JOURNAL OF INVESTMENT MANAGEMENT, Vol. 2, No. 1, (2004), pp. 79–85

www.joim.com

WORKING PAPERS

"Working Papers" provides a review of significant working papers in investment management. This section draws from recent research in order to highlight an area of topical interest. In selecting papers pertinent to a prominent topic, "Working Papers" acknowledges current trends in the investment management business, while simultaneously directing the reader to interesting and important recent work.

TECHNICAL ANALYSIS

Sanjiv Ranjan Das[a], Senior Editor *and David Tien*[a]

Technical analysis is one of the oldest ideas for stock selection, and in the opinion of some modern economists, too old. Even though the approach flies in the face of modern tenets of weak-form market efficiency, it has an appeal that has preserved it in an increasingly quantitative financial environment. Lo *et al.* (2000) categorize this old versus new world tension as a tussle between algebraic and geometric analysis within the finance pantheon. In their words,

These linguistic barriers underscore an important difference between technical analysis and quantitative finance: technical analysis is primarily visual, whereas quantitative finance is primarily algebraic and numerical. Therefore, technical analysis employs the tools of geometry and pattern recognition, and quantitative finance employs the tools of mathematical analysis and probability and statistics. In the wake of recent breakthroughs in financial engineering, computer technology, and numerical algorithms, it is no wonder that quantitative finance has overtaken technical analysis in popularity—the principles of portfolio optimization are far easier to program into a computer than the basic tenets of technical analysis. Nevertheless, technical analysis has survived through the years, perhaps because its visual mode of analysis is more conducive to human cognition, and because pattern recognition is one of the few repetitive activities for which computers do not have an absolute advantage (yet).

It seems that this debate is far from being concluded. In this short article, we examine a few of the working papers that represent the state of affairs in this area of research. The long-standing question asked of technical trading has been (a) whether this analysis yields profitable trading strategies (unconditional analysis), and (b) whether the various patterns that are tracked by technical analysts do indeed precede marked changes in the time series of stock returns (conditional analysis).

1 Methodological issues

Whether or not trading rules are able to beat the market often depends on the mode of

[a] Santa Clara University, Santa Clara, CA, USA.

analysis. In a short theoretical article Reitz (2002) presents an argument for technical analysis being within the realm of rational expectations based decision-making.

The argument is essentially quite simple and proceeds as follows. Exchange rate movements may emanate from a time-homogeneous statistical system punctuated by regime shifts. Within each regime the system sustains parameter stability. When a regime shift occurs, a new set of parameters controls the new process of exchange rates. When a regime shift occurs, it is not known to market participants who only infer the new regime after a period of observation of exchange rates. This learning is essentially Bayesian. Reitz's idea is that the use of a trading rule to determine the presence of a regime shift may be interpreted as one way to utilize more information in their trading strategies. The examination of past prices using technical analysis may be interpreted as a proxy for Bayesian learning, which is fully consistent with rational expectations, in which market forecasts are such that they minimize forecast errors.

Much of the literature does find that technical analysis evidences the power to make fairly accurate forecasts of prices, even though it is not always possible to profit from the predictions. The paper by Reitz is a consistent theoretical approach to explaining these findings.

Data snooping is an issue in empirically assessing the results of technical analysis. Snooping occurs when data are used more than once for model selection or for statistical inference. The historical rules we see today are basically the ones that we have after survival over a period of time. Therefore, conditioning on rule survivorship, they have above average returns, or evidence some degree of predictive ability. Sullivan *et al.* (1999) argue that the finding of predictability in all 26 trading rules investigated by Brock *et al.* (1992) may be attributed partially to survivorship. Their approach to correct for data snooping uses the "Reality Check Bootstrap" estimator of White (1997). They found that, while the correction did not invalidate the trading rules in-sample, it eliminated all profits on an out-of-sample basis. This is usually the argument against data-snooping—incessant searching for rules in-sample, results in overfitting, and failure of the rules out-of-sample. Further, if the period of data is very large (100 years in the Sullivan *et al.* (1999) study), the possibility of rejection in-sample is remote. This evidence of bootstrapping is revisited by Ready (2002) who suggests that the results of Brock *et al.* (1992) are attributable to data snooping, and after accounting for transactions costs, no profitability remains.

A more recent paper by Jacquier and Yao (2002) adopts a genetic algorithm approach (see Allen and Karjalainen, 1995) to mimicking trading rules, and assesses their performance after controlling for the data-snooping bias and transactions costs. This paper shows that there is a horizon effect, i.e. the performance of the rule depends on the measurement period—rules based on a one-year measurement do poorly out of sample, but those using a ten-year period do rather well. This leads to the conclusion that to benefit from technical trading, investors need to be long-run oriented in their portfolio strategies. To assess data-snooping, they use a statistic that is corrected for snooping, compared to one where no correction is applied. Over short trading horizons, it turns out that the impact of data snooping is very high, and if adjusted for, leads to unprofitable trading results. In general, out-of-sample tests find the profit spigot to be dry.

To summarize, the main methodological issues pertaining to the empirical performance of technical trading rules are as follows. First, trading rules are inherently Bayesian in their underpinnings, and are measurable within rational expectations settings. Second, the performance of technical rules must be assessed after accounting for data-snooping biases and transactions costs. Third, performance is horizon dependent, with performance being positively biased with longer evaluation horizons. Finally, trading rules may provide evidence of return forecastability, but need not result in better than benchmark-adjusted performance. We take a closer look at this next.

2 Are technical trading rules profitable?

This is the canonical question in this literature. Bessembinder and Chan (1998) examine technical sell signals and find little evidence that these signals forecast negative returns, except in periods prior to the 1940s. Transactions costs are found to eliminate any material profits based on technical signals. While technical traders usually use the Dow Jones index for analysis, Bessembinder and Chan also find that an index constructed from CRSP returns does just as well. While technical signals evidence mild forecastability, their results provide no ability to reject a hypothesis of market efficiency.

The standard approach to technical analysis is that of point and figure charting. The main idea in this approach is that time is not important, and only changes in price space need to be examined for trading indications. There are many different types of point and figure signals, such as heads-and-shoulders patterns, broadening tops and bottoms, double-tops, open-high-low-close charts, etc., and most of these have been studied in great detail over various economic epochs, beginning from the early 1900s. These methods are usually implemented as data-filtering tools, and as discussed earlier, provide Bayesian updating for portfolio decisions. The evidence on these rules has been fairly mixed. Hauschild and Winkelmann (1985) found that aggregated across firms, the point and figure rule failed to produce profitable results. But there are many other early papers that suggest just the opposite (for an example: see Davis, 1965). Anderson (2001) undertakes a detailed re-examination of the point and figure technique. The unique feature of this study is the use of ultra-high-frequency data on the S&P 500 index. He also provides a set of useful appendices in which Boolean logic is used to express the point and figure patterns in a rigorous way. Altogether, eight trading rules are examined. An interesting finding is that trading profits are positively linked to trading volume, which in turn is related to price volatility. While these trading rules are found to be profitable, the paper does not correct for transactions costs or data snooping. Nor is there a comparison to a conventional benchmark. Nevertheless, there appears to be forecastability, although profits seem to decline over time, suggesting that markets have become more efficient in recent years.

Can excess returns accruing to technical trading rules be caused by incorrect specification of the asset price process? This issue is considered in the paper by Feng and Smith (1997). Using the same trading rules as in Brock *et al.* (1992), this paper enhances the asset price specification by considering a jump-diffusion model as well as an ARCH(1) plus jump model. Using an appropriate bootstrapping methodology, they find that trading profits are reduced or disappear completely when these extended models are used. Hence, the residuals from the standard asset pricing models may be spuriously leading to conclusions of successful trading rules. In other words, the profits found for trading rules may simply be fair compensation for jump risk premia.

There are two main conclusions that emanate from the literature on the empirical evaluation of technical trading rules. First, most of the rules that are commonplace today are found to be reasonably good at forecasting stock returns. This evidence is consistent with the occurrence of a technical pattern presaging a structural shift in the returns process. Second, even though there is forecasting power, the results are tinged with the inability to beat realistic benchmarks. Moreover, these results are also dependent on correctly accounting for data snooping, as well as risk premia for non-standard features of the stochastic processes used for asset prices.

We note that, empirically, there is evidence (see Lo *et al.*, 2000) that conditional on a technical pattern, there is a structural shift in the data. This is more prevalent in the realm of economic time series that are more susceptible to regime shifts, such as in the foreign exchange markets. Therefore, it is not surprising that many of the papers in this literature tend to be focused on currency markets. It is to this evidence that we turn next.

3 Foreign markets

Though the evidence on the success of technical trading rules in US equity markets is mixed, there is a substantial body of research that shows that technical trading rules generate economically significant returns in the foreign exchange markets. Levich and Thomas (1993) find that simple moving average-based trading rules generate significant (using bootstrapped standard errors) returns in the currency futures market. Many researchers have sought to refine technical trading profits by employing more sophisticated techniques to generate trading rules. Neely *et al.* (1997) show that trading rules generated using genetic programming techniques yield highly significant out-of-sample excess returns. Fernández-Rodriguez *et al.* (2000) show that Nearest Neighbor techniques, essentially a computational pattern-recognition method, significantly improves upon the performance of standard moving average rules.

The depth and liquidity of foreign currency markets also lends itself to studying the performance of high-frequency technical trading rules. Dempster *et al.* (2001) develop technical models using both genetic algorithms and reinforcement learning and analyze the performance of these trading rules at horizons as short as 15 min. They find that after accounting for modest transaction costs, their intraday trading strategies can generate annual returns in the range of 10–20%. Not surprisingly, they find that performance is typically best at the 15 minute trading frequency and falls off dramatically as one moves to trading once per day. Furthermore, they also find that reinforcement learning, essentially a computational learning method that is very close in spirit to dynamic programming, tends to overfit in sample, thus leading to poor performance relative to a genetic algorithm.

Though the papers cited above show that technical analysis can generate positive, statistically significant returns, it is still not clear that technical trading systems actually improve the welfare of investors. Dewachter and Lyrio (2003) address this question by applying a novel welfare decomposition to returns from simple technical trading models. Their decomposition attributes the utility cost of a technical trading strategy to either (1) expectational errors arising from the technical model or (2) the sub-optimality of the portfolio allocation resulting from the model (the sub-optimality arises from the fact that they only consider bang-bang trading rules that go fully long for positive signals and fully short for negative signals). They find that the cost of expectational errors generated by the technical models range from 0.6% to 2% in certainty equivalent terms for a log-utility investor. The contrast between the large costs of expectational errors presented in this paper and previous results showing the significant forecasting ability of technical analysis highlights the importance of implementation issues when considering the results of trading models. Though the results presented by Dewachter and Lyrio should give prospective technical analysts pause, the use of relatively naive trading rules in their analysis makes it difficult to pass judgment on technical analysis in general.

While most research on technical analysis in a foreign context has focused on currencies, there is an increasing amount of work on applying technical trading models to international equity markets. Detry and Grégoire (2001) test a variety of moving average rules on market indexes in Europe. Their results largely mimic those found using US data. While technical analysis does appear to generate positive returns, the returns do not appear to be large enough to overcome the probable transaction cost of implementing the strategies tested. Along similar lines, Fernández-Rodriguez *et al.* (1999) study the profitability of moving average and break out technical rules on the General Index of the Madrid Stock Exchange. Again, though they find evidence of positive returns, the returns do not appear to be large enough to compensate for trading costs. These results are not totally surprising given that these papers test well-known technical rules that are likely being used by traders in the markets studied.

The evidence from foreign markets supports much of what is known from domestic studies. First, it appears to be very difficult to profitably use technical analysis in equity markets. Though technical rules do appear to convey some information, the implementation costs of technical trading strategies negate the positive returns to this information. Second, technical trading rules appear to perform much more robustly in currency markets, perhaps due to a higher frequency of regime shifts relative to equity markets. Lastly, recent developments in computational learning appear to offer superior returns to traditional technical analysis.

4 Why does technical analysis work?

The papers discussed above focus on the performance and implementation of technical trading models in a variety of settings. The more fundamental issue, however, is why does technical analysis appear to enhance one's information set given that the data inputs are publicly available. As noted above, the "success" of technical analysis is often attributed to its ability to predict structural changes in the data generating process. This fact makes the currency markets a natural venue for employing technical strategies because central bank interventions are just the type of external shocks that can spark regime shifts. LeBaron (1999) showed that profits from simple moving average trading rules exhibited a high correlation with daily US official interventions. Several papers have tried to examine the causal relationship between technical trading profits and interventions. Neely (2000) confirms the results in LeBaron (1999) by noting that technical trading rule performance degrades substantially when intervention periods are excluded. He explores the relationship more closely by employing high frequency exchange rate data to examine trading model performance immediately before and after intervention. Contrary to the hypothesis that regime shifts induced by intervention drive the returns to technical analysis, he finds that the profits to technical trading actually *precede* interventions. He goes on to show that central bank interventions tend to be negatively correlated with deviations from purchasing power parity exchange rate levels and that central banks tend to trade against recent price trends, indicating a bias to "lean against the wind."

In closely related work, Frenkel and Stadtmann (2000) replicate the results in Neely (2000) focusing exclusively on the US Dollar/Deutschemark market over a longer horizon. Importantly, they find that technical trading profits Granger cause interventions with no significant reverse causality. This series of results is interesting because it speaks to the destabilizing nature of technical traders in the short run. The results suggest that technical traders tend to push prices to extremes, inducing central bank intervention. One must be cautious not to read too much into these results in that central bank interventions (with the exception of Japanese Yen currency pairs) have been extremely rare in recent years.

Osler (2003) examines another hypothesis for the usefulness of support and resistance levels, a commonly used technical indicator. Her analysis shows that market participants in the currency markets tend to cluster stop-loss and take-profit (price-contingent) orders at round numbers, i.e. 125.00 versus 125.07 for the Japanese Yen/US Dollar market. Once these support/resistance levels are hit, pent up demand hits the market and tends to cause price cascades. Thus, technical analysis is useful in identifying points where clusters of orders are located and profiting from the cascades that occur when these orders are filled.

In a similar vein, Kavajecz and Odders-White (2002) examine the correlation between technical indicators and the depth of the limit order book at the NYSE. Again, the basic idea is that methods to identify support and resistance levels in fact reveal information about the limit order book and the level of liquidity provided in the market. Interestingly, they find some evidence that variation in the depth of the limit order book Granger causes changes in technical analysis indicators.

Identifying the underlying factors behind the predictive success of technical analysis is arguably the key objective of research on technical analysis. As the research presented above shows, we appear to be advancing toward a better understanding of the forces that technical analysis sheds light on. The work by Osler (2003) and Kavajecz and Odders-White (2002) both indicate that technical analysis proxies for information about supply/demand conditions in the financial markets. The fact that supply/demand matters is controversial in itself given that most asset pricing models portray individual securities as perfect substitutes. The work in Neely (2000) and Frenkel and Stadtmann (2000) forcefully rejects the hypothesis that central bank interventions are the driving force behind the success of technical analysis in currency markets.

5 Summary

Our understanding of technical trading has certainly improved over the years. We are much more aware of how to measure the performance of trading rules, subject to data-snooping biases and transactions costs. The papers we have looked at highlight the need to separate the issue of forecastability of trading rules from that of their profitability. Whereas, trading rules show reasonable forecasting performance, in that they presage changes in the price process, their ability to generate profits is far weaker, and not that persistent. Technical trading seems to be prevalent in markets impacted by macroeconomic shifts such as those of foreign exchange. Recent work at the microstructure level provides a very promising set of insights into why trading rules seem to work. With the increasing facility for automated trading, allowing for microlevel capture of trading profits, we are positioned well for an explosion in research in, and the use of, sophisticated technical trading rules.

References

Allen, F. and Karjalainen, R. (1995). "Using Genetic Algorithms to Find Technical Trading Rules." Working Paper, Wharton School, Rodney L. White Center for Financial Research Working Paper #20-95.

Anderson, J. (2001). "Point and Figure Charting: A Computational Methodology and Trading Rule Performance in the S&P 500 Futures Market." Working Paper, Queensland University of Technology—School of Economics and Finance, Discussion Paper No. 01-01.

Andrew, L., Mamaysky, H. and Wang, J. (2000). "Foundations of Technical Analysis: Computational Algorithms, Statistical Inference, and Empirical Implementation." *Journal of Finance* 55(4), 1705–1765.

Bessembinder, H. and Chan, K. (1998). "Do the Profits From Technical Trading Rules Reflect Inefficiencies?" Working Paper, University of Utah.

Brock, W., Lakonishok, J. and LeBaron, B. (1992). "Simple Technical Trading Rules and the Stochastic Properties of Stock Returns." *Journal of Finance* **47**, 1731–1764.

Davis, R. E. (1965). *Profit and Probability*. West Lafayette, Indiana, USA: Dunn & Hargitt's Financial Service.

Dempster, M., Payne, T. and Romahi, Y. (2001). "Intraday FX Trading: Reinforcement vs. Evolutionary Learning." Working Paper, University of Cambridge, Judge Institute of Management Studies Working Paper 23/2001.

Detry, P. and Grégoire, P. (2001). "Other Evidences of the Predictive Power of Technical Analysis: The Moving Averages Rules on European Indexes." Working Paper, Université Catholique de Louvain, CeReFiM.

Dewachter, H. and Lyrio, M. (2003). "The Cost of Technical Trading Rules in the Forex Market: A Utility-Based Evaluation." Working Paper, Erasmus Research Institute of Management, ERS-2003-052-F&A.

Feng, C. and Smith, S. D. (1997). "Jump Risk, Time-Varying Risk Premia, and Technical Trading Profits." Working Paper, Federal Reserve Bank of Atlanta 97-17.

Fernández-Rodriguez, F., Sosvilla-Rivero, S. and Andrada-Félix, J. (1999). "Technical Analysis in the Madrid Stock Exchange." Working Paper, FEDEA Working Paper 99/05.

Fernández-Rodriguez, F., Sosvilla-Rivero, S. and Andrada-Félix, J. (2000). "Technical Analysis in Foreign Exchange Markets: Linear Versus Nonlinear Trading Rules." Working Paper, FEDEA Working Paper 00/02.

Frenkel, M. and Stadtmann, G. (2000). "Trading Rule Profitability and Interventions in the Dollar-Deutschemark Market." Working Paper, WHU Otto Beisheim Graduate School of Management.

Hauschild, K. and Winkelmann, M. (1985). "Kapitalmarketeffizienz und Point and Figure Analyse." *Kredit Und Kapital*, **18**.

Jacquier, E. and Yao, T. (2002). "Evaluating Dynamic Strategies: The Free Lunch was no Banquet." Working Paper, Boston College.

Kavajecz, K. and Odders-White, E. (2002). "Technical Analysis and Liquidity Provision." Working Paper, Wharton School, Rodney L. White Center for Financial Research Working Paper #11-02.

LeBaron, B. (1999). "Technical Trading Rule Profitability and Foreign Exchange Intervention." *Journal of International Economics* **49**, 125–143.

Levich, R. and Thomas, L. (1993). "The Significance of Technical Trading-Rule Profits in the Foreign Exchange Market: A Bootstrap Approach." *Journal of International Money and Finance* **12**, 451–474.

Neely, C. (2000). "The Temporal Pattern of Trading Rule Returns and Central Bank Intervention: Intervention Does Not Generate Technical Trading Rule Profits." Working Paper, Federal Reserve Bank of St. Louis Working Paper 2000-018C.

Neely, C., Weller, P. and Dittmar, R. (1997). "Is Technical Analysis in the Foreign Exchange Market Profitable? A Genetic Programming Approach." *Journal of Financial and Quantitative Analysis* **32**(4), 405–426.

Osler, C. (2003). "Stop-Loss Orders and Price Cascades in Currency Markets." Working Paper, Brandeis University.

Ready, M. (2002). "Profits from Technical Trading Rules." *Financial Management* **31**(3).

Reitz, S. (2002). "Why Technical Trading Works—A Simple Illustration." Working Paper, University of Giessen.

Sullivan, R., Timmermann, A. G. and White, Jr., H.L. (1999). "Data-Snooping, Technical Trading Rule Performance, and the Bootstrap." *Journal of Finance* **54**(5), 1647–1692.

White, H. (1997). "A Reality Check for Data Snooping." Technical Report, NRDA, San Diego, CA.