Information in Stock Price Patterns*

Steven Jens

October 5, 2009

Abstract

In their paper of 2000 [5], Lo, Mamaysky, and Wang tested whether ten widely-used stock price chart patterns provided information about future one-day returns. They found that seven of them did, but made no attempt to explain why this might be. In this paper, I compare one-day returns conditional on technical indicators to one-day returns on days with similar characteristics. I find that this explains the abnormal returns behavior for the rectangle bottom pattern.

1 Introduction

In 2000, Lo, Mamaysky, and Wang [5] attempted to bridge the gap between academic finance and the technical analysis side of the financial industry by surveying five pairs of the chart patterns most frequently used by technical analysts, giving them precise definitions, and testing whether they provide information regarding future returns. They did not test any specific trading rules, simply whether the distribution of one-day returns three days after the completion of a pattern was significantly different from the unconditional distribution of one-day returns. Of the ten patterns they tested, they found that seven produced distributions of one-day returns significantly different from the unconditional distributions. They did not attempt to explain this phenomenon.

*Thanks to WRDS, of the Wharton School at the University of Pennsylvania, for providing CRSP's data sets of stock prices and to the Apache project for their chi-squared library.

Lo et al. normalized their returns data by subtracting off the mean daily returns and dividing by the standard deviation of daily returns for each stock over a five-year period. However, some market characteristics, including momentum [4] and volatility [7], are known to change over long periods of time but persist over shorter periods – periods shorter than five years. It seems reasonable to suppose that a pattern might occur more frequently in, for example, high-volatility environments than in low-volatility environments and that subsequent returns would be better compared to returns in high-volatility environments than to unconditional returns.

I follow the procedures of Lo et al. in comparing daily returns three days after each technical pattern to unconditional returns, but I also compare the conditional returns to the daily returns on a proxy day. The proxies are chosen based on two characteristics. First, the proxies are roughly 35 days before the end of the pattern, which is generally well before the start of the pattern but within a period of time over which volatility and momentum are likely to persist. Second, the proxies are three days after a minimum for a pattern ending in a minimum and three days after a maximum for a pattern ending in a maximum.

Using far less data than Lo et al., I only find three out of eight patterns to provide significant information relative to unconditional returns, while they found significant results from six of the eight patterns that we both tested. Of the three patterns for which I find significant information relative to unconditional returns, none of these patterns provide significant information relative to returns on their proxies. For two of the three, I believe this to be due to a limitation of my data, but I conclude that the rectangle bottom pattern is not informative about the future so much as it is characteristic of the environments in which it occurs.

In the next section I explain how Lo et al. and I identify patterns and the standard we use for informativeness. In section 3 I discuss our data collection and results. In section 4 I wrap up with a brief discussion of how my inquiries in this area might proceed.

2 Methods

Most of my methods are copied from Lo et al.

Their methods addressed two questions: first, what defines presence of a technical pattern? Second, given the presence or lack of the pattern, what does it mean for the pattern to be informative?

2.1 Identifying Patterns

The first problem is to formulate a precise definition of the patterns being tested. Lo et al. defined five pairs of patterns, one member of each pair being the upside-down version of the other. Four of these pairs of patterns were based on the characteristics of five consecutive extrema; the other pair was the "double top" and "double bottom". As the "double bottom" was not found to be informative and six of the other eight were informative, I focus on those eight patterns.

Four of the patterns are characterized by three consecutive maxima separated by two minima, with the following additional characteristics:

- In a **Head and Shoulders** pattern,
  - the middle maximum is the highest
  - the outer two maxima are each within 1.5% of their average
  - the two minima are each within 1.5% of their average
- In a **Broadening Top**,
  - the last maximum is higher than the middle maximum, which is higher than the first maximum
  - the first minimum is higher than the second minimum
- In a **Triangle Top**,
  - the last maximum is less than the middle maximum, which is less than the first maximum
  - the first minimum is lower than the second minimum
- In a **Rectangle Top**,
  - All maxima are within 0.75% of their average

- Both minima are within 0.75% of their average
- All maxima exceed all minima

The other four patterns are identical, but upside-down, and are called, respectively, the **Inverse Head and Shoulders**, the **Broadening Bottom**, the **Triangle Bottom**, and the **Rectangle Bottom**. Lo et al. required that patterns be completed in a 35 trading-day window; I used a 35 calendar-day window. Lo et al. reported no tests on how frequently this requirement was binding, and I performed none, but most of the windows I looked at had at least a dozen extrema. Because a chart pattern can't be identified on the day the pattern completes, the conditional one-day returns were measured three days after the end of the pattern.

The five consecutive extrema based on which the patterns are identified are not from the raw stock price data, but from a smoothed version. Like Lo et al., I run my stock prices through a kernel regression estimator, which produces a new data series each element of which is a weighted average of the points in the old series. The weight on each point in the average is proportional to a function of the distance between the dates. Generally, closer dates are weighted more heavily than dates well into the past of the future; in our case, we use a normal PDF with a mean of zero and a standard deviation characterized by a "bandwidth" parameter, the choice of which will be explained later. For example, if a bandwidth of 1 day is selected, the smoothed stock price for Wednesday will be a weighted average of Monday's actual closing stock price with a weight proportional to $\phi(-2)$, Tuesday's actual closing stock price with a weight proportional to $\phi(-1)$, Wednesday's actual closing stock price with a weight proportional to $\phi(0)$, and so on, with the constant of proportionality chosen such that the weights add up to 1.

The "optimal" bandwidth is the bandwidth which minimizes the sum of the squared differences between each point's actual value and the value predicted for the point by a kernel regression using the rest of the points. When Lo et al. used optimally-smoothed data to detect patterns and presented their findings to technical analysts in industry, the practitioners told them the data sets had been too smoothed. By trial and error, Lo et al. settled on a bandwidth of 30% of the optimal value. Their paper lamented the lack of a theoretical basis for settling on 30%, but in the identification of technical patterns which are not themselves theoretically justified, an objective standard that passes the Turing test is more important than a theoretically-justified standard. What *does* require formal justification is the basis on which the technical patterns will be judged, and Lo et al. successfully defined two such standards, of which I use one.

2.2 Measuring Informativeness

Lo et al. used two standards to measure whether one conditional distribution is significantly different from another. One is the Kolmogorov-Smirnov test, which was designed for comparing probability distributions. The one I use splits the unconditional distribution into deciles and counts how many conditional one-day returns fall in each tenth of the distribution. This is less elegant – why ten buckets? – but is easier to use and makes it easier to visualize in what manner two distributions differ in addition to determining whether the difference is statistically significant.

According to Lo et al. equation (17), if n_j observations fall into decile j out of n observations total, and both distributions are actually the same,

$$\sum_{j=1}^{10} \frac{(n_j - 0.10n)^2}{0.10n} \overset{a}{\sim} \chi_9^2 \tag{1}$$

3 Results

Lo et al. collected, over seven separate five-year spans, daily data for each of ten stocks in each of five market capitalization segments (i.e., they broke up their universe of stocks into quintiles and took ten from each group). For this inquiry, I settle for five stocks in each of the five segments from the years 2003 to 2008, inclusive. Consequently, my results are less significant than theirs, but I still get some significant results which match theirs.

As mentioned before, the use of deciles allows us to examine how a conditional distribution differs from the unconditional distribution of returns. The "rectangle bottom" is the pattern which generates the distribution most significantly different from the unconditional distribution. Figure 1 shows how returns conditional on a rectangle bottom pattern fall into the deciles of the unconditional distribution. By inspection, the main difference between the rectangle bottom returns and unconditional returns is that returns conditional on a

Pattern	N
Head & Shoulders (HS)	459
Inverse Head & Shoulders (IHS)	789
Broadening Top (BTOP)	414
Broadening Bottom (BBOT)	759
Triangle Top (TTOP)	203
Triangle Bottom (TBOT)	351
Rectangle Top (RTOP)	427
Rectangle Bottom (RBOT)	693
Unconditional	26253

Table 1: Total number of occurrences of each pattern

rectangle bottom are less likely to be very large or very small. In other words, the rectangle bottom pattern seems to occur in low-volatility environments.

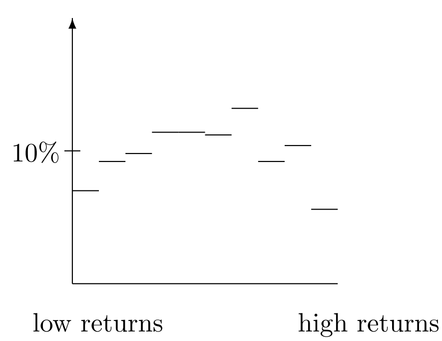

Figure 1: post-RBOT returns vs. unconditional returns

To determine this more formally, I create a proxy return distribution for each pattern with significant results. Rather than the one-day return three days after the end of the 35-day window, this collects the one-day return after the first maximum in the window in the case of the broadening top and the rectangle top and the first minimum in the case of the rectangle bottom. The conditional returns and their proxy returns are each found three

Pattern	p-value
Head & Shoulders (HS)	7.23 %
Inverse Head & Shoulders (IHS)	47.94 %
Broadening Top (BTOP)	2.87 %
Broadening Bottom (BBOT)	95.09 %
Triangle Top (TTOP)	57.58 %
Triangle Bottom (TBOT)	24.08 %
Rectangle Top (RTOP)	0.75 %
Rectangle Bottom (RBOT)	0.0477 %

Table 2: Significance of the differences of return distributions conditional on patterns from the unconditional distribution of returns

days after an extremum of the same kind, and each occur in similar market environments to the extent that those environments persist for more than a month.

As table 3 indicates, the one-day returns conditional on the broadening top pattern are significant at the 5% level. The difference between returns conditional on the broadening top pattern and returns conditional on the broadening top proxy are not significant at the 5% level, but there are over 26000 returns in the unconditional sample and only 414 data points in the broadening top proxy return set. The comparison of the broadening top proxy return distribution to the unconditional distribution indicates that the environment a month before broadening top patterns form is not significantly different from the environment on a random day. We see similar results for the rectangle top.

The rectangle bottom, on the other hand, is almost surely a marker of its environment rather than an independent source of information about the future distribution of returns. It predicts the future quite well, but it predicts the past even better, and the proxy returns – anticipating the pattern – are not distinguishable from the rectangle bottom conditional returns using the amount of data I have.

target distribution	comparison distribution	significance
BTOP	unconditional	2.87%
BTOP	BTOP proxy	9.00%
BTOP proxy	unconditional	62.56%
RTOP	unconditional	0.746%
RTOP	RTOP proxy	9.38%
RTOP proxy	unconditional	34.84%
RBOT	unconditional	0.0477%
RBOT	RBOT proxy	23.84%
RBOT proxy	unconditional	0.0286%

Table 3: Comparison of results to proxies

4 Future Directions

Much about this paper is more in the manner of an initial inquiry than a publishable result. I could add a lot of power to my tests by adding more data and I would feel more comfortable with the statistics if a model were more precisely specified. The results so far suggest that different patterns may need different treatments. The fact that patterns ending in maxima are consistently identified roughly 70% more frequently than their inverted counterparts is suspicious. Before I increase the amount of data, I will add some tests to my code to explain or remove this anomaly and verify that my results are not due to bugs in my software.

I had originally intended to see whether any of the information about one-day returns contained in these technical patterns is explained by the Fama-French three-factor model [2], but, given my results to date, inquiries related to volatility (e.g., French, Schwert, and Stambaugh 1987 [3] or Ang et al 2006 [1]) or liquidity risk (e.g., Pastor and Stambaugh 2003 [6]) seem more likely to be fruitful.

References

[1] Andrew Ang, Robert J. Hodrick, Yuhang Xing, and Xiaoyan Zhang. The cross-section of volatility and expected returns. *The Journal of Finance*, 61(1), 2006.

[2] Eugene F. Fama and Kenneth R. French. Common risk factors in the returns on bonds and stocks. *Journal of Financial Economics*, 33(1):3–53, 1993.

[3] Kenneth R. French, G. William Schwert, and Robert F. Stambaugh. Expected stock returns and volatility. *Journal of Financial Economics*, 19(1):3 – 29, 1987.

[4] Narasimhan Jegadeesh and Sheridan Titman. Returns to buying winners and selling losers: Implications for stock market efficiency. *The Journal of Finance*, 48(1):65–91, 1993.

[5] Andrew W. Lo, Harry Mamaysky, and Jiang Wang. Foundations of technical analysis: Computational algorithms, statistical inference, and empirical implementation. *The Journal of Finance*, 55(4):1705–1765, 2000.

[6] Ľuboš Pástor and Robert F. Stambaugh. Liquidity risk and expected stock returns. *Journal of Political Economy*, 111(3):642–685, 2003.

[7] G. William Schwert. Why does stock market volatility change over time? *The Journal of Finance*, 44(5):1115–1153, 1989.