Foundations of Technical Analysis: Computational Algorithms, Statistical Inference, and Empirical Implementation: Discussion

Author(s): Narasimhan Jegadeesh

Source: *The Journal of Finance*, Vol. 55, No. 4, Papers and Proceedings of the Sixtieth Annual Meeting of the American Finance Association, Boston, Massachusetts, January 7-9, 2000 (Aug., 2000), pp. 1765-1770
Published by: Blackwell Publishing for the American Finance Association
Stable URL: http://www.jstor.org/stable/222378
Accessed: 24/11/2009 12:16

Osler, Carol, and Kevin Chang, 1995, Head and shoulders: Not just a flaky pattern, Staff Report No. 4, Federal Reserve Bank of New York.
Poggio, Tomaso, and David Beymer, 1996, Regularization networks for visual learning, in Shree Nayar and Tomaso Poggio, eds.: *Early Visual Learning* (Oxford University Press, Oxford, U.K.).
Press, William, Brian Flannery, Saul Teukolsky, and William Vetterling, 1986, *Numerical Recipes: The Art of Scientific Computing* (Cambridge University Press, Cambridge, U.K.).
Pruitt, Stephen, and Robert White, 1988, The CRISMA trading system: Who says technical analysis can't beat the market? *Journal of Portfolio Management* 14, 55–58.
Riesenhuber, Maximilian, and Tomaso Poggio, 1997, Common computational strategies in machine and biological vision, in *Proceedings of International Symposium on System Life* (Tokyo, Japan), 67–75.
Rouwenhorst, Geert, 1998, International momentum strategies, *Journal of Finance* 53, 267–284.
Simonoff, Jeffrey, 1996, *Smoothing Methods in Statistics* (Springer-Verlag, New York).
Smirnov, N., 1939a, Sur les écarts de la courbe de distribution empirique, *Rec. Math.* (*Mat. Sborn.)* 6, 3–26.
Smirnov, N., 1939b, On the estimation of the discrepancy between empirical curves of distribution for two independent samples, *Bulletin. Math. Univ. Moscow* 2, 3–14.
Tabell, Anthony, and Edward Tabell, 1964, The case for technical analysis, *Financial Analyst Journal* 20, 67–76.
Treynor, Jack, and Robert Ferguson, 1985, In defense of technical analysis, *Journal of Finance* 40, 757–773.

Discussion

NARASIMHAN JEGADEESH*

Academics have long been skeptical about the usefulness of technical trading strategies. The literature that evaluates the performance of such trading strategies has found mixed results. This literature has generally focused on evaluating simple technical trading rules such as filter rules and moving average rules that are fairly straightforward to define and implement. Lo, Mamaysky, and Wang (hereafter LMW) move this literature forward by evaluating more complicated trading strategies used by chartists that are hard to define and implement objectively.

Broadly, the primary objectives of LMW are to automate the process of identifying patterns in stock prices and evaluate the usefulness of trading strategies based on various patterns. They start with quantitative definitions of 10 patterns that are commonly used by chartists. They then smooth the price data using kernel regressions. They identify various patterns in the smoothed prices based on their definitions of these patterns. Their algorithms allow them to recognize patterns objectively on the computer rather

* University of Illinois at Urbana-Champaign.

than visually. They then compare the "postpattern" distribution of stock returns with the unconditional return distributions to evaluate the usefulness of charting techniques.

I. Concerns about Charting

Many academics and practitioners have serious concerns with investing based on charting techniques. These concerns seem to run deeper than the "linguistic barriers" emphasized in LMW. Perhaps the most important reason why charting techniques have not been more widely accepted is that they are built on weak foundations. For instance, chartists believe that selected patterns in the history of stock prices tend to repeat. However, there does not seem to be a plausible explanation as to why these patterns should indeed be expected to repeat.

Of course, from a strictly efficient markets perspective, there should be no predictable pattern in stock returns. But in recent times many academics are open to the idea that markets may not be fully efficient and prices may be affected by behavioral biases documented in the psychology literature. For instance, Barberis, Shleifer, and Vishny (1998) and Daniel, Hirshleifer, and Subrahmanyam (1998) appeal to conservatism bias and investor overconfidence to explain evidence of market underreaction or overreaction to information documented by DeBondt and Thaler (1985) and Jegadeesh and Titman (1993), among others. It is hard to see, however, the nature of behavioral biases that will lead to return predictability based on the charting patterns evaluated in LMW.

"Foundations of Technical Analysis" in the title led me to believe that this paper would provide theoretical justification for the use of charting techniques. This issue, however, is not addressed in the paper. LMW do attempt to rationalize the use of these techniques by pointing out that patterns in prices and trading volumes such as that in their Figure 1 contain information. Most readers will have little difficulty in accepting that prices and trades are driven by information, but this information pertains to what happened in the past. What is sorely needed is a reason to believe that any such information will help us predict future prices.

II. Pattern Recognition

One of the main motivations behind the algorithm developed here for pattern recognition is that the techniques currently used by chartists are extremely subjective. What may appear to be a head-and-shoulders pattern to one chartist, for example, may appear to be noise or an altogether different pattern to another chartist. Although the techniques in LMW are closer to objective pattern recognition, they also contain several layers of subjectivity.

First, the definitions of various patterns in terms of local extrema and their relative positions are subjective. Second, the choice of a particular smoothing technique is also subjective. Section I.A of LMW lists a variety of smoothing estimators that are available in the literature, and for this paper, the authors have chosen the kernel regression estimator. Finally, the choice of bandwidth parameter h used in the kernel regression is also subjective.

To an extent, subjective choices of this nature are the prerogatives of the researchers and cannot be entirely avoided. The main issue, however, is whether the results are robust with respect to different algorithms. A priori, I expected that robustness would not be a serious issue here, but the results in the paper suggest otherwise. The paper reports that the bandwidths for the kernel regressions obtained by minimizing the cross-validation functions yielded fitted values that the authors deemed to be too smooth based on their discussions with practicing chartists. So the paper uses bandwidths that are 30 percent of the optimal (in a statistical sense) bandwidths. A potentially interesting avenue for future research would be to examine the extent to which the results are sensitive to various choices discussed in the last paragraph.

A more basic question here is whether in the first place it is necessary to fit a smoothed function to objectively identify patterns. Naturally, there will be many more local extrema in the raw data than in the smoothed data. However, it is possible that some of these extrema can be screened out by appropriate filters (such as discarding local maxima that are less than x percent above the nearest local maximum). This is another issue that can be addressed in future work to convince the readers that the price data need to be smoothed to identify patterns on the computer.

III. Evaluating the Informativeness of Patterns

LMW evaluate the information content of the technical trading rules by comparing the return distribution conditional on the occurrence of a particular pattern with the corresponding unconditional return distribution. The extant literature typically judges the usefulness of trading rules by evaluating their profitability because this is the metric that is of most direct interest to investors. Although the paper does not discuss the profitability of these trading rules, the empirical results in the paper allow us to address this issue.

Table I presents the standardized mean returns (the difference between conditional and unconditional mean returns divided by the unconditional standard deviation) conditional on the occurrence of various patterns for the NYSE/AMEX sample and the Nasdaq sample. The standardized mean returns are remarkably close to zero, and the largest magnitude of standardized mean is only -0.062 (for BBTOP in the NYSE/AMEX sample). Clearly, these results do not suggest that any of these trading rules yields profits

Table I

This table presents test statistics under the hypothesis that the conditional one-day mean standardized returns, three days following the conclusion of an occurrence of one of the 10 technical indicators in LMW, equal zero. Standardized return is computed as the difference between the raw return and the unconditional mean return divided by unconditional standard deviation. The data are taken from the tables in LMW indicated in parentheses. See Table I in LMW for descriptions of the technical indicators.

	HS	IHS	BTOP	BBTOP	TTOP	TBTOP	RTOP	RBTOP	DTOP	DBTOP
					NYSE/AMEX Stocks (1962–1996)					
Standardized Mean (Table III)	−0.038	0.040	−0.005	−0.062	0.021	−0.009	0.009	0.014	0.017	−0.001
Occurrences (Table I)	1,611	1,654	725	748	1,294	1,193	1,482	1,616	2,076	2,075
t-statistics	−1.525	1.627	−0.135	−1.696	0.755	−0.311	0.346	0.563	0.775	−0.046
					Nasdaq Stocks (1962–1996)					
Standardized Mean (Table IV)	−0.016	0.042	−0.009	0.009	−0.020	0.017	−0.052	0.043	0.003	−0.035
Occurrences (Table II)	919	817	414	508	850	789	1,134	1,320	1,208	1,147
t-statistics	−0.485	1.200	−0.183	0.203	−0.583	0.478	1.751	1.562	0.104	−1.185

that are of any significance from an economic standpoint. The t-statistics in the table indicate that none of the conditional mean returns are reliably different from zero.

These results indicate that the technical trading rules cannot be independently used for identifying profitable investment opportunities. However, my discussions with practitioners suggested that chartists tend to use signals from patterns in price history in conjunction with other information to time their trades. For instance, a chartist may not buy all stocks that exhibit a particular pattern, but when he has a favorable view about a particular stock he will initiate his trade after observing certain patterns in prices. The idea here is that after the observation of one of these patterns the chartist feels relatively more confident that his favorable information will eventually get incorporated into stock prices.

The usefulness of such conditional trading strategies is hard to verify objectively (short of directly evaluating a chartist's actual trades) because empiricists typically do not have access to the private information used by chartists. However, this notion is consistent with the intuition behind Brown and Jennings's (1989) model. In that model, technical analysts use the history of prices in conjunction with their private information to disentangle the noise in prices induced by supply uncertainty. When an agent in their model does not have access to private information (empiricists are typically in this situation), technical analysis will not help in predicting future prices.

Although the conditional mean returns following various patterns are not different from unconditional mean returns, LMW reject the hypothesis that the conditional return distributions are the same as the unconditional distributions for seven out of 10 patterns in the NYSE/AMEX sample (Table IV in LMW) and for all 10 patterns in the Nasdaq sample (Table VI in LMW). These rejections, therefore, are not due to differences in means but are due to differences in the higher moments. The paper does not present statistical tests of the hypothesis that any specific moment of the conditional distribution equals its unconditional counterpart. The results in Table III and Table IV in LMW, however, suggest that some of the point estimates of unconditional variance, skewness, and kurtosis may be different from the corresponding conditional statistics, but I was not able to visually detect any consistent pattern across various samples and subperiods. I would encourage the authors to explore this issue more closely in future work.

Even if we grant that some higher order unconditional moments are indeed different from their conditional counterparts, at least two more issues need to be addressed before we can conclude that these differences are useful from a practical standpoint. First, the paper needs to address the issue of how investors can benefit from any ability to predict higher moments. It is possible that investors may be able to exploit conditional variance forecasts in options-based investment strategies, but it is also possible that these conditional forecasts are already incorporated in option prices. Applications of forecasts of conditional skewness and kurtosis are not quite as obvious and merit further discussion.

The second issue, which is perhaps more important, is whether the use of these trading rules is the most efficient way to forecast these higher moments. For instance, it is well known that conditional variances can be predicted reasonably well using ARCH and related techniques. Are the trading patterns in this paper better predictors of conditional variances than the ARCH models? Similarly, if an investor is interested in other conditional moments there may be better techniques to predict these moments directly than these trading rules.

IV. Concluding Remarks

LMW is an important addition to the literature on technical trading. One of the stated goals of LMW is to "extend the reach of technical analysis by augmenting its tool kit with some modern techniques in pattern recognition." They accomplish this goal by providing algorithms that could potentially remove some of the subjective elements in pattern recognition and automate the task. The evidence provided here, however, does not support the hypothesis that technical analysis can be used as the basis for profitable trading strategies. Therefore, I do not think that the results of this paper will move academic critics any closer to technical analysts.

My discussion here highlights several dimensions along which this paper can be extended. Additionally, the approach in this paper can potentially be modified to identify profitable trading rules. This paper selects certain trading rules and examines whether these rules can predict future return distributions. In future work, one may let the algorithm choose the patterns in prices that best predict future returns. In related work, Allen and Karjalainen (1999) report that genetic algorithms are not able to consistently predict Standard and Poor's composite index returns. However, the results with the nonparametric approach in this paper applied to individual stock returns may well be more promising. Of course, it is important to exercise sufficient caution so that any extension of this research along this path is not viewed as pure data mining.

REFERENCES

Allen, Franklin, and Risto Karjalainen, 1999, Using genetic algorithms to find technical trading rules, *Journal of Financial Economics* 51, 245–271.

Barberis, Nicholas, Andrei Shleifer, and Robert Vishny, 1998, A model of investor sentiment, *Journal of Financial Economics* 49, 307–343.

Brown, David, and Robert Jennings, 1989, On technical analysis, *Review of Financial Studies* 2, 527–551.

Daniel, Kent, David Hirshleifer, and Avanidhar Subrahmanyam, 1998, Investor psychology and security market under- and overreactions, *Journal of Finance* 53, 1839–1886.

DeBondt, Werner F. M., and Richard Thaler, 1985, Does the stock market overreact?, *Journal of Finance* 40, 793–805.

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