Is the Hurst Exponent better than Technical Analysis indicators for detecting trends?

Working Paper · February 2017

DOI: 10.13140/RG.2.2.21618.09928

1 author:

Tushar Chande

Is the Hurst Exponent better than Technical Analysis indicators for detecting trends?

Abstract

This paper is an empirical exploration at the intersection of mathematics and technical analysis using the Hurst Exponent, which seems firm in theory, but is squishy in practice. First, its value is hard to pin down. Second, the critical values designating trend onset change with methodology and are rarely exceeded in market data. Third, major trends can go unnoticed, and other quirks give false readings. Fourth, adding it as a filter to technical trend-following systems substantially degraded the absolute and risk-adjusted performance in our tests. I conclude that the Hurst Exponent is not necessarily better than conventional technical indicators.

Key words: Hurst Exponent; rescaled range analysis; detrended fluctuation analysis; technical analysis; trend-following; system testing

I. Introduction

The maxim "Buy low, sell high" assumes the market is trending higher, and symbolizes the importance to detecting trends in achieving investment goals. Two distinct approaches have developed to accomplish this goal: mathematical analysis and technical analysis.

In mathematics, trend detection takes the form of a search for autocorrelation in time series. A large body of literature has developed starting with the work on river flows (Hurst (1951)), extending from rescaled range analysis (Anis and Lloyd (1963), Mandelbrot and Wallis (1969)) to the fractal and chaos formulation (Peters (1991), (1995)) on detecting the presence of trends in time series data. This line of research continued with a new method of estimating the Hurst exponent called detrended fluctuation analysis (Peng et al (1994)), though variations of the rescaled range have also been proposed (Sanchez Granero , Segovia and Garcia Perez(2008)). The mathematics of fractional Brownian motion and fractal analysis is also very well developed (see Kantelhard et al (2002)).

Technical analysis comprises another set of tools also dedicated to determining trends, often predating the evolution from Hurst. Though these methods were considered more art than science, they have gradually gained academic attention, also putting them on a firmer footing (see Lo, Mamayasky and Wang (2000), Jagdeesh and Titman (2001)).

I focus on the intersection of the two methods because most of the literature does not do so, in order to understand how the two approaches can reinforce one another.

II. Literature Review

Mitra (2011, 2012) offers a good review of the prior work and explores the values of Hurst coefficient and how they relate to profits from a technical trading system. He calculated the Hurst exponent over 30-bar (consecutive trading days) periods and found them to be correlated to simple-moving average models. He did not directly filter the moving average models with the Hurst exponent. In his second study he used 60-bars to calculate the Hurst exponent via DFA method for a dozen different market indexes, and compared the profitability of a 10-day simple moving average rule to Hurst exponent values. He showed that profitability was generally greater when the Hurst exponent was greater than 0.55. However, there is much variability in the results; the Hurst exponent was not directly used as a filter, and because the data length is quite short, the calculated values of the Hurst exponent may be unstable. Thus, it is difficult to rely on these results.

Karamchandani and Jain (2013) studied daily data for Indian Rupee (USD/INR) exchange rate and tried to combine the Hurst exponent with technical analysis. They used N=1226 and found a Hurst Exponent of 0.499, concluding that daily log returns were randomly distributed. They tested a trend-following strategy using moving-average cross-over models with averages of 5 and 151 days, and 51 and 201 days. They also tested a counter-trend strategy using Bollinger bands, i.e., the standard deviation of daily closes added to a 21-day or 31-day simple moving average. Both their trend-following and counter-trend strategies were profitable, with the trend-following strategy being more profitable. Thus, the “random” price series prediction of Hurst exponent (Hurst exponent = 0.49) had no effect on the profitability of either approach, and had no predictive value, even though we might have expected counter-trend strategies to be more profitable.

Eom et al (2008) used DFA-driven Hurst exponent calculations to measure the predictability of the direction of 1-day ahead returns from 60-months of return data and found data with higher Hurst exponent were more predictable. However, they do not show any profitability measures or test any trading systems.

Qian and Rasheed (2004) make the case for using a time interval of 1024 daily time periods and show that the critical value of Hurst Exponent can be estimated from analytical equations (Anis and Lloyd (1963), Peters (1991, 1994)) to be 0.5436 or 0.5607 for random data. Since the two values differ, they used Monte Carlo simulations with Gaussian data to show that the estimated value is, on average, 0.5454, close to the Anis and Lloyd analytical estimate. However, due to the variability in the numerical simulation data, they recommend using a Hurst exponent greater than 0.65. They calculated Hurst exponent data for the Dow Jones Industrial Average and found it ranged from 0.42 to 0.68. They trained a neural network for predictability but did not test any trading systems. Unfortunately, the data as reported in their Table 3.2 is erroneous, so we redo it in a later section.

Overall, the literature does not compare the two technologies (mathematics and technical analysis) side-by-side calculated on the same underlying data. The promise of the Hurst Exponent is clear: if it exceeds some theoretical critical level, then that time series is trending, and then trend following strategies will be successful. From a trader's point of view, this raises many questions. One, can the Hurst Exponent be calculated uniquely, or do different techniques give different readings? Two, can the critical value be calculated uniquely, or do different techniques give different estimates? Three, are trend-following strategies always profitable if the Hurst Exponent is greater than some critical level? Four, are trend-following strategies always unprofitable when the Hurst Exponent is below the critical level? Five, can the Hurst Exponent fail to rise above the critical value during a major trend, i.e., is it possible to miss a major trend? Six, can conventional trend-following strategies identify trends when Hurst Exponent does not? Seven, does the Hurst Exponent used as a filter improve the returns from conventional trend-following strategies?

This paper systematically addresses the above questions. I revisit the calculations in the literature (Sanchez Granero et al (2008), Qiang and Rasheed (2004), Rzeszotko (2012)), and use them to calculate the theoretical and numerical values of Hurst Exponent to confirm the presence of trends in samples 1024-periods long. Armed with these values, I tested them against conventional technical analysis indicators and built trend-following systems. I directly compared Hurst Exponent calculations from the R/S method and DFA method to technical trend-following indicators, and explored their relative profitability and effectiveness at detecting tradeable trends. These analyses extend the literature through the side-by-side empirical exploration of mathematical and technical indicators of trend detection on the same underlying data.

The paper is organized as follows. The details of my implementation of the rescaled range method and the detrended fluctuation analysis are placed in the appendix. The methods section (III) details the approach to the calculations. The results of the testing are discussed in section IV, followed by a discussion and summary in section V.

III. Methods

I calculated the Hurst Exponent using the rescaled range (R/S) method and Detrended Fluctuation Analysis (DFA) method using 1024 days of data. This length was chosen for computational stability and in order to cross-check my calculations with the literature. The detailed steps of my approach are reported in the appendix.

The trend-following system testing on futures data is done with proprietary software that has been battle tested over twenty years of real-time trading and is in excellent agreement with actual trading results on diversified portfolios. A portfolio of 40 major US and international futures contracts was used covering all major sectors of the futures markets. It assumed $100 as slippage and commission per round turn and used continuous back-adjusted futures contracts.

The trend-following models were chosen arbitrarily without any attempt at optimization. A 20-period and 100-period simple moving average combination was used for the moving average trend-following model. A 100-period price channel was used for the break-out style model. The difference between the two is that the latter requires a new 100-period high or low for entry, while the former does not. The Hurst Exponent was used as a filter. This means that a new 100-day high or low was traded only if the Hurst Exponent was greater than a reference value. In the case of the moving average system, a long (or short) entry was taken only if the 20-period average was above (below) the 100-period average and the Hurst Exponent was greater than a reference value. The 100-day channel used a 60-period channel for exits, so that a 100-period long entry exited on a 60-period low (or short entries exited on a new 60-period high). The moving average model exited on a crossing over of the two averages.

I also arbitrarily picked two technical trend measures. To be consistent with the 100-day breakout model, I calculated the Channel-100 score that measured the percentage of time a market's close was within 15% of new 100-day highs or lows. For example, during a major trend, a stock or futures market will make a series of new highs or lows, and the Channel-100 score will increase.

I also calculate a trend-following score using six exponential moving averages with lengths arbitrarily chosen from the Fibonacci series starting with 21 days and ending with 233 days to set up a set of four comparison pairs. For example, I arbitrarily compare the 21-day exponential average of the close to the 55-day exponential moving average of the close. If the 21-day average is greater than the 55-day average, 1 is added to the score. If the 21-day average is below the 55-day average, then 1 is subtracted. We compare the following pairs: 21-55, 34-89, 55-155, 89-233, where each number corresponds to the exponential average of that length. If the shorter moving average is above the longer moving average in every pair, then the total score will be +4. Conversely, if every shorter average is below its corresponding long average then the score will be -4. I call this the moving average cross over (MAXO) score.

I then compared the trend-following scores to Hurst Exponent values on both, futures and stocks, since stocks can sometimes have longer lasting trends with larger amplitude than futures contracts.

IV. Results

IV. A - Theoretical Estimates of Hurst Exponent for Random Data

It is useful to have a benchmark for the Hurst Exponent for random data in order to decide if the actual data have a significant trend or not. I redo the Table 3.1 from Qian and Rasheed (2004) because the published data in that table are erroneous. Qian and Rasheed (2004) used equations for estimated Hurst Exponent from Anis and Lloyd and Peters. To this I added the Purczynski formula from Rzeszotko (2012). The formulas are not reproduced here, but are in the cited literature.

We note that the three sets of formulas, which correct for small sample effects, essentially give the same answer: the expected value of Hurst Exponent of 0.55~0.56 for random data (see Table 1). Thus, we conclude that for a data sample of 1024 days (or about 4 years) the calculated Hurst Exponent must exceed at least 0.56, to suggest a potentially significant trend in the sample at that moment. If one adds the errors in the regression estimates, we can bump up the value 0.58 to give us a cushion.

IV.B - Hurst Exponent for Random Data by Numerical Simulations

Again following Qian and Rasheed (2004) I used a Gaussian random number generator to generate 10,000 sequences of 1024 daily returns at random, and then calculated the Hurst Exponent using both rescaled range method and Distributed Fluctuation Analysis method, using the same set of random data. In this sense we go beyond Qian and Rasheed, by using both methods.

From Table 2, first observe that for random data the DFA method (column 4, last row) gave an expected Hurst Exponent of 0.5012, which is very close to the expected value of 0.50 from large samples. On the other hand, the rescaled range method (column 2, last row) reports a Hurst Exponent of 0.5591, about 11.5% higher than the expected value of 0.50. So the simulations showed that the DFA method had a greater standard deviation but, on average, reported a lower value for the Hurst Exponent than the rescaled range method.

IV. C - Critical Values of the Hurst Exponent

If one adds 1.96 times the standard deviation observed in the numerical data computed over 100,000 points to the average value, then the upper limit for random data Hurst Exponent (at 5% confidence level) is 0.64 for the rescaled range method, and 0.61 for the DFA method. These calculations agree closely with Qian and Rasheed who report 0.64 as the upper bound via the (R/S) method. My calculations also agree with Sanchez Granero et al (2008, Table 1) even though my implementation of the DFA method is different. They report a value of 0.60+/-0.03 for a series of length 1000 and minimum length n=8, i.e., their upper bound is 0.66. Thus, we can be confident that for all practical purposes, Hurst Exponent values greater than 0.65 on samples with 1024 time bars should indicate the presence of trends.

We then calculate the theoretical estimates for sample sizes smaller than 1024 time periods and find that the smaller the sample, the higher the theoretical Hurst Exponent estimated from that sample for random data (see Figure 1). As the sample size length decreases to 128 time intervals, a practical value for traders, the theoretical R/S estimate required to meet the non-random threshold increases by about six percent. So, we conclude that raising the theoretical reference value to 0.60 will cover most common situations of data length and randomness. These results are broadly consistent with Sanchez Granero et al (2008).

IV. D - Comparing Hurst Exponent and Technical Trend Indicators on Futures Data

The Hurst Exponent almost never exceeded 0.65 for any of the currency futures contracts I tested (see Table 3) implying that there were "no trends" in these markets for 23 years, which is simply not supported by the real-time experience of commodity trading advisors.

The 0.58 threshold for Hurst Exponent was exceeded between 28% and 51% of the time (see column 2), and the longest streak of continuous readings above 0.58 ranged from 307 to 651 periods. These markets met MAXO the trend-following criteria of fully long or fully short between 64% and 70% of the time. The futures markets were near the edges of a 100-day channel between 61% and 73% of the time. These data show that technical indicators more readily picked up trending markets than the Hurst Coefficient did. Even if we use the lower 0.58 threshold for the Hurst Exponent, the Hurst Exponent trend detection frequency was only about two thirds of the frequency of technical analysis indicators.

IV. E - Comparing Hurst Exponent and Technical Trend Indicators for Dow 30 Stocks

Table 4 shows that the Hurst coefficient never rose above 0.65 when calculated by either method for stocks in the Dow 30 Industrials Average. Thus, the Hurst Exponent detected no trends at the 0.65 level over the 12.5 years or so of data. When the Hurst value is set to 0.58, the theoretical 'maximum' for random data, we get signals of trends more often. However, there is little correlation between the Hurst Coefficient and the technical analysis trend measures.

In Figure 2 we show that the technical analysis measurements in Table 4 are correlated, i.e., when the market is making new 100-day highs or lows consistently over many months, the moving average score is +4 or -4, i.e., the shorter averages are above the corresponding longer averages (and vice versa). We should expect that when the technical analysis is showing definite trends, the Hurst Exponent will at least exceed 0.58, and ideally, stay there throughout the trend. However, Figure 3 shows that the Hurst Exponent is uncorrelated to the MAXO score data from Table 4. This means that the Hurst Exponent does not pick up trends as consistently as the technical analysis measures do.

IV. F - Rescaled Range Method Fails To Detect Major Trends

I hand-checked the Hurst Exponent against the price action of many stocks and found that it routinely missed major trends. I illustrate this vulnerability by using Nvidia (NVDA) stock data form 12/31/2014 to 1/30/2017 to show that the Hurst Exponent missed an explosive, unrelenting rally during all of 2016. In Figure 4 we show that technical indicators responded quickly to the emerging trend. Whereas in Figure 5 we show that the Hurst Exponent was stuck below 0.58 for most of the rally, with the DFA method able to push its exponent above 0.65 only at the latter stages of the move. Thus, the combination of long data period and power law requirement means that this indicator could miss major moves in the market.

IV. G - Hurst Exponent via (R/S) vulnerable to Cyclic Fluctuations

I generated synthetic data over 1024-day interval with a linear move from 25 to 30, or sinusoidal fluctuations around 25 (between 30 and 20) with 1, 2, 4, 8, 16 and 32 cycles. The (R/S) Hurst Exponent values are shown in Table 5, and Figure 6 has the contrived data with 4 cycles over the entire time period.

As Figure 6 shows, most technical analysts would say that this stock was not trending. It is surprising that even data with 8 cycles can produce relatively high Hurst Exponent values above 0.58. Thus, high values on the Hurst Exponent do not necessarily correspond to rising or falling prices over a long period, but could include fluctuating markets as well. I repeated these calculations with triangular waves in the data with similar results (not shown).

IV. H - Filtering Trend-following systems with Hurst Exponent

I filtered a 100-period breakout-style system and a moving average system using 20-period and 100-period simple moving averages. The Hurst Exponent base line is with a value of 0, when the filter has no effect. We then jump to a Hurst of 0.40, which still has no effect, which means that virtually all the values were greater than 0.40. After that, Hurst Exponent values increased in steps of 0.025 from 0.4 to 0.65, since 0.65 was the critical level for a 1024-period calculation.

I show the gross returns for a period ranging from 01/1993 through 8/2015, a period of 220 months in Figure 7 and vary the Hurst Exponent from 0.4 to 0.65 (since about 4 years of data are needed to seed the Hurst calculations). Since the critical values are above 0.56, it is clear that adding the Hurst Exponent filter massively reduces profitability. Risk adjusted measures of performance also declined once the Hurst Exponent started rising above 0.50 (Figure 8). Adding the Hurst Exponent filter to a 20-100 moving average system also decreased profitability with increasing Hurst Exponent (see Figure 9). Adding the Hurst filter merely delays the start of the trade relative to the exit, reducing the time in trade or missing it entirely. Both occurrences will lower profitability.

V. Discussion and Conclusion

The Hurst Exponent is widely accepted as an accurate guide to the presence of trends. However, there are many practical problems with its application. In answer to the first question we posed, there is no unique method to calculate the Hurst Exponent, each method produces a different value, and the literature does not have consistent descriptions of the methods. For each method, the Hurst exponent calculation gives a different value as the time period used increases, and as the length of underlying sub-intervals increases. Hence, it is not possible to find a unique value of the Hurst exponent.

Different methods also yield different answers for the critical value of the Hurst Exponent. Thus, it is difficult to pin down precisely how to set the critical levels. I updated the calculations reported in the literature to find 0.56 as the theoretical Hurst Exponent for a sample size of 1024, and 0.64 as its numerically calculated counterpart. I then showed that the level 0.65 rarely occurs in historical data using futures and US stocks. Hence, we have to use a lower value, such as 0.56 or 0.58 to even get measurable data. Shortening the length of the calculation period raises the level above which trends should theoretically exist. Numerical simulations with random data generally lead to estimates of the critical level for Hurst exponent that are greater theoretical calculations. Hence, though shorter data lengths are more practical, the reference value is higher, making it even less likely that these values will be exceeded in market data. So the answer to our second question is that the critical value cannot be calculated uniquely.

My calculations found that the 1024-period Hurst Exponent does not correlate with other trend-following indicators from technical analysis such as moving averages and 100-period price channels. The different time intervals used in the Hurst exponent calculations and the technical indicator calculations should not have any bearing on whether the two correlate or not. In addition my calculations show that Hurst calculations can fail to detect major trends in a time series, and are vulnerable to sinusoidal or triangular cycles in the data. Finally, adding the Hurst Exponent as a filter to a trend-following system severely reduced system profitability, which we tested using both breakout-style and moving-averaged based systems. The risk-adjusted measures of performance also suffered.

So the answer to question three is that trend-following strategies are not always profitable when the Hurst Exponent is greater than some critical level. In answer to question four, trend-following systems are not always unprofitable when the exponent is below the critical level. The answer to question five is that is certainly possible to miss major trends. For question six, conventional technical indicators easily identified trends when the exponent does not. Lastly, for question seven, using the Hurst exponent as a filter may not necessarily improve the returns from conventional trend-following strategies.

Thus, I conclude that it is possible that the Hurst Exponent is not necessarily better than conventional technical indicators for detecting and following trends.

References

Anis, A. A. and E. H. Lloyd. "The expected value of the adjusted rescaled Hurst range of independent normal summands.", Biometrika, 63(1963) 111-116.

Barbulescu, A.; C. Serban; and C. Maftei. "Evaluation of Hurst exponent for precipitation time series." In Proc. WSEAS International Conference on Computers, Corfu, Greece (2010).

Eom, C.; S. Choi; G. Oh; and W.S. Jung. "Hurst exponent and prediction based on weak-form efficient market hypothesis of stock markets." Physica A 387(2008), 4630-4636.

Hurst H.E. "Long-term storage of reservoirs: an experimental study." Transactions of the American Society of Civil Engineers, 115(1951) 770-799.

Jegadeesh, N. and S. Titman. " Profitability of Momentum Strategies: An Evaluation of Alternative Explanations." Journal of Finance, 56(2001) 699–720.

Kantelhardt J.W.; S. A. Zschiegner; E. Koscielny-Bunde; S. Havlin; A. Bunde; and H. E Stanley. "Multifractal detrended fluctuation analysis of nonstationary time series." Physica A 316(2002) 87-114.

Karamchandani, M. and S Jain. "Indian forex market: Revisited through Hurst exponent and technical analysis." International Journal of Commerce, Business and Management, 2(2013) 185-192.

Lo A.; H. Mamayasky; and J. Wang. "Foundations of technical analysis: computational algorithms, statistical inference, and empirical implementation." J. of Finance, LV(2000) 1705-1765.

Mandelbrot B.B. and J. R. Wallis. "Robustness of rescaled range R/S in the measurement of noncyclic long-run statistical dependence." Water Resources, 5(1969) 967-988.

Mitra, S. M. "Trends in stock prices and range to standard deviation ratio." International Journal of Business and Management, 6(2011) 223-234.

Mitra, S.M. "Is Hurst exponent value useful in forecasting financial time series?" Canadian Center of Science and Education, 8(2012) 111-120.

Peng C.K.; S. V. Buldyrev; S. Havlin; M. Simons; H. E. Stanley; and A. L. Goldberger. "Mosaic organization of DNA nucleotides." Physical Review E, 49(1994) 1685-1689.

Peters E.E. "Chaos and Order in the Capital Markets: a new view of cycles, prices and market volatility." Wiley (1991).

Peters E.E, "Fractal market analysis: applying chaos theory to investment and economy.", Wiley (1995).

Qian, B. and K. Rasheed. "Hurst exponent and financial market predictability." In Proceedings of the 2nd IASTED, International Conference on Financial Engineering and Applications, Cambridge, MA: ActaPress(2004).

Sanchez Granero M. A.; J. E. Trinidad Segovia; and J. Garcia Perez. " Some comments on Hurst exponent and the long memory process on capital markets." Phsyisca A 387(2008) 5543-5551.

Rzeszotko, Z. " Estimation of the Hurst exponent for market indices by use of rescaled range method and CAS Mathematica. Case Study: S&P500." The Electronic Journal of Mathematics and Technology, 6(2012) 256-267.

Appendix : Calculating the Hurst Exponent

To calculate the Hurst Coefficient with the (R/S) method, start with 1025 daily data points and calculate daily log returns d_n, $d_n = \ln(c_n/c_{(n-1)})$, where n=2,3, … 1025 and c_n is today's close and $c_{(n-1)}$ is yesterday's close. This gives a total N=1024 values.

Next, divide the 1024 data points into equal, non-overlapping sub-intervals of length 2^s, s= 3,4 … 10, yielding 128 sub-intervals of length t= 8, 64 sub-intervals of length t=16, 32 sub-intervals of length t=32 and so on.

Within each sub-interval of length *t*,

a) Calculate the mean, $\mu_{s,}$ as $(1/t)\sum d_{s,}$ s=1,2..t.

b) Subtract the mean from each daily log return d_s, to create $z_s = d_s - \mu_s$, s=1,2… t.

c) Cumulate each individual z_s into Z_s where if t=1 $Z_s = z_s$, and for t>1 $Z_s=\sum z_{s,}$ s=1,2,..t.

d) Calculate the maximum (max Z_s) and minimum (min Z_s) values of Z_s

e) Calculate the standard deviation of t daily returns d_s (S_s) as the square root of

$S^2{}_s= (1/(t-1))\sum(d_s-\mu_s)^2$, s=1,2..t.

f) Calculate rescaled range as R/S for the sub-interval of length t as $(\max Z_s - \min Z_s)/S_s$

Within each set of sub-intervals, say 128 of length t=8, 64 of t=16 and so on, average the calculated R/S value from each sub-interval, to get the representative scaled range for all intervals, $(R/S)_t$.

Lastly, convert the $(R/S)_t$ and length-t pairs into base-2 logarithms and calculate the slope of the ordinary least squares line as the Hurst Exponent value.

The central idea of Detrended Fluctuation Method (DFA) of Peng et al (1994) is to use ordinary least squares to fit a regression to sub-intervals of varying length, and calculate the standard deviation of the residual sum of squares, or the sum of squared due to error. This standard deviation, called a fluctuation, is then expected to follow a power law relationship to the length of sub-intervals like the rescaled range calculations.

We start with N=1024 daily closes (c_n, n = 1, 2, .. 1024), and the entire range N is sub-divided into equal, non-overlapping sub-intervals of length 2^s, s= 3,4 … 10, yielding 128 sub-intervals of length t= 8, 64 sub-intervals of length t=16, 32 sub-intervals of length t=32 and so on.

Within each sub-interval of length *t*,

a) Use an ordinary least squares regression to fit a straight line to the $c_{s,}$ s = 1, 2… t.
b) For each daily close, calculate the estimate at each point from the regression, Y_s for each c_s as $Y_s = m^* c_s + K$ where *m* is the slope and K is the constant from the regression
c) Calculate the residual sum of squares as $E_t = \sum(c_n – Y_s)^2$, where s = 1, 2.. t.
d) Calculate the fluctuation as the standard deviation of the residuals as $F_t = \text{sqrt}(E_s/t)$

Now average the F_t over all intervals of a particular length *t*. We then plot a log-log graph of the logarithm of the fluctuations averaged F_t form all the sub-intervals versus the logarithm of the length of sub-intervals, *t*, looking for the power law $F_t = Kt^H$ where H is the Hurst coefficient and K is a constant.

The literature suggests calculating DFA by subtracting the mean price from all N data and cumulating the deviations (see Barbulescu (2010)), before calculating fluctuations with the cumulative deviations. However, it did not give meaningful results comparable to R/S analysis. Using daily log returns also did not produce meaningful estimates for the Hurst Exponent (see Mitra (2011), (2012)).

Table 1
Theoretical (R/S) Expected Values (EV) for Hurst Exponent of 'Small' Random Data Sample

Log2 t	Log2 Anis EV	Log2 Peters EV	Log2 Purczynski EV
4	1.942176	1.896373	1.966955
5	2.566090	2.543370	2.588641
6	3.147310	3.135995	3.158332
7	3.701969	3.696322	3.707370
8	4.239360	4.236540	4.241999
9	4.765211	4.763802	4.766491
10	5.283210	5.282505	5.283823
Hurst Exponent	0.554050	0.560707	0.549642

Author's calculations using formulas in Qian and Rasheed (2004) (from Annis and Lloyd, Peters) and separately in Rzeszotko(2012) (from Purczynski)

Column 1: Logarithm with base 2 of the length of the sub-intervals used to calculate the Hurst exponent. So a 16-day interval would have Log2(16) = 4. Column 2 has the logarithm to base 2 of the expected values calculated via the formulas of Anis and Lloyd(1963) as reported by Qian and Rasheed(2004). Column 3 has the logarithm to base 2 of the expected values with the Peters correction as reported by Qian and Rasheed(2004). Column 4 has the logarithm to base 2 of the expected values calculated from the formula of Purczynski as reported by Rzeszotko (2012). The last row has the Hurst Exponent values obtained by a linear regression of the values in columns 2, 3 and 4 against the values in column 1. These are the theoretically expected values of the Hurst Exponent for a 1024-day sample using rescaled range analysis.

(Place in section IV A)

Table 2: Hurst Coefficient Estimates From Numerical Simulations for Random Data

Run #	(R/S) Average	(R/S) Standard Deviation	DFA Average	DFA Standard Deviation
1	0.5624	0.0389	0.5048	0.0457
2	0.5466	0.0482	0.4785	0.0643
3	0.5576	0.0356	0.5013	0.0517
4	0.5496	0.0320	0.4867	0.0385
5	0.5549	0.0337	0.5030	0.0484
6	0.5668	0.0356	0.5139	0.0489
7	0.5488	0.0502	0.4923	0.0644
8	0.5785	0.0459	0.5231	0.0577
9	0.5827	0.0416	0.5224	0.0560
10	0.5429	0.0269	0.4865	0.0383
Over all Data points	0.5591	0.0415	0.5012	0.0541

Each row reports the average Hurst Exponent using 10,000 samples of length 1024 each using the rescaled range (R/S) and detrended fluctuation analysis (DFA) method on the same data.

(Place in section IV B)

Table 3: Measuring Trends in Continuous Futures Currency Contracts
with Hurst Coefficient and Technical Analysis

Currency Futures	**% HE> 0.58**	**Max Streak > 0.58**	**% HE > 0.65**	**MAXO Score**	**CH100 Score**
Australian Dollar	42.09%	448	0	68.98%	60.98%
Canadian Dollar	28.78%	651	0	69.54%	69.05%
Euro	51.28%	307	0	68.51%	66.11%
Japanese Yen	44.22%	533	0	67.29%	72.66%
Swiss Franc	32.67%	992	1.09%	64.82%	69.20%

A total of 4659 intervals 1024 bars long starting from 1/4/1993 through 9/9/2015

Column 2 gives the percentage of intervals that the Hurst Exponent exceeded 0.58 in the data. Column 3 gives the longest streak of consecutive days with Hurst Exponent greater than 0.58. Column 4 gives the percentage of time the Hurst Exponent was greater than 0.65. The MAXO Score is the percentage of time the score was either +4 (all four moving average pairs showed shorter average above the longer exponential moving average) or -4 (all four moving average pairs showed the shorter average below the longer exponential moving average). The MAXO score is measure of trending behavior when the moves are sufficiently long lasting to align the moving average one above the other or one below the other. The last column is the percentage of time prices were within 15% of new 100-period highs or lows. During major trends prices will tend to near rising new highs-channel or falling new-lows channel. The technical trading indicators showed the market was trending more often than did the Hurst Exponent.

(Place in section IV D)

Table 4: Hurst Exponent and Technical Indicators for Dow Jones 30 Stocks (10/04-1/17)

Symbol	% HX>0.65	% HX>0.58	% HX<0.51	MAXO Score	CH100 score	% HX_DFA >0.65	MaxStreak
AAPL	0%	67.88%	0%	75.47%	51.14%	0%	676
AXP	0%	43.11%	23.60%	69.78%	48.37%	0%	472
BA	0%	47.45%	0%	61.02%	40.05%	0%	765
CAT	0%	56.93%	0%	68.56%	29.98%	0%	999
CSCO	0%	9.44%	6.08%	65.64%	41.31%	0%	136
CVX	0%	14.21%	41.80%	67.10%	42.14%	0%	108
DD	0%	59.12%	1.02%	72.21%	45.11%	0%	1015
DIS	0%	9%	2.97%	76.16%	58.25%	0%	63
GE	0%	50.80%	8.71%	63.80%	44.38%	0%	1000
GS	0%	54.26%	0.54%	63.89%	36.35%	0%	367
HD	0%	6.72%	25.94%	74.31%	73.14%	0%	94
IBM	0%	7.49%	39.51%	69.25%	49.83%	0%	49
INTC	0%	1.90%	0.19%	65.06%	29.39%	0%	11
JNJ	0%	0.73%	3.99%	68.71%	54.01%	0%	8
JPM	0%	0.15%	15.13%	69.88%	44.96%	0%	3
KO	0%	9.10%	41.36%	64.28%	41.85%	0%	39
MCD	0%	0%	40.58%	67.93%	50.17%	0%	-5
MMM	0%	37.13%	2.77%	76.20%	63.80%	0%	469
MRK	0%	27.74%	7.45%	70.66%	43.26%	0%	181
MSFT	0%	12.60%	39.56%	67.10%	43.60%	0%	226
NKE	0%	0%	41.12%	72.75%	60.54%	0%	-5
PFE	0%	0.24%	9.68%	66.33%	47.25%	0%	3
PG	0%	6.76%	10.66%	64.82%	47.40%	0%	40
TRV	0%	0%	63.31%	69.20%	56.98%	0%	-5
UNH	0%	3.89%	54.65%	68.91%	62.58%	0%	22
UTX	0%	14.40%	10.41%	70.90%	43.07%	0%	126
V	0%	0%	88.68%	78.46%	77.20%	0%	-5
VZ	0%	1.31%	11.29%	62.68%	40.58%	0%	8
WMT	0%	16.01%	38.88%	61.02%	32.70%	0%	278
XOM	0%	0.05%	47.98%	61.75%	33.28%	0%	1

Columns 2, 3 and 4 show the percentage of time the Hurst Exponent exceeds the value shown in the heading. Column 5 shows the percentage of time the stock had a Maxo score of +4 or -4 (see details at the bottom of Table 3). Column 6 shows the percentage of time the price was within 15% of the 100-day high or low. Column 7shows the percentage of time the Hurst Exponent calculated via the DFA method exceeded 0.65. The last column show the longest streak of consecutive days with R/S Hurst Exponent greater than 0.58.

(Place in section IV E)

Table 5: Hurst Exponent via (R/S) Method for Various Degrees of Data Cyclicality

Synthetic data (Range:20-30)	Hurst Exponent Via (R/S)
Straight Line Up	1
1 Cycle	1
2 cycles	0.94
4 cycles	0.787
8 cycles	0.621
16 cycles	0.43
32 cycles	0.26

The synthetic data and the corresponding Hurst Exponent are shown to show the vulnerability of the rescaled range method to cycles in the data.

(Place in section IV G)

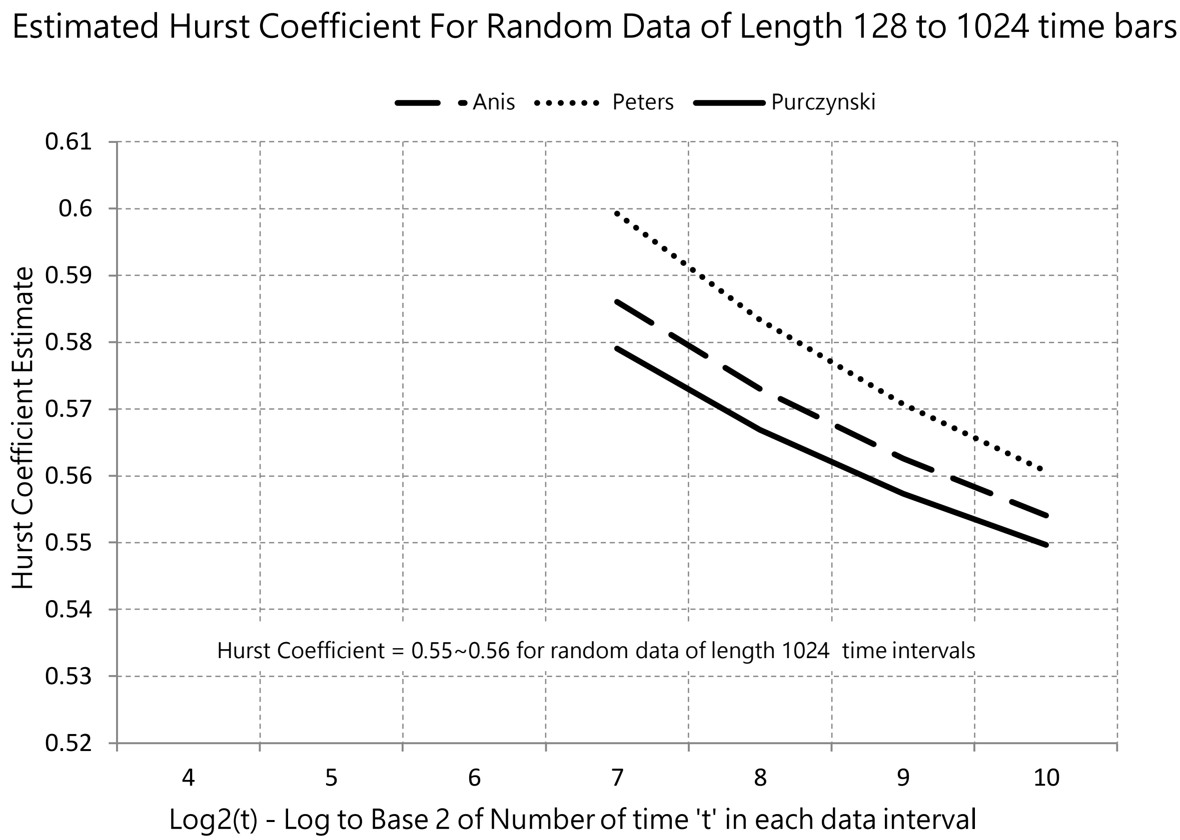

Figure 1: We estimate Hurst Exponent values for random data, using theoretical formulas corrected for small sample size, and time intervals ranging from 128 to 1024 time intervals. The three different formulas (from Lloyd and Anis, Peters and Purczynski as obtained from papers by Qian and Rasheed and Rzeszotko) show that for a sample of size 1024 time intervals, the expected Hurst Exponent is about 0.55 of 0.56.

(Place in section IV C)

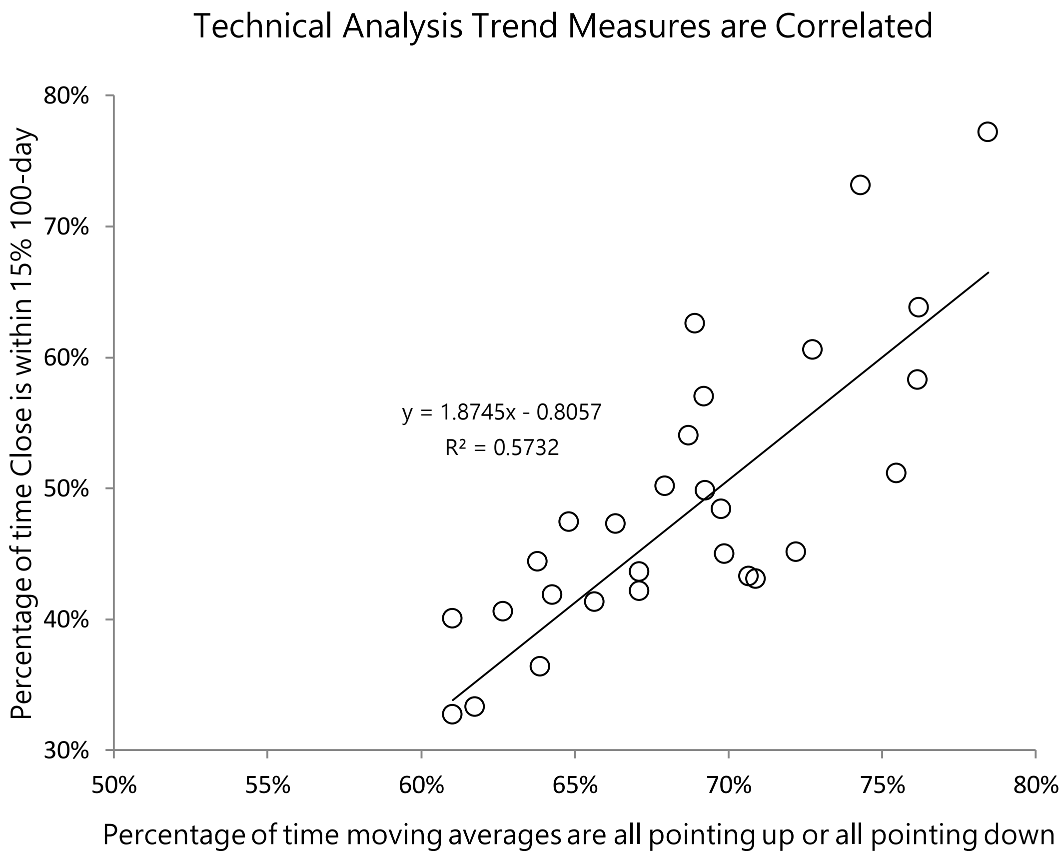

Figure 2: Technical analysis measures of trend are correlated, using data for all Dow 30 stocks and data from 10/04 to 1/17. Thus, stocks making new highs or lows usually have all of their moving averages aligned, such that shorter moving averages are above longer moving averages (or vice versa).

(Place in section IV E)

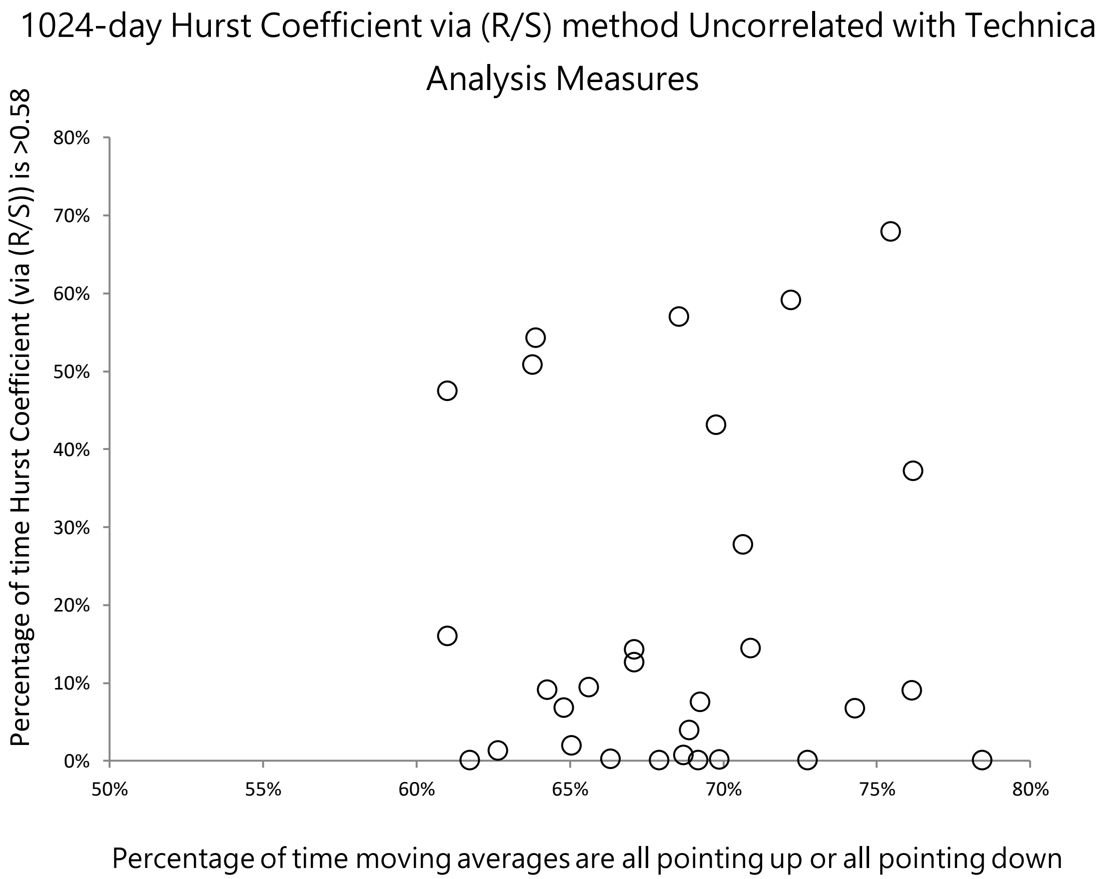

Figure 3: The Hurst Exponent over 1024-day time interval is uncorrelated with the percentage of time all moving averages are pointing up or down, i.e., the stock is trending.

(Place in section IV E)

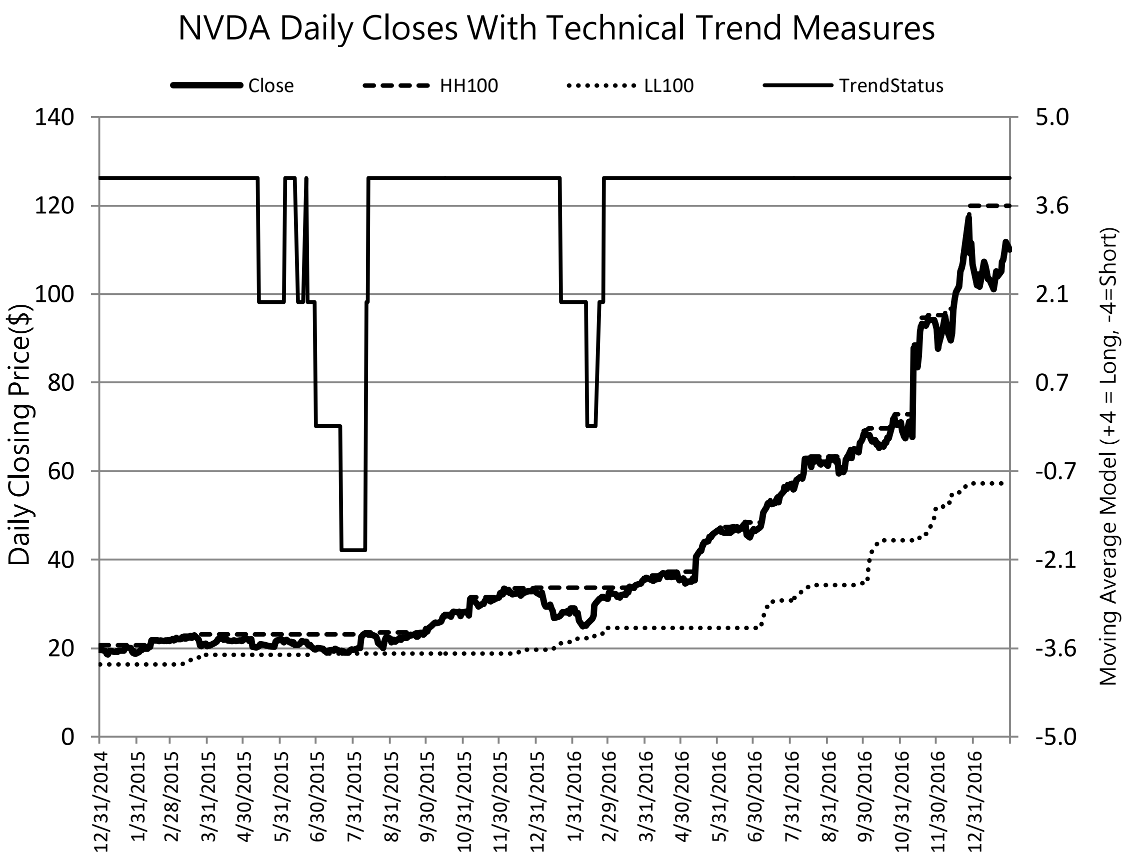

Figure 4: We show Nvidia (NVDA) stock data during 2016 which saw an impressive rally. The technical indicators showed a series of new highs on the 100-day channel, and the moving-average score was pinned at +4 for major portions of the move. Thus, technical indicators were able to identify the uptrend in this case.

(Place in section IV F)

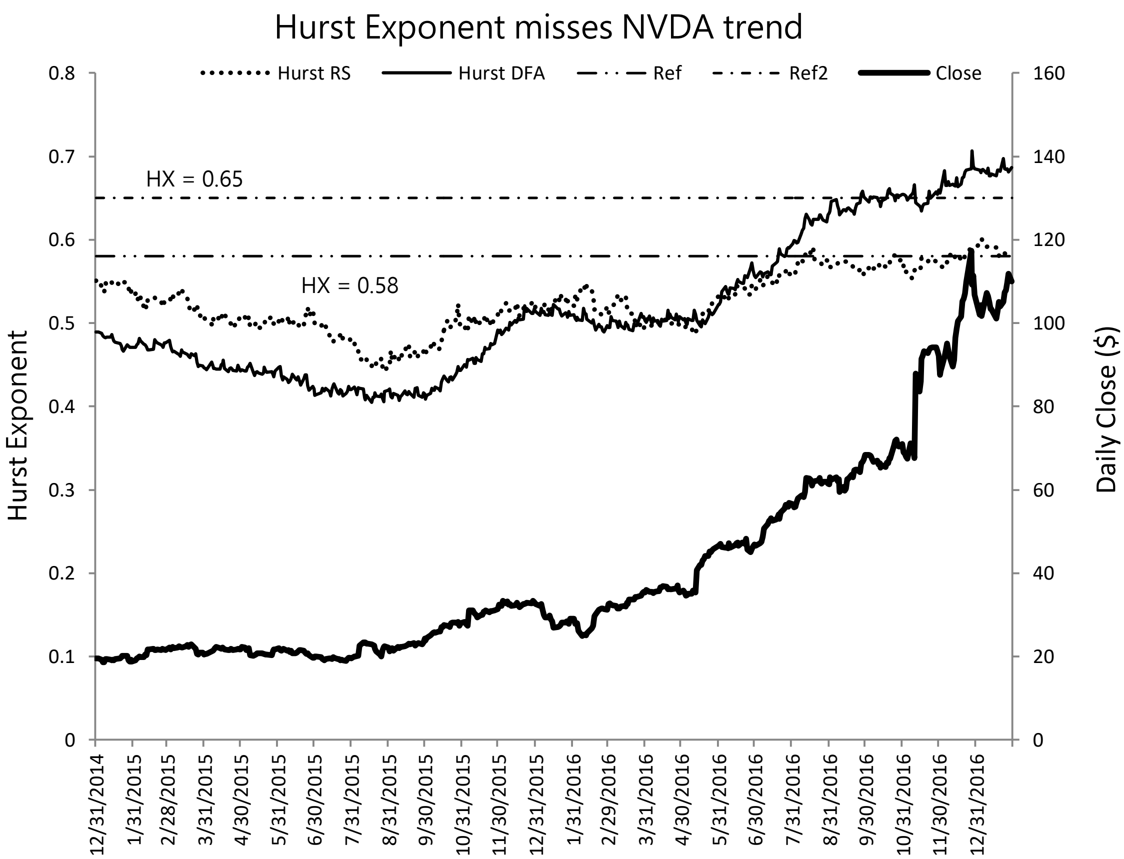

Figure 5: We show the daily close of Nvidia (NVDA) stock and the daily calculations of Hurst Exponent on a rolling 1024-day window using the rescaled range (Hurst RS) and de-trended fluctuation method (Hurst DFA). The R/S method never pushed the Hurst Exponent above 0.65, and barely rose above 0.58. The DFA Hurst Exponent reached 0.65 only toward the latter stages of the move.

(Place in section IV F)

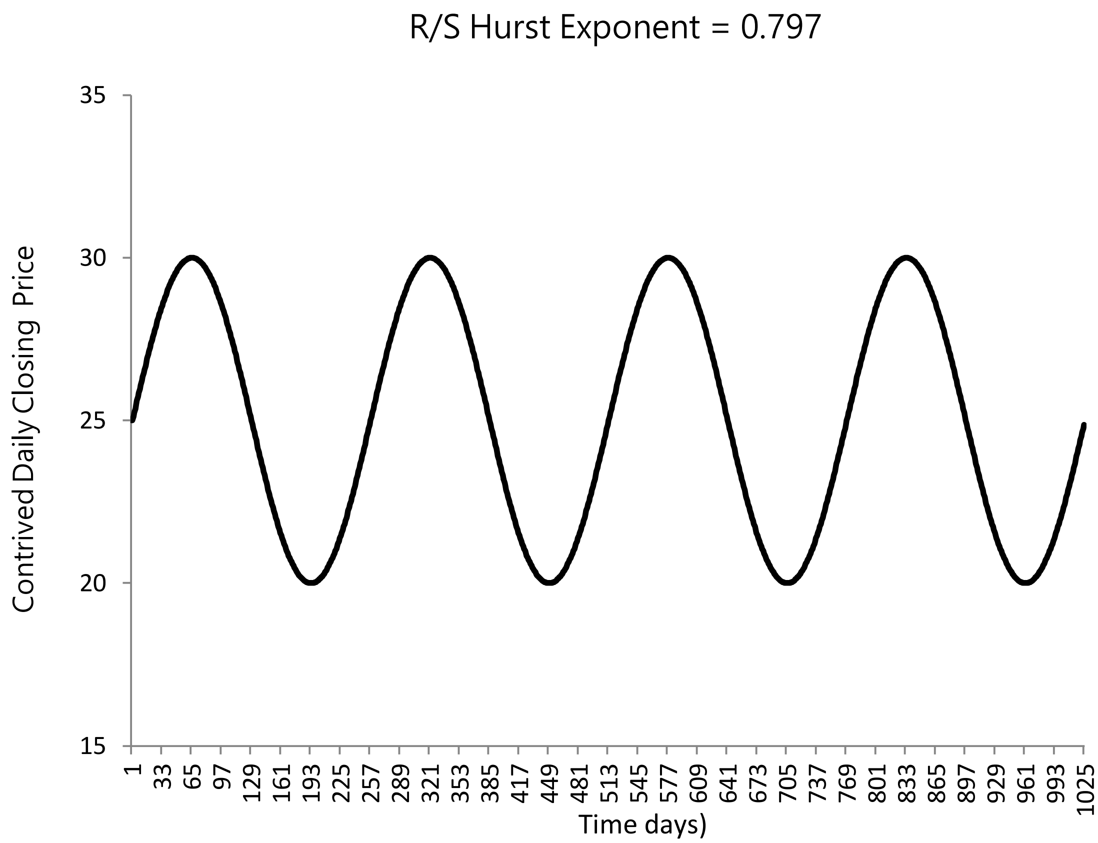

Figure 6: Synthetic data with built-in cyclicality can also produce high Hurst Exponent values, even though they do not have directional price moves.

(Place in section IV G)

Simulated Gross Profits for Diversified Futures Portfolio with Hurst filter (Jan-93 to Aug-15)

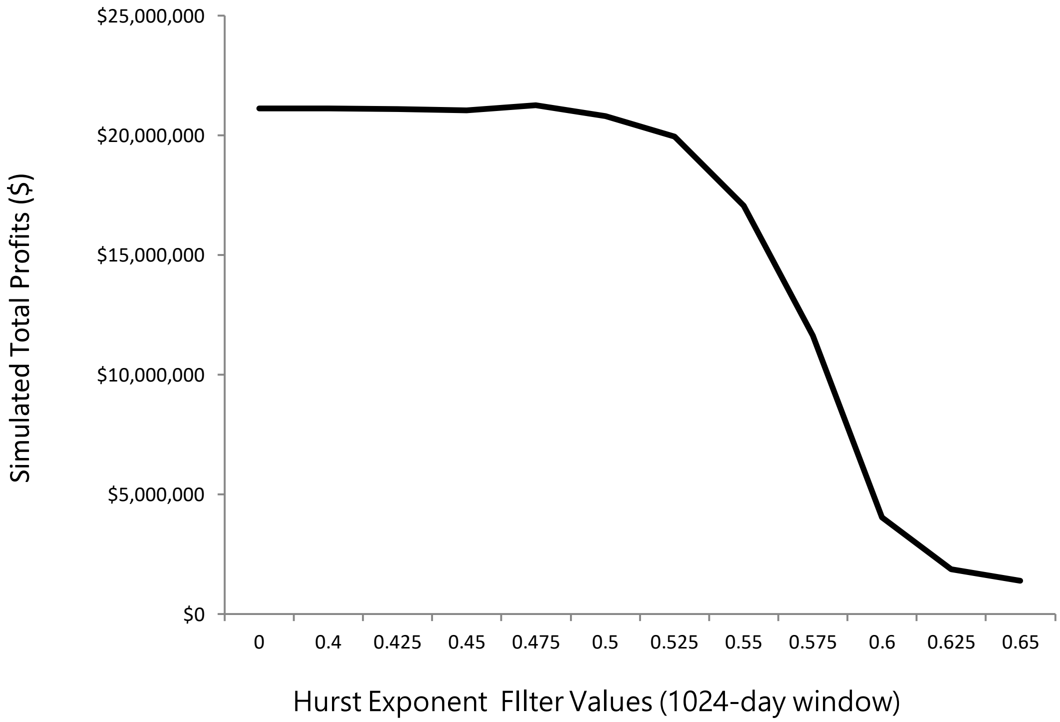

Figure 7: Adding a Hurst Exponent filter to a 100-bar channel breakout system reduces profitability. For example, between 0 and 0.45 the Hurst Exponent has very little effect on system performance since most of the 1024-day values are greater than 0.45. Once the Hurst Exponent increases past 0.5, the entry point gets delayed and profitability decreases. As the Hurst Exponent rises to 0.65, very few entry signals are accepted and systems profitability is greatly reduced. Thus, adding a Hurst Exponent Filter at or above 0.56 (the critical values from theoretical analysis) massively erodes system profitability.

(Place in section IV H)

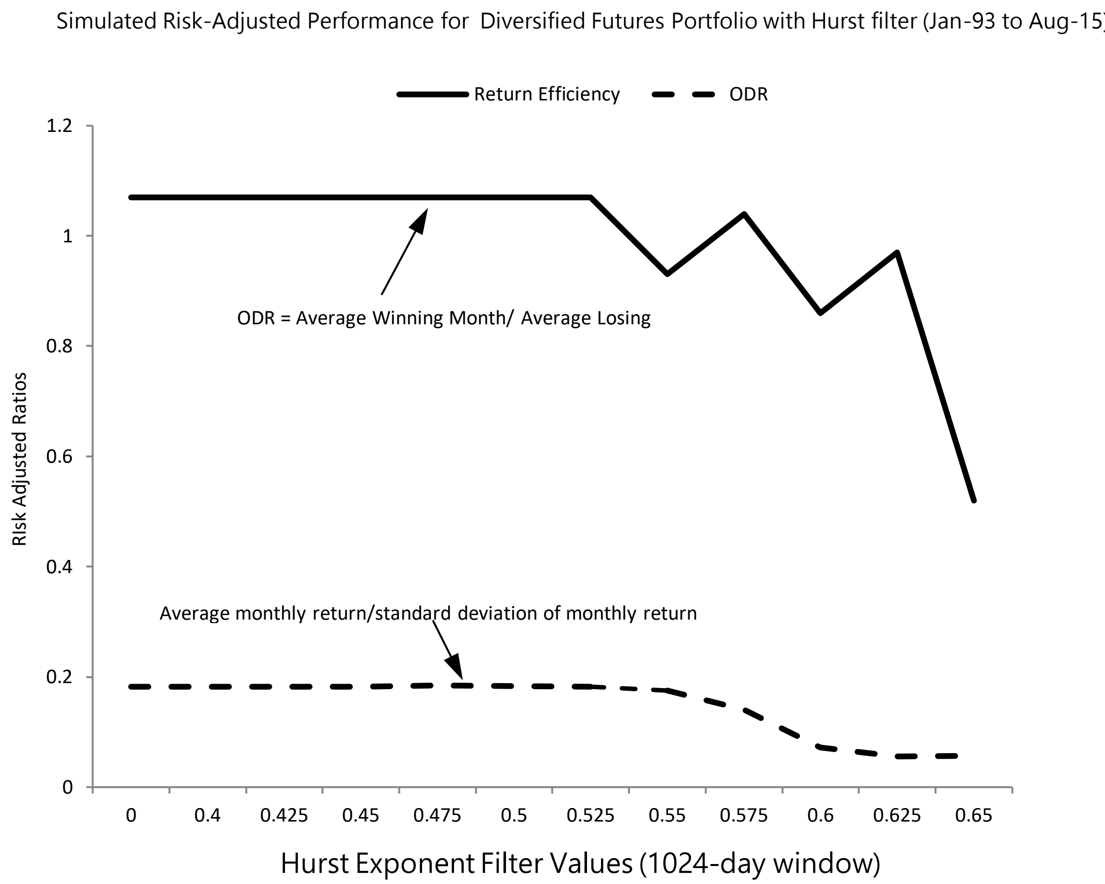

Figure 8: Risk-adjusted performance is measures do not improve with heavier filtering as the value of the Hurst exponent used to filter trades increases.

(Place in section IV H)

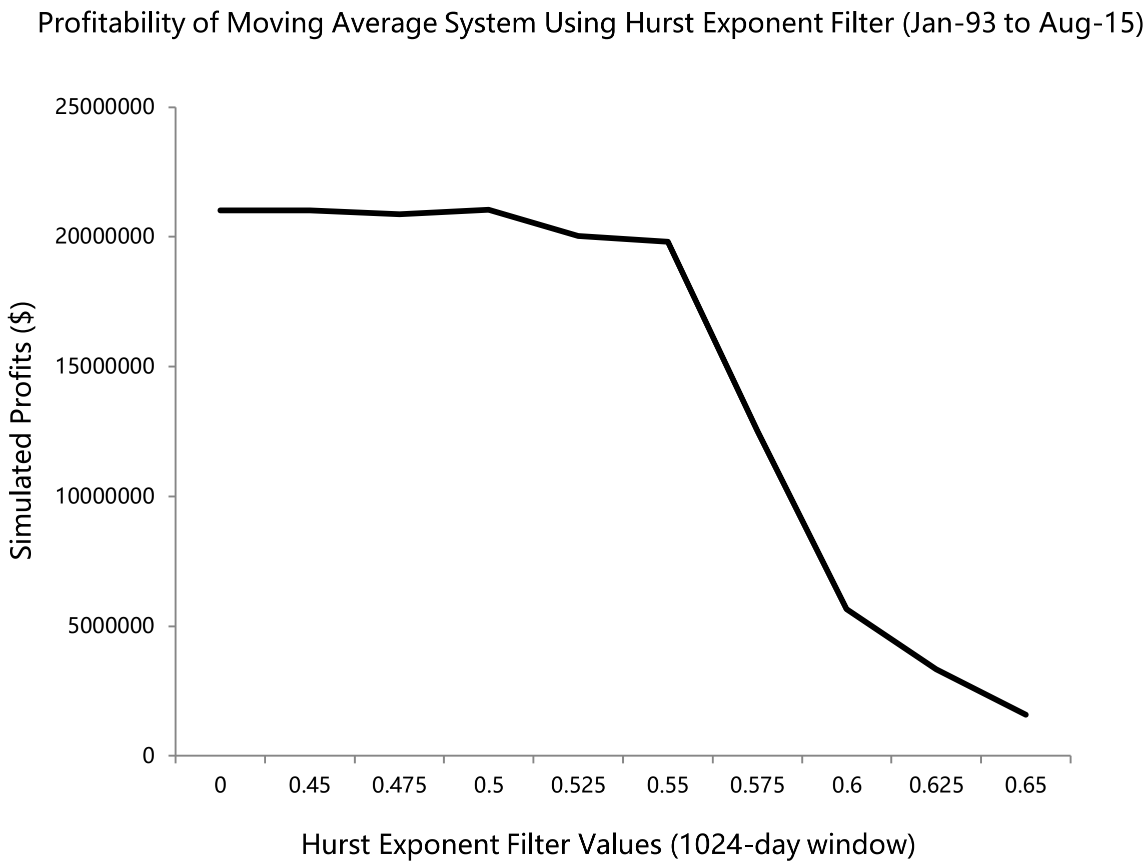

Figure 9: Profitability of moving average trend following system (using 20-bar shorter average and 100-bar longer average) on a diversified portfolio of futures markets decreased as the level of Hurst exponent used for filtering increased. For a sample size of 1024, the range of critical values ranges from 0.56-0.65, i.e., if the Hurst Exponent is greater than this range, then the data show the presence of trends. The Hurst Exponent was virtually never greater than 0.65. Simulated gross profits were essentially unchanged between say 0 and 0.50, i.e., the filtering had no effect. This implies there virtually no values of the Hurst Exponent below 0.50 or so.

(Place in section IV H)