Abstract

The purpose of this paper is to provide a usable framework for detecting, measuring and exploiting trends in financial markets. Using technical analysis (TA) indicators we challenge Efficient Market Hypothesis (EMH) that says that markets are random and that is not possible to regularly outperform a passive investment strategy.

If a fair coin is flipped a thousand times, it is bound to land on one side at least a few times in a row. Similarly, stocks may go in one direction for a few days in a row. Occasionally, such runs repeat a few times in a month. When looking at a chart of such stock, one would say there was a trend. This is in line with Efficient Market Hypothesis. Indeed, if every up day was followed by a down day, markets would be perfectly predictable. In other words, existence of trends is an expected consequence of EMH.

Trends in financial markets have been researched by many trend followers as well as academics. Traders simply find trends and trade according to the old saying “cut losses short and let winners ride”. Academics, on the other hand, say trends can be found even in random walk data, like a coin that landed on heads a few times in a row. Therefore, they say, any profits arising from trend trading must be explained by luck. Both parties seem to agree that trends exist, but there is no consensus on how to measure market trends and whether success of trend followers is due to luck or skill.

The first part of this paper presents research of market trends and shows reasons why trend following should be a profitable approach to trade markets. We point out Hurst exponent's flaws that render it useless to detect long term memory in market prices. As an alternative, we propose to use trading systems and TA indicators. We introduce a new TA indicator called “Min-Max” that can be used to read prevailing market direction. We follow trends using this indicator alongside with Simple Moving Average and compare results against trend following on various types of random walk data. If markets are mean reverting by nature, trends would be shorter and weaker than those in random walk data. In this case trend following using TA indicators would be unprofitable. On the other hand, if markets are trending by nature, trends found in historical prices would be stronger and longer. As a

consequence trend following would show profitability. Our tests confirm that historical prices trend more than random walk data. By detecting this difference, this paper will show that trend following with TA indicators regularly outperforms passive investment strategies and generates out-sized returns.

For this research we used historical market data that include stocks, commodities, currencies and indices dating back to 1971 where possible. Such a wide range and a long period of time allows us to measure reactions to events such as the crashes of '87 and '97, invention of the Internet, the Internet stocks bubble, attacks of 9/11, Russia's default, the Argentine economic crisis, the forming of the European currency and many more that influenced markets around the world. For random data we used a computer simulated fair coin toss where the price has an equal probability of going up and down. We also test random data with artificially introduced random trend elements. Such data has weak, but unpredictable, trends that can be exploited by a good trading system. For the last type of random data, we used scrambled historical prices. Scrambling removes price's long term memory and therefore makes data more similar to a fair coin toss data.

The second part of this paper is an example of a quantitative trend following strategy, that combines SMA with the Min-Max indicator. The strategy is fully mechanical and is very simple to implement and follow in live trading. This strategy beat the S&P 500 index by a wide margin on years 1991 to 2010. We also backtested this strategy on random data, showing that hypothetical profit earned on historical market prices is significantly larger than profit earned on random data.

In the last part of this paper, we use two hypothesis testing methods to show how the results of our trading strategy could not be achieved simply by luck. The first method is a P-value measure on the difference between profitable and unprofitable days. In the second method, we use random buy and sell orders to generate 75000 random historical results, each with the number of trades equal to our historical testing. Both methods show how unlikely it is to achieve positive results by luck, and thus confirm this strategy's ability to exploit a market inefficiency.

Theoretical basis and a practical example of trend following

Łukasz Wojtów

Zamość University of Management and Administration

*) Łukasz Wojtów can be contacted at lukasz.wojtow@gmail.com

Table of Contents

1. Introduction

There are two schools of thought when it comes to financial markets: one is Efficient Market Hypothesis (EMH) that claims markets and their participants are rational, never make mistakes and price changes are randomly distributed. According to EMH believers, success in trading is only due to luck. Because markets are random, the probability of beating the stock market in any given period is 50 %. Therefore, out of 1000 investors, there will be 500 that beat the market in one year. Out of those 500, there will be 250 that beat the market in two years in a row and so on. Out of one million investors there will be one who beats the market 20 years straight. But is probability of beating stock market really 50 % and what does this number actually mean? Surely not all strategies are equal. For example, selling highly leveraged naked options is a certain way to go bankrupt. If chances depend on approach and leverage, perhaps there is a way to make them higher than 50 %? For example, if markets have tendency to trend, then trading in the direction of main trend should put the odds in traders' favour.

On the other side of the argument there are traders who claim markets are not rational, to a small extent predictable and inefficiencies can be exploited to generate out-sized returns. Traders who know this, trade only when probabilities are on their side. But who is on the other side of the trade? If all market participants are rational and never make mistakes, how come one is buying and the other one is selling at exactly the same price and time? Does successful trading boil down to finding those mistakes and taking the other side of the trade? Is trading against a trend a mistake?

This paper provides empirical research and statistics of long term market trends in order to answer these questions. In our research we use mostly traded stocks, indices, currency pairs and commodities. Chapter 2 introduces various types of random walk data that will be used to compare trends characteristics with real market data. Chapter 3 shows different measures of trends, showing that real prices trend more than random walk data. Chapter 4 tries to provide reasons for trends in financial markets while Chapter 5 is an example of a trend following method. Chapter 6 discusses the statistical significance of the achieved results.

2. Random Data

Technical Analysis (TA) is based on the assumption that past prices can influence the future, that is, decision making is based solely on past prices. Of course, such an approach will never be profitable with random data. Although it is still possible to find similar TA signals (chart patterns, moving averages crossovers and so on), making trading decisions is futile, as future prices are not influenced by the past. It is interesting to see how various trading systems and measures of trends perform with different types of random data compared to real market data. By finding quantitative differences, we can identify inefficiencies that can be utilized to generate alpha.

For the purpose of this paper, we introduced three types of random data:

1. Data simulating a fair coin toss. For each head a security's price rises 1.010101010101 %, for each tail, a security's price falls 1 %. The difference between the rise and fall values is necessary to compensate for so called "drag", i.e. adding 1 % and subtracting 1 % yields price equal to 99.99 % of the original price. If this were repeated a sufficient number of times, the price would eventually fall to zero. Twenty different securities were generated, each with 10400 prices ("days"), simulating 40 years of history.
2. Biased coin toss. For each head, a security's price rises 1.010101010101 %; for each tail, security's price falls 1 %. Only one coin is tossed at a time, but there are two coins in simulation: one has a 51 % chance to land on heads and the other a 51 % chance to land on tails. Each coin is tossed a random number of times (between 0 and 1000) and then replaced by the other coin. We don't know which coin is being used at any given moment. This type of data will have weak trends, but it is not possible to predict the price much further in advance because of its random nature. Again, 20 securities were generated, 10400 days each.
3. Scrambled real data. We used markets' historical daily changes to create new time series, but days and securities were chosen in random order. For example, we started quotes at

$1000. To simulate next day we added DJIA's percentage change on 10th September 1982. For the next day we used Bank of America's change on 23rd August 2009, for another day we used gold's change on 5th February 2002, and so on. As each day is is used, it is removed from the pool of available changes. The total number of available days was 187514, out of which 20 different securities were generated. This type of data will have daily changes exactly the same as real market data, but any long term price memory (i.e. past to future influence) is removed.

3. Trend measures

By measuring the length, strength and likeliness of trends in financial data, we can answer the question of whether trend following is a valid trading methodology leading to out-sized returns, or just guesswork with any profits being a fluke. Unfortunately, there is no consensus on how to measure trends and each method has its own weaknesses. In the next chapter we present a few methods of detecting trends and show their results on various type of data.

3.1 The Hurst exponent

The Hurst exponent is often used as a measure of long term memory (autocorrelation), i.e. whether past data influences future data. It has been used to find trends in financial data in [1],[2] and [3]. We believe it has some weaknesses that disqualify it as a tool to measure market trends. For example, let's imagine a security where at the beginning we have a strong up trend, and then a long consolidation period with prices remaining relatively flat. Chart 1 is an example of such data (a simple moving average of 200 periods is shown in red):

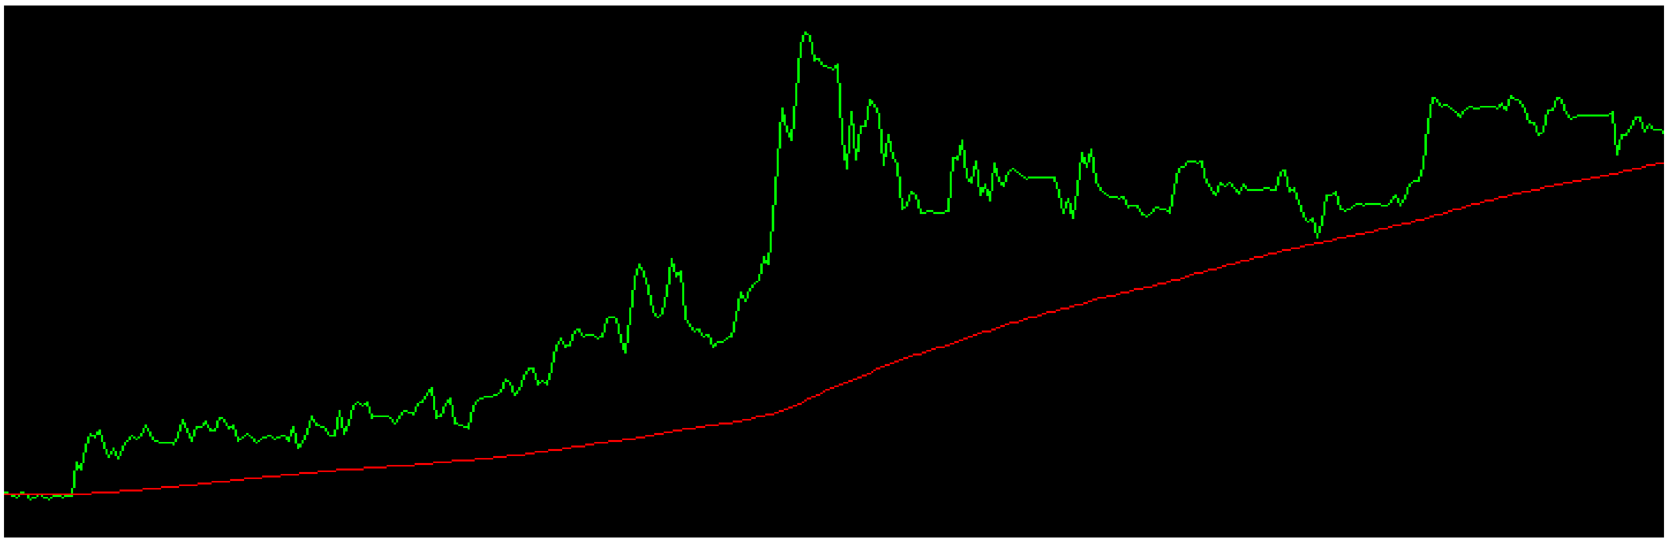

Chart 1. Hypothetical market prices with SMA

When calculating the Hurst exponent, the flat period will bring an end result closer to 0.5 (meaning no trends). However, for a trader using a moving average as a trading signal, the flat period will not mean losses as no new signal is generated. Whilst the Hurst exponent looks at all data, trend following systems can wait out flat periods with either no positions or a position previously established and will not be negatively affected. Also, it is possible that markets sometimes trend and sometimes mean revert, averaging each other out. For example, breaking a trend could occur because the price is too far from its fundamental value. In such case, the Hurst exponent would be close to 0.5, yet a skilled trader would be able to make a profit. For this reason, we believe there are better, simpler and more practical ways to measure market trends. In some way, even complete trading systems themselves can be used to analyse time series [4].

3.2 Simple Moving Average as a measure of trends

Let's consider a simple trading system which goes long at the next day's open if today's close is above the simple moving average of specified length, and goes short if the price is below it. Such a system will achieve profitability if there are strong trends in traded instruments and will incur losses where no trends are present. In some way, the system's total profit can inform us about the nature of data. If there is a profit, the data seems to be trending, if there is a loss, the data seems to be mean reverting.

We checked every SMA's length between 130 and 390 days (6 – 18 months). The years tested are 1971-2010, and no commissions are taken into account. Every time a signal is generated we opened a position worth 10 % of current account value. Chart 2 shows the profit factor achieved by such a system for different SMA's lengths.

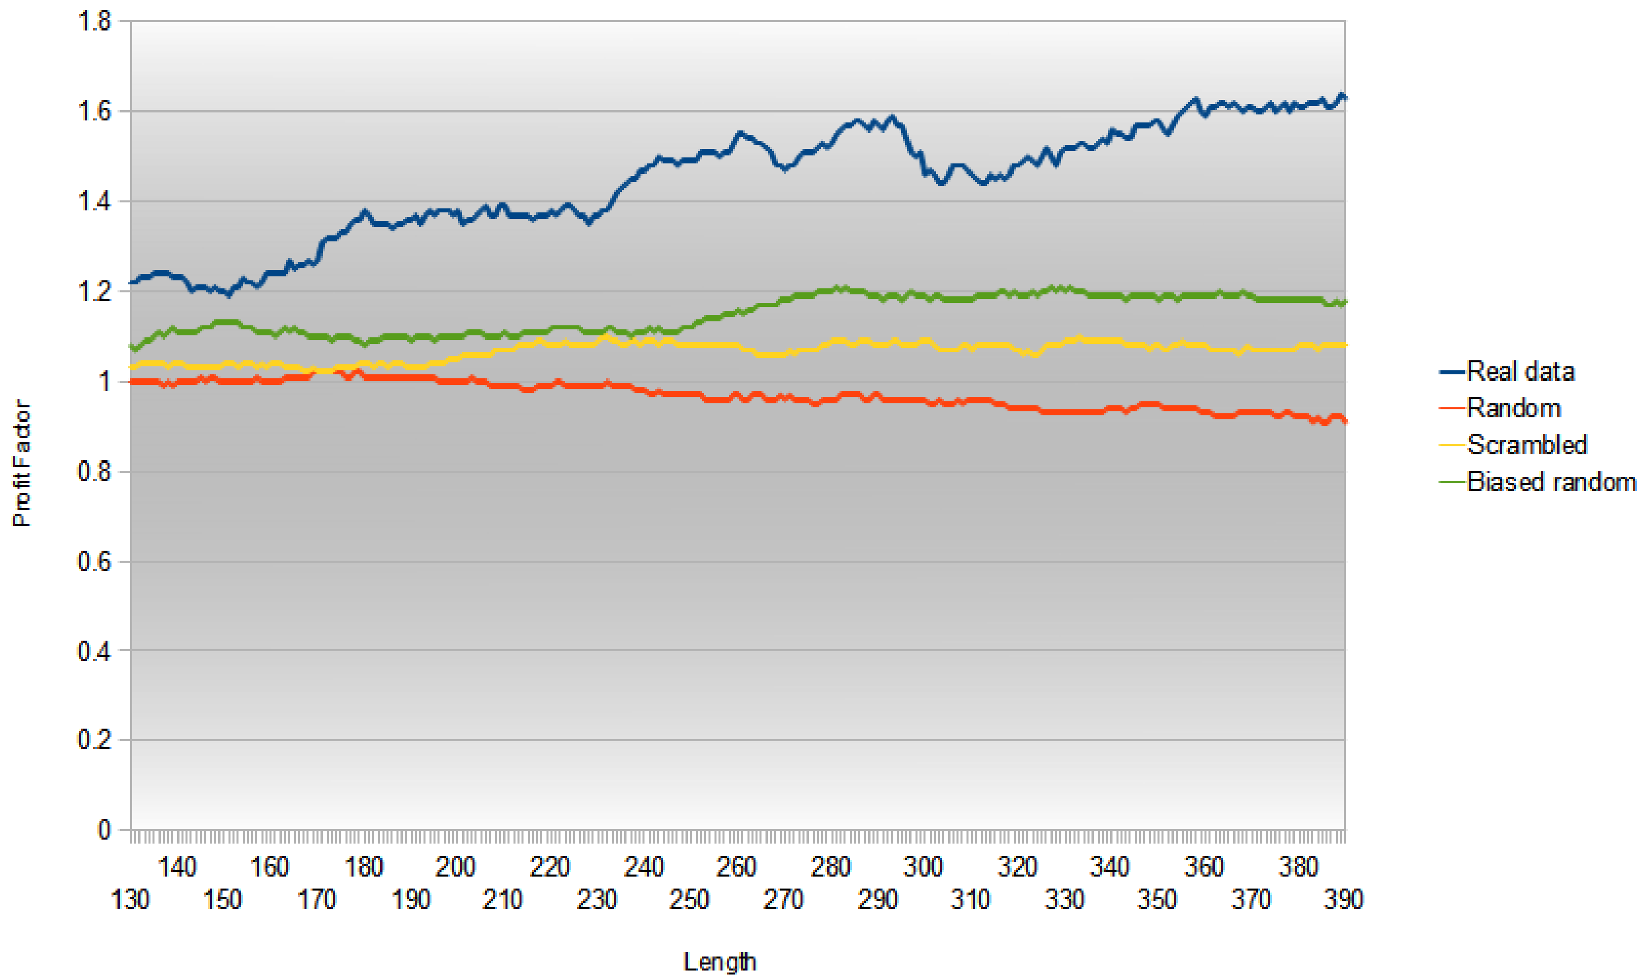

Chart 2. Profit factor of Simple Moving Average

As expected, the profit factor for real market data is above the profit factor for any other type of data, for all SMA's lengths. Interestingly, the profit factor stays above 1.0 (which indicates profitability and thus market trends) for scrambled data. This could be explained by the fact that some markets (namely stocks and commodities) are prone to long term appreciation due to inflation. Simply put, if there are more up days than down days, scrambling the data does not remove long term "inflationary" trends.

3.3 The Min-Max Indicator

At this point we introduce a new Min-Max indicator, which we believe can be used to read long-term market trends. In order to calculate indicator's value one needs to perform following steps:

- decide the length of the indicator, for example 260 days,
- find a maximum price during last “length” of days
- count how many days ago was the maximum, this will be “maximum index”
- find a minimum price during last “length” of days
- count how many days ago was the minimum, this will be “minimum index”
- calculate difference: minimum index – maximum index

If the result is positive, the trend is up, if the result is negative, the trend is down. Visually, one can draw a straight line from the maximum to the minimum index. The slope of the line indicates the direction of the main trend as shown below:

Chart 3. Silver's chart with Min-Max indicator

We tested the Min-Max indicator in the same way as we tested SMA with opening a long position if trend is up, and going short if the trend is down. The profit factor achieved will tell us about data's nature. We checked every length of the Min-Max indicator between 130 and 390 days, with no commissions, from 1971 to 2010.

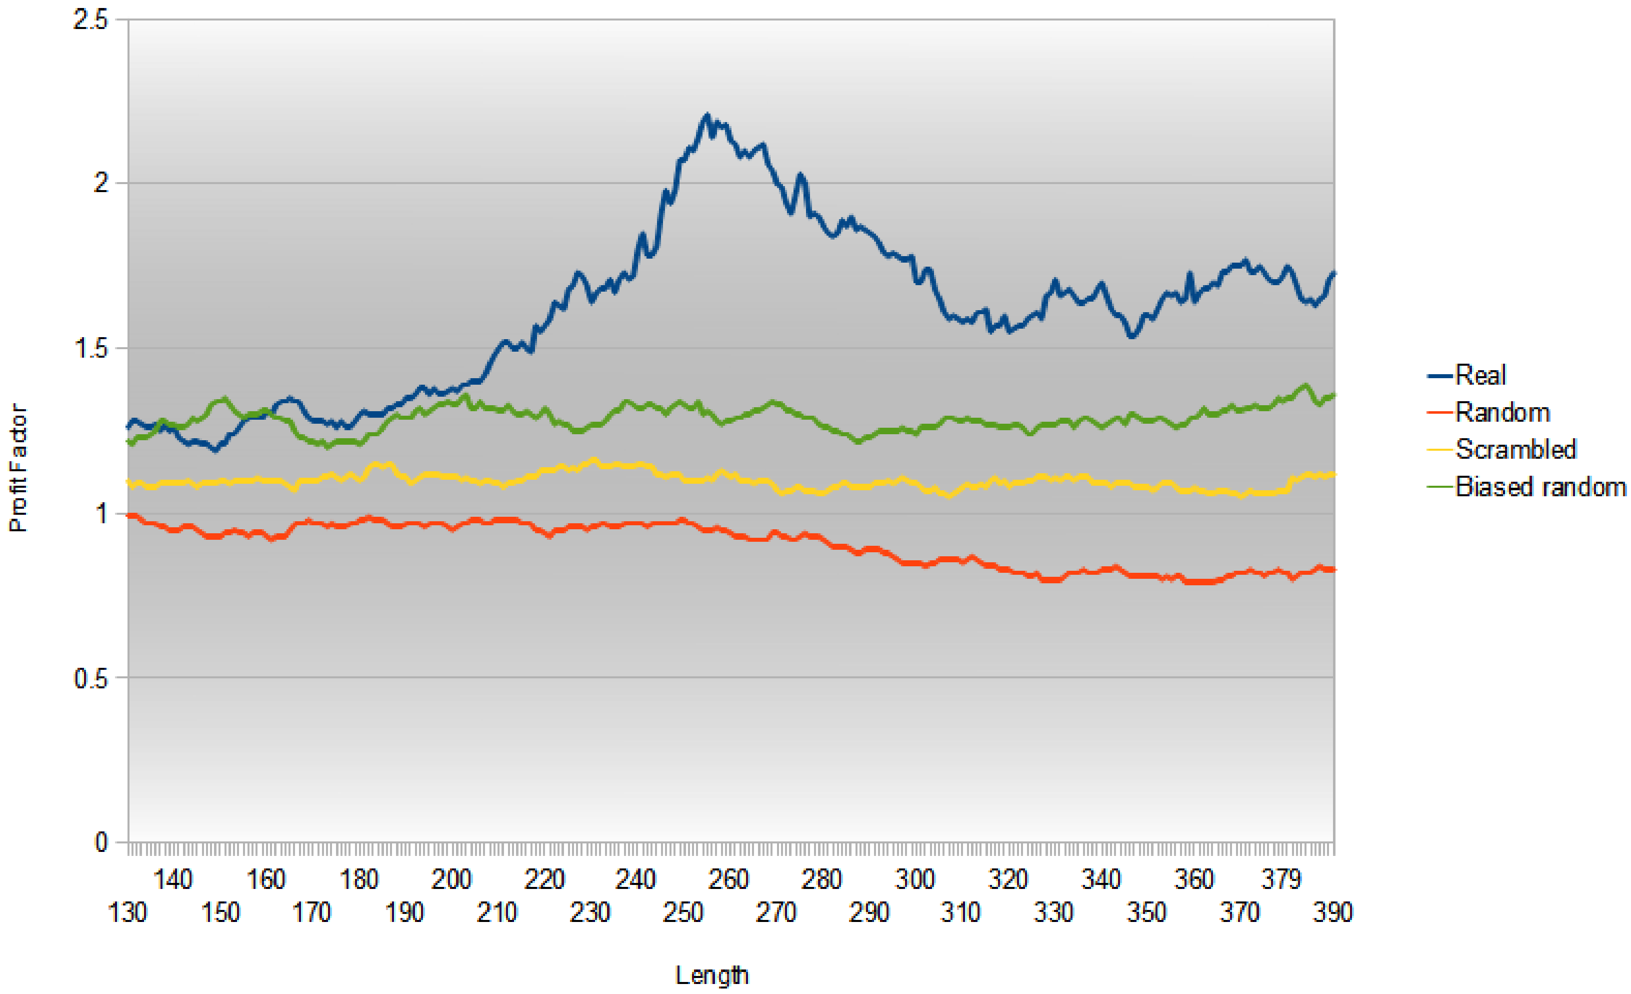

Chart 4. Profit factor of Min-Max indicator for various types of data

As shown in chart 4 above, again, real market data shows the biggest profit factor, thus suggesting trends are longer and stronger than trends in random data. Table 1 below compares the average profit factor for both indicators.

	Real Data	**Random**	**Scrambled**	**Biased Random**	**Average row**
SMA's Profit Factor	1.44	0.96	1.06	1.14	1.15
Min-Max Profit Factor	1.61	0.9	1.09	1.28	1.22
Average column	1.52	0.93	1.07	1.21	

Table 1. Summary of profit factors for various indicators and data types.

3.4 Conclusions

The above results show that historical prices are prone to trending more than any type of random data. The profit factor for market historical prices is significantly higher, meaning following long term trends is a valid trading methodology and generates alpha if done skilfully. In the Min-Max example, the average profit factor for scrambled data equals 1.10 with a standard deviation of 0.023, meaning that the average result for real data (1.61) was 22 sigma away from the mean and the best result (2.20) was 47 sigma away from the mean. That shows historical market trends cannot be explained by luck or coincidence.

4. Reasons for trends in financial markets

There could be many reasons that explain long term trends such as margin calls or stop losses, constant leverage and hedge funds.

Margin calls or stop losses. When a trader opens a long position with leverage 2:1, the trader's profit grows twice as fast as the appreciation of that traded security. This is also true for losing. If the underlying market falls 50 %, trader's loss is 100 %. In such a situation, the broker is responsible for closing the trader's position (in fact the broker will close this position just before the 100 % loss mark). In order to close a long position, the broker has to sell and often does so at market price, thus pushing the price even lower. The same is true if a trader wants to close a losing position. In other words, falling prices cause some traders to sell.

Constant leverage. Imagine a situation where a trader opens a long position with leverage 2:1 and would like to keep leverage constant during the time the position is opened. If security rises 50 %, the trader has 100 % profit, but that means that leverage falls to 1.5:1. In order to bring the leverage to 2:1, the trader needs to buy more, thus pushing the price higher. In other words, rising prices cause some traders to buy. It works other way round when losing; if the price goes against a trader, the leverage becomes higher. In order to keep the leverage constant, one needs to reduce the position.

Hedge funds [5]. This concept is based on the observation that investors commit their money to funds that have performed well recently and funds invest new money into the same assets. This is also true the other way round; investors withdraw money from funds performing badly. In order to meet cash requirements, funds close positions in their assets (where prices are going against the fund).

There could be other, psychological reasons that generate trends in financial markets, including investors' fear and greed that prompt them to buy in rising markets and sell when prices fall. However this is beyond the scope of this research.

5. Example trading system

For this research we created a mechanical trend trading system based on objective rules that can be backtested on historical prices. Simple moving average is used as a signal indicator. System opens a long position at tomorrow's open if today's close is above SMA and opens a short position if below. Close a long position at tomorrow's open if today's close is below SMA and vice versa for short. Trade only in the direction given by the Min-Max indicator, i.e. if Min-Max shows an uptrend – open long positions only, if Min-Max shows a downtrend – open short positions only. Use average daily change to decide stop loss distance and use fixed fractional as position sizing rules.

Choosing a right set of parameters is often the "make or break" of live trading results. There are numerous ways to deal with this problem, one of which is walk forward optimization. Every year, we optimize the system on last 5 years. Then, we find the best MAR ratio in order to choose parameters to trade for the next year. The parameters were found by an exhaustive search using the following values:

Simple moving average: from 60 days to 260 days, step 5

Min-Max: from 130 days to 390 days, step 10

Stop loss size: from 3 to 5 times 20-days average daily change (close-to-close range), step 1

No commissions are taken into account when optimizing. The best parameters are shown in Table 2:

Optimization Years	Best Min-Max	Best Sma	Best StopLoss
1985 - 1990	260	90	3
1986 - 1991	260	60	3
1987 - 1992	260	60	3
1988 - 1993	260	75	5
1989 - 1994	260	65	3
1990 - 1995	260	85	4
1991 - 1996	340	85	5
1992 - 1997	310	70	4
1993 - 1998	310	90	4
1994 - 1999	310	90	4
1995 - 2000	300	100	5
1996 - 2001	300	135	4
1997 - 2002	270	135	4
1998 - 2003	270	135	4
1999 - 2004	260	260	5
2000 - 2005	270	190	4
2001 - 2006	270	60	3
2002 - 2007	300	60	3
2003 - 2008	300	60	3
2004 - 2009	260	175	4

Table 2: Best parameters for the strategy found by optimizing 5-year rolling periods.

And so, during year 1991 we used parameters 260, 90 and 3 for Min-Max, SMA and Stoploss respectively. In year 1992 we used 260, 60 and 3, and so on.

5.1 Risk adjustment

One of the problems facing an active manager who decides to trade a mechanical system is to choose the right amount of risk. This is often expressed as a percentage of the capital risked per trade. Choosing risk that is too small may result in small drawdowns, but also low profits and the manager may not beat the benchmark. On the other hand, choosing risk too large could result in bankruptcy even with a profitable system. We decided that we would like the maximum drawdown in the walk forward test to be 48 %. Although too big for institutions, this should be acceptable for individual investors. Also, this is S&P 500's maximum drawdown during the years 1971 – 1990 and so it seems fair to require active managers to be able to weather-out drawdown at least as deep as “Buy and Hold” investors.

Normally, a trader can optimize the risk taken during backtests and choose a required value. This is not possible in walk forward optimization as the parameters (especially stop loss size) keep changing and so is an optimal risk level. To work around this problem, we backtested every set of parameters from Table 2 for years 1971 to 1990 with a risk equal to 0.1 % of capital, and then calculated risk required to bring the maximum drawdown to the desired 48 %. Table 3 summarizes the risk used in the walk forward test:

Walk forward year	Parameters			Maximum Drawdown 1971 – 1990 with risk 0.1 %	Risk used in walk forward test (% of capital)
	Min-Max	SMA	Stop Loss		
1991	260	90	3	-6.81	0.9
1992	260	60	3	-10.31	0.6
1993	260	60	3	-10.31	0.6
1994	260	75	5	-6.41	1
1995	260	65	3	-10.12	0.6
1996	260	85	4	-7.53	0.8
1997	340	85	5	-7.48	0.8
1998	310	70	4	-7.9	0.8
1999	310	90	4	-7.87	0.8
2000	310	90	4	-7.87	0.8
2001	300	100	5	-6.97	0.9
2002	300	135	4	-10.48	0.6
2003	270	135	4	-10.72	0.6
2004	270	135	4	-10.72	0.6
2005	260	260	5	-11.57	0.5
2006	270	190	4	-12.4	0.5
2007	270	60	3	-7.37	0.9
2008	300	60	3	-7.38	0.8
2009	300	60	3	-7.38	0.8
2010	260	175	4	-6.73	0.9

Table 3: Risk and parameters used by strategy in walk-forward test.

5.2 Results

Chart 5 shows the equity graph for this trading system with an account starting balance set to 330.20 (S&P 500 opening value on 2nd January 1991). The yellow line shows the account value, the red line shows S&P 500 for the same period. Chart 6 shows the experienced drawdown. Commissions are taken into account.

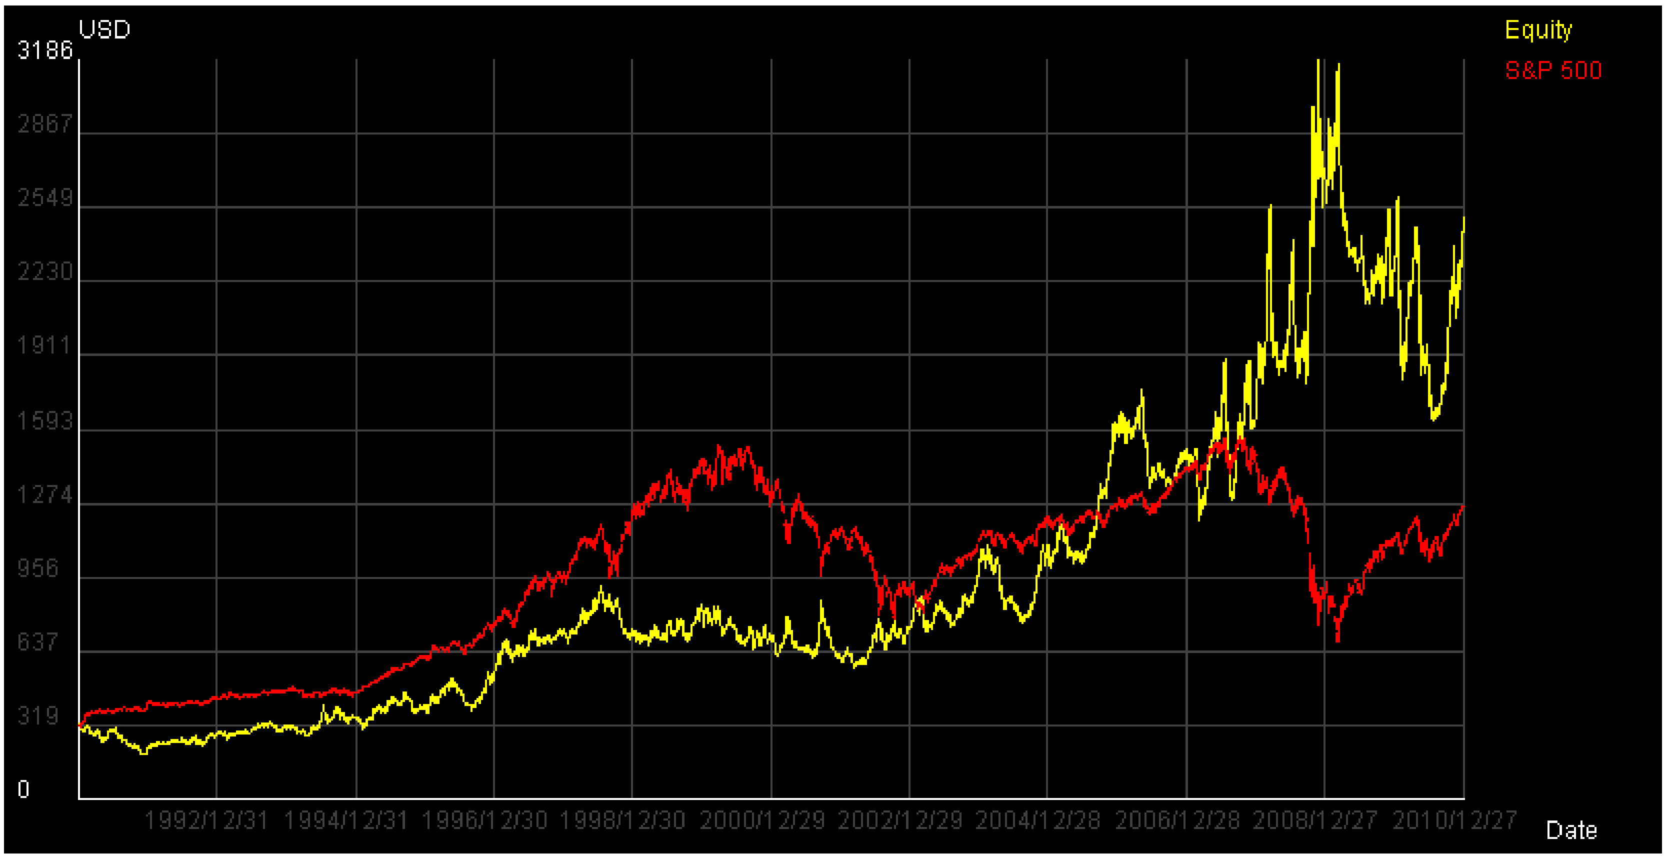

Chart 5. Strategy's results compared with S&P 500, years 1991 – 2010

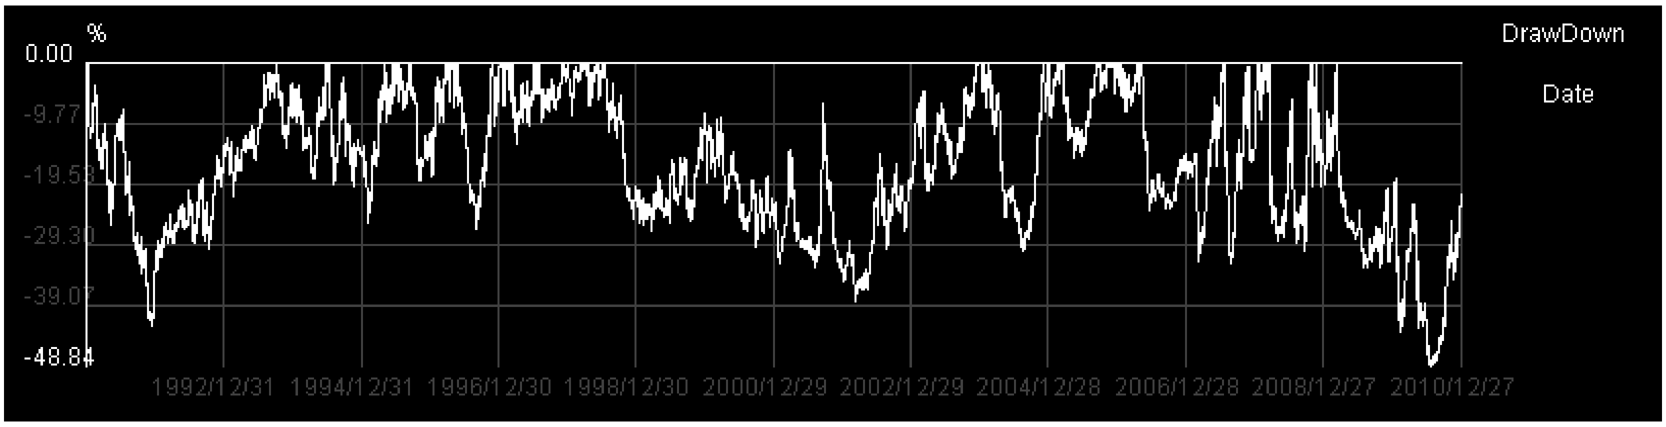

Chart 6. Strategy's drawdown. Maximum value was about 49 %.

During the 20 years of the walk-forward test, this strategy achieved 659 % profit compared to 285 % made by S&P 500. The maximum strategy's drawdown was 49 %, S&P 500's maximum drawdown was 57 %. The strategy's profit to drawdown ratio was 13.4, while for S&P 500 the ratio was 5.0.

It is worth noting, that during the 1990s this strategy did not beat the benchmark. This could be because 1990 to 2000 was the best decade in S&P 500's history, with total growth at 316 %, meaning annual growth of an unsustainable 12 % (all time average annual growth is 3.5 %). For comparison, chart 7 shows the account's growth from year 2000:

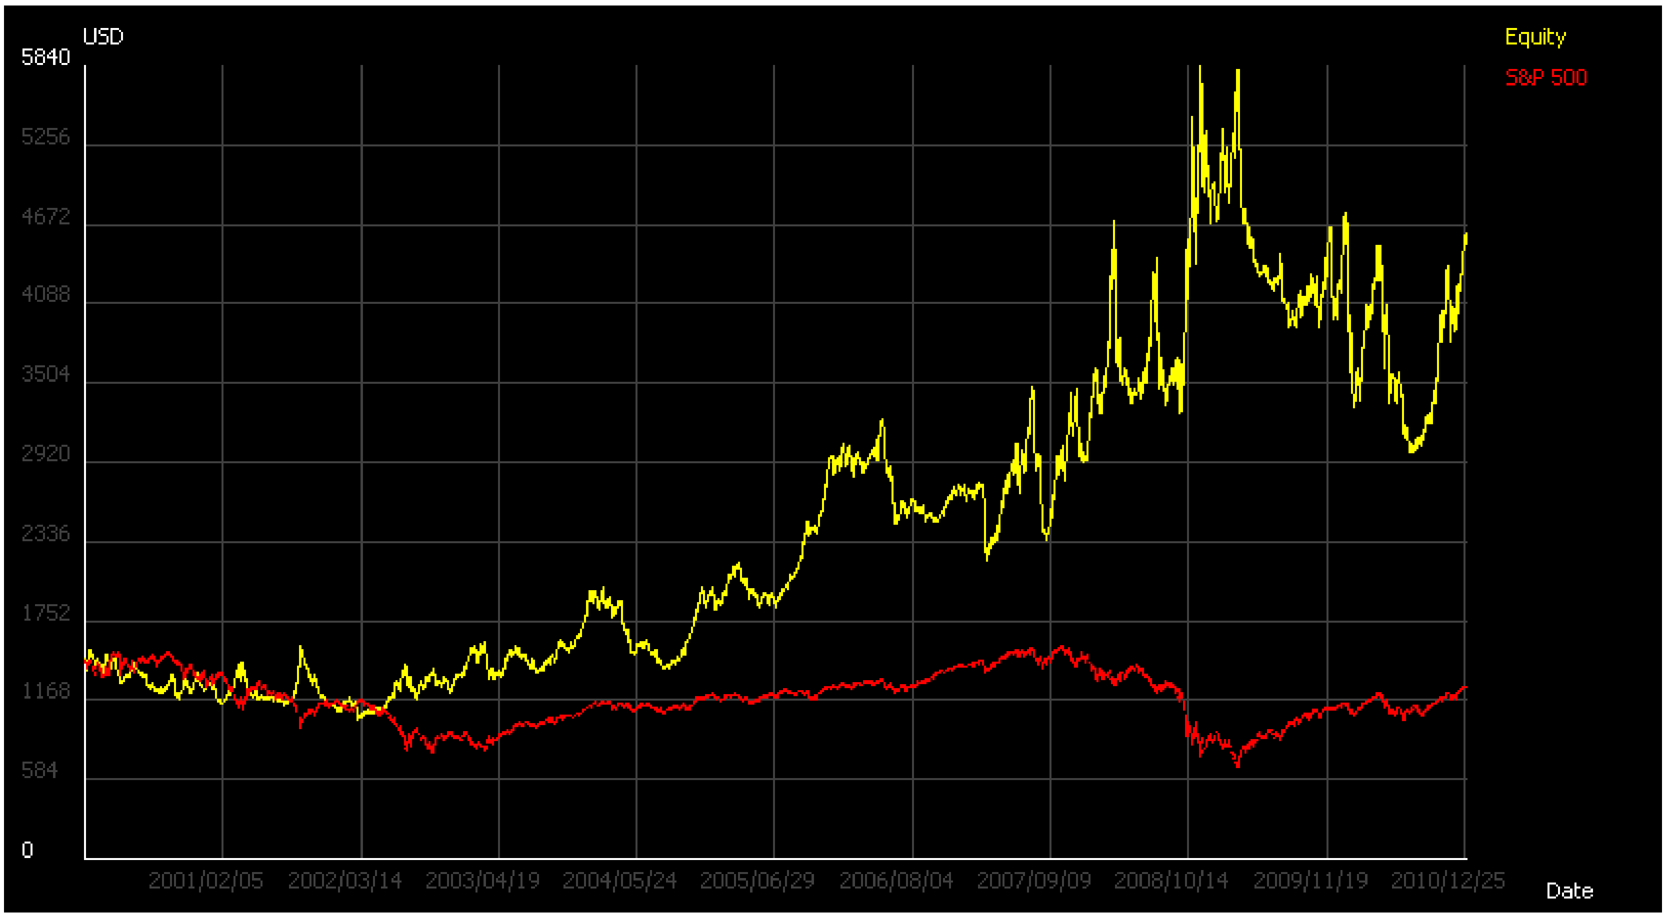

Chart 7. Strategy results compared with S&P 500, 2000 – 2010

Results for each market are shown in Table4 (years 1991-2010):

Name	Total profit (%)	Number of transactions
Apple	72.3	133
AudUsd	7.30	157
Bank of America	-1.50	149
Coffee	32.00	163
DJIA	16.16	182
EurUsd	13.47	138
Ford	3.33	172
FTSE 100	19.47	159
Natural gas	-12.02	160
GbpUsd	11.30	164
General Electric	21.8	151
Gold	14.06	168
J. P. Morgan	8.84	169
NASDAQ	44.64	121
Nikkei	26.97	127
Silver	-45.46	204
Sugar	-7.22	165
UsdCad	-11.05	176
UsdChf	17.81	130
UsdJpy	38.29	148

Table 4. Strategy's results for each market individually

The strategy was profitable on the vast majority of securities (15 out of 20, or 75 %) proving its robustness, necessary from a trading system. Chart 8 below shows the profit for each transaction. Winning transactions are evenly spread out through the whole testing period, also confirming the system's robustness. The biggest profit (+37.1 %) was made by a "Buy" transaction on Apple, bought on 6th February 2004 at $11.22 and sold on 29th March 2006 at $59.13. The biggest loss (-2.7 %) was made by a "Buy" transaction on UsdCad, bought on 19th July 1991 at 1.1600, and sold (stopped out) on 26th July 1991 at 1.1497.

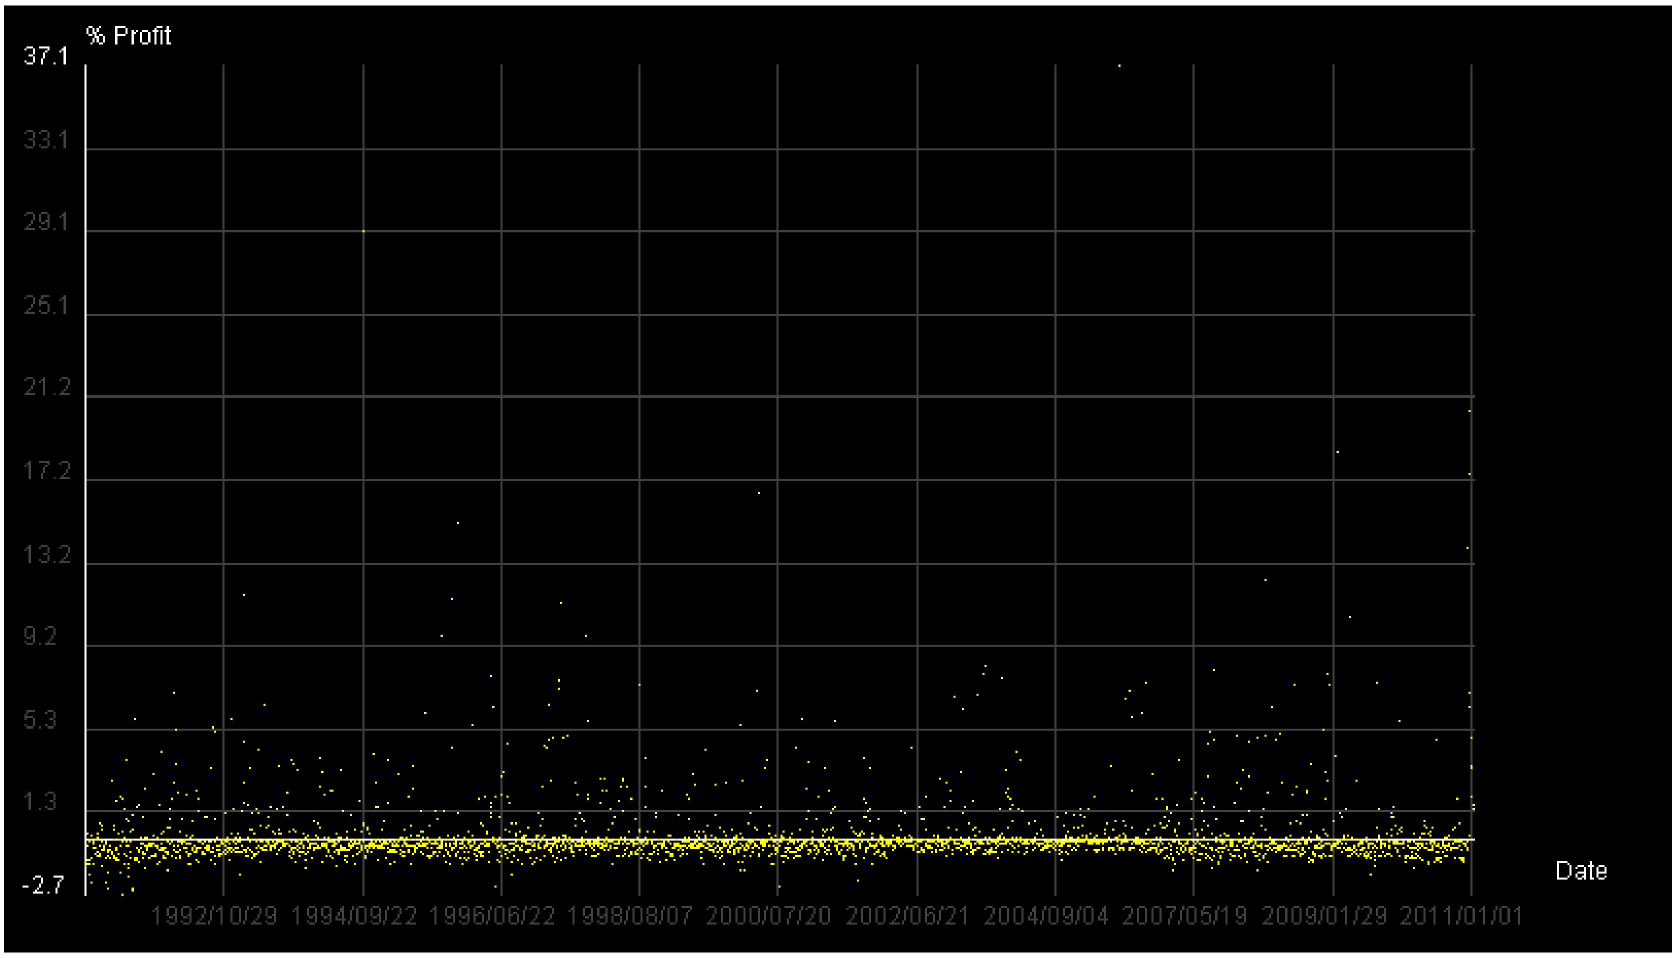

Chart 8. Transactions' profits made by system

So far, when talking about risk, we mean percentage of capital risked per transaction and maximum historical drawdown. One more, often overlooked, measure of potential risk taken by a strategy is leverage, which can have a devastating effect on an account. This can happen even if historical backtests show the maximum drawdown is acceptable. Recent events with MF Global and earlier LTCM, illustrate that even the most intelligent and knowledgeable investors can fail if things get out of hand. Such events cannot be discovered by backtesting simply because they have never happened before. Nevertheless, they all have one thing in common – a high leverage killed a seemingly safe strategy. It is highly recommended not to exceed a certain value of leverage, values in the low tens are commonly considered as a reasonable maximum. We believe, every investor should monitor his/her own leverage on a day-to-day basis. Chart 9 shows leverage used by our strategy.

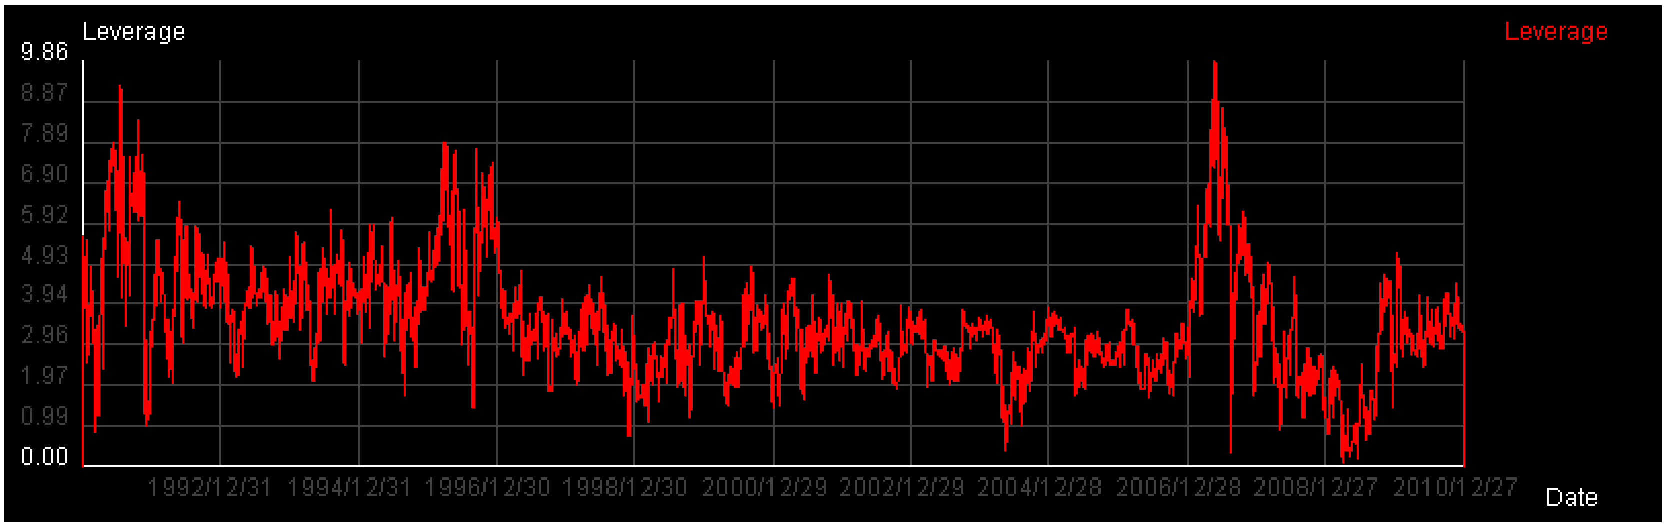

Chart 9. Leverage used by strategy

Table 5 shows our strategy's results. For comparison, results on random data are also shown. As can be seen, real data results are better by all metrics except for the percentage of profitable transactions. Another startling difference is high average profit to average loss ratio. For some reason, our strategy closes losing trades the quickest on real data. This in itself, could be another market inefficiency worth looking into.

Metric	Real data	Fair coin	Biased coin	Scrambled
Total profit	659 %	20 %	80 %	-32 %
CAGR	10.7 %	0.94 %	3.01 %	-1.96 %
Maximum drawdown	-48.84 %	-40.05 %	-17.95 %	-63.53 %
MAR ratio	0.22	0.02	0.17	-0.03
Profit factor	1.24	1.03	1.10	0.93
Expectancy	0.19	0.02	0.07	-0.05
Total transactions	3136	2626	2610	1318
Percent profitable transactions	23.4 %	34.2 %	34.4 %	29.0 %
Average profit	15.15	13.59	14.35	16.68
Average loss	3.73	6.90	6.82	7.35

Table 5. Strategy's performance

5.3 Price shocks

The attacks of 9/11 caused price shocks in financial markets around the world. It is important to know, that strategy used in real life is shock proof. Our strategy opened short DJIA position on 7th August 2001. The position was held during the 9/11 attacks and was closed on 16th September, with a profit of +1.37 %. All together, our strategy made +21.2 % profit during September 2011. Another price shock was the flash crash of 6th May 2010. Our strategy lost only -6 % on that day.

6. Statistical significance of achieved results

Luck plays a huge role in short term results, but is irrelevant in the long term. To illustrate this point, let's think of trades as a fair coin toss. For example, after 10 tosses we might get 7 heads and only 3 tails. If every head represents a winning trade and every tail represents a losing trade, the trader's result would be 7:3, or 4 wins net. Depending on the risk taken, this could mean large profits on the account. A novice trader would be ecstatic, but more experienced traders know that 10 trades is not statistically significant. The probability of getting 7 or more heads in 10 tosses is 17 %, too high to exclude an element of luck. In other words, by measuring a system result's statistical significance, we can see if the system has an ability to read markets correctly or if the results can be explained by luck alone. If a system can read markets, we assume that it will be able to do so for at least some time and the system will make money in the future.

Statisticians know many different tests that enable us to estimate if a particular result is due to chance or not. Yet many traders use rule of thumb, like a certain minimum number of trades or a minimum value of profit factor. Such rules are often based on intuition or prior experience and can judge systems incorrectly. In order to reliably estimate a system's probability of making profit in the future, traders have to employ more formal statistical hypothesis testing. One of many possible tests is P-value, a probability of obtaining a result at least as extreme. In the example above it is the sum of probabilities of getting 7, 8, 9 and 10 heads out of 10 coin flips. If P-value is smaller than the required significance level, the null hypothesis is rejected. Rejecting the null hypothesis means the result is statistically significant.

Our null hypothesis is: the system's trades are not better than random, and no real trades should be placed on system's signals. If we can reject this hypothesis at a reasonable confidence level (which we set at demanding 99.5 %), we can use the system in real world trading. For the purpose of this test we will treat every day as one fair coin flip, looking at every market separately. Our walk-forward test trades 20 different markets, some of them with almost 20 years of history. All together, this gives almost 100000 days to analyse, but due to the fact that the system is often out of the market, this number falls to 63309 days, or 63309 coin flips. A fair coin has equal chances to come up heads or tails, so on average we should get 31654.5 of each and probability of getting 31654 or more heads is obviously 50 %. We will now calculate probability of getting as many heads as our system's number of winning days. If the system cannot read markets, its result should be about as good as a coin flip (50 %). To do this, every day when a position in a market was held for at least a portion of the day (a position could be stopped out during the day), we checked the market's change at the day's close relative to a previous day's close. If close-to-close change was positive – we counted the day as winning if the system held a long position, and as a losing day if the system held a short position. If the close-to-close change was negative, we counted the days the other way round. Out of 63309 days, 32543 (51.4 %) of them were counted as positive. Probability of getting 32543 or more heads in 63309 flips is $8.39*10^{-13}$ or 0.0000000000839 %. We are interested in two-sided P-value, which is $1.7*10^{-12}$ and is much smaller than the required significance level (0.005 for confidence 99.5 %). Therefore, we reject our null hypothesis at 99.99999999983 % level of confidence and decide that real trades can be placed according to the system's signals.

Another way to judge the system's ability to make profit in the future is to compare its result with a system that opens and closes trades randomly. In this test, on every day for every market we drew a random boolean variable. If it comes up “false” no action is taken until tomorrow. If it is “true” and currently a position is held, we close current position. If “true” and no position is held - a new position is opened, another random variable is drawn to decide if the new position should be long or short. Every position's value equals 10 % of the current account value. No stops are used and we take 3136 trades. Commissions are taken into account.

We repeated the above procedure 75000 times to gather a sample big enough to make conclusions. Table 6 shows the average values for different metrics and compares them with our walk-forward results.

Metric	**Random System Mean**	**Random System Standard Deviation**	**Real System Walk-forward result**	**Sigma away from the mean**
Total profit	-40.9 %	7.55	659 %	92
CAGR	-2.63 %	0.62	10.7 %	21
Maximum drawdown	42.74 %	6.87	49 %	1
MAR ratio	-0.06	0.006	0.22	44
Profit factor	0.78	0.04	1.24	9.72

Table 6. Random system's results compared with real system's walk-forward

It is worth noting that out of 75000 iterations only two yielded net profits (equal to 3.15 % and 0.27 % total profit). The average drawdown is close (one sigma away) to the walk-forward test drawdown, which means that the random system's risk was more or less representative of the walk-forward test.

It is safe to assume, therefore, that results obtained in the walk-forward test are sufficiently positive to claim that they are impossible to achieve simply by luck alone. For this reason, we believe the system has an edge and can read markets to generate out-sized returns.

7. Conclusions

In this paper we presented research that confirms trend following as an acceptable methodology upon which to to base investment decisions. We showed that real market historical prices have a tendency to trend and that these trends are longer and stronger than in random walk data. We used this knowledge to create a trend following system that beats passive investment.

This is in contrast with EMH, that claims market changes are random and no system can lead to regular profits. By understanding market trends, one can position a portfolio to profit in rising and falling markets. Our system can be used by all types of investors and does not require access to high computational power.

We were not able to reproduce the market environment down to the very detail. For example, trading volume was not taken into account when backtesting trading systems. The purpose of the test was to confirm or deny the positive system's expectancy, not to simulate an actual account. Also, we have not calculated the difference in interest rates when holding foreign exchange positions, or roll-over costs when simulating futures contracts. These can be positive or negative and are too low to have significant impact on the end results.

Our research used only two technical analysis indicators to detect market trends, but we are sure there are many more waiting to be discovered. Other indicators may be more reliable and this should translate into bigger profits on traders' accounts. Also, replacing signal or filter indicators with others may increase our system's performance. The system itself could be improved, for example changing the position sizing algorithm or stop loss management could drastically improve results. Our system is just an example of what can be done with trend following.

8. Appendix I – Historical market prices

In this research we use end-of-day data for the most popular markets.

Currency pairs

AudUsd: 7 January 1971 – 31 December 2010

EurUsd: 5 January 1971 – 31 December 2010 (European currency was introduced in 1999, for quotes before that date Euro Currency Unit was used.)

GbpUsd: 5 January 1971 – 31 December 2010

UsdCad: 5 January 1971 – 31 December 2010

UsdChf: 5 January 1971 – 31 December 2010

UsdJpy: 5 January 1971 – 31 December 2010

Commodities

Gold: 5 January 1971 – 31 December 2010

Coffee: 17 August 1973 – 31 December 2010

Natural gas: 3 April 1990 – 31 December 2010

Silver: 5 January 1971 – 31 December 2010

Sugar: 5 January 1971 – 31 December 2010

Indices

DJIA: 5 January 1971 – 31 December 2010

FTSE100: 5 January 1971 – 31 December 2010

NASDAQ COMP: 5 January 1971 – 31 December 2010

NIKKEI: 5 January 1971 – 31 December 2010

Stocks

Apple: 7 September 1984 – 31 December 2010

Bank of America: 29 May 1986 – 31 December 2010

Ford Motors: 3 January 1977 – 31 December 2010

General Electric: 5 January 1971 – 31 December 2010

J.P. Morgan: 5 January 1971 – 31 December 2010

9. Appendix II – Trading cost

Trading costs (commissions, spread and slippage) can have significant impact on live trading results. Wherever in this research we talked about taking into account trading costs, we included the following spreads (as Ask – Bid difference):

Apple: $0.01, AudUsd: 0.0003 points, B.O.A: $0.01, Coffee: 0.01 points, DJIA: 1.0 point, EurUsd: 0.0004 points, Ford: $0.01, FTSE100: 2.5 point, Natural gas: 0.01 point, GbpUsd: 0.0003 point, G.E.: $0.01, Gold: $0.5, J.P.M.: $0.01, NASDAQ: 1.0 point, Nikkei: 2.0 point, Silver: $0.05, Sugar: 0.01 point, UsdCad: 0.0004 point, UsdChf: 0.0003 point, UsdJpy: 0.02 point.

10. References

[1] Ian Kaplan, May 2003, "Estimating the Hurst Exponent"

[2] Roman Racine, April 2011, "Estimating the Hurst Exponent"

[3] Bo Qian, Khaled Rasheed "Hurst exponent and financial market predictability"

[4] Futures Magazine, March 2011 "William Eckhardt: The man who launched 1,000 systems"

[5] Dimitri Vayanos, Paul Wooley, December 2008, "An Institutional Theory of Momentum and Reversal"

[6] http://www.wolframalpha.com/input/?i=32543+heads+63309+flips