

Economics Department

Analysing Technical Analysis

SPYROS SKOURAS

ECO No. 97/36

EUI WORKING PAPERS



WP
330
EUR

EUROPEAN UNIVERSITY INSTITUTE

European University Institute



3 0001 0036 6990 2

EUROPEAN UNIVERSITY INSTITUTE, FLORENCE

ECONOMICS DEPARTMENT

EUI Working Paper ECO No. 97/36

Analysing Technical Analysis

SPYROS SKOURAS

**WP 330
EUR**



BADIA FIESOLANA, SAN DOMENICO (FI)

All rights reserved.
No part of this paper may be reproduced in any form
without permission of the author.

© Spyros Skouras
Printed in Italy in December 1997
European University Institute
Badia Fiesolana
I – 50016 San Domenico (FI)
Italy

Analysing Technical Analysis

Spyros Skouras*

email: skouras@datacomm.iue.it

This version, November 1997

Abstract

This paper proposes an expected utility framework for the treatment of theoretical and empirical issues in technical analysis. Circumstances are found in which a technical analysts' behaviour is 'rationalisable' and schemes for learning such rational behaviour are considered. The decision theoretic perspective developed is shown to be useful in formalising a measure of trading rule optimality, in creating improved rules and in using these rules to create powerful model specification tests. The framework is shown to be useful in relating the efficacy of technical analysis to the efficiency of financial markets and an empirical analysis of this relationship is provided.

*This paper was previously circulated with the title 'An (Abridged) Theory of Technical Analysis'. The research leading to this paper was supported by a grant from I.K.Y.. Thanks are due to Dave Cass, Alex Gumbel, Mark Salmon and particularly Ramon Marimon for useful comments and suggestions on previous versions of this paper.

1 Introduction and Motivation

Economists, until recently, have viewed technical analysis (or chartism) with a great degree of skepticism. Indeed, typically technical analysis was treated as an irrational mode of behaviour that can be approximated sufficiently well by the conceptual construct of a 'noise trader' (Black, 1986). This is somewhat paradoxical given the size and importance of the institutions utilising technical analysis: there are investment consultants who sell exclusively technical services and all major investment banks employ technical analysts. There are also private institutions which collect and sell data for analysts; and there is a huge range of expensive computer software as well as dozens of magazines and books on the subject.

The paradox is particularly striking when we reflect that in most other applications of economic modelling, purely irrational agents are anathema. Furthermore, it is surprising that technical analysis lacks a behavioural foundation given that technical analysts are driven by clearly economic motives to make decisions which are inherently quantifiable.

Part of the reason why technical analysis was dismissed as an irrational mode of behaviour was because it was inconsistent with a naive but popular reading of the efficient market hypothesis according to which prices followed a random walk (see for example Malkiel, 1996). However, it is now clear that the random walk is not a satisfactory model for prices and is neither a necessary nor a sufficient condition for market efficiency (see Fama, 1991 *inter alia*).

Apart from being inconsistent with the vogue in financial economics, a further offence committed by technical analysis was that the rules constituting it were fuzzy and subjective. As put by Tewels, Harley and Stone: "Chart patterns are almost completely subjective. No study has yet succeeded in mathematically quantifying any of them. They are literally in the mind of the beholder..."¹. Thus a divide grew between financial theorists and analysts: in the words of a famous analyst,

¹Quoted in Murphy (1986), p.17

they believed that "Chart reading is an art"² whereas theorists saw this activity as the latest branch of astrology³. Recently, however, Neftci (1991) has developed a methodology which has allowed economists to clearly distinguish which rules are well-defined in the sense of generating Markov-times. This has given economists a criterion according to which they can judge the formalisability of a rule and has therefore made it possible to study those aspects of technical analysis which are not artistic (i.e. are formalisable).

In conjunction with the provision of increasing evidence that technical trading rules are able to detect nonlinearities in financial time series (see e.g. Brock, Lakonishok and LeBaron, 1992 (henceforth Brock *et al.*), LeBaron (1992a,b), Levich *et al.* (1993) and the references therein), these recent developments have generated a significant interest in technical analysis. In particular, the recent literature has focused on applying and extending Brock *et al.*'s innovatory use of trading rules to characterise the distributions of financial time series.

This paper is primarily concerned with determining what the empirical observation of the use of a trading rule implies for the rationality and preferences of its user. In particular, in Section 2, we seek to find conditions under which technical analysis is a 'rationalisable' activity, in the sense of being consistent with expected utility maximisation.

In Section 3, we examine what becomes of technical analysts when they have limited information about the environment. If this is the case, analysts must learn their optimal actions, and we show how a decision theoretic approach to learning can be used to model this type of behaviour. This leads to the concept of an 'artificial technical analyst', which formalises the loose notion of what it means for a rule to be good or 'optimal' (see Allen F. and Karjalainen 1996, Neely *et al.* 1996, Pictet *et al.* 1996, Taylor 1994, Allen P. and Phang 1994, Chiang 1992, Pau 1991).

²This is taken from Murphy (1986), which is considered a 'classic' book on technical analysis.

³Investment analysts, who offer predictions based on the movement of the stars are becoming quite popular in the last few years, e.g. Weingarten H., 1996, *Investing by the Stars: Using Astrology in the Financial Markets*, McGraw-Hill.

This formalisation is important because it indicates that an explicit measure of rule optimality can and should be derived from a specific utility maximisation problem and that a rule which is optimal for all agents will not usually exist.

In the next two sections, we turn to empirical applications of our framework. Section 4 investigates the importance of using an optimal rule when trading rule returns are applied to characterise financial series. We find that it is 'easy', for any specific rule to lead to misleading results when it is chosen in an ad hoc manner, but that this problem can be mitigated by using an optimal rule in such applications.

Section 5 constitutes an investigation of the (hitherto ignored) relationship between market efficiency and technical trading at both a theoretical and an empirical level. Concepts are introduced which are helpful in calculating the level of transaction costs necessary for past prices to be, in a specified sense, 'useless' to risk-averse investors.

Section 6 closes this paper with a synopsis of its conclusions.

2 Technical analysis formalised

At a certain level of abstraction, technical analysis is the selection of rules determining (conditional on certain events) whether a position in a financial asset will be taken and whether this position should be positive or negative. One important difference between an analyst and a utility maximising investor is that the rules the analyst follows do not specify the magnitude of the positions he should take.

This leads us to the following definition of technical analysis:

Df.1: Technical Analysis is the selection of a mapping d which maps the information set I_t at t to a space of investment decisions Ω .

Ass. 1: The space of investment decisions Ω consists of three events $\{\text{Long, Short, Neutral}\} \equiv \Omega$. We will use a more convenient integer representation of these events, so $\Omega \equiv \{1, -1, 0\}$.

The event space on which these conditions are written are usually some quantifiable variables such as the time series of prices, volatility, or the volume of trading (e.g. Blume *et al.* 1994) of an asset. Here, we focus our attention on rules which are based solely on the realisation of a finite history of past prices. Restricting the information set of technical traders to past prices rather than, say, past volume is justified by the fact that in order to judge the effectiveness of any rule, prices at which trade occurs must necessarily be known. Hence, the restriction we will make allows an examination of technical analysis when the minimum information set consistent with its feasibility is available. This is made explicit in the assumption below.

Ass. 2: $I_t = P_t = \{P_t, P_{t-1}, P_{t-2}, \dots\}$.

We use the standard notation $E_t(\cdot)$ to refer to $E(\cdot/I_t)$.

Technical Trading Rules and Rule Classes.

Observing the practice of technical analysis, we are able to offer a sharper characterisation of the form that mappings $P_t \rightarrow \{1, -1, 0\}$ actually take. It is the case that rules actually used differ over time and amongst analysts, but are often very similar and seem to belong to certain 'families' of closely related rules, such as the 'moving average' or 'range-break' family (see Brock *et al.*). These families belong to even larger families, such as those of 'trend-following' or 'contrarian' rules (see for example Lakonishok *et al.* 1993). Whilst it is difficult to observe widespread use of any particular rule, certain 'families' are certainly very widely used. When we choose to analyze the observed behaviour of technical analysts, we will therefore need to utilise the concept of a rule family, because empirical observation of a commonly used type of mapping occurs at the level of the family rather than that of the individual rule. We formalise the distinction between a rule and a family by defining and distinguishing technical trading rules and technical trading rule classes.

Df. 2: A Technical Trading Rule Class is a single valued function $D : P_t \times x \rightarrow \Omega$ where x is a vector of parameters.

Df. 3: A Technical Trading Rule is a single valued function:

$$d_t = D(P_t, x = \bar{x}) : P_t \rightarrow \Omega$$

which determines a unique investment position for each history of prices⁴.

2.1 The revealed preferences of Technical Analysts

Having defined the main concepts required to describe technical analysis, we now attempt to identify investors who would choose to undertake this activity. In particular, we find restrictions on a rational (in the von Neumann-Morgenstern sense) agents' preferences that guarantee he will behave (i.e. will be) a technical analyst.

For this purpose, consider the following simple but classic investment problem. An investor i has an investment opportunity set consisting of two assets: A risky asset paying interest R_{t+1} (random at t) and a riskless asset (cash) which pays no interest. He owns wealth W_t and his objective is to maximise his next period expected utility of wealth by choosing the proportion of wealth θ invested in the risky asset. We will assume $\theta \in [-1, 1]$ reflecting the assumption that borrowing is not allowed but that shortselling of the risky asset is possible to a value determined by current wealth. His expectations E_t are formed on the basis of past prices P_t , as dictated by A2.

Formally, the problem solved is:

$$\begin{aligned} \max_{\theta \in [-1, 1]} E_t U^i(W_{t+1}) \\ \text{s.t. } W_{t+1} = \theta W_t(1 + R_{t+1}) + (1 - \theta)W_t \end{aligned}$$

Or equivalently,

⁴Notice that any set of rule classes $\{D_i\}_{i=1}^n$ can be seen as a meta-class itself, where the parameter vector $x = \{i, x_i\}$ determines a specific technical trading rule.

$$\max_{\theta \in [-1,1]} E_t U^i(W_t(1 + \theta R_{t+1})) \quad (1)$$

The solution to (1) is obtained at:

$$\theta^* = \arg \max_{\theta \in [-1,1]} E_t U^i(W_t(1 + \theta R_{t+1})) \quad (2)$$

and assuming that expectations are conditional on past prices only, this implies:

$$\theta^* : P_t \rightarrow [-1, 1]$$

So we see that an investor will not in general use trading rules as defined above, since he is interested not only in whether he should take a long or a short position but also what the **size** of this position is. An exception to this is the risk-neutral investor, whose maximisation problem is:

$$\max_{\theta \in [-1,1]} E_t(W_t(1 + \theta R_{t+1})) \quad (3)$$

Which simplifies to:

$$\max_{\theta \in [-1,1]} \theta E_t(R_{t+1}) \quad (4)$$

And in this case, θ will optimally take bang-bang solutions. Letting θ_n^* be the risk-neutral investor's optimal choice, and assuming that $E_t(R_{t+1}) = 0$ results in $\theta_n^* = 0$ (so that θ_n^* is a single-valued function), then:

$$\theta_n^* : P_t \rightarrow \Omega$$

Hence, we see that a risk-neutral investor conditioning on past prices, will choose technical trading rules. This result is summarised in the following proposition:

Proposition I: The risk-neutral investor solving (1) is an expected utility maximising agent who always uses technical trading rules. Thus, expected utility maximisation and technical analysis are compatible.

On the basis of proposition I, we can define a technical analyst as a risk-neutral investor:

Df. 4. A Technical Analyst, is a risk-neutral investor who solves:

$$\max_{d(P_t) \in D} dE_t(R_{t+1}) \quad (5)$$

Where D is a function space including all functions with domain $R^{\dim(P_t)}$ and image $\{-1, 0, 1\}$.

The returns accruing to an analyst when he uses a rule d are dR_{t+1} and will be denoted R_{t+1}^d .

Obviously, different trading rules will be associated with different returns.

3 Technical Analysts and learning.

Let us assume henceforth that the technical analyst does not know $E_t(R_{t+1})$ but has a history of observations of P_t on the basis of which he must decide his optimal action. This decision is a standard problem in learning theory where an agent must learn his optimal response in a game played against the (market) environment. Whilst these learning problems are conceptually simple, the key in solving them is inferring the correct conditioning of the data and this may often prove difficult. As a practical matter, any solution method can only be expected to give an approximation to the true solution. There are two approaches to modelling the analysts' learning problem given a learning sample of past prices, and they each give approximations which are valuable in different contexts.

They are chosen according to different metrics on the basis of which potential solutions may be judged, which are derived from:

i) *'Statistical' loss functions* : In this case, the objective is to learn $E_t(R_{t+1})$, typically involving the selection of an estimator \hat{R}_{t+1} for R_{t+1} through the choice of parameters from a set of models, (e.g. GARCH models) according to some goodness of fit criterion, such as least squares or maximum likelihood. It is then straightforward for the analyst replace $E_t(R_{t+1})$ with \hat{R}_{t+1} , and hence determine d^* . We will refer to this method as the 'Econometric approach' to learning, as this techniques involved are of an econometric nature.

ii) *Context dependent loss functions*: In this case, the objective is to learn d^* directly. This is achieved by choosing \hat{d}^* , an estimate for d^* which has been found to give in-sample optimal solutions to (5) from a specified function space D . This method may be termed the 'Decision theoretic approach', since in decision theory learning is not focused on determining the underlying stochastic environment, but in determining an action which is an optimal decision for a specified agent.

The two approaches differ in their focus, since the latter does not even involve the formation of an explicit expectation for R_{t+1} . The use of different loss functions for learning implies that we do not expect the two methods to yield the same solutions unless the 'true' solution is contained in both postulated models. Even then, the two approaches give the same solution only asymptotically and under certain regularity conditions. The advantage of the decision theoretic approach in the context of this paper is that it leads to rules which are optimal with respect to the technical analysts' loss function rather than with respect to a statistical criterion⁵ the properties of which may be irrelevant for the objective at hand.

Empirical studies tend to confirm that this distinction is important. For example, Leitch and Tanner (1991) show that standard measures of predictor performance are bad guides for the ability of a predictor to dis-

⁵An effort to develop a methodology for constructing econometric models based on general loss functions is under way (see e.g. Christofferson and Diebold (1995) and the references therein). A fully operational methodology of this form should bridge the gap between the econometric and decision theoretic approach to learning.

cern sign changes of the underlying variable⁶. Also, Taylor (1994) finds that trading based on a channel trading rule outperforms a trading rule based on ARIMA forecasts because the former is able to predict sign changes more effectively than the latter. This is likely to be due to the fact that the ARIMA forecast is chosen according to the 'wrong' criterion, (i.e. to minimise least squares)⁷. Finally, Kandel and Stambaugh (1996) show that statistical fitness criteria are not necessarily good guides for whether a regression model is useful to a rational (Bayesian) investor. These theoretical and empirical considerations suggest that any reasonable model of analysts' learning must be based on a decision theoretic perspective. Thus, this paper gives a decision theoretic treatment of the analysts' learning problem.

A crucial ingredient in forming a decision theoretic approach to technical analysis is an appropriate functional space from which to choose trading rules d . Just as the econometrician has a set of models for the distribution $F(R_{t+1}/\mathbf{P}_t)$, (e.g. GARCH, random walk models, etc.) from which he chooses a member, the analyst needs a set of rules from which

⁶The technical analyst is clearly interested in the sign of R_{t+1} , not its magnitude. In particular, the magnitude of R_{t+1} is irrelevant for his decision problem if $\text{sign}(R_{t+1})$ is known. Hence, he seeks a predictor $E^i(R_{t+1})$ which takes account of his loss function by being accurate in terms of a sign-based metric. Satchell and Timmerman (1995) show that, in general, standard least square error predictors do not have this property. Their proof derives from the fact that unless the distribution of $F(R_t)$ is restricted, there is a non-monotonic relationship between a predictor's squared errors and the probability of it correctly predicting $\text{sign}(R_{t+1})$.

⁷A number of studies of technical trading implicitly or explicitly assume away the possibility that there exists a nonmonotonic relationship between the accuracy of a prediction in terms of a metric based on Euclidean proximity and a metric based on the probability of predicting a sign change correctly. Examples are Taylor 1989a,b,c, Allen and Taylor 1989, Curcio and Goodhart 1991 and Arthur et al. 1996, who reward agents in an artificial stockmarket according to traditional measures of predictive accuracy. When the assumption is made explicit its significance is usually relegated to a footnote, as in Allen and Taylor 1989, fn. p.58, '*our analysis has been conducted entirely in terms of the accuracy of chartist forecasts and not in terms of their profitability or 'economic value' although one would expect a close correlation between the two*'. As we have argued however, the preceding statement is unfounded and the results of the above studies must be treated with extreme caution.

to choose his optimal rule. Ideally, we would like to have a procedure for checking whether a family of rules or a set of families contain the universally optimal rule, but it is doubtful if such a process exists. Hence we instead choose families which are empirically observed, reflecting the belief that they have been developed by analysts of the real world to serve them (according to their loss function) in a way analogous to that in which econometric models serve econometricians. The next section illustrates with an example how an analyst learns an optimal rule from a rule class.

3.1 Learning the Optimal Moving Average Trading Rule

The moving average rule class is one of the most popular rule classes used by technical analysts and has appeared in most studies of technical analysis published in economics journals. For these reasons, we will use it to illustrate how a technical analyst would learn the optimal rule within the moving average class. Let us begin with a definition⁸ of this class:

Df 6: The Moving Average rule class $MA(P_t, x_t)$ is a trading rule class s.t.:

$$MA(P_t, x) = \left[\begin{array}{l} 1 \text{ if } P_t \geq (1 + \lambda) \frac{\sum_{i=0}^n P_{t-i}}{n+1} \\ 0 \text{ if } (1 - \lambda) \frac{\sum_{i=0}^n P_{t-i}}{n+1} \leq P_t \leq (1 + \lambda) \frac{\sum_{i=0}^n P_{t-i}}{n+1} \\ -1 \text{ if } P_t < (1 - \lambda) \frac{\sum_{i=0}^n P_{t-i}}{n+1} \end{array} \right] \quad (6)$$

where $P_t = [P_t, P_{t-1}, \dots, P_{t-N}]$,

$x = \{n, \lambda\}$,

$X = \{N, \Lambda\}$,

$N = \{1, 2, \dots, N\}$, this is the 'memory' of the MA

$\Lambda = \{\lambda : \lambda \geq 0\}$ this is the 'filter' of the MA

⁸As defined, the moving average class is a slightly restricted version of what Brock *et al.* (1991, 1992) refer to as the 'variable length moving average class' (in particular, the restriction arises from the fact that the short moving average is restricted to have length 1).

Suppose a technical analyst postulates this as a set of models from which to choose his trading rule at time t . Then he effectively substitutes his objective function (5) with

$$\max_{n,\lambda} MA(\mathbf{P}_t, n, \lambda) E_t(R_{t+1}) \quad (7)$$

Since in any practical application $E_t(R_{t+1})$ is unknown, the optimal solution (n^*, λ^*) needs to be learned on the basis of m past prices. When there is little information about the distribution of \mathbf{P}_t , a natural estimate $(\hat{n}^*, \hat{\lambda}^*)$ for the optimal solution is⁹:

$$(\hat{n}^*, \hat{\lambda}^*) = \arg \max \sum_{i=t-m}^{t-1} MA(\mathbf{P}_i, n, \lambda) R_{i+1} \quad (8)$$

Let us fix ideas with an example. Suppose an analyst solving (8) wanted to invest on the Dow Jones Industrial Average and had access to $m = 250$ daily observations of this index at t , so that $I_t = \{P_t, P_{t-1}, \dots, P_{t-249}\}$. What rule would he choose? Figure I plots the answer to this question¹⁰ repeated 6157 times from $t=1/6/1962$ till $31/12/1986$.

Insert Figure I Here

The above figure plots a sequence of rules, denoted $\{d(n_t^*, \lambda_t^*)\}_{t=1}^{6157}$, which are optimal in a recursive sample of 250 periods. Such rules d_t^* which are optimised in the period before t , yield out of sample returns which shall be denoted $R_{t+1}^{d_t^*}$. It may be interesting to note that these rules

⁹If $\frac{1}{m} \sum_{i=t-m}^{t-1} MA(\mathbf{P}_i, n, \lambda) R_{i+1}$ converges uniformly to $E(MA(\mathbf{P}_t, n, \lambda) R_{t+1})$ as $m \rightarrow \infty$, then $(\hat{n}^*, \hat{\lambda}^*) \rightarrow (n^*, \lambda^*)$.

¹⁰The moving average parameters were restricted so that $N = \{1, 2, \dots, 200\}$ and Λ was discretised to $\Lambda = \{0, 0.005, 0.01, 0.015, 0.02\}$. This discretisation allowed us to solve (8) by trying all $\dim(N) \cdot \dim(\Lambda) = 1000$ points composing the solution space.

¹¹This data corresponds to the third subperiod used by Brock *et al.* and to most of the data used by Gencay (1996).

correspond to what Arthur (1992) terms 'temporarily fulfilled expectations' of optimal rules. It is difficult to interpret the sharp discontinuities observed in these expectations.

These optimal rules are important in investigations of any aspect of empirically observed rule classes in the same way that an estimated GARCH model is necessary for the evaluation of the usefulness of GARCH models in describing financial series. The reason for this will be illustrated in the next section.

4 The advantages of basing investigations of Technical Analysis on optimally learned rules.

Much of the economic literature on technical trading rules has asked whether popular types of rules such as the moving average class, will yield returns in excess of what would be expected under some hypothesized distribution of stock returns (e.g. Brock *et al.*, Levich and Thomas (1993), Neftci (1991), etc.). A serious criticism levied against this type of analysis arises from the fact that it involves a testing methodology which is not closed. The methodology is not closed in the sense that the choice of rules to be tested is *ad hoc* since it is made according to non-rigorous and often implicit criteria. Whilst this is clearly unappealing from a theoretical viewpoint, one might argue that at a practical level, a closed methodology in which optimally learned rules are used would yield very similar results. If so, the *ad hoc* approach to rule choice might seem justified (at least as a basis for empirical tests) because of its simplicity.

The objective of this section is to empirically refute this argument by showing that the results drawn from empirical investigations of rule returns depend crucially on the choice of rules and that therefore the received *ad hoc* approach to testing trading rule returns entertains two grave lacunae. These are:

- i. 'Striking' or 'anomalous' results may be coincidental. The likelihood of such coincidences appearing in the literature is augmented by the fact that published research is biased in favour of reporting 'anomalies' over 'regularities'. In Section 4.1, we show that analysis based on a small sample of *ad hoc* rules is subject to the possibility of leading to spurious conclusions since the distribution of rule returns in a class is very diverse and hence small samples of rules are unrepresentative.
- ii. Results may be less striking than they ought to be, because *ad hoc* selection is by definition suboptimal. Hence, analysis based on *ad hoc* rules is likely to be weaker than that based on optimally chosen rules. This is shown in Section 4.2.

While the effect of the two problems on the content of the reported conclusions work in opposite directions, they by no means cancel out; rather, the two effects compound the overall uncertainty regarding the validity of conclusions drawn when they are present¹². However, by applying the decision theoretic approach to technical analysis proposed above, we can overcome the problem of arbitrariness and close the methodology for testing hypotheses on return distributions of rule classes; this is achieved by restricting the economists' choice of rules to be the same as the rationalisable rules which are learned by a technical analyst¹³. This not only eliminates the arbitrariness, but also makes the hypothesis we are testing more precise, since we can characterise the rule the returns of which we are testing as an optimal action of a specific agent.

¹²The seriousness of these pitfalls should be expected to increase with the size of the rule classes from which a rule is arbitrarily chosen.

¹³Of course, a degree of arbitrariness remains in the economists' selection of the rule class to be tested. However, we have already mentioned that there exists much stronger empirical evidence on the basis of which to choose a rule class than for any specific rule.

4.1 The distributional diversity of ad hoc Technical Trading Rules

As has been remarked, the specific rules chosen for analysis by economists examining trading rule returns belong to classes which are known to be used widely; the same is not true for the specific rules and it has proved difficult to argue on any *a priori* grounds that these rules were representative of the rule classes to which they belong. This could however be supported on empirical grounds if the distributions of rule returns belonging to the same class were sufficiently similar¹⁴. In this case, any rule could be used as a proxy for the class as a whole¹⁵.

Unfortunately, as we shall illustrate below, these distributions typically do differ significantly. Figure 2 shows the returns accruing to each rule belonging to the moving average class if it were consistently applied during the period 1/6/1962-31/12/1986 on the DJIA.

Insert Figure II here

The figure above is meant to illustrate the problem with the *ad hoc* approach whereby specific rules are used as proxies for the distribution of expected returns of a whole class. The highest mean return from a rule in this class is much higher than the lowest mean return. To be more

¹⁴Most of the literature has focused only on the distribution of the first moments of rule returns. Notice that an analysts' optimal choice of rule is likely to produce rule returns which are most 'striking' in terms of first moment, because it is only in terms of this moment that his choice is optimised.

¹⁵That this is the case is suggested by Brock *et al*, who say that '*Recent results in LeBaron (1990) for foreign exchange markets suggest that the results are not sensitive to the actual lengths of the rules used. We have replicated some of those results for the Dow index*', p1734, fn. The 'recent results' to which Brock *et al* refer are a plot of a certain statistic of 10 rules. Apart from the fact that 10 rules constitute a small sample, the minimum statistic is almost half the size of the maximum statistic - so it is not entirely clear that these results support the claim made.

On the other hand, the conclusions Brock *et al.* draw are valid because it so happens that the rules they chose did not display extreme behaviour and in fact are valid *a fortiori* since they generated sub-average returns.

precise, the mean return of the best rule is 1270 times larger than that of the worst rule. Since the means are taken from samples with more than 6000 observations, the conclusion is very strong. This difference indicates a very large variance in returns accruing to rules within the same class. Furthermore, the morphology of the best and worst rules is very similar: The best rule is the three period moving average with no filter $MA(3, 0)$ and the worst is the four period moving average with a 2 percent filter $MA(4, 0.02)$.

The conclusions we draw from these results are the following:

- i. It is easy to ex post find a rule that will have 'unusual' expected returns. Rule returns have large variance.
- ii. The expected returns of rules display significant variance even within small areas of the classes' parameter space. This is important because some authors choose to calculate returns for a few rules sampled evenly from the space of all rules, reflecting the implicit assumption that rules are 'locally' representative. However, this assumption is unfounded.

4.2 Optimal vs. Representative returns distributions

Having illustrated that a small sample of rules from a class is insufficient for an analysis of the class as a whole, we now turn to a different issue. The purpose of this section is to show that even when a sufficiently large sample of rules are used for inferences about the mean performance of a class, this is still a bad way of judging the returns accruing to a user of a trading rule class. The reason for this, is that a technical analyst who at t chooses d^* from D , should be expected to have learned to make a better-than-average choice of d . Imposing the use of an 'average' rule is like estimating a GARCH model for a time series by choosing the GARCH specification which has average rather than 'minimum' least squared errors. We therefore conclude that 'representative' choices of

rules cannot be expected to be as good as the rules chosen by an agent who bases his decision on past experience. What this implies is that the trading rule returns obtained from following optimally learned trading rules $R_{t+1}^{d^*}$ should be expected to be greater than R_{t+1}^d for any d which is fixed with at the beginning of time.

The following table which utilises some of the information in Brock *et al.*¹⁶ (1991, Table V) is intended to show that indeed, the results reported there on the basis of various fixed rules d are much weaker than those which can be drawn by using the time-varying optimal rule d^* derived in section 3.

Insert Table I here

The table indicates that all t-ratios are much higher for the optimal rule we have developed. Hence, this table allows us to reject the hypothesis that the returns of the DJIA are normally, identically and independently distributed¹⁷ with much greater confidence than that offered by Brock *et al.*'s analysis¹⁸. Taken together, the results of this section constitute a strong case for the selection of trading rules according to an explicit criterion, such as the one we obtain in Section 3 by teaching analysts to choose rules which have performed well in the past.

5 Market efficiency and technical trading

It is often heard that 'If markets are efficient, then (technical) analysis of past price patterns to predict the future will be useless', (Malkiel, 1992).

¹⁶Note that Brock *et al.* (1992) reproduce only a part of this table

¹⁷The table also contains information which is sufficient to show that the Cumby-Modest (1987) test for market timing would, if the riskless interest rate were zero, confirm the ability of a technical analyst learning optimal rules to conduct market timing.

¹⁸It is expected that the optimal rule will be equally powerful as a specification test for other hypothesised distributions including those considered by BLL (AR, GARCH-M, EGARCH). However, we must leave confirmation of this for future research.

In this section, we attempt to develop a way of analysing the relationship between the efficiency of markets and the efficacy of technical analysis. However, at present there seems to be little consensus as to what an efficient market is (see LeRoy 1989 and Fama 1991) consequentially to the lack of an accepted model of financial markets. Most of the non-tautological definitions that have been proposed seek in their weak forms to incorporate the idea that profitable intertemporal arbitrage is not possible¹⁹ (Ross, 1987). In its very weakest forms, this is interpreted as meaning that once transaction costs are included, no risk-averse agent can increase his utility by attempting to 'time' the market. This statement is so weak that some authors (for example LeRoy, 1989, p1613 fn.) consider this notion of market efficiency to be intestable.. However we shall show below how this test can be conducted if we assume that the time series of prices is the market clearing equilibrium of an economy with a single risky asset.

We will refer to the version of the efficient market hypothesis that we have described as the Lack of Intertemporal Arbitrage (LIA) Hypothesis and discuss its implications for technical trading rules. We assume there exist agents in the market who are involved in solving (1), which we repeat here for convenience:

$$\max_{\theta \in [-1,1]} E_t U^i(W_t(1 + \theta R_{t+1})) \quad (1)$$

We will say that LIA is confirmed if knowledge of past prices does not affect the optimal actions of any market participant solving (1).

Df 4.1. The Lack of Intertemporal Arbitrage (LIA) Hypothesis holds in a market in which there exist agents who solve (1) and have utility functions $u_i \in U$ if $\forall i$

$$\arg \max_{\theta \in [-1,1]} EU^i(W_t(1 + \theta R_{t+1})/P_t) = \theta^* \quad \forall P_t \quad (9)$$

¹⁹A notable exception is Olsen et al. 1992, who propose a definition according to which 'efficient markets...are a requirement for relativistic effects and thus for developing successful forecasting and trading models'.

where:

$$\theta^* \equiv \arg \max_{\theta \in [-1, 1]} EU^i(W_t(1 + \theta R_{t+1})) \quad (10)$$

Where $E(\cdot/P_t)$ is the true expectation conditional on P_t and $E(\cdot)$ is the expectation of the true marginal distribution.

What this definition implies is that the true joint distribution $F(R_{t+1}, P_t)$ is such that knowledge of P_t in no way affects the actions of any market participant; it does not mean that $F(R_{t+1}) = F(R_{t+1}/P_t)$. For example, suppose P_t is only useful for predicting third and higher order moments of the distribution. Then in a market with mean-variance agents, actions will not be affected by knowledge of P_t although in a market populated with other types of agents this may be the case. Hence, according to our definition, a market is efficient with respect to a class of agents and the efficiency of a market can be viewed as a function of the wideness of this class. Formally, efficiency is determined by the wideness of the space of utility functions U^i for which we accept the null hypothesis that LIA holds when we test:

$$H_0(\text{LIA}) : \arg \max_{\theta \in [-1, 1]} EU^i(W_t(1 + \theta R_{t+1})/P_t) = \theta^* \text{ for some } P_t \quad (11)$$

versus

$$H_1(\text{Not LIA}) : \arg \max_{\theta \in [-1, 1]} EU^i(W_t(1 + \theta R_{t+1})/P_t) \neq \theta^* \text{ for some } P_t \quad (12)$$

5.1 Technical Trading Rules and LIA

5.1.1 A sufficient condition on rule returns for the rejection of LIA

If the distributions $F(R_{t+1})$ and $F(R_{t+1}/P_t)$ are known, then testing LIA is straightforward. When this is not the case, the usual approach for this

type of test is to use an estimated model for the unknown distributions. Here we take an alternative testing approach based on technical trading rule returns.

In particular, noticing that H_1 necessarily holds if there exists a trading rule $d(\mathbf{P}_t)$ s.t.:

$$EU^i(W_t(1 + \theta^* d(\mathbf{P}_t) R_{t+1})) > EU^i(W_t(1 + \theta^* R_{t+1})) \quad (13)$$

We conclude that a test for LIA based on trading rules, can be obtained by replacing H_1 with:

$$H_{1'} : \exists d \text{ s.t. } EU^i(W_t(1 + \theta^* R_{t+1}^d)) > EU^i(W_t(1 + \theta^* R_{t+1})) \quad (14)$$

Whilst this alternative hypothesis is weaker than (12), we shall see that it is still powerful. We show this below where we in turn test H_0 vs. $H_{1'}$ under risk-neutral, mean-variance and risk-averse specifications of U^i .

5.1.2 The Risk-Neutral Case

In this case, by the linear structure of U^i , $H_{1'}$ becomes (assuming θ^* is positive, i.e. $E(R_{t+1}) > 0$):

$$H_{1'}^{rn} : \exists d \text{ s.t. } E(R_{t+1}^d) > E(R_{t+1}) \quad (15)$$

Suppose we use the optimal moving average rules $\{d(n_t^*, \lambda_t^*)\}_{t=1}^{6157}$ derived in section 3 and the corresponding returns $R_{t+1}^{d^*}$ to test $H_{0'}^{rn}$. Then referring to the table below, we conclude that the probability that H_0 (LIA) is accepted is extremely low.

Rules	Mean Return	St. Dev.	$\Pr(R_{t+1}^{d*} \leq R_{t+1})$
R_{t+1}	0.0002334	0.008459	-
R_{t+1}^{d*}	0.000801	0.008335	8,887e-5

Table II *Note that the last column was calculated by assuming normality of both R_{t+1} and R_{t+1}^{d*}*

Hence we can conclude with great confidence that there exist intertemporal arbitrage opportunities (LIA is rejected) for risk-neutral agents investing in the market for the DJIA index.

5.1.3 The Mean-Variance Case

Suppose now that U^i is not linear, but instead is such that i has mean-variance utility. Is it still the case that LIA is violated? The reason this might not be the case is that although there exists a rule satisfying (16) it involves greater variance than R_{t+1} and hence is not preferred by mean-variance agents. For example, LeRoy (1989) argues that:

...even though the existence of serial dependence in conditional expected returns implies that different formulas for trading bonds and stock will generate different expected returns, because of risk, these alternative trading rules are utility-decreasing relative to the optimal buy-and-hold strategies.

In order to take account of this possibility when testing for LIA, $H_{1'}^{mv}$ is the relevant alternative hypothesis where $H_{1'}^{mv}$ is exactly as in (15), only U^i is a quadratic utility function and hence i is interested only in the mean and variance of R_{t+1} and R_{t+1}^d .

However, the following proposition shows that if there exists a rule that mean-dominates a long position, then it will also variance dominate it and hence the case LeRoy describes can never occur.

Proposition II: If market returns R_{t+1} are mean dominated by the distribution of a rules' returns R_{t+1}^d but are positive, then the market returns will also be variance-dominated.

Proof:

Notice that we are interested in unconditional distributions²⁰, i.e. in the case when I_t is unknown and hence $d(\mathbf{P}_t)$ is a random variable. This is implied by the notation since $E(R_{t+1}^d) = E(E_t(R_{t+1}^d)/I_t)$.

Ex hypothesi, $E(R_{t+1}^d) > E(R_{t+1}) \geq 0$

$$\Rightarrow E(d_t \cdot R_{t+1}) > E(R_{t+1}) \geq 0$$

$$\Rightarrow [E(d_t \cdot R_{t+1})]^2 > [E(R_{t+1})]^2 \quad (16)$$

And clearly,

$$[d_t]^2 \cdot R_{t+1}^2 \leq R_{t+1}^2 \Rightarrow E([d_t]^2 \cdot R_{t+1}^2) \leq E(R_{t+1}^2) \quad (17)$$

Together the two above inequalities imply (using the fact that $Var(X) = E(X^2) - E(X)^2$):

$$Var(d_t \cdot R_{t+1}) < Var(R_{t+1}) \quad (18)$$

Or equivalently,

$$Var(R_{t+1}^d) < Var(R_{t+1}) \blacksquare$$



Corollary II.1: $H_{1'}^{rn}$ is a sufficient condition for $H_{1'}^{mv}$ when $E(R_{t+1}) \geq 0$, where $H_{1'}^{m-v}$ is the mean-variance version of (14), i.e.:

$$H_{1'}^{mv} : \exists d \text{ s.t. } EU^i(W_t + W_t \theta^* R_{t+1}^d) > EU^i(W_t + W_t \theta^* R_{t+1})$$

for every $EU^i(x)$ which is increasing w.r.t $E(x)$, decreasing w.r.t $Var(x)$

²⁰The result holds a fortiori (and also much more trivially) if I_t is known.

This follows almost trivially from Prop. II. To show it, notice that ex hypothesi:

$$E(R_{t+1}^d) > E(R_{t+1}) \geq 0$$

So by Prop. II:

$$Var(R_{t+1}^d) < Var(R_{t+1})$$

These two inequalities imply also that:

$$\begin{aligned} E(W_t + W_t \theta^* R_{t+1}^d) &> E(W_t + W_t \theta^* R_{t+1}) \\ Var(W_t + W_t \theta^* R_{t+1}^d) &< Var(W_t + W_t \theta^* R_{t+1}) \end{aligned}$$

Since $EU^i(x)$ is a mean-variance function, it directly follows that:

$$H_{1'}^{mv} : \exists d \text{ s.t. } EU^i(W_t + W_t \theta^* R_{t+1}^d) > EU^i(W_t + W_t \theta^* R_{t+1})$$

and hence a weaker form of $H_{1'}^{mv}$ is:

$$H_{1'}^{mv} : E(R_{t+1}^d) > E(R_{t+1}) \geq 0 \blacksquare$$

Hence it follows that the risk-neutral case implies the mean-variance case and that there existed arbitrage opportunities for mean-variance agents in the market for the DJIA. Indeed, note that Proposition II is confirmed in Table II.

5.1.4 The Risk-Averse Case

In this case, H_0 is a great deal more complicated to test. An exception arises when R_{t+1} and R_{t+1}^d are normally distributed. Then:

Proposition III: If R_{t+1} and R_{t+1}^d are normally distributed and $E(R_{t+1}^d) > E(R_{t+1}) \geq 0$, then R_{t+1}^d stochastically dominates R_{t+1} (and hence all risk-averse agents will prefer R_{t+1}^d).

Proof:

In normal environments mean-variance domination and stochastic domination are equivalent (Hanoch and Levy, 1969). This together with Proposition II yield the desired conclusion. ■

If the assumptions of Proposition III are not known to be satisfied, we can reformulate (14) in terms of a stochastic domination criterion of R_{t+1}^d over R_{t+1} . This is shown in Proposition IV below:

Proposition IV: A sufficient condition for (14) is that $E(R_{t+1}) \geq 0$ and $\exists d$ s.t. $M(\gamma) \geq 0 \forall \gamma$ and $M(\gamma) > 0$ for at least one γ , where:

$$M(\gamma) = \int_{-\infty}^{\gamma} R_{t+1} dF(R_{t+1}) - \int_{-\infty}^{\gamma} R_{t+1}^d dG(R_{t+1}^d)$$

Proof

As is well known, the condition of Proposition IV is a sufficient condition for:

$$EU^i(R_{t+1}^d) > EU^i(R_{t+1}) \forall \text{ concave } U^i$$

Notice now that when $E(R_{t+1}) \geq 0$ then $\theta^* \geq 0$ and so $U^i(W_t(1 + \theta^*x))$ is also concave in x , since:

$$\frac{\partial}{\partial x} U^i(W_t(1 + \theta^*x)) = W_t \theta^* U' \geq 0$$

$$\frac{\partial^2}{\partial x^2} U^i(W_t(1 + \theta^*x)) = (W_t \theta^*)^2 U'' < 0$$

Therefore it must also be that

$$EU^i(W_t(1 + \theta^* R_{t+1}^d)) > EU^i(W_t(1 + \theta^* R_{t+1})) \blacksquare$$

Hence, (14) can be replaced with:

$$H_1^{ra} : \exists d \text{ s.t. } \int_{-\infty}^{\gamma} R_{t+1} dF(R_{t+1}) - \int_{-\infty}^{\gamma} R_{t+1}^d dG(R_{t+1}^d) \geq 0 \forall \gamma \quad (19)$$

and the inequality is strict for at least one γ

Whilst a formal statistical test of (16) is feasible, it is incredibly cumbersome computationally (see Tolley and Pope, 1988) especially when, as here, there are many observations on R_{t+1} and R_{t+1}^d . Here, we offer a casual evaluation of whether H_0 can be rejected by inspecting a plot of the sample version of $\hat{M}(\gamma)$ for the optimal moving average returns R_{t+1}^d .

Insert Figure 3 Here

Observing figure 3, we notice that for small γ , $M(\gamma) < 0$. This indicates that the minimum returns from the optimal trading rule resulted in smaller returns than the long position. Hence, for example, an agent with a minimax utility function would prefer **not to use** the trading rule. Therefore, it is unlikely that H_0^a can be rejected and thus we are unable to show that the use of trading rules or conditioning on past prices is utility increasing for all risk-averse agents.

5.2 Efficiency with Transaction Costs

So far we have shown that without transaction costs, there existed an arbitrage opportunity for agents in the DJIA index who had mean-variance utility. We now turn to see how the inclusion of transaction costs affect these results. First of all, transaction costs will alter the analysts' learning problem; hence, we replace (8) with:

$$\begin{aligned}
 (\hat{n}^*, \hat{\lambda}^*) &= \arg \max_{n \in \mathbb{N}, \lambda \in \Lambda} \{ \sum_{i=t-m}^{t-1} MA(\mathbf{P}_i, n, \lambda) R_{t+1} \\
 &\quad - c |MA(\mathbf{P}_t, n, \lambda) - MA(\mathbf{P}_{t-1}, n, \lambda)| \} \quad (20)
 \end{aligned}$$

c are proportional transaction costs

And derive $\left\{ d(\hat{n}_t^*, \hat{\lambda}_t^*) \right\}_{t=1}^{6157}$ and R_{t+1}^d for various levels of c . Our ob-

jective will be to determine the level of transaction costs²¹ c for which H_0 can be rejected in favour of H_1^{rn} at the 95% confidence level²².

In Table III below we have, amongst other things, tabulated the returns from a rule used by the analyst who solves (17). The level of costs at which H_0 can be rejected under the assumption that $R_{t+1}, R_{t+1}^{d*} \sim N$ *i.i.d.* is represented by the line dividing Table III. Notice that this table incorporates the special case $c = 0$, as described in Table II.

Rules	Mean Return	St. Dev.	$\Pr(R_t^d \leq R_t)$	$\Pi_{t=1}^{6157}(1 + R_t^d)$
Market R_t	0.0002334	0.008459	-	2.378
Opt. TTR				
$c=0$	0.000801	0.008335	8.887e-05	110.7
$c=0.0001$	0.0007304	0.008328	0.0005109	71.41
$c=0.0002$	0.0006911	0.008321	0.001238	55.88
$c=0.0003$	0.0006196	0.008318	0.00533	35.63
$c=0.0004$	0.0005508	0.008311	0.01788	22.99
$c=0.0005$	0.0004893	0.008305	0.0452	15.44
$c=0.0006$	0.0004707	0.00829	0.058	13.67
$c=0.0007$	0.0003741	0.00828	0.1755	7.101
$c=0.0008$	0.0003099	0.008273	0.3061	4.456
$c=0.0009$	0.0002763	0.008205	0.3876	3.453
$c=0.001$	0.0002191	0.008215	0.5379	2.13

²¹Note that as defined, the cost of switching from a long to a short position and vice versa is $2c$.

²²It is important to note that Proposition II can be extended to the case of transaction costs if these are small enough. The same is not true for Proposition I if transaction costs are proportional. For a more extensive discussion of technical analysis with transaction costs see Skouras, 1997.

Table III: *The first column indicates which level of costs is under consideration. The next two columns indicate the empirical mean and the standard deviation of the rules' returns (note that Prop. II is confirmed). The fourth column shows the probability (under the assumption of normal distributions) that the mean returns from a specific rule were smaller or equal to those of a non market timer. The final column shows the cumulative returns from each strategy during the whole time period.*

The table indicates that at the 5% level of significance, LIA will be accepted for $c \geq 0.06\%$. The mean return of the optimal rule remains larger for $c \leq 0.09\%$ (but not for the usual margin of confidence). Whilst these levels of c are probably high enough to guarantee that in today's cost conditions LIA might be violated²³, costs were certainly larger at the beginning of the sample we have considered. How large the decrease in transaction costs has been and how it has affected different types of investors is a question which is beyond the scope of this paper, so we do not attempt to answer it. We must add the warning that the time-series used is not adjusted for dividends and hence our results are likely to be biased against LIA.

6 Conclusions

This paper has been organised around the objective of developing definitions and assumptions which would allow technical analysis to be approached in a utility maximisation framework. Its starting point is the illustration that technical analysis is consistent with expected utility maximisation in a typical investment problem when preferences are risk-neutral.

Viewing technical analysis as the decision problem of an agent learning to maximise his expected utility, formalises the notion of 'optimal

²³An investor with access to a discount broker, e.g. via email, can purchase 1000 shares of a company listed on the NYSE for \$14.95.

technical analysis' hitherto loosely alluded to in various papers. The formalisation is revealing because it clearly illustrates that a rule can only be optimal with respect to a specific decision problem and hence a specific class of rules, a level of transaction costs, a position in the market and most importantly a utility function. This indicates that a generally optimal technical analysis is a chimera and that when rules are chosen, it is useful to base this choice on an explicit criterion based on a decision problem.

In a more empirical vein, we show that using learned rules can lead to inferences which are more powerful than those based on arbitrary rules as well as subject to fewer data-mining problems. Consequently, we suggest that model specification tests based on rule returns as pioneered by Brock *et al.* (1992) should be augmented by use of 'artificial' technical analysts in the spirit of Sargent (1993).

Finally, we have tried to investigate the relationship between trading rule returns and market efficiency. This investigation has limited itself to developing and applying a way of empirically rejecting LIA, the hypothesis that past prices do not affect investment decisions. The conclusion drawn from our empirical test is that if the DJIA is the only risky asset in an economy, the hypothesis can be rejected for agents with mean-variance utility facing low enough transaction costs; however, the same is not true for all risk-averse agents. We interpret the magnitude of transaction costs for which this hypothesis is rejected as a measure of market inefficiency.

Natural extensions of this work lie mainly in empirical applications of the developed framework. Firstly, further development of the idea that technical analysis is an effective form of prediction for certain types of loss functions is warranted, since it is likely that this could lead to the development of useful classes of trading rules. Secondly, a more detailed application of the 'optimal' rule to model specification tests could yield significant insights as to the nature of financial time-series. Finally, it is quite easy to extend the framework so as to allow the technical analyst to choose rules conditional on variables other than past prices. This indicates that, in principle, even fundamental analysis is not beyond the

scope of this 'theory of technical analysis'.

References

- [1] Allen F., Karjalainen R., 1996, 'Using Genetic Algorithms to find Technical Trading Rules', working paper, University of Pennsylvania.
- [2] Allen P. M., Phang H. K., 1994, Managing Uncertainty in Complex Systems: Financial Markets, in *Evolutionary Economics and Chaos theory*, in L. Leyesdorff, P. Van den Besselaar, (eds) 1994, *Evolutionary Economics and Chaos Theory*, Pinter, London.
- [3] Allen P. M., Phang K., 1993, Evolution, Creativity and Intelligence in Complex Systems, in H. Haken and A. Mikhailov (eds.), *Interdisciplinary approaches to Nonlinear Complex Systems*, Springer-Verlag, Berlin, 1993.
- [4] Allen H., Taylor M. P., (1990), "Charts, Noise and Fundamentals in the London Foreign Exchange Market", *Economic Journal*, 100, pp.49-59
- [5] Allen H., Taylor M. P. (1989), The Use of Technical Analysis in the Foreign Exchange market., *Journal of International Money and Finance* 11: 304-314.
- [6] Arthur W. B., 1992, 'On learning and adaptation in the economy', Santa Fe Institute wp 92-07-038.
- [7] Arthur W. B., Holland J. H., LeBaron B., Palmer R., Tayler P., 1996, Asset Pricing Under Endogenous expectations in an Artificial Stock Market, Santa Fe Institute working paper.
- [8] Black, F., 1986, 'Noise', *Journal of Finance* 41, 529-43.
- [9] Blume L., Easley D., O'Hara M., 1994, Market Statistics and Technical Analysis: The role of Volume, *Journal of Finance*, Vol. XIIX, No. 1.
- [10] Brock W., Lakonishok J., LeBaron B., 1991, Simple Technical Trading Rules and the Stochastic properties of Stock Returns, Santa Fe Institute working paper.

- [11] Brock W., Lakonishok J., LeBaron B., Simple Technical Trading Rules and the stochastic properties of Stock Returns, *Journal of Finance*, 47(5): 1731-1764.
- [12] Chiang T. F., 1992, Technical Trading Rules Based on Classifier Systems: A New Approach to Learn from Experience, Unpublished Dissertation UCLA.
- [13] Conroy R., Harris R., 1987, Consensus Forecasts of Corporate Earnings: Analysts' Forecasts and Time Series Methods." *Management Science* 33 (June 1987):725-738.
- [14] Christofferson P. F., Diebold F. X., (1995), 'Further results on Forecasting and Model Selection Under Asymmetric Loss', Working Paper.
- [15] Cumby R. E., D. M. Modest, 1987, 'Testing for Market Timing Ability: A framework for forecast evaluation', *Journal of Financial Economics* 19, pp169-189.
- X [16] Curcio R., Goodhart C., 1991, Chartism: A Controlled Experiment., Discussion Paper #124, Financial Markets Discussion Group Series, LSE.
- X [17] Fama E., 1991, Efficient Capital Markets: II, *Journal of Finance*, XLVL, 5, pp.1575-1617.
- X [18] Goldberg M. D., Schulmeister S., (1989), "Technical Analysis and Stock Market Efficiency", Discussion Papers...
- [19] Hanoch G., Levy H., 1969, "The efficiency analysis of choices involving risk", *Review of Economic Studies*, 36, 335-346.
- [20] Jensen M., 1978., 'Some anomalous evidence regarding market efficiency', *Journal of Financial Economics* 6(2-3), 107-147.
- [21] Kandel S., Stambaugh R. F., 1996, 'On the predictability of stock returns: An asset-allocation perspective', *Journal of Finance*, LI, 2.
- [22] Kaufman, 1978, *Commodity Trading Systems and Methods*, John Wiley & Sons.
- [23] Lakonishok J., Vishny R. W., Shleifer A., 1993, 'Contrarian Investment, Extrapolation and Risk', NBER, wp 4360.

- [24] LeBaron B, 1991, 'Technical Trading rules and Regime shifts in Foreign Exchange', SFI WP 91-10-044.
- [25] LeBaron B, 1992, 'Do Moving Average Trading Rule Results Imply Non-linearities in the Foreign Exchange Markets?', U. Wisconsin w.p..
- [26] Leitch G. and J.E. Tanner, 1991, 'Economic Forecast Evaluation: Profits Versus the Conventional Error Measures', AER 81(3), pp. 580-590.
- [27] Levich R., and L. Thomas, 1993, 'The significance of technical-trading rules profits in the foreign exchange market: A bootstrap approach', Journal of International Money and Finance, 56, 269-290.
- [28] Markowitz, 1987, Mean Variance Analysis in Portfolio Choice and Capital Markets, McMillan.
- [29] Malkiel B. G., 1992, 'Efficient Market Hypothesis', in Eatwell J., Milgate M. and P. Newman (eds.), The New Palgrave Dictionary of Banking and Finance, McMillan.
- [30] Malkiel B. G., 1996, A RandomWalk Down Wall Street, , W. W. Norton & Company. Inc.
- [31] Murphy J. J., 1986, Technical Analysis of the Futures Markets, New York: New York Institute of Finance.
- [32] Neely C., Weller P., Dittmar R., 1996, 'Is Technical Analysis in the Foreign Exchange Market Profitable? A Genetic Programming Approach', CEPR d.p. #1480.
- [33] Neftci S. N., 1991, 'Naive Trading Rules in Financial Markets and Wiener-Kolmogorov Prediction Theory: A Study of 'Technical Analysis' ', Journal of Business, 64(4).
- [34] Olsen R. B., M. M. Dacorogna, U. A. Müller, O. V. Pictet, 1992, 'Going Back to the Basics -Rethinking Market Efficiency', O&A Research Group discussion paper.
- [35] Pau L. F., 1991, Technical analysis for portfolio trading by symmetric pattern recognition, Journal of Economics Dynamics and Control, Journal of Economic Dynamics and Control 15, 715-730.

- [36] Pictet O. V., Dacorogna M. M., Dave R. D., Chopard B., Schirru R., Tomassini M., 1996, 'Genetic Algorithms with collective sharing for robust optimization in financial applications', *Neural Network World*, 5(4), pp.573-587.
- [37] Ross, S.A., 1987, 'the Interrelations of Finance and Economics; Theoretical Perspectives', *AER*, 77(2), 29-34.
- [38] Sargent T. J., 1993, *Bounded Rationality in Macroeconomics*, Oxford: Clarendon, Oxford University Press.
- [39] Satchell S., Timmerman A., 1995, "An Assessment of the Economic Value of Nonlinear Exchange Rate Forecasts", WP UCSD, March 1995.
- [40] Skouras S., 1997, *A Theory of Technical Analysis (Unabridged)*, Mimeo., European University Institute.
- [41] Taylor S. J., 1994, 'Trading Futures using a channel rule: A study of the predictive power of technical analysis with currency examples', *Journal of Futures Markets*, 14(2), pp.215-235.
- [42] Teweles R.J., Harlow C.V., Stone H.L., 1977, *The Commodity Futures Game*, McGraw-Hill.
- [43] Treynor J. L., Ferguson R., 1985, In *Defence of Technical Analysis*, *Journal of Finance*, XI, 3.
- [44] Tolley H. D., Pope R. D., 1988, *Testing for Stochastic Dominance*, *Journal of Agricultural Economics*.
- [45] White H, 1992, *Artificial Neural Networks: Approximation and Learning Theory*, Blackwell.

Rule	N(buy)	N(sell)	Buy	Sell	Buy>0	Sell>0	buy-sell
Long $\forall t$	6157	0	0.00023	-	-	-	-
Brock et al (n_t, λ_t) (50,0)	3468	2636	0.00036 (0.90076)	-0.0004 (-1.16108)	0.5167	0.4879	0.00041 (1.78607)
(50,0.01)	2782	1985	0.00053 (1.64014)	0.00003 (-0.70959)	0.5230	0.4861	0.00049 (1.89872)
(150,0)	3581	2424	0.00037 (0.94029)	-0.00012 (-1.49333)	0.5205	0.4777	0.00049 (2.11283)
(150,0.01)	3292	2147	0.00035 (0.80174)	-0.00018 (-1.67583)	0.5216	0.4742	0.00052 (2.13824)
(200,0)	3704	2251	0.00037 (0.92753)	-0.00016 (-1.64056)	0.5173	0.4780	0.00053 (2.23379)
(200,0.01)	3469	2049	0.00038 (0.96907)	-0.00018 (-1.66579)	0.5189	0.4763	0.00056 (2.26328)
Average			0.00037	-0.00011			0.00048
Opt. TTR, $\{d(n_t^*, \lambda_t^*)\}_{t=1}^{6157}$	3313	2650	0.00095 (3.95033)	-0.00067 (-4.57949)	0.5337	0.4675	0.00162 (7.34848)

Table I. The first row of this table indicates which rule is being used. The rules in parentheses represent members of the moving average class, as they constitute specifications of pairs of (n_t, λ_t) . The second and third rows indicate the number of buy and sell signals generated. The fourth and fifth indicate the mean return on days on which buy and sell signals have occurred. The next two columns report the proportion of days in which buy or sell signals were observed in which returns were greater than zero. Finally, the last column reports the difference between buy and sell signals. The numbers in parentheses report results of t -tests testing whether the numbers above them are different to zero²⁴ For the exact tests, see Brock et al. (1992)

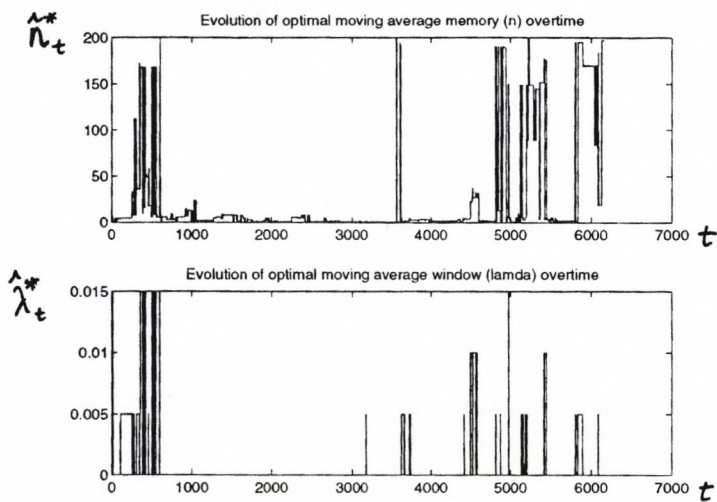


Figure 1: Evolution of each optimal parameter n_t^* and λ_t^* respectively, during $t \in [T^S, T^F]$.




Figure 2: Mean Returns of each rule (n, λ) . The mean is taken over $t \in [T_s, T_f]$.




Figure 3: This is $\widehat{M}(\gamma)$ the sample version of $M(\gamma)$



EUI WORKING PAPERS

EUI Working Papers are published and distributed by the
European University Institute, Florence

Copies can be obtained free of charge
– depending on the availability of stocks – from:

The Publications Officer
European University Institute
Badia Fiesolana
I-50016 San Domenico di Fiesole (FI)
Italy

Please use order form overleaf

Publications of the European University Institute

To The Publications Officer
 European University Institute
 Badia Fiesolana
 I-50016 San Domenico di Fiesole (FI) – Italy
 Telefax No: +39/55/4685 636
 e-mail: publish@datacomm.iue.it
 <http://www.iue.it>

From Name
 Address.

- ☐ Please send me a complete list of EUI Working Papers
- ☐ Please send me a complete list of EUI book publications
- ☐ Please send me the EUI brochure Academic Year 1998/99

Please send me the following EUI Working Paper(s):

No, Author
Title:
No, Author
Title:
No, Author
Title:
No, Author
Title:

Date

Signature



Working Papers of the Department of Economics
Published since 1996

ECO No. 96/1

Ana Rute CARDOSO
 Earnings Inequality in Portugal: High
 and Rising?

ECO No. 96/2

Ana Rute CARDOSO
 Workers or Employers: Who is Shaping
 Wage Inequality?

ECO No. 96/3

David F. HENDRY/Grayham E. MIZON
 The Influence of A.W.H. Phillips on
 Econometrics

ECO No. 96/4

Andrzej BANIAK
 The Multimarket Labour-Managed Firm
 and the Effects of Devaluation

ECO No. 96/5

Luca ANDERLINI/Hamid
 SABOURIAN
 The Evolution of Algorithmic Learning:
 A Global Stability Result

ECO No. 96/6

James DOW
 Arbitrage, Hedging, and Financial
 Innovation

ECO No. 96/7

Marion KOHLER
 Coalitions in International Monetary
 Policy Games

ECO No. 96/8

John MICKLEWRIGHT/ Gyula NAGY
 A Follow-Up Survey of Unemployment
 Insurance Exhausters in Hungary

ECO No. 96/9

Alastair McAULEY/John
 MICKLEWRIGHT/Aline COUDOUEL
 Transfers and Exchange Between
 Households in Central Asia

ECO No. 96/10

Christian BELZIL/Xuelin ZHANG
 Young Children and the Search Costs of
 Unemployed Females

ECO No. 96/11

Christian BELZIL
 Contiguous Duration Dependence and
 Nonstationarity in Job Search: Some
 Reduced-Form Estimates

ECO No. 96/12

Ramon MARIMON
 Learning from Learning in Economics

ECO No. 96/13

Luisa ZANFORLIN
 Technological Diffusion, Learning and
 Economic Performance: An Empirical
 Investigation on an Extended Set of
 Countries

ECO No. 96/14

Humberto LÓPEZ/Eva ORTEGA/Angel
 UBIDE
 Explaining the Dynamics of Spanish
 Unemployment

ECO No. 96/15

Spyros VASSILAKIS
 Accelerating New Product Development
 by Overcoming Complexity Constraints

ECO No. 96/16

Andrew LEWIS
 On Technological Differences in
 Oligopolistic Industries

ECO No. 96/17

Christian BELZIL
 Employment Reallocation, Wages and
 the Allocation of Workers Between
 Expanding and Declining Firms

ECO No. 96/18

Christian BELZIL/Xuelin ZHANG
 Unemployment, Search and the Gender
 Wage Gap: A Structural Model

ECO No. 96/19

Christian BELZIL
 The Dynamics of Female Time Allocation
 upon a First Birth

ECO No. 96/20

Hans-Theo NORMANN
 Endogenous Timing in a Duopoly Model
 with Incomplete Information

*out of print

ECO No. 96/21

Ramon MARIMON/Fabrizio ZILIBOTTI
'Actual' Versus 'Virtual' Employment in
Europe: Is Spain Different?

ECO No. 96/22

Chiara MONFARDINI
Estimating Stochastic Volatility Models
Through Indirect Inference

ECO No. 96/23

Luisa ZANFORLIN
Technological Diffusion, Learning and
Growth: An Empirical Investigation of a
Set of Developing Countries

ECO No. 96/24

Luisa ZANFORLIN
Technological Assimilation, Trade
Patterns and Growth: An Empirical
Investigation of a Set of Developing
Countries

ECO No. 96/25

Giampiero M.GALLO/Massimiliano
MARCELLINO
In Plato's Cave: Sharpening the Shadows
of Monetary Announcements

ECO No. 96/26

Dimitrios SIDERIS
The Wage-Price Spiral in Greece: An
Application of the LSE Methodology in
Systems of Nonstationary Variables

ECO No. 96/27

Andrei SAVKOV
The Optimal Sequence of Privatization in
Transitional Economies

ECO No. 96/28

Jacob LUNDQUIST/Dorte VERNER
Optimal Allocation of Foreign Debt
Solved by a Multivariate GARCH Model
Applied to Danish Data

ECO No. 96/29

Dorte VERNER
The Brazilian Growth Experience in the
Light of Old and New Growth Theories

ECO No. 96/30

Steffen HÖRNIG/Andrea LOFARO/
Louis PHILIPS
How Much to Collude Without Being
Detected

ECO No. 96/31

Angel J. UBIDE
The International Transmission of Shocks
in an Imperfectly Competitive
International Business Cycle Model

ECO No. 96/32

Humberto LOPEZ/Angel J. UBIDE
Demand, Supply, and Animal Spirits

ECO No. 96/33

Andrea LOFARO
On the Efficiency of Bertrand and
Cournot Competition with Incomplete
Information

ECO No. 96/34

Anindya BANERJEE/David F.
HENDRY/Grayham E. MIZON
The Econometric Analysis of Economic
Policy

ECO No. 96/35

Christian SCHLUTER
On the Non-Stationarity of German
Income Mobility (and Some Observations
on Poverty Dynamics)

ECO No. 96/36

Jian-Ming ZHOU
Proposals for Land Consolidation and
Expansion in Japan

ECO No. 96/37

Susana GARCIA CERVERO
Skill Differentials in the Long and in the
Short Run. A 4-Digit SIC Level U.S.
Manufacturing Study

ECO No. 97/1

Jonathan SIMON
The Expected Value of Lotto when not all
Numbers are Equal

ECO No. 97/2

Bernhard WINKLER
Of Sticks and Carrots: Incentives and the
Maastricht Road to EMU

ECO No. 97/3

James DOW/Rohit RAHI
Informed Trading, Investment, and
Welfare

*out of print

ECO No. 97/4

Sandrine LABORY
 Signalling Aspects of Managers'
 Incentives

ECO No. 97/5

Humberto LÓPEZ/Eva ORTEGA/Angel
 UBIDE
 Dating and Forecasting the Spanish
 Business Cycle

ECO No. 97/6

Yadira GONZÁLEZ de LARA
 Changes in Information and Optimal Debt
 Contracts: The Sea Loan

ECO No. 97/7

Sandrine LABORY
 Organisational Dimensions of Innovation

ECO No. 97/8

Sandrine LABORY
 Firm Structure and Market Structure: A
 Case Study of the Car Industry

ECO No. 97/9

Elena BARDASI/Chiara MONFARDINI
 The Choice of the Working Sector in
 Italy: A Trivariate Probit Analysis

ECO No. 97/10

Bernhard WINKLER
 Coordinating European Monetary Union

ECO No. 97/11

Alessandra PELLONI/Robert
 WALDMANN
 Stability Properties in a Growth Model

ECO No. 97/12

Alessandra PELLONI/Robert
 WALDMANN
 Can Waste Improve Welfare?

ECO No. 97/13

Christian DUSTMANN/Arthur van
 SOEST
 Public and Private Sector Wages of Male
 Workers in Germany

ECO No. 97/14

Søren JOHANSEN
 Mathematical and Statistical Modelling of
 Cointegration

ECO No. 97/15

Tom ENGSTED/Søren JOHANSEN
 Granger's Representation Theorem and
 Multicointegration

ECO No. 97/16

Søren JOHANSEN/
 Ernst SCHAUMBURG
 Likelihood Analysis of Seasonal
 Cointegration

ECO No. 97/17

Maozu LU/Grayham E. MIZON
 Mutual Encompassing and Model
 Equivalence

ECO No. 97/18

Dimitrios SIDERIS
 Multilateral Versus Bilateral Testing for
 Long Run Purchasing Power Parity: A
 Cointegration Analysis for the Greek
 Drachma

ECO No. 97/19

Bruno VERSAEVEL
 Production and Organizational
 Capabilities

ECO No. 97/20

Chiara MONFARDINI
 An Application of Cox's Non-Nested
 Test to Trinomial Logit and Probit
 Models

ECO No. 97/21

James DOW/Rohit RAHI
 Should Speculators be Taxed?

ECO No. 97/22

Kitty STEWART
 Are Intergovernmental Transfers in
 Russia Equalizing?

ECO No. 97/23

Paolo VITALE
 Speculative Noise Trading and
 Manipulation in the Foreign Exchange
 Market

ECO No. 97/24

Günther REHME
 Economic Growth, (Re-)Distributive
 Policies, Capital Mobility and Tax
 Competition in Open Economies

*out of print

ECO No. 97/25

Susana GARCIA CERVERO
A Historical Approach to American Skill
Differentials

ECO No. 97/26

Susana GARCIA CERVERO
Growth, Technology and Inequality: An
Industrial Approach

ECO No. 97/27

Bauke VISSER
Organizational Structure and Performance

ECO No. 97/28

Pompeo DELLA POSTA
Central Bank Independence and Public
Debt Convergence in an Open Economy
Dynamic Game

ECO No. 97/29

Matthias BRUECKNER
Voting and Decisions in the ECB

ECO No. 97/30

Massimiliano MARCELLINO
Temporal Disaggregation, Missing
Observations, Outliers, and Forecasting:
A Unifying Non-Model Based Procedure

ECO No. 97/31

Marion KOHLER
Bloc Formation in International Monetary
Policy Coordination

ECO No. 97/32

Marion KOHLER
Trade Blocs and Currency Blocs: A
Package Deal?

ECO No. 97/33

Lavan MAHADEVA
The Comparative Static Effects of Many
Changes

ECO No. 97/34

Lavan MAHADEVA
Endogenous Growth with a Declining
Rate of Interest

ECO No. 97/35

Spyros VASSILAKIS
Managing Design Complexity to Improve
on Cost, Quality, Variety, and Time-to-
Market Performance Variables

ECO No. 97/36

Spyros SKOURAS
Analysing Technical Analysis