Journal of Financial Economics 106 (2012) 473–491

Contents lists available at SciVerse ScienceDirect

Journal of Financial Economics

journal homepage: www.elsevier.com/locate/jfec

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

Pierre Bajgrowicz [a], Olivier Scaillet [b,*]

[a] Université de Genève, and Litasco SA, 40 Bd du Pont d'Arve, 1211 Geneva, Switzerland
[b] Université de Genève and Swiss Finance Institute, 40 Bd du Pont d'Arve, 1211 Geneva, Switzerland

ARTICLE INFO

Article history:
Received 23 November 2010
Received in revised form
5 December 2011
Accepted 6 December 2011
Available online 16 June 2012

JEL classification:
C12
C15
G11
G14

Keywords:
Technical trading
False discovery rate
Persistence
Transaction costs

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.

1. 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.

☆ 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 EC^2 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).

* Corresponding author. Tel.: +41 22 379 88 16; fax: +41 22 379 81 04.
E-mail addresses: pbajgrowicz@litasco.ch (P. Bajgrowicz), olivier.scaillet@unige.ch (O. Scaillet).

0304-405X/$ - see front matter
http://dx.doi.org/10.1016/j.jfineco.2012.06.001

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 $\text{FDR}^{+/-}$ 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