Journal of Financial Economics
Volume 106, Issue 3, December 2012, Pages 473-491

# Technical trading revisited: False discoveries, persistence tests, and transaction costs ☆

Pierre Bajgrowicz [a], Olivier Scaillet [b]

Show more

Share Cite

https://doi.org/10.1016/j.jfineco.2012.06.001
Get rights and content

## Abstract

We revisit the apparent historical success of technical trading rules on daily prices of the Dow Jones Industrial Average index from 1897 to 2011, and we use the false discovery rate (FDR) as a new approach to data snooping. The advantage of the FDR over existing methods is that it selects more outperforming rules, which allows diversifying against model uncertainty. Persistence tests show that, even with the more powerful FDR technique, an investor would never have been able to select ex ante the future best-performing rules. Moreover, even in-sample, the performance is completely offset by the introduction of low transaction costs. Overall, our results seriously call into question the economic value of technical trading rules that has been reported for early periods.

## Introduction

Whether technical trading rules can consistently generate profits, as opposed to just being lucky every now and then, is the subject of an ongoing debate. Practitioners have devoted significant resources to technical trading, which uses past price and volume data to infer future prices. A substantial segment of the investment industry employs indicators that include moving averages, support and resistance levels, and other filter rules. Technical indicators are as ubiquitous on professional information systems as on popular finance websites and online retail brokers. In spite of its popularity among practitioners, academics have long been skeptical about the merits of technical analysis. They argue that it is inconsistent with the theory of market efficiency, which states that all available information must be reflected in security prices. In hopes of resolving this conflict, researchers have undertaken numerous empirical studies of technical trading rules. Some have found results in favor of the ability of trading rules to deliver superior returns, e.g., Neftci (1991), Brock, Lakonishok, and LeBaron (BLL, 1992) Neely, Weller, and Dittmar (1997), Sullivan, Timmermann, and White (STW, 1999) Lo, Mamaysky, and Wang (2000), and Kavajecz and Odders-White (2004). Other studies conclude that trading rules cannot be used to predict future prices. For example, Fama and Blume (1966), Bessembinder and Chan (1998), Allen and Karjalainen (1999), and Ready (2002) show that transaction costs outweigh the predictive power of trading rules. In addition to the impact of transaction costs, researchers have warned against the danger of data snooping which raises the possibility that the reported results are spurious. Menkhoff and Taylor (2007) provide an extensive review of the literature on the use of technical analysis in foreign exchange markets.

In this paper, we revisit the apparent historical success of trading rules during early time periods found in previous studies, including studies reaching an overall negative conclusion such as Ready (2002). In particular we examine the performance of the 7,846 trading rules of STW on daily prices of the Dow Jones Industrial Average (DJIA) index between January 1897 and July 2011. The first contribution is to apply the false discovery rate (FDR) methodology developed by Barras, Scaillet, and Wermers (2010) in the context of mutual funds selection, as a new approach to select outperforming rules while accounting for data snooping. We show that the FDR approach has numerous advantages compared with existing methods. The second contribution is to perform a rigorous analysis of the economic value of the trading rules. We focus on two issues that have been only partly addressed in the literature: the impact of transaction costs and the question of whether investors could have reasonably selected the future outperforming rules without the benefit of foresight. Equipped with the more powerful FDR approach to detect rules with true predictive ability and accounting for transaction costs ex ante, we perform persistence tests in which we measure the out-of-sample performance of the selected rules. We are the first to carry out such a comprehensive persistence analysis of trading rules. Only by combining all these relevant factors can the economic value of the strategies be truly assessed.

To illustrate the problem of data snooping, imagine you put enough monkeys on typewriters and that one of the monkeys writes *The Iliad* in ancient Greek. Because of the sheer size of the sample, you are likely to find a lucky monkey once in a while. Would you bet any money that he is going to write *The Odyssey* next? The same principle applies to trading rules. By looking long enough and hard enough on a given set of data, an investor always finds a trading rule parameterization that works, even if it does not genuinely possess predictive power. For a discussion of the dangers of data snooping, see Lo and MacKinlay (1990), White (2000), and the references therein. Diebold (2006) also warns against the danger of in-sample overfitting. Kosowski, Naik, and Teo (2007) study the impact on detecting hedge fund performance.

In this paper we propose a new methodology to select superior trading rules while accounting for data snooping based on the FDR. We employ the $FDR^+$ and the $FDR^-$, developed by Barras, Scaillet, and Wermers (2010). The $FDR^{+/-}$ gives the proportion of false discoveries—rules with no genuine performance, separately among the rules selected as delivering statistically significant positive and negative performance. As we show in a Monte Carlo experiment, the FDR approach has advantages compared with statistical methods used in previous studies, e.g., the bootstrap reality check (BRC) of White (2000) employed by STW, and its stepwise extension by Romano and Wolf (RW, 2005). The BRC indicates only whether the rule that performs best in the sample beats the benchmark, after accounting for data snooping. It provides no information on the other strategies. In practice, investors prefer not to base their investment decision on a single strategy. Though potentially able to detect further outperforming rules, the RW method relies on the conservative familywise error rate (FWER), which results in a lack of power; see Romano, Shaikh, and Wolf (2008b) for a discussion. One further problem with methods derived from the BRC such as the RW method is that they do not select further strategies once they find a rule whose performance is due to luck, even if there remain an important number of true outperforming rules in the population. The FDR approach by tolerating a certain (small) proportion of false discoveries, does not suffer from the problem. We run a Monte Carlo study calibrated to the setting of our empirical work and taking into account the cross-sectional dependence of trading strategies. The Monte Carlo simulations illustrate that situations in which a rule with no genuine predictive power achieves one of the highest performance are common in practice. They also show that the FDR approach greatly improves the chances of detecting all true outperforming rules and behaves well even if the rules are not independent. Using the FDR method, an investor can construct a portfolio of rules on which to base his investment decision and, hence, diversify against model risk.

With the help of our new more powerful rules selection approach, we investigate whether the trading rules can make money. BLL show examples of historical performance and consider them as proof of the usefulness of the trading rules. STW argue that the findings of BLL are not spurious as the best rule passes the BRC data snooping test. However, although it can be the case that we are able to find rules that perform well historically, no indication exists that we can select these rules ex ante. Another important issue not addressed ex ante in BLL and STW is the impact of transaction costs. The rules selected before transaction costs produce very frequent trading signals, and their predictive power is likely to be offset by transaction costs. Previous studies do not treat transaction costs as endogenous to the selection process. Hence, the relevant question is: Could investors reasonably have anticipated which rules would generate performance outweighing transaction costs? To answer this question, we perform persistence tests, adding a transaction cost each time a buy or sell signal is generated. Specifically, we measure the out-of-sample performance of a portfolio of rules selected using our new FDR approach and updated every month using data from the previous month. Rebalancing the portfolio monthly has the further advantage of being closer to what is done in practice than previous studies. Investors never get the chance to trade over multiple-year periods before being evaluated, and they update their trading rules regularly in an attempt to adapt to the changing economic environment. The persistence analysis is a major contribution of this paper. STW qualify as out-of-sample the results for the period after the original BLL study but, in fact, they always measure performance in-sample. Persistence analysis has been applied to mutual funds, e.g., Carhart (1997). To our knowledge, however, this is the first time this type of persistence tests are performed on technical trading rules. Jacquier and Yao (2002) implement another approach to persistence analysis also inspired by the mutual fund literature. They follow Brown and Goetzmann (1995) and estimate the probability that a trading rule beats the benchmark over consecutive periods. Their study is limited to the ten moving average rules of BLL and finds that the performance is not persistent at horizons shorter than five years.

Our tests show that, even with our new FDR rule detection approach, the reason for choosing the rules with future superior performance is clear only to researchers examining the price data ex post. Contrary to the mutual fund literature, we conclude that there is no hot hands phenomenon. In addition, even the in-sample historical performance is canceled already with the inclusion of low (conservative) transaction costs. Again, it is only by considering all the relevant aspects—performance persistence, transaction costs, and data snooping—together that we can correctly assess the economic value of the strategies. Our study confirms the results of Ready (2002) and Allen and Karjalainen (1999), who also deal with data snooping and rule selection, though in a very different fashion based on a genetic algorithm.

Our analysis indicates that the past period of predictability reported by numerous studies is not really a puzzle. The BLL results should be viewed as a statistical anomaly, discovered ex post by extensive data snooping. In any case, they should not be viewed as an episode of market inefficiency, as the hypothetical predictability could not have been exploited. Although we provide evidence against the usefulness of the simple trading rules of STW to deliver superior returns when applied in a blue-chip investment environment (DJIA index), our results say little about the existence of profitable trading strategies in other markets or using different trade frequencies. The growing number of institutions getting involved in highfrequency trading hints that profitable algorithmic strategies can be found. Our results do, however, indicate that investors should be wary of the common technical indicators present on any investment website or professional information system and advertised as obvious money-making tools.

Section 2 reviews existing methods to account for data snooping and presents the FDR based approach. Section 3 describes the universe of 7,846 technical trading rules, the performance measurement, and the data. Section 4 illustrates the advantage of the FDR approach by applying it in the same framework as STW. Section 5 presents the persistence analysis, while simultaneously accounting for transaction costs. It also investigates the impact of short sale constraints. Section 6 gathers concluding remarks. Appendices contain technical details on the implementation of the FDR approach and results of Monte Carlo experiments showing the advantages of the FDR method. We also review the literature on gauging total transaction costs and their evolution over time, and we provide up-to-date data for current market conditions. An Appendix with supplementary empirical and simulation results as well as files with the data set and programs used in the paper are posted on the Journal of Financial Economics web page.

---

## Section snippets

### Data snooping measures

In Section 2.1, we review existing data snooping methods. We present our new approach based on the false discovery rate in Section 2.2, before discussing in Section 2.3 how we can use it to construct a portfolio of trading rules.…

### Trading rules, data, and performance measurement

We describe the universe of trading rules and the data in Section 3.1, and explain how we measure performance in Section 3.2.…

### Long-term in-sample performance

For each of the sample periods, the columns on the right-hand side of Table 1 display the in-sample performance of the best rule in the sample and the corresponding BRC $p$-value, as reported in STW. The columns on the left-hand side present the performance and size of the portfolio obtained using the RW method to control the FWER at the 5% level. Based on such in-sample evidence discovered ex post, BLL and STW conclude that technical rules can be used to generate profits. These results have no…

### Persistence analysis

The question addressed in this section is simple but essential to evaluate the economic value of the trading rules: Could investors reasonably have anticipated which rules would generate superior returns after transaction costs? It is important to ask what information could have been used to select the outperforming rules. If the answer is that the prediction could have been made based on an analysis of investment flows, fiscal policy, or market psychology, then price data alone are not…

### Conclusion

Previous studies, e.g., BLL and STW, have reported examples of technical trading rules generating superior returns, at least during early time periods. Based on such results observed ex post, they have concluded that trading rules were useful to deliver profits. In our paper, we reassess this apparent historical success.

First, we propose a new approach to select outperforming rules while accounting for data snooping based on the false discovery rate. The FDR method is designed to control false…

---

## References (55)


F. Allen *et al.*
Using genetic algorithms to find technical trading rules
Journal of Financial Economics (1999)

C. Blanchet-Scalliet *et al.*
Technical analysis compared to mathematical models based methods under parameters mis-specification
Journal of Banking and Finance (2007)

M. Brennan *et al.*
Strategic asset allocation
Journal of Economic Dynamics and Control (1997)

G. D'Avolio
The market for borrowing stock
Journal of Financial Economics (2002)

D. Duffie *et al.*
Securities lending, shorting, and pricing
Journal of Financial Economics (2002)

C. Geczy *et al.*
Stocks are special too: an analysis of the equity lending market
Journal of Financial Economics (2002)

P.-H. Hsu *et al.*
Testing the predictive ability of technical analysis using a new stepwise test without data snooping bias
Journal of Empirical Finance (2010)

C. Jones *et al.*
Short-sale constraints and stock returns
Journal of Financial Economics (2002)

D. Keim *et al.*
Transactions costs and investment style: an inter-exchange analysis of institutional equity trades
Journal of Financial Economics (1997)

R. Kosowski *et al.*
Do hedge funds deliver alpha? A Bayesian and bootstrap analysis
Journal of Financial Economics (2007)



## Cited by (191)

A False Discovery Rate approach to optimal volatility forecasting model selection
2023, International Journal of Forecasting

Show abstract

Technical analysis, spread trading, and data snooping control
2023, International Journal of Forecasting

Show abstract

No pain, no gain: You should always incorporate trading costs for a bias-free evaluation of trading rule overperformance
2022, Economics Letters

*Citation Excerpt :*
…In particular, authors typically adjust for either trading costs or data snooping, but rarely for both. The relatively few counterexamples find substantially reduced evidence in favor of “abnormal” predictability, see, e.g., Bajgrowicz and Scaillet (2012) for speculative trading rules, Chen and Velikov (2021) for trading strategies based on cross-sectional pricing anomalies, or Barras et al. (2010) for mutual fund performance (fund manager skill). A more popular, alternative approach is to estimate and discuss break-even trading costs after applying MTPs.3…

Show abstract

Testing technical trading strategies on China's equity ETFs: A skewness perspective
2022, Emerging Markets Review

Show abstract

Network herding of energy funds in the post-Carbon-Peak Policy era: Does it benefit profitability and stability?
2022, Energy Economics

Show abstract

Machine Learning and the Stock Market ↗
2023, Journal of Financial and Quantitative Analysis

## Recommended articles (6)

Research article

### Trading rules, competition for order flow and market fragmentation

Journal of Financial Economics, Volume 115, Issue 2, 2015, pp. 330-348

Show abstract

Research article

### Illusory profitability of technical analysis in emerging foreign exchange markets

International Journal of Forecasting, Volume 30, Issue 2, 2014, pp. 192-205

Show abstract

Research article

### Profitability of simple technical trading rules of Chinese stock exchange indexes

Physica A: Statistical Mechanics and its Applications, Volume 439, 2015, pp. 75-84

Show abstract

Research article

### Technical analysis and sentiment embeddings for market trend prediction

Expert Systems with Applications, Volume 135, 2019, pp. 60-70

Show abstract

Research article

### Modeling asset returns under time-varying semi-nonparametric distributions

Journal of Banking & Finance, Volume 118, 2020, Article 105870

Show abstract

Research article

### Systemic co-jumps

Journal of Financial Economics, Volume 126, Issue 3, 2017, pp. 563-591

Show abstract

☆ We would like to thank the editor and the referee for constructive criticism and numerous suggestions that have lead to substantial improvements over previous versions of the paper. We are grateful to M. Franscini-Scaillet for helping us to get the data on fund structure costs and futures trading costs via her industry contacts. We thank L. Barras, I. Chaieb, M. Dubois, B. Dumas, L. Frésard, E. Ghysels, R. Gibson, R. Gonzalez, P. Hsu, M. Rockinger, A.Timmermann, and A. Treccani, as well as seminar participants at the finance seminar of the University of Athens (2009), the Econometric Society European Meeting (2008), the European Finance Association conference (2008), the Bachelier Finance Society World Congress (2008), the European Financial Management Association annual meeting (2008), the Computational and Financial Econometrics workshop (2008), the Society for Financial Econometrics (SoFiE) inaugural conference (2008), the Swiss Society for Financial Market Research conference (2008), the Meeting of the Swiss Society of Economics and Statistics (2008), the Midwest Finance Association meeting (2008), the Brown Bag Seminar of the University of Zurich (2008), the University of Rome Workshop on Quantitative Finance (2008), the EC2 conference (2007), the Swiss Doctoral Workshop in Finance (2007), and the Association Française de Finance conference (2007) for helpful comments. We received financial support from the Swiss National Science Foundation through the National Center of Competence in Research: Financial Valuation and Risk Management (NCCR FINRISK).

View full text