Identifying Patterns in Financial Time Series

James A. Shine, US Army Corps of Engineers (retired) and George Mason University
James E Gentle, George Mason University
Charles R. Perry, US Department of Agriculture (retired)
JSM 2014, Boston, MA, USA
7 August 2014

Abstract

One important application in financial and other time series is the construction of trendlines (both positive and negative) for specific time periods. Such trendlines have been done visually and manually for decades, but computational methods for automatic or semi-automatic trendline generation is a much newer science. We have developed an algorithm for generating trendlines from specific time series data observations using minimum values, moving averages and constrained optimization. We will show results of this algorithm on various sets of time series data. This tool could prove useful for future analysis of financial and other time series.

Key Words: visualization, prediction, time series, financial data

Introduction

An important application in time series analysis is the visualization of time series, and in particular the construction of trendlines for specific time periods. Trendlines can visualize patterns that are difficult to see by just examining the numerical values of the time series. One example of such a time series is stock prices.

We have been looking at various stock price time series for some time. We have attempted analysis by both visual and nonvisual approaches. We present some of the visual approaches here.

Related Work

Construction of stock price trendlines is not a new area of research; analysts have been producing them for decades. The most well-known efforts date back to the mid-1900s. Pioneering work was presented by Edwards and Magee, and more recent work has been performed by Bassetti and others [1], [2].

In the pre-computer era, trendline construction usually involved graph paper (semi-log was popular) and rulers. In the last 30 years or so, newer approaches have been possible using computer graphics and visualization. These newer approaches can obviously be a lot more sophisticated than the old graph paper and ruler techniques, but all approaches have had mixed predictive results at best.

A wide range of stock market models have also been developed over this period. One of the earliest was the Dow Theory. In particular the concept of bull and bear markets has been around for some time. In a bull market, stock prices tend to increase over time, and in a bear market, they tend to decrease [1]. Finding a trendline algorithm that can accurately show where one market stops and the other starts is obviously a highly desired result.

There are many other finer points of stock price modeling, such as basing points procedures, support and resistance, volume anomalies, and trendline scales [1]. In addition, there are certain signature patterns such as “head and shoulders” that are postulated to indicate shifts in the market. A thorough listing of these patterns is documented by Achelis among others [3].

Algorithm

As with some of our previous work, we wished to produce some relatively simple open source code to demonstrate our approach. There are multiple proprietary approaches to trendline creation such as pitstock, stratasearch and optionshouse. Trendlines are also available on financial web sites such as finance.yahoo.com. However, these approaches are proprietary and the algorithm is generally not readily available to the average user. We wanted to test ideas and document which worked best, so we wrote our own code in the open source statistical language R.

We start by taking stock data from the finance.yahoo.com web site, using an existing R function entitled “getStockPrice” (author: Dr. James Nolan). We can choose different start and finish dates and different stocks, depending on the input parameters to this function. We then take the obtained data (in the form of a time series) and look for points where the direction (up or down) of the stock prices change. We begin at the start date and check data values successively at a specific lag forward from the previous value, an adjustable parameter we call step size. If the data values continue to increase or decrease, we move to a new value using the step size, until there is a change in direction (a decrease after increases, or an increase after decreases). This becomes a break point, and the process continues with the new trend until another (break) point is found that switches the trend direction again. After the entire data set has been processed, we are left with a series of break points. These points can be used for multiple purposes, and in the future we hope to use them in classification experiments, but for the work discussed here, we plot the trend lines on a time series graph in red.

The two primary adjustable parameters in the work discussed here are the step size and the length (in days, months or years) of the input data set.

Results

We examined many different stocks and time periods to test our algorithm. We found that it works reasonably well on stock price time series. As might be expected, adjusting the step size changes the results; a shorter step size shows more short-term trends, while a longer step size shows just the big trends.

Figure 1 shows trendlines for General Electric stock in 2007 with a step size of 5, and Figure 2 shows trendlines for the same stock and time period with a step size of 40.

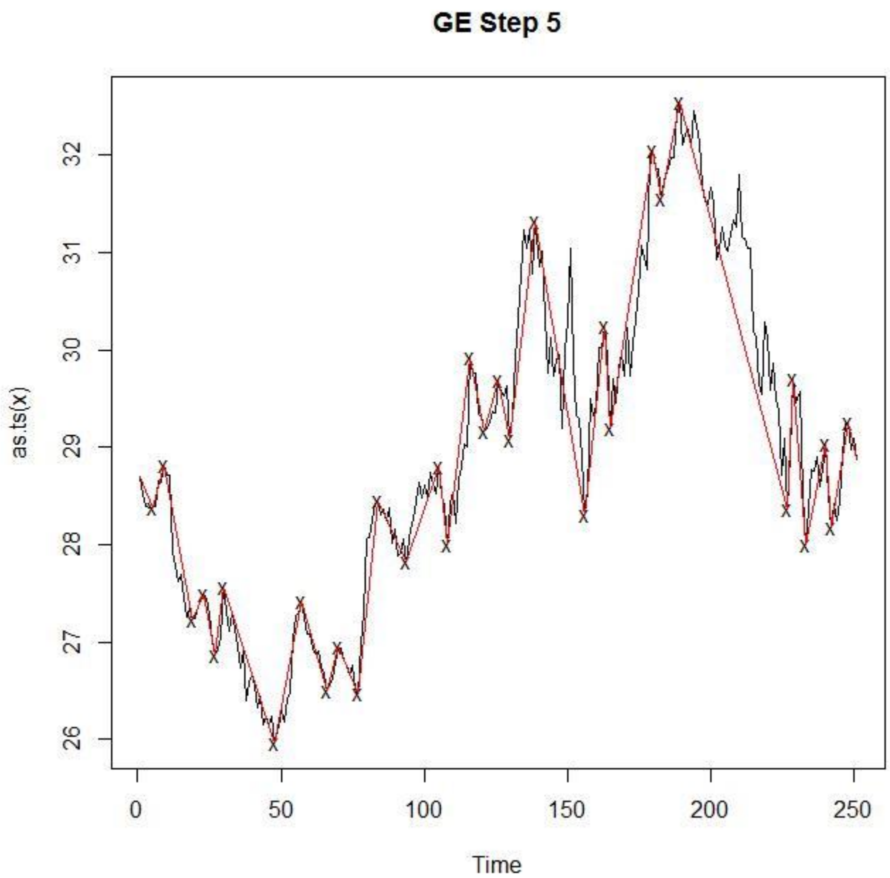

.

Figure 1: Trendlines for General Electric stocks, year 2007, step size 5

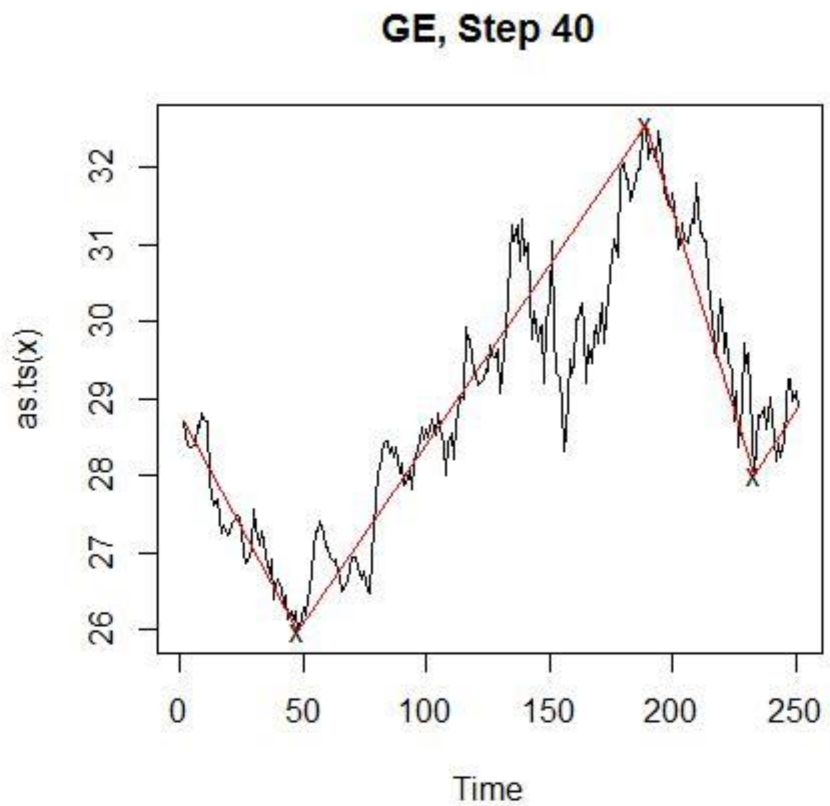

Figure 2: Trendlines for General Electric stocks, year 2007, step size 40

We also examined the effects of a moving average smoother on trendline analysis. Stock moving averages tend to produce the same trendlines as the raw data. Figure 3 shows trendlines for a moving average (5 points) of General Electric stock in 2007; the step size is 40, as in Figure 2. Note that the trendlines in Figure 2 and Figure 3 are identical.

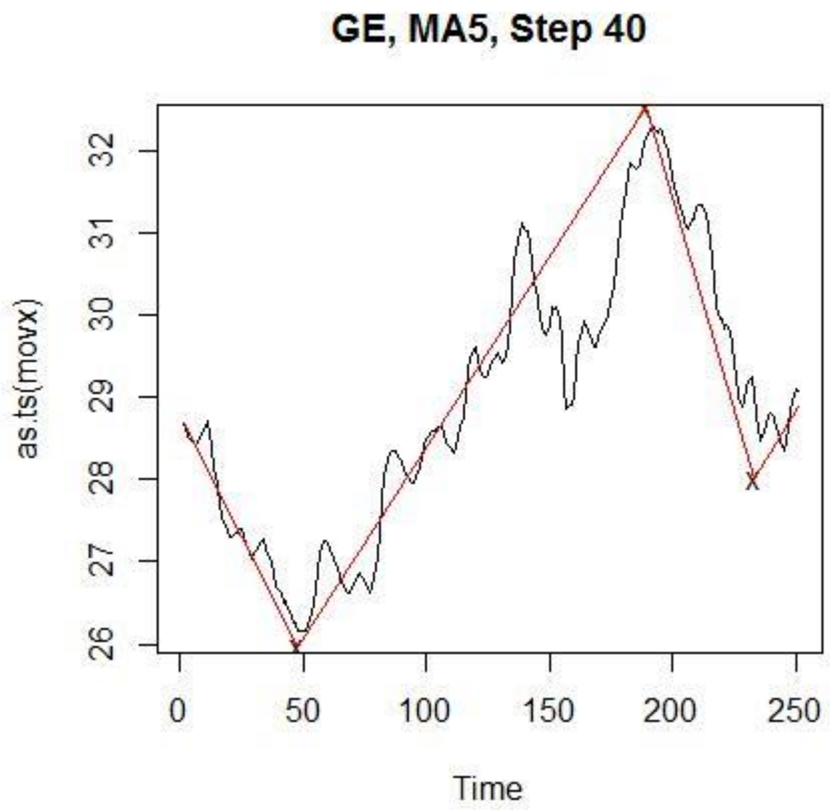

Figure 3: Trendline for General Electric stocks, moving average (5 points), in 2007 with a step size of 40.

Conclusions and Future Work

We have developed a simple tool in R for creating trendlines that works pretty well on stock price time series. Performing this trendline analysis on moving averages of stock data does not seem to improve the results significantly, at least in tests so far. Parameter adjustment allows trendlines on different scales to be discovered. This tool could prove useful for future analysis of financial and other time series.

For future work, we hope to study in depth a supervised classification approach with variables such as line slopes, length of lines, and the location of the points where the slope changes. We also hope to do further analysis of a comparison of moving average results to those of the raw stock data. In addition, we hope to add a third type of line, representing a “mule market” when the overall trend is neither significantly up or significantly down.

References

[1] R.D. Edwards, J. Magee and W.H.C. Bassetti, “Technical Analysis of Stock Trends”, 10th Edition, CRC Press, 2013.

[2] W.H.C. Bassetti, “Zen Simple: Beat the Market with a Ruler”, MaoMao Press, 2012.

[3] S.B. Achelis, “Technical Analysis from A to Z”, 2nd Edition, McGraw Hill, 2001.