Financial Returns and Efficiency as seen by an Artificial Technical Analyst*

Spyros Skouras
European University Institute
skouras@datacomm.iue.it

First Version: December 1996
This version, March 1998

Abstract

We introduce trading rules which are selected by an artificially intelligent agent who learns from experience - an Artificial Technical Analyst. It is shown that these rules can lead to the recognition of subtle regularities in return processes whilst reducing data-mining problems inherent in simple rules proposed as model evaluation devices. The relationship between the efficiency of financial markets and the efficacy of technical analysis is investigated and it is shown that the Artificial Technical Analyst can be used to provide a quantifiable measure of market efficiency. The measure is applied to the DJIA daily index from 1962 to 1986 and implications for the behaviour of traditional agents are derived.

*The research leading to this paper was supported by a grant from I.K.Y.. Thanks are due to Buz Brock, Dave Cass, Alex Gumbel, Blake LeBaron, Mark Salmon and particularly Ramon Marimon for useful comments and suggestions on previous versions of this paper. Blake LeBaron is also thanked for providing the necessary data.

1 Introduction

In the last few years, increasing evidence that technical trading rules can detect nonlinearities in financial time series has renewed interest in technical analysis (see e.g. Brock, Lakonishok and LeBaron, 1992 (henceforth Brock *et al.*), Neftci (1991), Levich *et al.* (1993) and LeBaron (1998)). Based on this evidence, much research effort has also been devoted to examining whether trading rules can be used to evaluate and create improved time-series and theory driven returns models (e.g. Hudson *et al.* (1996), Kho (1996), Gencay (1996)).

The term technical analysis is used in these empirical studies to refer to the practice of investing according to well-known technical trading rules. However, in certain areas of financial theory (particularly asymmetric information models) technical analysis is defined to be any conditioning of expectations on past prices (e.g. Brown and Jennings (1989), Treynor and Ferguson (1985)). Indeed, noisy asymmetric information models in which rational agents condition on past prices reflecting (a noisy signal of) each others information provide one explanation of why technical analysis is observed. Unfortunately though, theoretical models which lead to conditioning consistent with the *precise form* of observed technical trading rules are currently unavailable: there is as yet no positive model of investment behaviour which leads to actions similar to those of real Technical Analysts.

This paper models Technical Analysts as agents whos' actions are *de facto* consistent with observed technical trading rules. Our terminology therefore will be consistent with that of researchers examining empirical aspects of technical analysis and as such will be more narrow than that of theoretical models. The objective however is not the modelling of Technical Analysts *per se*; rather, it is to use our model of a Technical Analyst to derive a more sophisticated approach to examining the properties of trading rules. It is somewhat surprising that some studies have found a single arbitrarily selected rule to be "effective" over long periods (e.g. Brock *et al.*) given that real Technical Analysts use different rules in different times and in different markets. In order to truly evaluate the effectiveness of technical analysis *as implemented* we need a model of how analysts adapt to the market environment.

We provide such a model by introducing Technical Analysts who are artificially intelligent agents (see e.g. Marimon *et al.* 1990). In Section 2 technical analysis is introduced in the simple case where agents are fully

informed and circumstances in which it may be a rational activity are derived. This is a necessary step so that in Section 3 we can introduce a model of a Technical Analyst who learns from his environment - an Artificial Technical Analyst. This agent chooses amongst technical trading rules and his actions are the outcome of an explicit decision problem which formalises the loose notion of what it means for a rule to be "good" or optimal (examples of applications in which a metric for comparing rules is implicit are Allen F. and Karjalainen 1996, Neely *et al.* 1996, Pictet *et al.* 1996, Taylor 1994, Allen P. and Phang 1994, Chiang 1992, Pau 1991). 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 different types of investors (in terms of risk-aversion and budget constraints) will not usually exist.

A standard application of artificially intelligent agents is to design them so that their actions can reveal interesting aspects of the environment in which they are placed (see Sargent 1993, pp. 152-160). In this vein, we will use our Artificial Technical Analysts to reveal certain regularities in financial data. In particular, in Section 4 they will be used to characterise financial series as in Brock *et al.,* showing that they can provide characterisations which are more convincing and powerful than those previously attainable.

In Section 5 we propose a notion of weak market efficiency which is defined almost exclusively as a time series property on prices and show that empirical tests of this condition can be based on the returns obtained by the Artificial Technical Analyst. This is a step in addressing the relationship between market efficiency and the profitability of technical analysis, an issue that has appeared in some of the theoretical literature (e.g. Brown and Jennings, 1989) but is absent from many recent empirical investigations of technical analysis.

Section 6 closes this paper with a synopsis of its conclusions. The main finding is that the Artificial Technical Analyst can provide evidence corroborating the view many Technical Analysts hold of econometric returns models and market efficiency: that the former are inadequate and that the latter is not always present.

2 Technical analysis with full information

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 description of technical analysis:

Df.1: Technical Analysis is the selection of a function d which maps the information set $\mathbf{I}_t$ at t to a space of investment decisions $\mathbf{\Omega}$.

Ass. 1: The space of investment decisions $\mathbf{\Omega}$ consists of three events $\mathbf{\Omega} \equiv$\{Long, Neutral, Short\}. It is a central feature of technical analysis that the magnitude of these positions are undetermined and hence must be treated as constant.

It is crucial to stress that this definition of technical analysis captures the features of technical analysis *as practiced.* As mentioned in the introduction, it is consistent with what is referred to as technical analysis in the literature on empirical properties of trading rules. It is *not* consistent with the definition of technical analysis used in the literature on financial equilibrium with asymmetric information where any agent who conditions on past prices is often called a technical analyst (e.g. Brown and Jennings 1989, Treynor and Ferguson (1985)). The reason we do not allow a more general definition is because we wish to examine the properties of observed trading rules which at first sight seem very different to the investment behaviour we would expect from utility maximising agents. Technical trading rules as defined here have also been referred to as "market timing" rules for which an equilibrium analysis has been developed by Merton (1981).

The event space on which these conditions are written are usually the realisations of some random variable such as prices, volatility measures, or the volume of trading (e.g. Blume *et al.* 1994) of an asset. Here, we focus our attention on rules which are defined on the realisation of a history of past prices. Restricting the information set of Technical Analysts 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.

Ass. 2: $\mathbf{I}_t = \mathbf{P}_t = \{P_t, P_{t-1}, P_{t-2}, ..., P_{t-N+1}\}$.
Where P_t refers to prices at t; henceforth, we use $E_t(\cdot)$ to refer to $E(\cdot/\mathbf{I}_t)$.

Technical Trading Rules and Rule Classes.

We now impose some structure on the form that the functions $\mathbf{P}_t \rightarrow \mathbf{\Omega}$ take, based on observation of how technical analysis is actually conducted. 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 argue that use of any particular rule is widespread, 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 function 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 $D : \mathbf{P}_t \times \mathbf{x} \rightarrow \mathbf{\Omega}$ where $\mathbf{x}$ is a parameter vector in a parameter space B ($\mathbf{x} \in B \subseteq \mathbb{R}^k, \mathbf{P}_t \in \mathbb{R}_+^N$).

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

$$d_t = D(\mathbf{P}_t, \mathbf{x} = \overline{\mathbf{x}}) : \mathbf{P}_t \rightarrow \mathbf{\Omega}$$

which determines a unique investment position for each history of prices[1].

2.1 Technical Analysis and rationality

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 rational (i.e. expected utility maximising) agents' preferences that guarantee they behave as if they were

[1] Notice that any set of rule classes $\{D_i\}_{i=1}^n$ can be seen as a meta-class itself, where the parameter vector $\mathbf{x}' = \{i, \mathbf{x}\}$ determines a specific technical trading rule. Loosely speaking, a rule class can be thought of as an analogue to a parametric model in econometrics and a meta-class as a semi-parametric model.

Technical Analysts. The purpose of this exercise is to clarify the meaning of "optimal technical analysis" in a full information setting. This concept can then be applied to the more interesting case where optimal behaviour must be learned from experience.

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 expected utility of wealth at the end of the next period by choosing the proportion of wealth θ invested in the risky asset. We will assume $\theta \in [-s, l]$, $(s, l \in \mathbb{R}_+)$ so that borrowing and short-selling are allowed but only to a finite extent determined as a multiple of current wealth. His expectations E_t are formed on the basis of past prices $\mathbf{P}_t$, as dictated by $A2$.

Formally, the problem solved is:

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

or equivalently,

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

the solution to which is obtained at:

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

and is the solution of the first order condition:

$$E_t\{R_{t+1} U'(W_t(1 + \theta R_{t+1}))\} = 0$$

In general therefore, $\theta^* : \mathbf{P}_t \to [-s, l]$ is a function that depends on the joint distribution of the random variables R_{t+1} and $U'(W_t(1+\theta R_{t+1}))$ conditional on $\mathbf{P}_t$ and hence indirectly also on W_t. Rational investment behaviour thus generally differs from technical analysis in that investment behaviour cannot be described by a function consistent with our definition of a trading rule which specifies that θ^* may only take three distinct values and cannot be a function of W_t.

In the special case that investors are risk-neutral however, the first order condition above is inapplicable and instead θ^* takes bang-bang solutions

depending on $sign(E_t\{R_{t+1}\})$. Denoting θ^*_{rn} the solution to the risk-neutral investor's problem and assuming that $\theta^*_{rn} = 0$ when $E_t\{R_{t+1}\} = 0$, then:

$$\theta^*_{rn} : \mathbf{P}_t \rightarrow \{-s, 0, l\}$$

which is compatible with the definition of a technical trading rule.

We have therefore shown that only 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.

A direct corollary of this proposition is that expected utility maximisation and technical analysis are compatible. This allows us to define a technical analyst as an expected utility maximising investor:

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

$$\max_{d \in D} d(\mathbf{P}_t) \cdot E_t(R_{t+1}) \tag{3}$$

where D is the space of all functions with domain $\mathbb{R}_+^{\dim(\mathbf{P}_t)}$ and image $\{-s, 0, l\}$.

We will let $R^d_{t+1} \equiv d(\mathbf{P}_t) \cdot R_{t+1}$ denote the returns accruing to an analyst who uses a rule d. Clearly, different trading rules lead to different expected returns.

3 Artificial Technical Analysts

Having specified what is meant by optimal technical analysis in the case of full information, 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 at t. This is a similar amount of information to that possessed by econometricians and hence a technical analyst with an effective mode of learning his optimal actions in this environment is an artificial intelligent agent in the sense of Sargent (1993) or Marimon *et al.* (1990). The term "effective" is not well defined in these circumstances due to the lack of an accepted model for how learning should be conducted. Models for learning can be justified as "reasonable" in a context-dependent way (i.e. depending on the interaction of available information and the decision in

which it will be used) but necessarily retain some *ad hockery*. In this section, we will propose a "reasonable" model for how a Technical Analyst might try and learn his optimal actions. The term Artificial Technical Analyst will refer to a Technical Analyst operating under conditions of imperfect knowledge of his probabilistic environment but equipped with an explicit "reasonable" mechanism for learning optimal actions.

3.1 Parametrising Analysts' learning

Typically, the learning technique of an artificial agent is similar to that of an econometrician: the agent attempts to learn the solution to (3) by selecting a predictor $\hat{R}_{t+1}$ for $E_t(R_{t+1})$ from some parametric model specification (e.g. GARCH-M). The predictor selection is made according to some statistical fitness criterion, such as least squares or quasi-maximum likelihood. The artificial adaptive agent then chooses the optimal action d^* which solves (3) where $E_t(R_{t+1})$ is replaced by $\hat{R}_{t+1}$.

Whilst for some applications this may be a useful approach, we are forced to depart from this methodology somewhat due to the fact that there is significant empirical evidence that statistical fitness criteria can be misleading when applied to decision problems such as that of the Technical Analyst. For example, 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. Taylor (1994) finds that trading based on a channel trading rule outperforms a trading rule based on ARIMA forecasts chosen to minimise in-sample least squares because the former is able to predict sign changes more effectively than the latter[2]. More generally, Leitch and Tanner (1991) show that standard measures of predictor performance are bad guides for the ability of a predictor to discern sign changes of the underlying variable[3].

[2] In some circumstances, the technical analyst is interested in the sign of R_{t+1} more than in its magnitude. In particular, the magnitude of R_{t+1} is irrelevant for his decision problem if $sign(R_{t+1})$ is known or if $sign(R_{t+1})$ and $|R_{t+1}|$ are independent. Hence, a prediction which is formulated to take into account the purpose for which it will be used is likely to be accurate in terms of a sign-based metric. Satchell and Timmerman (1995) show that, without severe restrictions on the underlying series, least square metrics are not directly related to sign-based metrics.

[3] A number of studies of technical trading implicitly or explicitly assume away the possibility that there exists a non-monotonic relationship between the accuracy of a prediction in terms of a metric based on least squares and a metric based on the profit maximisation.

These empirical considerations suggest that any reasonable model of analysts' learning must take his loss function into account. One way to achieve this would be to create Bayesian Artificial Technical Analysts but this would require specification of priors on $F(R_{t+1}/\mathbf{P}_t)$, which might be very difficult. Instead, we confine the Artificial Technical Analyst to a frequentist perspective and calculate an estimator $\hat{d}^*$ giving in-sample optimal solutions to (3) from a specified function space D; this is then used as an estimate for d^*. This type of inference is commonly employed in decision theoretic applications where learning is not focused on determining the underlying stochastic environment, but in determining an action which is an optimal decision for an agent with very little knowledge of this environment.

The main ingredients for creating an Artificial Technical Analyst who learns according to the decision theoretic approach are an in-sample analogue to (3) and an appropriate way of restricting the functional space D. The simplest in-sample analogue to (3) is:

$$\max_{\mathbf{x}\in B} \Sigma_{i=t-m}^{t-1} D(\mathbf{P}_i, \mathbf{x}) \cdot R_{i+1} \tag{4}$$

The motivation for this is that if $\frac{1}{m}\Sigma_{i=t-m}^{t-1} D(\mathbf{P}_i, \mathbf{x})R_{i+1}$ converges uniformly[4] to $E\left\{D(\mathbf{P}_t, \mathbf{x}) \cdot R_{t+1}\right\}$ as $m \to \infty$, it is also the case that the argument maximising the former expression converges to the maximiser of the latter. In other words, we are after a consistent estimator of the solution to (3).

Next we choose D so as to impose some restrictions on the solution to (4) that allow regularities of the in-sample period to be captured. Having no theory to guide us on how to make this choice we will use empirically observed rule classes D^e in the hope that the reason Technical Analysts use them is because they are useful in solving problems similar to (3). Ideally, we would like to have a "specification test" to check whether D^e contains the optimal solution to (3) but it is doubtful if a process for achieving this exists.

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), p.58 fn., *"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 such studies must be treated with caution.

[4]Conditions for this are provided in Skouras (1998).

These considerations lead us to define an Artificial Technical Analyst as follows:

Df. 6: An Artificial Technical Analyst is an agent who solves:

$$\max_{\mathbf{x}\in B} \Sigma_{i=t-m}^{t-1} D^e(\mathbf{P}_i, \mathbf{x}) R_{i+1} \tag{5}$$

where D^e is some empirically observed D.

Turn now to an example illustrating the mechanics of this agent that will be useful in subsequent sections.

3.1.1 Example: An Artificial Technical Analyst Learns the Optimal Moving Average 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 an Artificial Technical Analyst might operate by using it as a specification for D^e. Let us begin with a definition[5] of this class:

Df. 7: The Moving Average rule class $\mathbf{MA}(\mathbf{P}_t, \mathbf{x}_t)$ is a trading rule class s.t. :

$$\begin{aligned}
MA(\mathbf{P}_t, \mathbf{x}) &= \left[\begin{array}{c} -s \text{ if } P_t < (1-\lambda)\frac{\Sigma_{i=0}^{n} P_{t-i}}{n+1} \\ 0 \text{ if } (1-\lambda)\frac{\Sigma_{i=0}^{n} P_{t-i}}{n+1} \leq P_t \leq (1+\lambda)\frac{\Sigma_{i=0}^{n} P_{t-i}}{n+1} \\ l \text{ if } P_t > (1+\lambda)\frac{\Sigma_{i=0}^{n} P_{t-i}}{n+1} \end{array}\right] \\
\text{where } \mathbf{P}_t &= [P_t, P_{t-1}, ..., P_{t-N}], \\
\mathbf{x} &= \{n, \lambda\}, \\
\mathbf{X} &= \{\mathbf{N}, \mathbf{\Lambda}\}, \\
\mathbf{N} &= \{1, 2, ...N\}, \text{ this is the "memory" of the MA} \\
\mathbf{\Lambda} &= \{\lambda : \lambda \geq 0\} \text{ this is the "filter" of the MA}
\end{aligned} \tag{6}$$

Now if $D^e = MA(\mathbf{P}_t, \mathbf{x})$, (5) becomes:

$$\max_{n,\lambda} \Sigma_{i=t-m}^{t-1} MA(\mathbf{P}_i, n, \lambda) R_{i+1} \tag{7}$$

[5] As 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).

Let us fix ideas with an example. Suppose an Artificial Technical Analyst solving (7) wanted to invest in the Dow Jones Industrial Average and had m daily observations of $MA(\mathbf{P}_i, n, \lambda)R_{i+1}$ derived from $N+m+1$ observations of P_t. Let $N = 200$ and $m = 250$. Suppose also that $s = l = 1$ (i.e. position size cannot exceed current wealth). What do the solutions to (7) look like? Figure I plots the answer to this question[6] with observations from t=1/6/1962 till 31/12/1986[7] (i.e. 6157 repetitions).

Insert Figure 1 Here

Figure 1 is a plot of a sequence of rule parameters $\left\{(\widehat{n}_t^*, \widehat{\lambda}_t^*)\right\}_{t=1}^{6157}$ which are optimal in a recursive sample of 250 periods. It is difficult to interpret the sharp discontinuities observed in this sequence but they suggest there is additional structure in the series that a more intelligent Artificial Technical Analyst (one with a more sophisticated learning mechanism) might be able to identify. From the sequence of solutions $\left\{(\widehat{n}_t^*, \widehat{\lambda}_t^*)\right\}_{t=1}^{6157}$ optimal rules $\widehat{d}_t^* = MA(\mathbf{P}_t, \widehat{n}_t^*, \widehat{\lambda}_t^*)$ can be directly derived. These rules are optimised in the period before t, and yield out of sample returns[8] which shall be denoted $R_{t+1}^{d^*}$ and used in subsequent sections.

We now turn to applications of the Artificial Technical Analyst.

4 Artificial Technical Analysts and the distribution of returns in financial markets.

Much of the 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 null hypothesis on the distribution of

[6] $\mathbf{\Lambda}$ was discretised to $\mathbf{\Lambda} = \{0, 0.005, 0.01, 0.015, 0.02\}$. This discretisation allowed us to solve (7) by trying all $\dim(\mathbf{N}) \cdot \dim(\mathbf{\Lambda}) = 1000$ points composing the solution space in each of the 6157 recursions. More sophisticated search methods could lead to more intelligent Artificial Technical Analysts but such niceties do not seem necessary when D^e is as narrow as it is in this example. Furthermore it makes conditions under which (7) converges uniformly much weaker (see Skouras (1998)).

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

[8] It may be interesting to note that these rules correspond to what Arthur (1992) terms "temporarily fulfilled expectations" of optimal rules.

returns (e.g. Brock *et al.*, Levich and Thomas (1993), Neftci (1991), etc.). For example, under a null that returns are a random walk, trading rules can do no better than the market mean; if this is contradicted by the data it is evidence that the null can be rejected. A serious criticism levied against this practice arises from the fact that it involves a testing methodology which is not closed. The openness derives from the *ad hoc* choice of rules the returns of which are examined, since these are selected according to non-rigorous and often implicit criteria. We therefore propose restricting our choice of rules to those chosen by Artificial Technical Analysts; in this way we can overcome the problem of arbitrariness and close the methodology for testing hypotheses on return distributions of rule classes[9].

Artificial Technical Analysts should allow us to avoid a grave lacuna involved in the open *ad hoc* approach. This is that "anomalous" results may be coincidental. 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 mean returns across rules in a class is very diverse and hence small samples of rules are unrepresentative[10].

Furthermore, we show in section 4.2 that by using the more sophisticated rules of the Artificial Technical Analyst we can construct more powerful tests of the null hypotheses.

4.1 The variation of returns across rules

In this section we show that small samples of rules are insufficient to generate reliable conclusions about the behaviour of rule returns even within a relatively narrow class. It would be interesting to try and derive returns processes under which this is the case, but this has not yet been achieved. We must therefore rely on empirical evidence to see whether rule returns are correlated closely enough within a class to justify using a few rules as proxies

[9] Of course, a degree of arbitrariness remains in our 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. The arbitrariness involved in the specification of learning schemes may be an additional problem, but overall such choices are generally considered to be robust and are certainly more robust than choices of arbitrary rules.

[10] 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".

for the behaviour of the class as a whole[11].

Figure 2 shows the returns accruing to each rule belonging to the moving average class if it were applied to the data used in Section 3.1.1.

Insert Figure 2 here

Figure 2 illustrates the inadeqaucies of the *ad hoc* approach whereby specific rules are used as proxies for the distribution of expected returns of a whole class. Notice in particular two highly prominent facts evident in this figure.

Firstly, the highest mean return from a rule in this class is **1270** times larger than that of the worst rule. Since the means are taken from samples with more than 6000 observations, it is unlikely that sampling uncertainty can account for these differences. We must conclude that returns accruing to rules within the same class vary very significantly.

Secondly, the expected returns of rules display significant variance even within small areas of the classes' parameter space. 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 $MA(4,0.02)$. This is important because most researchers choose to calculate returns for a few rules sampled evenly from the space of all rules, reflecting the unfounded implicit assumption that rules are "locally" representative.

Taken together these two observations imply that *ad hoc* rules cannot be the basis for convincing tests of specifications of models for returns. Rules must be selected according to an explicit procedure which is justifiable on theoretical grounds. Artificial Technical Analysts provide such a procedure.

[11]That 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. These results can be found in LeBaron (1998).

On the other hand, the conclusions Brock *et al.* draw are valid (if only by coincidence) since the rules they chose generated slightly sub-average returns.

4.2 Artificial Technical Analysts' rules generate above average returns

We have illustrated the fact that the use of a small sample of rules from a class is an unreliable way of deriving conclusions regarding the behaviour of rule returns for the whole class. We cannot know on any *a priori* grounds how large this sample must be in order for the mean returns across sampled rules to be representative of mean returns across the class. However, we now show that "representative" mean returns are in any case likely to be less than those obtained by an Artificial Technical Analyst and hence will be less effective in rejecting hypotheses on returns distributions.

Intuitively speaking, the reason for this is that a technical analyst who at t chooses d^* from D^e, should have learned to make a better-than-average choice of d. Imposing the use of a "representative" rule by an Artificial Technical Analyst would be the analogue in a decision theoretic setting of estimating a parametric model for a time series by choosing the parameters which have 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.

The following table serves to empirically confirm this reasoning. Utilising some of the information in Brock *et al.*[12] (1991, Table V) it shows that 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_t^* derived in Section 3.1.1. That the rules chosen by Brock *et al.* are "average" across the space of all rules can be verified by inspection of their position in Figure 2.

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

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

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 occured. 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 zero. For the exact tests, see Brock et al.* (1992).

The table indicates that all t-ratios are much higher for the optimal rule we have developed. This means we can reject the null hypothesis that the returns of the DJIA are normally, identically and independently distributed[13] with much greater confidence than that offered by Brock *et al.*'s analysis. We expect that the optimal rule will be equally powerful as a specification test for other hypothesised distributions including those considered by Brock *et al.* (AR, GARCH-M, EGARCH). However, we must leave confirmation of this for future research.

It is on the basis of this evidence that we propose the use of Artificial Technical Analysts' rules for model specification tests. This decreases the probability of obtaining misleading results whilst at the same time delivering more powerful conclusions.

5 Market efficiency and technical trading

We often read that "If markets are efficient, then (technical) analysis of past price patterns to predict the future will be useless", (Malkiel, 1992). In this section, we attempt to analyse the relationship between the efficiency of markets and the efficacy of technical analysis, with a view to a formal assessment of this statement.

At present there seems to be little consensus as to what empirical properties an efficient market should display (see LeRoy 1989 and Fama 1991) consequently to the lack of an accepted equilibrium model for financial markets. In the context of a specific equilibrium model, successive refinements on the definition of Fama (1970) have led to Latham's (1989) definition according to which a market is E-efficient with respect to an information set

[13] The 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.

I_t iff revelation of I_t to all market participants would leave both prices and investment decisions unchanged.

We will now propose a definition of efficiency which has the advantage that its testability does not hinge on the assumption of a specific equilibrium model. Our definition is a necessary condition for a market to be E-efficient. It is consistent with the weak requirement implicit in almost all definitions that profitable intertemporal arbitrage should not be possible on the basis of some information set[14] (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, p.1613 fn.)) consider this notion of market efficiency to be non-testable. However we shall show below that under certain assumptions this test can be empirically implemented. For simplicity, the analysis is restricted to the case where past prices are the information set with respect to which we evaluate efficiency (weak-form).

We will refer to the version of the efficient market hypothesis that we have alluded to as the Lack of Intertemporal Arbitrage (LIA) Hypothesis and discuss its implications for technical trading rules. It will become evident that this efficiency notion is formulated so that it is consistent with the idea that if technical analysis "works", then markets must be inefficient. In this sense it formalises the efficiency notion to which many empirical analyses of trading rule returns allude, yet typically leave undefined (e.g. Hudson *et al.* 1996, Taylor 1992).

LIA is a sensible equilibrium notion if it is reasonable to assume that there exist some agents in the market who are involved in solving (1), which we repeat here for convenience:

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

We will say that LIA is confirmed if knowledge of past prices does not affect the optimal actions of any agent i solving (1). A desired property efficiency notions have failed to deliver is a way of quantifying near-efficiency. We will see that LIA allows precisely such a quantification.

[14] An 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'.

Df. 5.1. The Lack of Intertemporal Arbitrage (LIA) Hypothesis holds in a market in which there exist agents who solve (1) and have utility functions U in a set $\mathbf{U}$ **if for all** $U \in \mathbf{U}$

$$\theta_i^* (\mathbf{P}_t, W_t) = \overline{\theta}^* (W_t) \ \forall \ \mathbf{P}_t$$

where

$$\theta^* (\mathbf{P}_t, W_t) \equiv \arg \max_{\theta \in [-s,l]} EU(W_t(1 + \theta R_{t+1}) | \mathbf{P}_t) \tag{8}$$

$$\overline{\theta}^* (W_t) \equiv \arg \max_{\theta \in [-s,l]} EUW_t(1 + \theta R_{t+1})) \tag{9}$$

What this definition implies is that if LIA holds, the true joint density $f_{R_{t+1},\mathbf{P}_t}$ is such that knowledge that $\mathbf{P}_t = \overline{\mathbf{P}}$ in no way affects the investment behaviour of any market participant for all $\overline{\mathbf{P}}$; this is not a sufficient condition for $f_{R_{t+1}} = f_{R_{t+1}} | \mathbf{P}_{t+1}$, i.e. independence of R_{t+1} and $\mathbf{P}_t$ (or in the terminology of Clements and Hendry (1996) that R_{t+1} is not predictable by $\mathbf{P}_t$). For example, suppose $\mathbf{P}_t$ 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 $\mathbf{P}_t$ 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 size of this class. Formally, the degree of efficiency is determined by the size of the space $\mathbf{U}^E \equiv \left\{ U : \theta^* = \overline{\theta}^* \ \forall \ \mathbf{P}_t \right\}$.

We may interpret the size of $\mathbf{U}^E$ as a measure of near-efficiency with respect to $\mathbf{P}_t$: as the "size" of the set $\mathbf{U}^E$ increases, less agents find $\mathbf{P}_t$ useful and hence the market becomes more efficient with respect to $\mathbf{P}_t$. We must be careful in defining what we mean by increases in this setting. Here we will simply say that $\mathbf{U}^B$ is larger than $\mathbf{U}^A$ if $\mathbf{U}^A \subset \mathbf{U}^B$ in which case the proposed measure would imply the reasonable conclusion that market A is less efficient than B. Such comparisons are relevant if there exist two markets which may be treated as seperate on *a priori* grounds or if we wish to compare the efficiency of a single market during different time periods.

To decide whether a market is efficient with respect to a given U we may derive θ^* and $\overline{\theta}^*$ for that U and evaluate the following hypothesies:

$$H_0(\mathbf{LIA}) : \theta^* (\mathbf{P}_t, W_t) = \overline{\theta}^* (W_t) \ \forall \ \mathbf{P}_t \tag{10}$$

versus

$$H_1(\mathbf{Not\ LIA}) : \theta_i^* (\mathbf{P}_t, W_t) \neq \overline{\theta}^* (W_t) \ \text{some} \ \mathbf{P}_t \tag{11}$$

5.1 Technical Trading Rules and LIA

5.1.1 The Artificial Technical Analyst provides a condition on rule returns for testing LIA

It is one of the most important features of LIA that it can be directly tested without any assumptions on how prices are formed. To see this, notice that if the distributions $f_{R_{t+1}}$ and $f_{R_{t+1}}|\mathbf{P}_t$ are known, LIA can be directly evaluated, in some cases even analytically. If the distributions are unknown, one approach for empirically testing (10-11) would be to use an estimated model for the unknown distributions $f_{R_{t+1}}$ and $f_{R_{t+1}}|\mathbf{P}_t$. However, such a model may not be feasible because in practice it may be difficult to forecast R_{t+1} even when it is known to be predictable (for the distinction between predictability and forecastability, see e.g. Clements and Hendry, 1996). Here we propose an alternative testing approach based on the returns obtained by an Artificial Technical Analyst. This approach is based on the fact that a sufficient condition for LIA to be violated is that technical trading is preferred over always being long by some investors. We show this formally below:

Proposition II. Suppose an investor maximises (unconditional) expected utility by investing a fraction of wealth $\overline{\theta}^*$ in a risky asset (see(9)), but can increase his expected utility by investing $\overline{\theta}^*$ according to a trading rule $d(\mathbf{P}_t)$ (i.e. $EU^i(W_t(1+\overline{\theta}^* d(\mathbf{P}_t)R_{t+1})) > EU^i(W_t(1+\overline{\theta}^* R_{t+1}))$). Then LIA does not hold[15], i.e. $\theta_i^*(\mathbf{P}_t, W_t) \neq \overline{\theta}^*(W_t)$ for some $\mathbf{P}_t$.

Proof

The assumption, may be rewritten as:

$$E\{E[U^i(W_t(1+\overline{\theta}^* d(\mathbf{P}_t)R_{t+1}))|\mathbf{P}_t]\} > E\{E[U^i(W_t(1+\overline{\theta}^* R_{t+1}))|\mathbf{P}_t]\}$$

Which implies that $\exists$ a set $\mathbb{Q} \subseteq \mathbb{R}_+^k$ s.t. $\Pr(\mathbf{P}_t \in \mathbb{Q}) > 0$ and $\forall\ \mathbf{P}_t \in \mathbb{Q}$

$$E[U^i(W_t(1+\overline{\theta}^* d(\mathbf{P}_t)R_{t+1}))|\mathbf{P}_t] > E[U^i(W_t(1+\overline{\theta}^* R_{t+1}))|\mathbf{P}_t]$$

Now define:

$$\theta^{**} = \left\{ \begin{array}{c} \overline{\theta}^* d(\mathbf{P}_t) \text{ if } \mathbf{P}_t \in \mathbb{Q} \\ \theta^* \text{ otherwise} \end{array} \right\}$$

[15] Note that expectations are taken with respect to $f_{R_{t+1},\mathbf{P}_t}$ so that $d(\mathbf{P}_t)$ is a random variable unless it is made explicit that the expectation is conditional on $\mathbf{P}_t$.

Then clearly, $\exists\ \theta^{**} \neq \overline{\theta}^*$

$$s.t. E[U^i(W_t(1+\theta^{**}R_{t+1}))|\mathbf{P}_t] \geq E[U^i(W_t(1+\overline{\theta}^*R_{t+1}))|\mathbf{P}_t] \forall \mathbf{P}_t$$

and the inequality is strict if $\mathbf{P}_t \in \mathbb{Q}$ which directly implies (11).

qed■

We conclude that a test for LIA based on trading rules, can be obtained by replacing (10-11) with the stronger but (in some cases) empirically verifiable conditions:

$$H_{0'} : EU^i(W_t(1+\overline{\theta}^*R^d_{t+1})) \leq EU^i(W_t(1+\overline{\theta}^*R_{t+1}))\ \forall d \tag{12}$$

$$H_{1'} : \exists\ d \text{ s.t. } EU^i(W_t(1+\overline{\theta}^*R^d_{t+1})) > EU^i(W_t(1+\overline{\theta}^*R_{t+1})) \tag{13}$$

Whilst rejection of $H_{0'}$ is not a necessary but only a sufficient condition for rejection of H_0, we now show that it turns out that even $H_{0'}$ can be empirically rejected for some important utility function classes.

5.1.2 The Risk-Neutral Case

In this case, all $U \in \mathbf{U}$ are linear and the test (12-13) becomes (assuming $\overline{\theta}^*$ is positive, i.e. $E(R_{t+1}) > 0$):

$$H^{rn}_{0'} : E(R^d_{t+1}) \leq E(R_{t+1})\ \forall\ d \tag{14}$$

$$H^{rn}_{1'} : \exists\ d\ s.t.\ E(R^d_{t+1}) > E(R_{t+1}) \tag{15}$$

Suppose we use the rules $\{d(n^*_t, \lambda^*_t)\}_{t=1}^{6157}$ of the Artificial Technical Analyst as derived in Section 3.1.1 and the corresponding returns $R^{d^*}_{t+1}$ to test $H^{rn}_{0'}$. 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.	$\Pr(R^{d^*}_{t+1} \leq R_{t+1})$
$\mathbf{R}_{t+1}$	0.0002334	0.008459	-
$\mathbf{R}^{d^*}_{t+1}$	0.000801	0.008335	8.887e-5

Table II *Note that the last column was calculated conditional on the (false) assumption that* R_{t+1} *and* $R^{d^*}_{t+1}$ *are normal i.i.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 Quadratic Case

Suppose now that **U** includes quadratic funcions. Is it still the case that LIA is violated? The reason this may not be the case is that although there exists a rule satisfying (15) it involves greater variance than R_{t+1} and hence is not preferred agents with quadratic utility. 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, we formulate the quadratic version of (12-13) which assumes that

$$\mathbf{U}^{MV} = \left\{ U : U(W) = aW^2 + bW + c, a \leq 0, \frac{-b}{2a} \geq W \right\}$$

so for all $U \in \mathbf{U}^{MV}$ it is the case that $EU(x)$ is increasing w.r.t. $E(x)$, and decreasing w.r.t $Var(x)$.

The following proposition shows the somewhat surprising result that if there exists a rule that mean-dominates a long position, then it will also variance dominate it and hence LeRoy's statement is a logical impossibility in rather general circumstances. The proposition is also interesting in its own right, it is crucial for our purposes because it establishes general circumstances in which the quadratic case collapses to the risk-neutral case.

Proposition III: If the trading rule $d_1(\mathbf{P}_t)$ mean dominates trading rule $d_2(\mathbf{P}_t)$ and both lead to positive expected returns

$$E(d_1 R_{t+1}) > E(d_2 R_{t+1}) \geq 0$$

and (a): The second rule is always long, i.e. is the 'Buy and Hold' strategy and long positions are not smaller than short positions

$$\begin{aligned} d_2(\mathbf{P}_t) &= l \text{ all } \mathbf{P}_t \\ l &\geq s \end{aligned}$$

or (b): Trading rules have a binary structure and position sizes are symmetric

$$\begin{aligned} d_1, d_2 &\in \{-s, l\} \\ l &= s \end{aligned}$$

Then the returns from rule 1 have a smaller variance than the returns from rule 2

$$V(d_1 R_{t+1}) < V(d_2 R_{t+1})$$

Proof:
Part (a)
Let $\overline{d}_1 \equiv \frac{1}{l} d_1$
Then the assumption implies $E(d_1 R_{t+1}) > E(l R_{t+1}) \geq 0$

$$E(\overline{d}_1 R_{t+1}) > E(R_{t+1}) \geq 0$$

Notice that $\overline{d}_i = \left\{-\frac{s}{l}, 0, 1\right\}$ and by assumption $l \geq s$ hence $0 > -\frac{s}{l} \geq -1$ so $\left(\overline{d}_i\right)^2 \leq 1$.
Hence $(\overline{d}_i R_{t+1})^2 \leq (R_{t+1})^2$
Using the fact that $V(x) = E(x^2) - E(x)^2$ it follows that:

$$V\left(\overline{d}_1 R_{t+1}\right) < V(R_{t+1})$$

From which it also follows directly that

$$V(d_1 R_{t+1}) < V(R_{t+1})$$

Part (b)
Let $\overline{d}_i \equiv \left(\frac{1}{l}\right) d_i$, $i = 1, 2$
Then the assumption $E(d_1 R_{t+1}) > E(d_2 R_{t+1}) \geq 0$ implies

$$E(\overline{d}_1 R_{t+1}) > E(\overline{d}_2 R_{t+1}) \geq 0$$

Notice that $\overline{d}_i = \{-1, 1\}$ so $\left(\overline{d}_i\right)^2 = 1$.
Hence $(\overline{d}_i R_{t+1})^2 = (R_{t+1})^2$, $i = 1, 2$
Using the fact that $V(x) = E(x^2) - E(x)^2$ it follows that:

$$V\left(\overline{d}_1 R_{t+1}\right) < V\left(\overline{d}_2 R_{t+1}\right)$$

Hence also:

$$V(d_1 R_{t+1}) < V(d_2 R_{t+1})$$

qed■

This Proposition is useful for establishing the two following corollaries, but also because it provides a shortcut to ranking rules by performance in

terms of Sharpe Ratios - an exercise attempted by many researchers and practitioners as an alternative to ranking by mean returns (e.g. LeBaron 1998b, Sullivan, Timmerman and White 1997). In the circumstances indicated however, Sharpe Ratios are inversely related to mean returns and such exercises are often redundant[16].

Corollary III.1: $H_{1'}^{rn}$ is a sufficient condition for the rejection of LIA (i.e. (13) $\Rightarrow$ (15)) when $E(R_{t+1}) \geq 0$, if U^i is a mean-variance utility function and $l \geq s$.

Proof

This follows from Proposition IIIa. To show it, notice that ex hypothesi:

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

So by Proposition IIIa:

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

These two inequalities imply also that:

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

which implies that for all $U \in \mathbf{U}^{MV}$:

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

Hence it follows that as long as $l \geq s$ which is very plausible, the risk-neutral case implies the mean-variance case and therefore mean-variance investors in the DJIA would find knowledge of past prices useful. Indeed, note that Proposition IIIa is empirically confirmed in Table II.

[16]Special cases of this proposition have been previously established by Skouras (1997, 1996) for $l = s = 1$. This narrow circumstance seems to have been independently examined by Acar (1998) who shows that if rules are binary, the variance of a rule is inversely related to its mean (as in Part b). However, his proof is erroneous because as we have shown this holds only when $E(R_t) > 0$ (which he does not assume).

5.1.4 The Risk-Averse Case

In this case **U** is a concave class of functions and H_0 is a great deal more complicated to test. As is usually the case, an exception arises when R_{t+1} and R^d_{t+1} are normally distributed.

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

Proof:

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

If the assumptions of Proposition IV are not known to be satisfied, we can reformulate (12-13) in terms of a stochastic domination criterion of R^d_{t+1} over R_{t+1}. This is shown in Proposition V below:

Proposition V: A sufficient condition for (13) is that $E(R_{t+1}) \geq 0$, $l \geq s$ and $\exists\ d\ s.t\ M(\gamma) \geq 0\ \forall\ \gamma$ and $M(\gamma) > 0$ for at least one γ, where

$$M(\gamma) \equiv \int_{-\infty}^{\gamma} [f_{R_{t+1}}(x) - f_{R^d_{t+1}}(x)]dx$$

and $f_{R_{t+1}}$, $f_{R^d_{t+1}}$ are the marginal densities of $R_{t+1}$and R^d_{t+1} respectively.

Proof

As is well known, the assumption on $M(\gamma)$ is a sufficient condition for:

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

Notice now that when $E(R_{t+1}) \geq 0$ then $\overline{\theta}^*(W_t) \geq 0$ (see (9)) and so $U(W_t(1+\overline{\theta}^* x))$ is also concave in x, since:

$$\begin{aligned} \frac{\partial}{\partial x} U(W_t(1+\overline{\theta}^* x)) &= W_t \overline{\theta}^* U' \geq 0 \\ \frac{\partial^2}{\partial x^2} U(W_t(1+\overline{\theta}^* x)) &= (W_t \overline{\theta}^*)^2 U'' < 0 \end{aligned}$$

Therefore it must also be that

$$EU(W_t(1+\overline{\theta}^* R^d_{t+1})) > EU(W_t(1+\overline{\theta}^* R_{t+1}))$$

qed■

From this proposition we notice that (12-13) can be replaced with:

$$H_{0'}^{ra} : \nexists\ d \text{ s.t. } \int_{-\infty}^{\gamma} [f_{R_{t+1}}(x) - f_{R_{t+1}^d}(x)]dx \geqslant 0\ \forall\ \gamma \tag{16}$$

$$H_{1'}^{ra} \quad : \quad \exists\ d\ s.t. \int_{-\infty}^{\gamma} [f_{R_{t+1}}(x) - f_{R_{t+1}^d}(x)]dx \geqslant 0\ \forall\ \gamma \tag{17}$$

and the inequality is strict for some γ

Whilst a formal statistical test of (16-18) is feasible, it is incredibly cumbersome computationally (see Tolley and Pope, 1988) especially when there are many observations on R_{t+1} and R_{t+1}^d. This obstacle forces us to offer only a casual evaluation of whether H_0 can be rejected. Such an evaluation can be conducted 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, $M(\gamma) < 0$, indicating that the minimum returns from the Artificial Technical Analysts' rule resulted in smaller returns than the minimum market return. This implies that an agent with a utility function which greatly penalizes extremely low returns would prefer **not to use** the trading rule. Hence, even without taking account of sample uncertainty we are unable to reject LIA in the risk-averse case. Taking sample uncertainty into account using a formal statistical procedure cannot reverse this result, since it could only make rejection of the null more difficult. We conclude that there are risk-averse agents whose action cannot be proved to depend on knowledge of past prices using this data and the type of procedure we propose. Clearly however there may be more powerful tests that could lead to a different result.

5.2 Efficiency with Transaction Costs

So far we have shown that without transaction costs, there existed an arbitrage opportunity for investors in the DJIA index who had mean-variance utility. We now consider how the inclusion of transaction costs affects these results. First of all, transaction costs will alter Artificial Technical Analysts'

objective function. To take account of this, we let $s = l = 1$ as before and replace (7) with:

$$\begin{aligned} & (\hat{n}_t(q_{t-1}), \hat{\lambda}_t(q_{t-1})) \qquad (18) \\ = \ & \arg \max_{n \in \mathbf{N}, \lambda \in \mathbf{\Lambda}} \{\Sigma_{i=t-m}^{t-1} MA(\mathbf{P}_i, n, \lambda) R_{t+1} - c \left| MA(\mathbf{P}_i, n, \lambda) - q_{t-1} \right|\} \end{aligned}$$

where c are proportional transaction costs and q_{t-1} is the position at $t-1$, i.e. $q_{t-1} \in \{-1, 0, 1\}$. Beginning with $q_0 = 0$, we set $q_1 = MA\left(\mathbf{P}_t, \hat{n}_t\left(q_0\right), \hat{\lambda}_t\left(q_0\right)\right)$ and iterate so as to obtain the sequence $\{q_t\}$. From this, we next obtain $\left(\hat{n}_t^*, \hat{\lambda}_t^*\right) = (\hat{n}_t(q_{t-1}), \hat{\lambda}_t(q_{t-1}))$ with which we derive $\left\{MA(\mathbf{P}_t, \hat{n}_t^*, \hat{\lambda}_t^*)\right\}_{t=1}^{6157}$ and $R_{t+1}^{d^*}$for various levels of c which are proportional transaction costs. Our objective will be to determine the level of transaction costs[17] c for which LIA holds, i.e. $H_{0'}^{rn}$ cannot be rejected in favour of $H_{1'}^{rn}$ at the 95% confidence level[18]. This exercise is conceptually similar to that of Cooper and Kaplanis (1994) who try to estimate the level of deadweight costs that would explain the home bias in international equity portfolios. Allowing for transaction costs changes our near-equilibrium notion in that we now seek pairs $\{\mathbf{U}^E, c\}$ rather than just $\mathbf{U}^E$ for which LIA holds.

In Table III below we have, in the second column, tabulated the returns from a rule used by the analyst who solves (18). 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

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

[18] It is important to note that Proposition III 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.

special case $c = 0$ reported 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

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. 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 the mean market returns. 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 confidence margin)[19]. These levels of c make it tempting to argue that with todays' cost conditions[20] LIA is violated. However, 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 and which we do not attempt to answer. We must

[19] Note that Proposition IIIa can be extended to the case with transaction costs if these are "small enough". Table III is consistent with the results of Proposition IIIa.

[20] An investor with access to a discount broker, e.g. via email, can purchase 1000 shares of a company listed on the NYSE for a $14.95 fee. However, micro-structure frictions such as bid-ask spreads should also be taken into account.

add the warning that the time-series used is not adjusted for dividends and hence our results are likely to be biased against LIA.

In this section, we have derived conditions on $\{\mathbf{U}^E, c\}$ that ensure that the trading rule returns of an Artificial Technical Analyst are such that we cannot reject our version of the weak form of the efficient market hypothesis. The restrictions derived are not trivial and hence even though our efficiency notion is only a special case of other standard notions and even though we have only tested sufficient conditions for its' rejection, we have found a class of agents defined by their preferences and transaction costs who would optimally condition actions on past prices. Note that there are numerous asymmetric information models which generate equilibria for which LIA is violated (e.g. Hussman 1992, Brown and Jennings 1989, Treynor and Ferguson 1985) so some of the available models may describe the data generating process accurately. The methodology proposed is useful because it unties efficiency from a specific equilibrium notion and allows near-efficiency comparisons across time and markets by comparison of the generality of the conditions on $\{\mathbf{U}^E, c\}$ that would ensure LIA is not rejected. In thus provides a quantifiable measure of near efficiency. Equally importantly, it formalises a sense in which markets can be characterised as inefficient when empirical studies find trading rules to be profitable. It therefore captures an informal notion of efficiency used by many researchers.

6 Conclusions

Empirical investigations of financial series have previously found technical trading rules to be useful research tools. In this paper we develop a more sophisticated form of technical analysis and show empirically that it can be used to obtain powerful characterisations of financial series.

Our first step is to show that technical analysis is consistent with utility maximisation since in a full information environment risk-neutral investors will use technical trading rules. This formalisation is revealing because it clearly illustrates that the optimality of a trading rule can only be judged in the context of 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. A generally optimal technical trading rule is therefore a chimera and hence choices from any class of rules should be derived from explicit criteria based on some decision problem.

Relaxing investors' knowledge of the stochastic behaviour of returns to a realistic level, we develop a simple model for how Technical Analysts might behave in environments where they can learn from experience. An agent using a simple learning model to make choices among technical trading rules is called an Artificial Technical Analyst. This construct is useful because it provides a way to empirically evaluate claims of real-world technical analysts that financial time-series models are inadequate for their decisions and that their profits prove markets are not always efficient.

In this vein, we use the rules which are optimal for the Artificial Technical Analyst in empirical applications and find that in addition to being subject to fewer data-mining problems than arbitrary fixed rules used in previous studies, they also lead to more powerful inferences. Consequently, we suggest that bootstrap based model specification tests based on rule returns as pioneered by Brock *et al.* (1992) should be augmented with artificially intelligent agents in the spirit of Sargent (1993).

Finally, we investigate the relationship between trading rule returns and market efficiency. Our analysis has been based on developing and applying a way of empirically rejecting LIA, the hypothesis that past prices do not affect investment decisions. Whilst this may seem intuitively plausible, it is a new definition of efficiency that has the advantage of being testable independently of an equilibrium model. It relates to previous notions in that it captures the main intuition underlying empirical tests of weak-form efficiency and is a necessary condition for some standard theoretical notions of efficiency such as Latham's (1989) E-efficiency. It also has the advantage that it provides a quantifiable notion of near-efficiency.

Our empirical test uses the DJIA index as a proxy for the market portfolio. For the data considered (daily 1962-1986), LIA can be rejected for agents with mean-variance utility facing low enough transaction costs; however, there do exist other risk-averse agents who cannot be shown to find conditioning on past prices to be useful. We interpret the magnitude of transaction costs and the generality of the preferences for which this hypothesis is rejected as a measure of market efficiency. An interesting experiment which we must leave for future investigations is the comparison of the size of this measure to ones obtained from other financial series. We note that a by-product of our analysis is the observation that under general assumptions the variance of trading rule returns are inversely related to their mean.

There are many natural extensions of this work so we can only make a few indicative suggestions. Firstly, the Artificial Technical Analyst could

be made more intelligent by making his learning more sophisticated and by widening the space of trading rules from which he may choose. Secondly, it would be important to try and find models for returns that are consistent with returns obtained by the Artificial Technical Analyst. Finally, it is quite easy to extend the framework so as to allow for analysts who choose rules conditional on variables other than past prices. This indicates that the Artificial Technical Analyst is capable even of fundamental analysis.

References

[1] Acar E., 1998,'Expected returns of directional forecasters', in Acar E. and Satchell S. (eds), *Advanced Trading Rules*, Butterworth Heinemann.

[2] Allen F., Karjalainen R., 1996, "Using Genetic Algorithms to find Technical Trading Rules", working paper, University of Pennsylvania.

[3] 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.

[4] 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-Verlang, Berlin, 1993.

[5] Allen H, Taylor M. P.,(1990), "Charts, Noise and Fundamentals in the London Foreign Exchange Market", Economic Journal, 100, pp.49-59

[6] 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.

[7] Arthur W. B., 1992, 'On learning and adaptation in the economy', Santa Fe Institute wp 92-07-038.

[8] 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.

[9] Black, F., 1986, 'Noise', Journal of Finance 41, 529-43.

[10] Blume L., Easley D., O'Hara M., 1994, Market Statistics and Technical Analysis: The role of Volume, Journal of Finance, Vol. XIIX, No. 1.

[11] Brock W., Lakonishok J., LeBaron B., 1991, Simple Technical Trading Rules and the Stochastic properties of Stock Returns, Santa Fe Institute working paper.

[12] Brock W., Lakonishok J., LeBaron B., Simple Technical Trading Rules and the stochastic properties of Stock Returns, Journal of Finance, 47(5): 1731-1764.

[13] Brown David P, and Robert H. Jennings, 'On Technical Analysis', Review of Financial Studies, 2(4), 527-55.

[14] Clements M. P., Hendry D. F., 1996, 'Forecasting in Macroeconomics', in Cox D. R., Hinkley D. V. and Barndoff-Nielsen O.E. (eds.), Time Series Models in econometrics, finance and other fields,

[15] Cooper Ian and Evi Kaplanis, 1994, 'Home Bias in Equity Portfolios, Inflation hedging and International Capital Market Equilibrium", Review of Financial Studies, 7, 1, pp. 45-60.

[16] Chiang T. F., 1992, Technical Trading Rules Based on Classifier Systems: A New Approach ro Learn from Experience, Unpublished Dissertation UCLA.

[17] Conroy R., Harris R., 1987, Consensus Forecasts of Corporate Earnings: Analysts' Forecasts and Time Series Methods." Management Science 33 (June 1987):725-738.

[18] Christofferson P. F., Diebold F. X., (1996), 'Further results on Forecasting and Model Selection Under Asymmetric Loss', Journal of Applied Econometrics, 11, 561-571..

[19] Cumby R. E., D. M. Modest, 1987, 'Testing for Market Timing Ability: A framework for forecast evaluation', Journal of Financial Economics 19, pp169-189.

[20] Curcio R., Goodhart C., 1991, Chartism: A Controlled Experiment., Discussion Paper #124, Financial Markets Discussion Group Series, LSE.

[21] Fama E., 1970, 'Efficient Capital MArkets: A review of Theory and Empirical Work', Journal of Finance 25, May, 383-417.

[22] Gencay Ramazan, 1996, 'Non-linear Prediction of Security Returns with Moving Average Rules', Journal of Forecasting, 15, 165-174.

[23] Goldberg M. D., Schulmeister S., (1989), "Technical Analysis and Stock Market Efficiency", Wissenschaftszentrum Berlin fur Sozialforschung, Berlin.

[24] Hanoch G., Levy H., 1969, "The efficiency analysis of choices involving risk", Review of Economic Studies, **36**, 335-346.

[25] Hudson R., Dempsey M., Keasey K., "A note on the weak form efficiency of capital markets: The application of simple technical trading rules to UK stock prices - 1935 to 1994.", Journal of Banking and Finance 20, 1121-1132.

[26] Hussman John, 1992, 'Market Efficiency and Inefficiency in Rational Expectations Equilibria', Journal of Economics Dynamics and Control, 16, 655-680.

[27] Jensen M., 1978., 'Some anomalous evidence regarding market efficiency', Journal of Financial Economics 6(2-3), 107-147.

[28] Kandel S., Stambaugh R. F., 1996, 'On the predictability of stock returns: An asset-allocation perspective', Journal of Finance, LI, 2.

[29] Kaufman, 1978, Commodity Trading Systems and Methods, John Wiley & Sons.

[30] Kho B. C., 1996, Time-varying risk premia, volatility and technical trading rule profits: Evidence from foreign currency futures markets, Journal of Financial Economics 41:249-290.

[31] Lakonishok J, Vishny R. W., Shleifer A, 1993, 'Contrarian Investment, Extrapolation and Risk', NBER, wp 4360.

[32] Latham, Mark, 1986, 'Informational Efficiency and Information Subsets', Journal of Finance, XLI, 1., pp.39-52.

[33] LeBaron B., 1998, 'Technical Trading rules and Regime shifts in Foreign Exchange', in Acar E. and Satchell S. (eds), *Advanced Trading Rules*, Butterworth Heinemann.

[34] LeBaron B., 1998b, 'An evolutionary bootstrap method for selecting dynamic trading strategies', SSRI w.p. 9805.

[35] LeBaron B, 1992, 'Do Moving Average Trading Rule Results Imply Nonlinearities in the Foreign Exchange Markets?', U. Wisconsin w.p..

[36] Leitch G. and J.E. Tanner, 1991, 'Economic Forecast Evaluation: Profits Versus the Conventional Error Measures', AER 81(3), pp. 580-590.

[37] Levich R., and L. Thomas, 1993, 'The significance of technical-trading rulles profits in the foreign exchange market: A bootstrap approach', Journal of International Money and Finance, 56, 269-290.

[38] Markowitz, 1987, Mean Variance Analysis in Portfolio Choice and Capital Markets, McMillan.

[39] Marimon M., McGrattan E., and Sargent T., 'Money as a Medium of Exchange in an Economy with Artificially Intelligent Agents', Journal of Economic Dynamics and Control, 14, 329-74.

[40] 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.

[41] Merton R. C., 1981, 'On market timing and investment performance, I: An equilibrium theory of market forecasts', Journal of Business 54, 363-406.

[42] Murphy J. J., 1986, Technical Analysis of the Futures Markets, New York: New York Institute of Finance.

[43] Neely C., Weller P., Dittmar R., 1996, 'Is Technical Analysis in the Foreign Exchange Market Profitable? A Genetic Programming Approach',CEPR d.p. #1480.

[44] 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).

[45] 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.

[46] 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.

[47] 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.

[48] Ross, S.A., 1987, 'the Interrelations of Finance and Economics; Theoretical Perspectives', AER, 77(2), 29-34.

[49] Sargent T. J., 1993, Bounded Rationality in Macroeconics, Oxford: Claredon, Oxford University Press.

[50] Satchell Steve and Alan Timmerman, 1995, An Assesment of the Economic Value of Non-linear Foreign Exchange Rate Forecasts, Journal of Forecasting, Vol. 14, 477-497.

[51] Skouras S., 1997, 'Analysing Technical Analysis', European University Institute working paper 97-36.

[52] Skouras S., 1996, 'A Theory of Technical Analysis (Unabridged)', Mimeo. Draft 1, European University Institute.

[53] Sullivan Ryan, Timmerman Allan and Halbert White, 1997, 'Data-snooping, technical trading rule performance and the bootstrap', UCSD working paper.

[54] 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.

[55] Taylor S. J., 1992, 'Rewards Available to Currency Futures Speculators: Compensation for Risk or Evidence of Inefficient Pricing', Economic Record, Supplement.

[56] Treynor J. L., Ferguson R., 1985, In Defence of Technical Analysis, Journal of Finance, XL, 3.

[57] Tolley H. D., Pope R. D., 1988, Testing for Stochastic Dominance, Journal of Agricultural Economics.

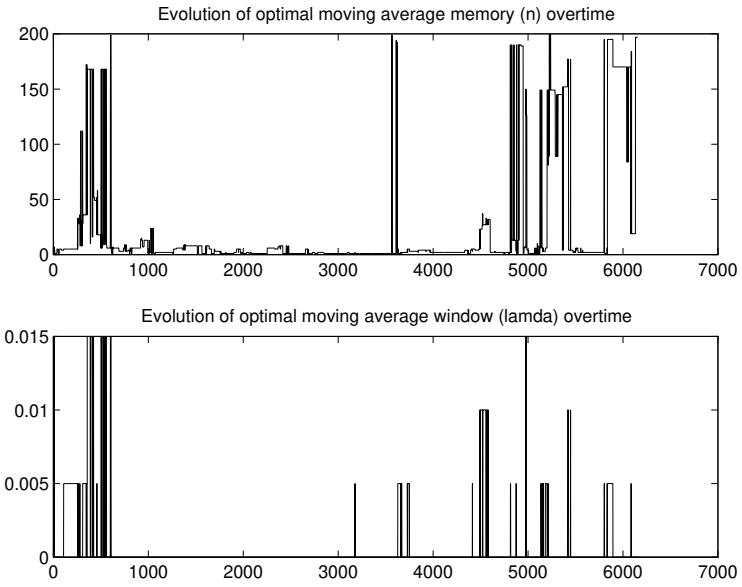

Figure 1: Evolution of each optimal parameter n_t^* and λ_t^* respectively, during $t \in \left[T^S, T^F\right]$.

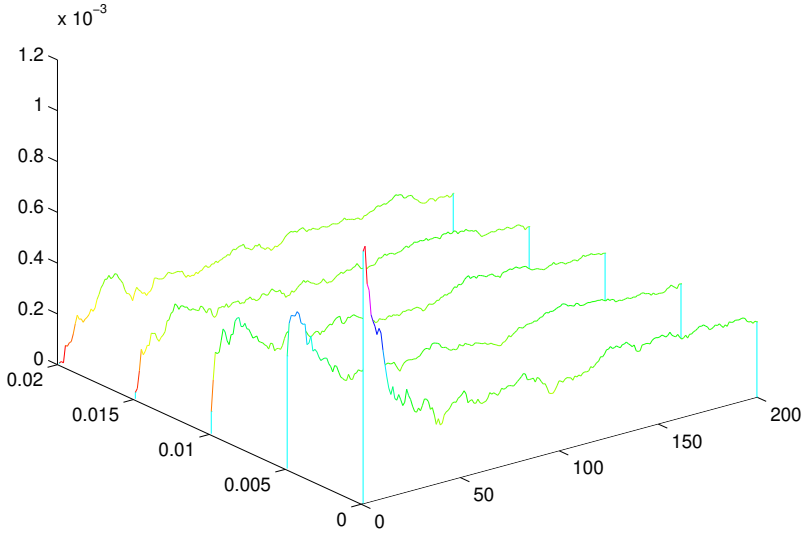

Figure 2: Mean Returns of each rule (n, λ). The mean is taken over $t \in [T_s, T_f]$.

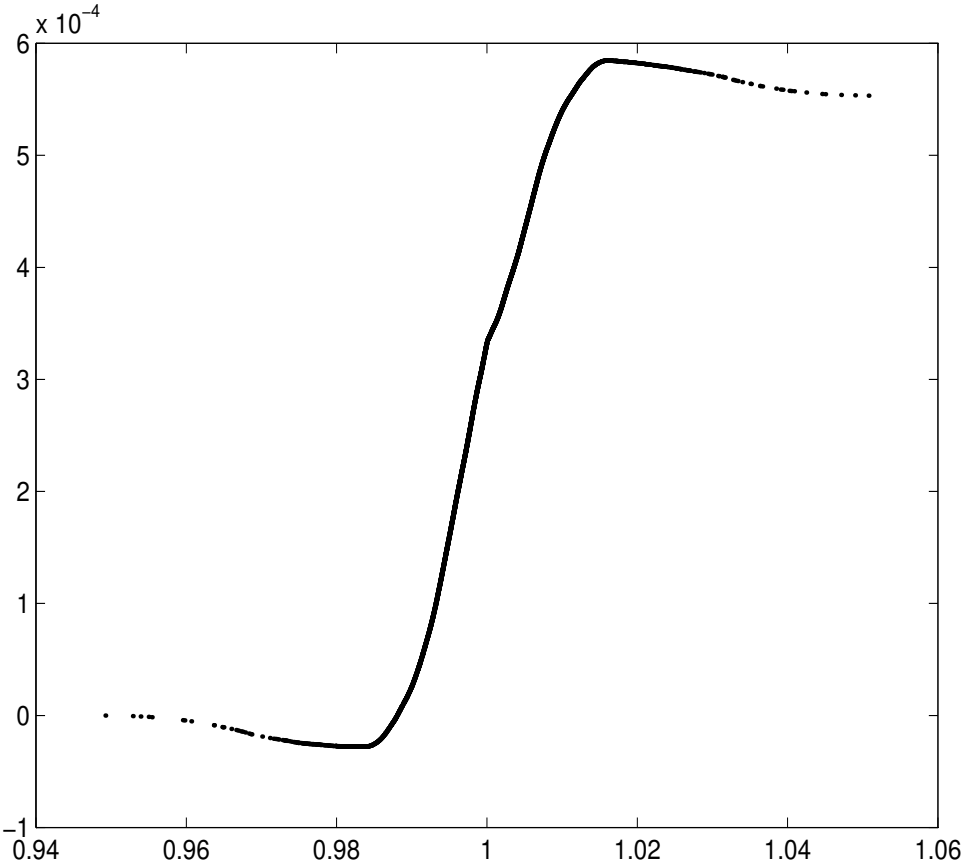

Figure 3: This is $\widehat{M}(\gamma)$ the sample version of $M(\gamma)$