports of its success are quickly brushed off with explanations that "even a broken clock is correct twice a day." Yet despite the harsh criticism, technical analysis has been widely used for as long as financial markets have existed.[1] Even today, many individual traders and institutions alike rely in some part on technical analysis to make trading decisions.

Discrediting prevailing financial theories and finding ways to earn excess profits has motivated many researchers to study technical analysis over the years. Yet despite thousands of studies, there has yet to be conclusive evidence one way or the other that technical analysis absolutely works. Many studies have been conducted on the predictive ability of individual indicators, with varying results. At times, these studies can be criticized for “data-snooping.” Data-snooping involves looking at past price data before conducting the experiment, which can imply that the researcher specifically designed the experiment to have biased positive results. To compensate for data-snooping, experiments should be conducted over multiple time periods to determine if an indicator or model consistently provides incremental information. Another issue with studying technical analysis stems from the large number and variation of indicators. Due to sheer volume, it would take substantial time and effort to completely study technical analysis as a whole. A further issue with studying technical analysis is that many aspects of technical analysis are extremely subjective. The same chart patterns or indicator values can be interpreted in different ways by individual traders. Many indicators also operate on a spectrum, never giving a clear “buy” or “sell” signal further contributing to the subjectivity of technical analysis. Professional technical analysts actively monitor multiple indicators to determine buying and selling opportunities, so studying the ability of one indicator on its own does not accurately reflect how it is used in practice. Lastly, just because one indicator is shown to have worked in one study, it does not definitively imply that it will work in the future or in different markets. Despite the difficulties, the opportunity to find new ways to increase profits and disprove important financial theories provides sufficient motivation for studying technical analysis.

I. LITERATURE REVIEW

Early studies on technical analysis focused on researching one indicator at a time. Researchers would back-test the indicator on an index to determine its accuracy in predicting price movements. While such studies have been conducted for over a hundred years, one of the earliest studies that remains relevant was conducted by Brock et al in 1992.[1] In it, they study twenty-six variations of the two simplest and most popular indicators: moving averages and trading range break. These indicators were backtested on the Dow Jones Industrial Average over 100 years, from 1887 to 1986. They found that returns following buy signals were consistently positive, albeit only slightly so, while returns following sell signals were consistently negative. Annual average returns following buy signals were 12%, while annual average returns following sell signals were -7%. These findings provide strong support for the two strategies tested, especially considering the long time frame they were tested on. A follow up study conducted by Sullivan et al in 1997 determined that Brock's findings were not guilty of data-snooping.[2] However, another follow up study performed by Bessembinder and Chan in 1998 determined that when compensating for trading costs, the strategies analyzed by Brock were not significantly more profitable than a buy and hold approach.[3] This implies that while the technical indicators might have some predictive ability, they were unable to disprove the weak form of the efficient market hypothesis.

Another study conducted by Lo et al in 2000 analyzed chart patterns over 31 years, from 1962 until 1996 on the NYSE.[4] Lo's study implemented a computer model to automatically detect common technical chart patterns. His results found that many patterns did in fact provide incremental information. While Lo stopped short of making any claims about their profitability, he did assert that chart patterns could be useful in the investing process. Studies on individual technical indicators have been so numerous that it is impossible to mention them all, but Irwin and Park wrote a thorough summary of such studies in 2007.[5] Compiling information from 95 modern studies on technical analysis, Irwin and Park found that 56 had positive results, 20 negative results, and 19 mixed. They do admit that many of the studies with positive results could be guilty of data-snooping and again stopped short of making claims on the profitability of these indicators. Overall, previous studies on individual indicators have yielded encouraging yet inconclusive results on the predictive power and profitability of technical analysis.

Other studies have taken a different approach and instead try to determine the profitability of technical analysis by studying the performance of traders that rely purely on technical strategies. One such study was performed by Murphy in 1986.[6] In it he looks at the performance of 11 futures funds that base trading decisions solely on technical signals. Looking at returns from 1980 to 1985, Murphy concluded that the 11 technical futures funds did not provide a significant improvement in profitability when accounting for fees and expenses. With arbitrarily low fees and expenses, however, Murphy concluded that the technical trading strategies could provide a statistically significant advantage. A 1994 study by Blume et al claimed that traders who used information contained in market statistics such as price and volume performed better than those that did not.[7] Some studies have even conducted trading simulations to research the performance of technical analysts. Yet no study that relies on performance can conclusively tell us about the profitability of technical analysis as many omitted variables exist.

Studies that combine indicators to develop buy and sell signals more accurately represent how technical analysis is used in practice. This is exactly what Elaine Loh did in her 2007 study. Loh used a model that consisted of two common indicators: moving averages and a stochastic oscillator.[8] This model was then backtested on 5 Asian Pacific stock exchanges from 1990 to 2005. She determined that while moving averages were only accurate 50% of the time, the combined method was accurate almost 75% of the time. This shows that combining indicators can reduce noise and result in a more accurate trading strategy. Loh stops short of making comments on the profitability of such as strategy, admitting that transaction costs vary from trader to trader, yet the findings still offer strong support for technical analysis. Questions still remain: Would these findings remain constant in American stock exchanges? Would adding more indicators further reduce noise and improve accuracy?

One intriguing method that can combine multiple indicators to study technical analysis is called fuzzy logic. Fuzzy logic is method of computing that operates on “degrees of truth.” Rather than having binary inputs and outputs of true and false, fuzzy logic operates on a spectrum. This is useful for studying technical analysis because many technical indicators do not give strict “buy” or “sell” signals but rather give signals that still contain uncertainty. Fuzzy logic models have the added benefit of being able to incorporate many indicators. This makes it possible to study technical analysis in a way that closely mimics how it is actually used. In a 2000 study, Dourra et al uses fuzzy logic to research technical analysis.[9] Their model used three indicators: rate of change, a stochastic oscillator, and support and resistance. In this fuzzy logic model, the strength of the three indicators were all considered and combined to create trading signals. Unfortunately, the model was not backtested well. Dourra’s study only tested the model on four individual stocks over a two year period. The results were astoundingly positive, with returns almost double a buy and hold approach. These results are not significant because of the small sample size. However, Dourra’s study did introduce an exciting new way to study technical analysis.

Other researchers have also utilized fuzzy logic to more effectively study technical analysis. In a 2004 follow up to Lo et al’s study of chart patterns, Zhou and Dong used a fuzzy logic model to more accurately identify chart patterns and test their profitability.[10] They claim their method shows chart patterns can be profitable, something Lo fell short of proving. Recently, fuzzy logic has been applied to technical analysis in a variety of ways. It has been used to create new technical indicators.[11] Naranjo et al in 2015 used a fuzzy logic model to not only define a trading strategy but also to determine optimal capital allocation.[12] It has even been used to increase efficiency in high-frequency trading.[13]

While many researchers have studied technical analysis, I believe that there is much more that is yet to be determined. Studies of individual indicators have at times produced optimistic results on their predictive ability. Yet technical analysts usually consider many indicators when determining buying and selling opportunities. Fuzzy Logic appears to be a promising tool for analyzing technical analysis because it allows researchers to determine the profitability of models incorporating multiple indicators and accounting for uncertainty in signals, effectively mimicking the traders' strategies. Yet until now, no fuzzy logic model has been tested long term on US stock markets.

II. MODEL

The model used in this study incorporates three of the most popular technical indicators: moving averages, stochastic oscillator, and Bollinger bands. These three indicators were chosen not only for their popularity but also because they each come from a different category, meaning they measure something unique. Technical analysts utilize multiple indicators like this in their trading strategies attempting to reduce noise and obtain clearer signals. This model imitates how an actual technical analyst makes trading decisions, resulting in a more realistic study of technical analysis.

A fuzzy logic model is comprised of four components. The first is inputs, which are derived from the three indicators chosen. These inputs are fed through membership functions which translate the numeric input into a linguistic variable that describes the relative value of each input.

Then the model has rules which determine how the relative values of the inputs affects the output. The output itself has membership functions which translate the combination of inputs and rules into a single number which can then be used to determine trading signals. By allowing us to combine multiple inputs and multiple rules associated with these inputs, fuzzy logic provides a platform to model a trader's perspective when analyzing multiple indicators.

A. INPUTS

The first indicator used in the model is a moving average. Moving averages are among the simplest and most popular technical indicators. They have also been studied extensively with somewhat encouraging results, most notably in the 1992 study by Brock et al. A simple moving average is just the average of the previous "n" closing prices, with the oldest closing price being replaced by the most recent closing price each day. In this way, moving averages slowly trace market movements. A shorter moving average, comprised of a smaller "n" value, traces the market closely. A longer moving average with a larger "n" value will adjust more slowly to current market movements. The most common trading strategy involving moving averages comes from using two moving averages together, one relatively short and one relatively long. When the shorter moving average "crosses-over" the longer moving average, it is thought to signal the beginning or the end of a trend. For example, when the shorter moving average crosses above the longer moving average, a buy signal is given and the market is thought to be trending up. Likewise, when the shorter moving average crosses below the longer moving average, a sell signal is given and the up-trend is considered over, replaced by a down-trend. In this way, moving averages fall into the category of trend indicators.

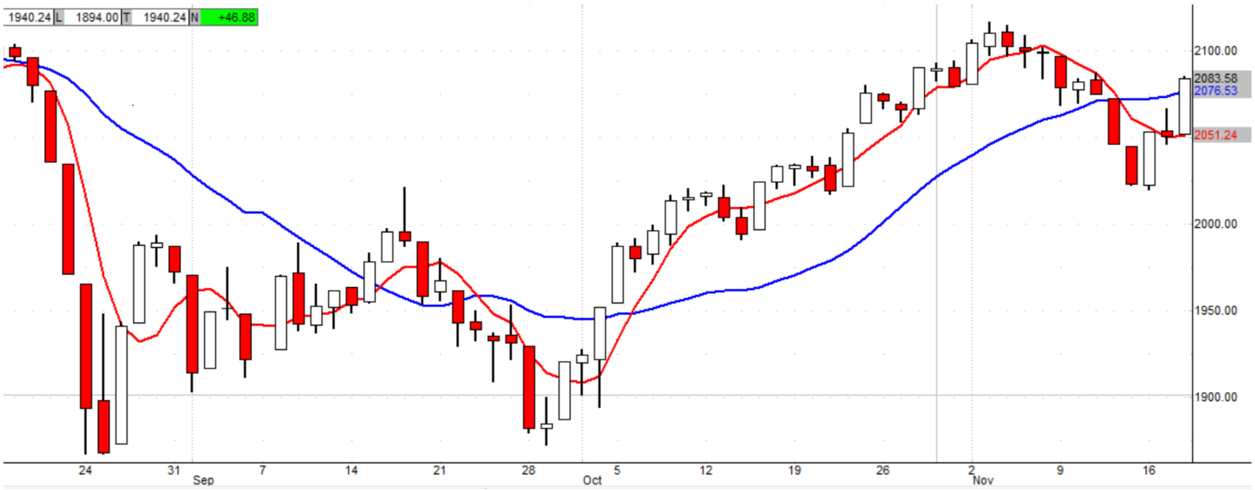

The above figure contains a short moving average (red) and a long moving average (blue).

Buy signals are given after the market has already started a move up, and sell signals are given after the market has started moving down. For this reason, moving averages lag behind the market. When the market is trending in one direction for a while, the moving average can create profitable buy and sell signals, as seen on the right half of the chart above. However, when the market moves sideways, or up and down without any clear direction, moving averages will provide misleading signals, as seen between September 14 and October 5 in the above chart. These false signals are called whipsaws. A whipsaw occurs when you have multiple crossovers in a short period of time, leading to buy and sell signals that are not accurately predicting a trend. Extending the length of the moving averages will lead to fewer whipsaws, as it will take longer for the short moving average to cross the long, however it will also delay the buy and sell signals allowing you to profit on a smaller portion of an actual trend. This tradeoff creates ambiguity in the effectiveness of moving averages.

This model uses a short moving average of 5 days and a long moving average of 20 days. They are calculated as follows:

$$5\ day\ moving\ average = SMA = \sum_{n-5}^{n} \frac{price(n)}{5}$$

$$20\ day\ moving\ average = LMA = \sum_{n-20}^{n} \frac{price(n)}{20}$$

In the above formulas, price(n) represents the closing price on day "n". The first input for the model is created by subtracting the short moving average from the long moving average. This indicates where the two moving averages are in relation to each other. This input was also price adjusted to normalize the distribution of values across many years of price data. Fuzzy logic requires us to define the values of inputs relative to their range. A relatively large divergence between the short and long moving averages in the 1980's when the S&P 500 was trading around 250 points would only appear to be a small divergence in 2015, when the S&P 500 was as high as 2100 points. Dividing by the current closing price fixes this issue and creates a normal distribution of values across a long time period. The formula for our first input is given below:

$$Y_1 = MA = 100 * \frac{SMA - LMA}{PRICE(n)}$$

This input, denoted MA, is positive when the short moving average is above the long, and negative when the short moving average is below the long.

The second indicator used is called a stochastic oscillator. A stochastic oscillator is a momentum indicator that compares current prices to the price range over a given time period, in this case the previous 20 trading days. This indicator was developed in the 1950's by George C Lane and has stood the test of time for many technical analysts, as it is still extremely popular today. The stochastic oscillator contains two components. The first, which is denoted by %K, is calculated as follows:

$$\%K = 100[\frac{C - L20}{H20 - L20}]$$

In the above formula, C represents the most recent closing price, while H20 and L20 denote the high and low of the previous 20 trading days respectively. The second component of the stochastic oscillator is a simple moving average of %K, denoted %D. In our model we use a 5 day moving average of %K to determine %D. The formula for %D, which is our second input, is given below:

$$Y_2 = \%D = \sum_{n-5}^{n} \frac{\%K(nT)}{5}$$

By computing the five day moving average we reduce the oscillator's sensitivity to short term market movements. We also derive a third input from this indicator which is simply the difference between %K and %D, indicating the direction the stochastic oscillator has been recently moving.

$$Y_3 = SO = \%K - \%D$$

The stochastic oscillator operates under the assumption that during an upward trending market, closing prices tend to be close to the high of the move, and vice versa for a downward trend. In this way the stochastic oscillator can be used to confirm a trend. Trading signals are given when %K converges or diverges from its 5 day moving average %D. When %K is above %D and Y_3 is positive, this signals upward momentum, and vice versa for when %K is below %D. For this reason, the stochastic oscillator is considered a momentum indicator. Because of the way it is calculated, the stochastic oscillator has a set range of 0 to 100, and it is not necessary to adjust for large variations in prices as it already has a normal distribution over time.

There are two trading rules derived from %D on its own. The first occurs when %D is over 80 or below 20. These levels signal when the market is "overbought" or "oversold". In other words, if %D reaches a value of 80 after a long up move, the market is thought to be "overbought" and prices are expected to soon decline. However, it is not uncommon for %D to take extremely high values for an extended period of time during a strong up move, so this signal needs to be combined with another confirming indicator to produce reliable signals.

The other commonly used trading rule associated with just %D occurs when %D diverges from the current prices. For example, if prices continue to move higher while %D stays flat or moves downward, momentum is thought to be slowing and it could signal an end to the current upward move. Normally, this type of signal is obtained by graphically comparing the levels of %D and prices. However, an input was created for this model that attempts to measure divergence between %D and price. It is calculated by subtracting a 10 day linear regression of %D from a 10 day linear regression of closing prices.

$$Y_4 = DPD = \frac{Price(n) - Price(n-10)}{10} - \frac{\%D(n) - \%D(n-10)}{10}$$

The third indicator used in our model is called Bollinger Bands. Developed by famous financial analyst John Bollinger in the early 80's, Bollinger Bands follow a common practice of using short term volatility to track the current range of prices. They quickly became a popular indicator. There are two Bollinger Bands, one upper and one lower. They are calculated by adding and subtracting two standard deviations from the 20 day moving average. The upper band is thought to act as resistance to an upward moving market, meaning once prices reach the upper band they are expected to stop moving upward. Likewise, the lower band acts as support, as prices are expected to rebound off the lower band after a long down trend. These signals are ambiguous, however, as often times prices remain close to the upper or lower band for a while during a strong trend. For this reason, trading signals need to be confirmed with another indicator to be seen as reliable. John Bollinger himself recommended this when he first introduced these bands.

The above figure shows trends reversing after hitting the upper or lower band. It is also apparent that at times prices can remain on one of the bands while continuing their trend.

Our final two inputs measure how close the current price is to the upper and lower bands respectively. These inputs were also price adjusted to create a normal distribution of relative levels over time. The formulas for these inputs are as follows:

$$Y_5 = RES = 100 * \frac{\left(\sum_{n-20}^{n} \frac{Price(n)}{20}\right) + 2 * STDEV}{Price(n)}$$

$$Y_6 = SUP = 100 * \frac{\left(\sum_{n-20}^{n} \frac{Price(n)}{20}\right) - 2 * STDEV}{Price(n)}$$

In the above formulas, STDEV denotes the standard deviation of the last 20 closing prices.

B. MEMBERSHIP FUNCTIONS

Membership functions are the critical component of fuzzy logic, allowing us to define the input and output values on a spectrum rather than in a binary true or false. Membership functions relate the numeric value of the inputs to a linguistic variable, or word. This word is then incorporated in the rules which determines how the input values affect the output value.

Membership functions are curves that associate the value an input takes with the signal generated by that value. Each input has its own unique membership functions that translate its values to trading signals. Most of the trading signals are either “Buy” or “Sell,” however there is one “Hold” signal in our model as well. The membership functions are really what make this model “fuzzy.” A certain value for an input does not need to be strictly interpreted as a “Buy” or a “Sell” signal, but rather could be somewhere in the middle. Some values could be a weak buy or a weak sell or not generate a signal at all. In this way, the model effectively mimics the uncertainty faced by traders’ using technical analysis. When traders view these indicators, they develop thoughts based on the indicators’ values that help them determine the opportunities for a trade. However these thoughts contain uncertainty, which is incorporated into this model. It is almost impossible to replicate and automate this mental process, but fuzzy logic allows a good approximation of this action. The membership functions of each input are shown below, along with a brief explanation.

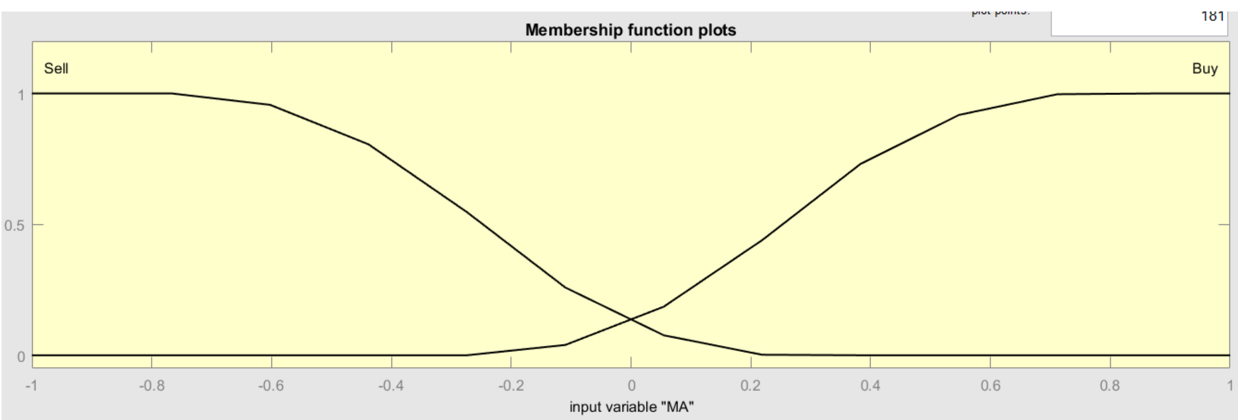

The figure above shows the two membership functions associated with our first input, MA. When MA is positive, the short moving average is above the long moving average, which is represented by the "buy" membership function. The membership function slowly increases as MA gets more positive, meaning that a definitive crossover will be a stronger buy signal than a very slight cross over. The "sell" membership function directly mirrors the "buy." These two membership functions extend to the high and low of the range of values MA takes. However after a certain point, extreme values are simply treated as definitive cross-overs.

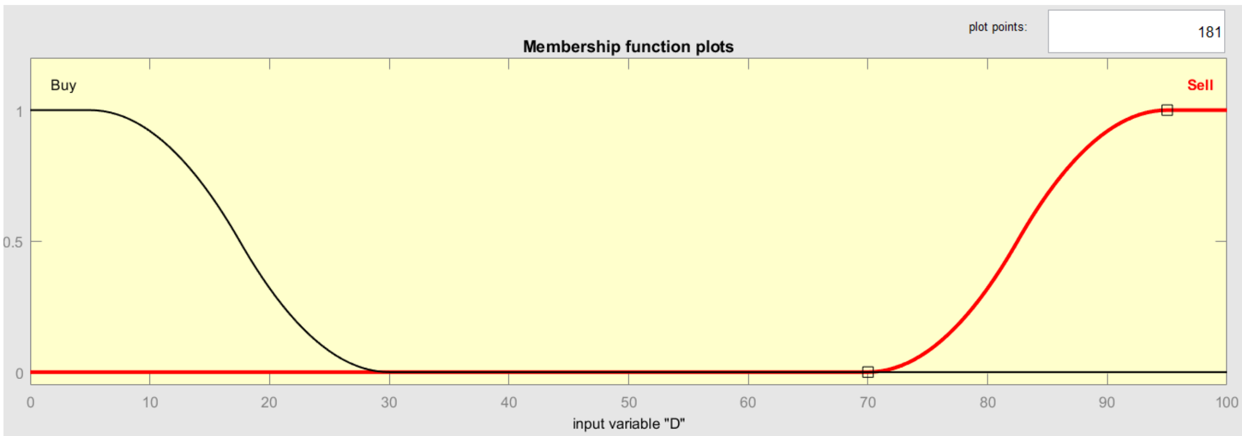

The above figure shows the membership functions for our second input, %D. Here we have buy signals generated as %D reaches extreme lows, and sell signals generated as %D reaches extreme highs. This allows us to incorporate the concept of %D indicating an "overbought" or "oversold" market.

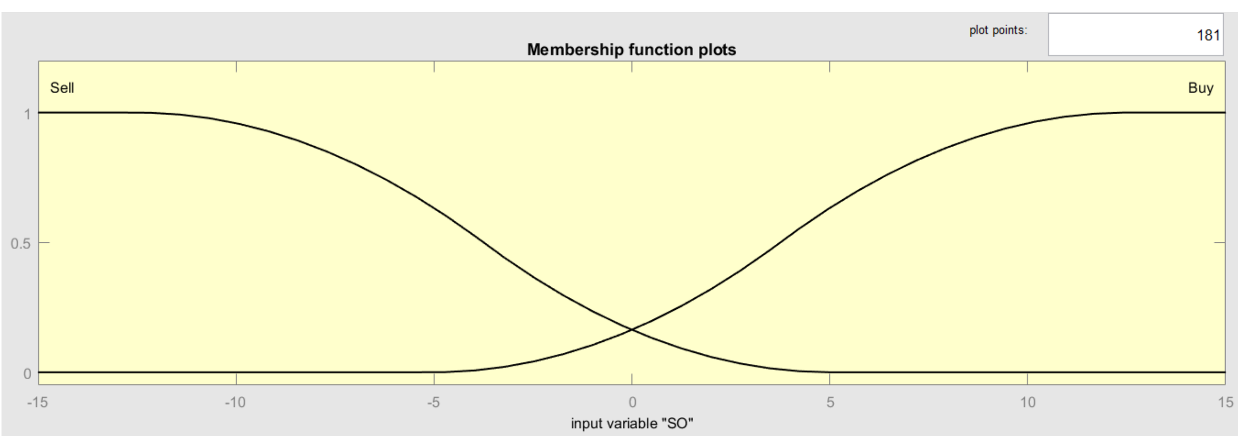

The above figure contains the membership functions associated with our third input, SO. Again we have buy and sell signals that mirror each other. These signals represent a positive or negative cross-over of %K and %D. The strength of the signals grows as the cross-over becomes more distinct.

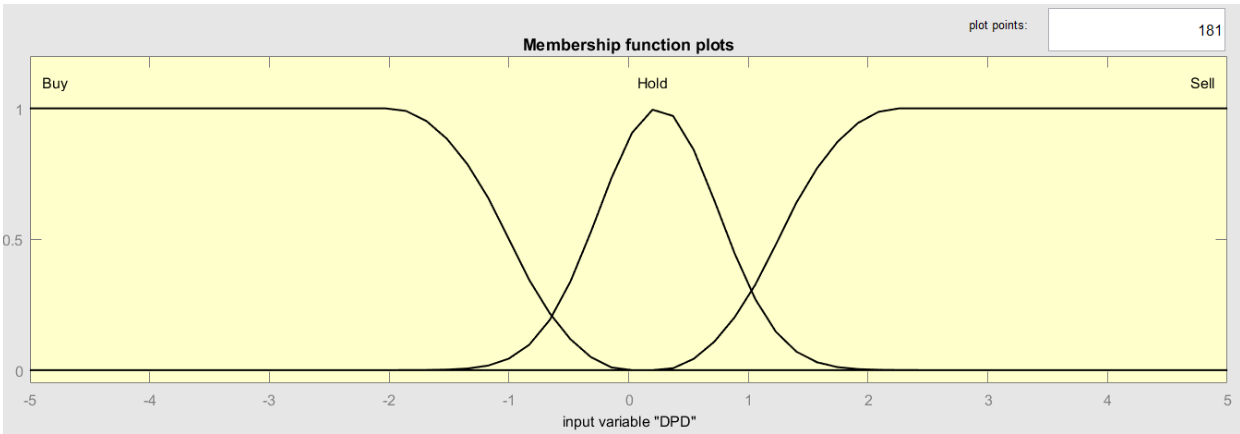

The membership functions of the fourth indicator, measuring divergence between %D and prices, are represented above. Buy and sell signals occur with increasing strength as the divergence between price and %D increases. A membership function for "hold" was also included for when there was no real divergence between prices and %D. This was included to help smooth out signals.

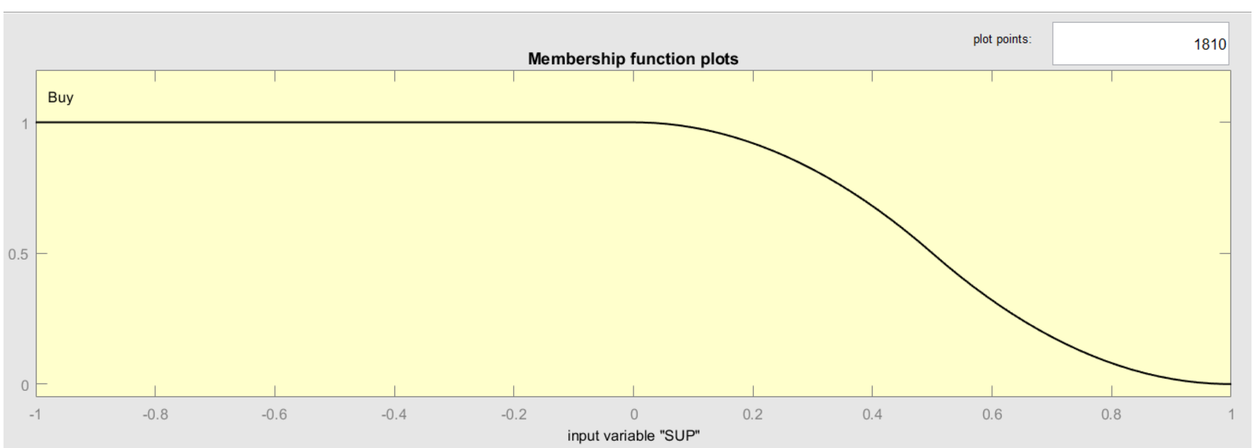

The membership function associated with the lower Bollinger band is shown above. Prices can be very far from the lower band, but relevant signals are only given when prices come close to touching the lower band, occurring at a value of 0 for our input. Prices moving to extreme values below the lower band were anomalies, and are treated the same as if they were simply touching the band. This membership function gives a stronger and stronger buy signal as prices get closer to the lower Bollinger band.

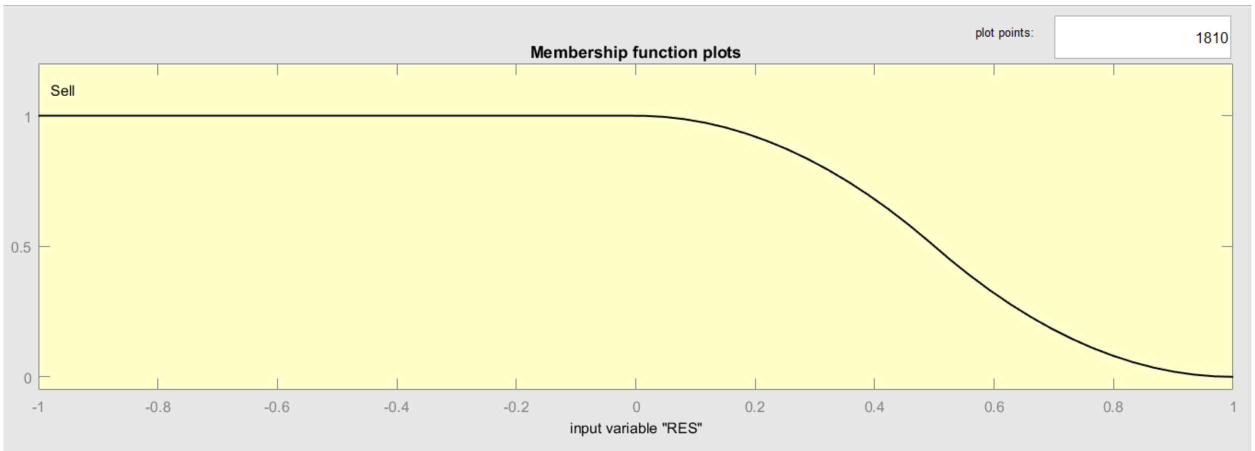

The membership function for our final input, which measures how close prices are to the upper Bollinger band, is shown above. It directly mirrors the membership function for SUP, but this time, a sell signal is generated as the price gets closer to the upper Bollinger band.

C. RULES

A fuzzy logic model uses rules in conjunction with the aforementioned membership functions to determine how the values of the inputs affects the trading signal given by the output. The rules for this model represent the aforementioned trading rules associated with each indicator.

Fuzzy logic also allows us to combine inputs in rules, allowing us to model traders' use of multiple indicators to determine signals. This model contains 11 rules, listed below with a brief explanation.

Rule 1: If (MA = "Buy") AND (SO = "Buy") THEN Output = "Strong Buy"

Rule 2: If (MA = "Sell") AND (SO = "Sell") THEN Output = "Strong Sell"

Rule 3: If (SO = "Sell") AND (%D = "Sell") THEN Output = "Sell"

Rule 4: If (SO = "Buy") AND (%D = "Buy") THEN Output = "Buy"

Rule 5: If (SO = "Sell") AND (RES = "Sell") THEN Output = "Sell"

Rule 6: If (SO = "Buy") AND (SUP = "Buy") THEN Output = "Buy"

Rule 7: If (%DPD = "Buy") THEN Output = "Buy"

Rule 8: If (%DPD = "Hold") THEN Output = "Hold"

Rule 9: If (%DPD = "Sell") THEN Output = "Sell"

Rule 10: If (MA = "Sell") AND (SO = "Buy") THEN Output = "Hold"

Rule 11: If (MA = "Buy") AND (SO = "Sell") THEN Output = "Hold"

Rules 1 and 2 are strong buy and sell signals that occur when both the moving average and the stochastic oscillator give the same signal. Combining the two inputs to create one signal was done to reduce noise from the signals on their own. Rules 10 and 11 indicate the uncertainty faced when moving averages and stochastic oscillators are giving different signals. Rule 3 indicates %D being overbought along with momentum slowing down. Commonly understood trading rules view this as a sell signal, meaning prices have climbed for a while, but momentum is slowing. Rule 4 directly mirrors this rule but on the buy side. Rules 5 and 6 use the stochastic oscillator to confirm buy and sell signals given when prices are close to either of the Bollinger bands. Rules 7, 8 and 9 are the signals given from the divergence between %D and prices. They were not combined with any other rule in this model.

These rules do not all hold the same weighting in our model. This is due to some being seen as more reliable than others. Rules 1, 2, 3, 4, 10 and 11 all have full weighting as moving averages and stochastic oscillators have held up in previous studies. Rules 7, 8 and 9 have only half weighting which limits their effect on the output. This was done because this input was created experimentally to measure a phenomenon normally observed graphically, and as such can't be taken as reliably. Rules 5 and 6 were given a weight of 0.75 to limit the affect the "SO" input has on our output. The stochastic oscillator is used to confirm all of our trading signals, and as such already has a significant impact on the output.

D. OUTPUT

The output from this model is given as a single numeric value for each set of input values. This number was arbitrarily given the range of 0 to 100. It can be seen as the combined strength of signals generated from all the trading rules. The output has its own membership functions that allow the rules to be translated to this output value.

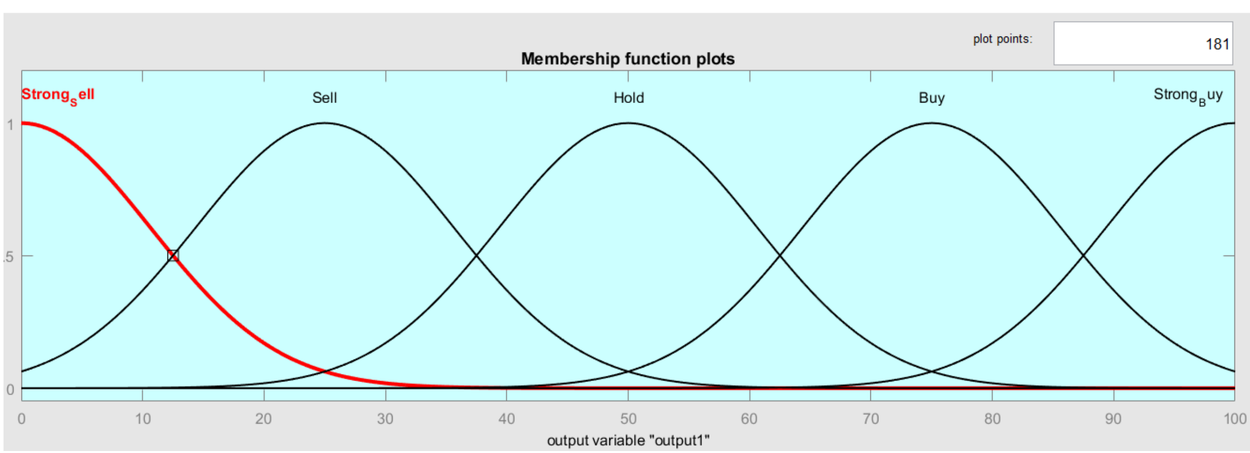

The membership functions associated with the output is shown above. As you can see, a larger output value corresponds to a stronger buy signal, and a lower output value corresponds to a stronger sell signal. This output variable can now be used to determine actual trading signals. Trading signals are determined by setting “trigger points” for the output. For example, if the buy trigger was set at 65, anytime the output reaches 65 a buy trade would be executed. If the sell signal was set at 35, this trade would be held until the output falls to 35. The values chosen for buy and sell trigger points plays a very important role in the results of our model. If the range between trigger points is set too small, we could expect many short term trades being executed, which would not be ideal. Likewise, if the range between trigger points was set too large, almost no trades would be executed. An ideal selection of trigger points would give us a moderate number of trades that might get us into a trade after a trend has formed but also reduces the number of whipsaws.

III. Data

This model was ran on daily price data from the S&P 500 over 30 years, from 1986 to 2015. The 30 year range is sufficiently long enough to give us a clear understanding of the model’s accuracy and consistency over time. I am using the S&P 500 because it is a good representation of the market as a whole. Market Portfolio Theory, which stands in antithesis to technical analysis, recommends a passive strategy of buying and holding the market. By using the S&P 500 I will be able to compare the results of this model with this ideal passive strategy. Data used to calculate the inputs was taken from daily high, low, and close prices. Opening prices on days following a buy or sell signal are used to simulate an executed trade.

IV. Results

Running the model over 30 years of daily S&P 500 data yielded daily output values ranging between 0 and 100. The output is distributed relatively normally around a median value of 48.6. The output is used to simulate trading strategies by setting "trigger points" for executing buy and sell trades. A higher output value means the model is recommending a buy trade. Likewise, a lower output value means the model is recommending a sell trade. Setting the buy trigger to 60 means that a buy trade is executed at the opening price on the day following an output value equal or greater to 60. If the sell trigger is 40, this trade will be held until an output value of 40 or lower is given.

Four unique sets of trigger points were chosen to simulate trading strategies using the model's output. Multiple sets of trigger points were used to ascertain a clearer understanding of the model's strengths and weaknesses. Since the model attempts to mimic a trader's ideas and actions, multiple sets of trigger points allow us to simulate different traders with potentially different risk preferences looking at the same information. The four sets of trigger points used were 60 and 40, 54 and 39, 54 and 37.75, and 69.5 and 36. These trigger points were chosen because they corresponded with critical areas in the output's distribution.

The trading simulation using these trigger points resulted in the following annual returns, shown below compared with market returns. Values in bold signify annual returns greater than the market.

Year	54-39	60-40	69.5-36	54-37.75	S&P 500
2015	-0.045	-0.072	-0.071	-0.061	-0.007

2014	**0.151**	**0.129**	-0.001	0.101	0.124
2013	0.101	0.087	0.081	0.117	0.296
2012	0.126	0.128	0.100	0.116	0.134
2011	**0.040**	-0.043	-0.044	-0.081	0.000
2010	**0.159**	0.125	**0.140**	**0.140**	0.128
2009	-0.100	-0.144	-0.085	0.036	0.235
2008	**-0.340**	**-0.257**	**0.008**	-0.406	-0.385
2007	**0.038**	**0.040**	-0.093	**0.042**	0.035
2006	-0.019	-0.030	0.041	-0.014	0.136
2005	**0.069**	**0.070**	0.002	**0.054**	0.030
2004	-0.018	-0.044	0.047	0.031	0.090
2003	0.081	-0.042	-0.024	0.181	0.264
2002	**-0.177**	**-0.043**	**-0.102**	**-0.163**	-0.234
2001	**0.078**	**0.050**	**0.000**	**-0.033**	-0.130
2000	**-0.033**	**0.010**	**0.023**	**-0.076**	-0.101
1999	**0.257**	**0.228**	0.005	0.191	0.195
1998	0.167	0.140	0.253	0.197	0.267
1997	**0.319**	0.280	0.297	**0.317**	0.310
1996	0.053	0.044	-0.097	0.044	0.203
1995	0.123	0.085	0.259	0.219	0.341
1994	**0.015**	**-0.015**	-0.030	**0.012**	-0.015
1993	-0.038	-0.042	-0.018	-0.031	0.071
1992	-0.010	-0.024	0.004	0.015	0.045
1991	0.199	0.169	0.099	0.151	0.263
1990	**-0.013**	**-0.023**	**0.012**	**-0.007**	-0.066
1989	0.091	0.062	0.098	0.108	0.273
1988	**-0.015**	**-0.090**	**0.012**	**-0.002**	-0.110

1987	**0.219**	**0.105**	**0.110**	**0.260**	0.020
1986	0.117	0.043	0.056	0.109	0.146
Average	0.053	0.031	0.036	0.052	0.085

As shown in the table above, no trading strategy came close to consistently beating market returns. The best performing set of trigger points, 54 and 39, only beat the market in 15 out of 30 years. It also yielded an average annual return of 5.3%, significantly lower than the average annual market return of 8.5%. This is before accounting for transaction costs and dividends. Dividends would likely be higher for a passive buy-and-hold approach as active trading would likely result in some foregone dividend payments. Transaction costs were not explicitly calculated due to the fact that they vary over time and between investors, but for transparency, the number of annual trades for each strategy was calculated and is shown below:

Year	54-39	60-40	69.5-36	54-37.75
2015	16	14	5	13
2014	14	12	3	12
2013	17	13	4	13
2012	15	13	3	12
2011	18	11	4	12
2010	18	11	4	14
2009	15	15	3	11
2008	13	14	3	9
2007	16	13	5	14
2006	14	13	5	13
2005	15	14	6	11

2004	19	18	5	11
2003	19	13	5	11
2002	15	13	4	13
2001	12	10	3	11
2000	14	14	3	9
1999	22	18	3	16
1998	13	10	3	13
1997	18	14	3	13
1996	15	12	9	13
1995	23	23	3	13
1994	15	16	6	9
1993	12	13	7	10
1992	14	12	5	11
1991	13	12	7	11
1990	13	13	5	10
1989	16	17	6	11
1988	17	15	6	12
1987	15	15	6	12
1986	15	13	6	11
Average	15.7	14	5	12
Total	471	414	140	354

As shown above, the number of annual trades was relatively consistent throughout the trading simulation for each strategy. As expected, a smaller range of trigger points resulted in more annual trades. If we had considered transaction costs and dividends, a passive buy-and-hold strategy would have been even more profitable than any of the technical strategies modeled.

Looking more closely at the model’s returns leads to one extremely interesting observation. The model’s trading strategies consistently beat the market in down years. Between 1986 and 2015, the S&P 500 had 7 years with losses of greater than 1%. In these 7 years, the model’s strategies consistently beat S&P 500 returns. The trigger points of 54 and 39 beat the market all 7 of these years, as did the trigger points of 60 and 40. The other two sets of trigger points beat the market in 6 of these 7 years. However, in years with large positive returns, the market almost always beat the model’s strategies. Between 1986 and 2015, the S&P 500 had 15 years with returns of greater than 10%. In these 15 years, even the best strategy, 54 and 39, only beat the market 4 times.

There are a few reasons why the model’s returns were overall significantly lower than market returns. The first, whipsaws, is a common problem in technical analysis. Since most technical indicators are trend following, there needs to be a trend for them to be profitable. When the market moves back and forth between trends frequently, trend following technical indicators can lead to inaccurate trading signals. This is referred to as whipsaws. While losses due to whipsaws are common in all technical trading strategies, this model had another less common failure. There were occasions when the market had huge up moves over days or even weeks without the model giving any buy signals. In a trend following technical strategy, you would expect to buy into up moves after they already started, missing out on some of the profits. You would not expect to miss 5% to 10% moves completely. It is not obvious why the model did not give buy signals during these times, as the underlying indicators were certainly giving buy signals. This particular failure does not appear to reflect any failure in technical analysis. While the model’s returns would have been higher had this not occurred, it did not occur often enough to significantly change the conclusions of this study. None of the model's strategies would have come close to beating market returns even if this did not occur.

V. Conclusions

The results of this study support the Efficient Market Hypothesis. The Efficient Market Hypothesis states that no publically available information can be used to consistently create returns greater than market returns. As we saw over a 30 year period of the S&P 500, the passive buy-and-hold approach recommended by the Efficient Market Hypothesis would have yielded significantly higher returns while taking considerably less effort. This is a sufficiently long enough time period to conclude that the model created in this study does not provide profitable information.

The results of this study can only lead to extremely limited conclusions. First, it is difficult to glean predictions for the future from this type of model. Past performance of a trading strategy does not indicate anything about future performance. Second, it is impossible to say anything about technical analysis as a whole due to the fact that an infinite number of technical trading strategies can be derived from different combinations of indicators. All we can say for certain is how this particular model did in this particular market over this time span.

The finding that technical strategies can consistently beat the market in down years is interesting and deserves further exploration. Does this imply that something inherent to technical strategies makes them more profitable during bear markets? Or is it simply that any active trading strategy allows you to avoid some losses during bear markets? This study leads to other unanswered questions. Would it be possible to create a different fuzzy logic model comprised of different indicators and rules that was able to beat the market? Would the model used in this study work better or worse on different securities such as large-cap or small-cap stocks or commodities? Would this model be profitable if you went short at every sell signal? Are there other, more effective methods of modeling technical analysis than fuzzy logic?

Even after answering these questions there will still likely be much uncertainty regarding the accuracy and profitability of technical analysis. Much of technical analysis is seen as an art form that is extremely challenging to methodically study. Despite the challenges and limitations, studying technical analysis is still important; even if it only serves to confirm the efficient market hypothesis.

VI. References

Brock, William, Josef Lakonishok, and Blake Lebaron. "Simple Technical Trading Rules and the Stochastic Properties of Stock Returns." *The Journal of Finance* 47.5 (1992): 1731-764. Web

Bessembinder, Hendrik. "Market Efficiency and the Returns to Technical Analysis."*Financial Management* 27.2 (1998): 5-17. *JSTOR*. Web. 30 Oct. 2015.

Sullivan, Ryan, Allan Timmermann, and Halbert White. "Data-Snooping, Technical Trading Rule Performance, and the Bootstrap." *J Finance The Journal of Finance*54.5 (1999): 1647-691. Web.

Lo, Andrew, Harry Mamaysky, and Jiang Wang. "Foundations of Technical Analysis: Computational Algorithms, Statistical Inference, and Empirical Implementation." (2000): n. pag. Web.

Park, Cheol-Ho, and Scott H. Irwin. "What Do We Know About The Profitability Of Technical Analysis?" *Journal of Economic Surveys* 21.4 (2007): 786-826. Web.

Murphy, J. Austin. "Futures Fund Performance: A Test of the Effectiveness of Technical Analysis." *Journal of Futures Markets* 6.2 (1986): 175-85. Web.

Blume, Lawrence, David Easley, and Maureen O'hara. "Market Statistics and Technical Analysis: The Role of Volume." *The Journal of Finance* 49.1 (1994): 153-81. Web

Loh, Elaine Y. L. "An Alternative Test for Weak Form Efficiency Based on Technical Analysis." *Applied Financial Economics* 17.12 (2007): 1003-012. Web

Dourra, Hussein, and Pepe Siy. "Investment Using Technical Analysis and Fuzzy Logic." *Fuzzy Sets and Systems* 127.2 (2002): 221-40. Web.

Zhou, Xu-Shen, and Ming Dong. "Can Fuzzy Logic Make Technical Analysis 20/20?"*Financial Analysts Journal* 60.4 (2004): 54-75. Web.

Escobar, Alejandro, Julian Moreno, and Sebastian Munera. "A Technical Analysis Indicator Based on Fuzzy Logic." *Electronic Notes in Theoretical Computer Science* 292 (2013): 27-37. Web. 30 Oct. 2015

Naranjo, Rodrigo, Albert Meco, Javier Arroyo, and Matilde Santos. "An Intelligent Trading System with Fuzzy Rules and Fuzzy Capital Management." *International Journal of Intelligent Systems* 30.8 (2015): 963-83. Web.

Kablan, A. "High Frequency Trading Using Fuzzy Momentum Analysis."*World Congress on Engineering* 1 (n.d.): n. pag. Web.