The Cost of Technical Trading Rules in the Forex Market: A Utility-based Evaluation

Hans Dewachter[a,b]*and Marco Lyrio[a]

[a] CES, Catholic University of Leuven
[b] RIFM and ERIM, Erasmus University Rotterdam

April 2004

Abstract

We compute the opportunity cost for rational risk averse agents of using technical trading rules in the foreign exchange rate market. Our purpose is to investigate whether these rules can be interpreted as near-rational investment strategies for rational investors. We analyze four different exchange rates and find that the opportunity cost of using chartist rules tends to be prohibitively high. We also present a method to decompose this opportunity cost into two parts: (*i*) a cost related to the misallocation of wealth, which increases with the investor's level of risk aversion (allocational cost); and (*ii*) a cost related to the investor's erroneous belief regarding the sign of the expected excess return (expectational cost). The results indicate that up to medium levels of risk aversion expectational costs are the principal component of the total opportunity cost of using technical trading rules. For higher levels of risk aversion, allocational costs become the most important component of the total opportunity cost.

Keywords: Technical trading rules, exchange rate.

J.E.L.: F31, G15.

*Corresponding author. Address: Center for Economic Studies, Catholic University of Leuven, Naamsestraat 69, 3000 Leuven, Belgium. Tel: +32(0)16-326859, e-mail: hans.dewachter@econ.kuleuven.ac.be. We are grateful for financial support from the FWO-Vlaanderen (Project No.:G.0332.01). We thank the editor, Michael Melvin, and an anonymous referee for valuable comments which improved this paper significantly. The authors are responsible for remaining errors.

1 Introduction

Despite the numerous studies reporting the pervasive use and profitability of technical trading rules (TTRs),[1] also called chartist rules, there is still a considerable amount of scepticism in the academic literature regarding their true value. Critics of chartist rules often point to the seemingly suboptimal nature of the portfolio composition implied by these rules (e.g. Skouras, 2001a). After all, investment strategies based on such rules (*i*) restrict the information set to a narrow group of pre-defined information variables, (*ii*) assume a positive relation of the sign of the signal with future expected excess returns[2], and (*iii*) imply a bang-bang type of investment strategy, i.e. a strategy where all wealth is invested either short or long. Each of these assumptions goes against the standard rational investor paradigm. The first two assumptions possibly go against the rationality of expectations formation, while the third is in general at odds with the assumption of risk aversion.

In this paper, we assess the value for rational risk averse investors of using TTRs in the foreign exchange rate market. The main motivation being that even if such rules turn out to be suboptimal, the observed practice of using these rules could still be considered as near rational for a large class of risk averse agents. More specifically, if the cost for risk averse agents of using TTRs is low, one may argue that following such rules of thumb comes close to the optimal trading strategy and could, therefore, be rationalized in terms of information cost type of arguments. Skouras (2001b), for instance, argues that such discrete rules are easier to learn than optimal decision rules.

The opportunity cost associated with the use of TTRs for risk averse agents can be decomposed into two components. We present a simple method to compute each of these components, which involves the introduction of a hypothetical, risk neutral agent. The first cost component originates from the suboptimality of the investment strategy due to the agent's level of risk aversion (allocational cost). It is associated with the misallocation of wealth and can be recovered by contrasting the portfolio positions of risk averse and risk neutral agents. Since both agents have identical expectations, the difference in their trading positions can be linked to the costs of risk averse agents investing according to bang-bang investment strategies. The second cost component relates to the potential error made by chartist traders due to the assumed relation between the chartist signal and the expected excess return (expectational cost). Risk neutral liquidity constrained agents have in common with chartist traders that investment strategies will typically be bang-bang solutions. As a result, differences in their trading positions isolate the costs associated with expectational errors in the relation between the technical trading signal and the expected return. If this relation is always positive, chartist trading strategies are equivalent to those of rational risk neutral liquidity constrained agents. In this case, chartist rules are, therefore, also rational (Skouras, 2001a). Combining these cost components results in a total opportunity cost for a

[1] See, among others, Gençay (1999), LeBaron (1992, 1999, 2002), Neely et al. (1997), and Taylor (1980).

[2] We assume here a standardized rule where a positive (negative) technical trading signal corresponds to a long (short) investment position.

rational risk averse agent of using TTRs. We use this technique to identify possible classes of risk averse agents for which these opportunity costs are limited. In this case, one could perhaps rationalize the use of TTRs in terms of near-rational behavior.

Computing the costs of chartist trading rules implies both the identification of technical trading signals and the design of a statistical model to relate the conditional moments of the excess returns to the technical trading signal. In this paper, we restrict the analysis to the class of moving average signals, or rules. This is the most widely used class of TTRs in the foreign exchange market and it has been shown to be robust in their profit generating capacities. We also opt for a relatively simple model relating return moments to the trading signal. While more advanced techniques such as the nonparametric regression technique of Brandt (1999), or nonlinear models such as neural nets (Gençay, 1999) or Markov switching models (e.g. Dewachter, 2001) could be used, we try to strike a balance between generality and computational costs. We, therefore, use a regression approach to estimate possible time-varying parameters of a Taylor expansion of the relation between return moments and trading signals. This approach is sufficiently flexible to allow for nonlinearities in the signal-return moment relation while at the same time is computationally tractable so as to allow for continuous updating of the parameters.

The remainder of the paper is organized in three main sections. In section 2, we discuss the proposed decomposition of the costs associated with the use of TTRs. The empirical results are presented in section 3. In this section, we first analyze the statistical models relating trading signals to return moments. We do find evidence of a nonlinear relation for both the conditional expected return and for the variance. Using these models to construct the optimal portfolio rule for classes of risk averse agents, we subsequently analyze the value of technical trading signals and the costs associated with the use of TTRs. We summarize the main findings of the paper in the concluding section.

2 The opportunity cost of TTRs

TTRs are typically rules of thumb that relate a certain information variable, the technical trading signal, to a trading position. In other words, a TTR specifies a mapping from the time t signal z_t to an advised trading position $\alpha_{CH}(z_t)$, expressed as a percentage of the agent's initial wealth. A typical feature is the discontinuity in the mapping $\alpha_{CH}(z_t)$. We assume that the trading signal has been standardized such that the trading rule can be described as:

$$\alpha_{CH}(z_t) = \begin{cases} b_L \text{ if } z_t > 0 \\ 0 \text{ if } z_t = 0 \\ -b_S \text{ if } z_t < 0 \end{cases} \tag{1}$$

where $b_S, b_L > 0$ denote the liquidity constraint faced by the chartist trader on a short and long position, respectively.[3]

[3] Although not considered here, trading rules can also include bands of inaction.

Obviously, the optimal portfolio composition might differ from the one implied by the TTR. As noted by Skouras (2001a), utility-based optimal trading rules typically depend on various factors, including the level of risk aversion and rational expectations about the conditional return distribution. Adopting a standard mean-variance approach, the optimal trading strategy for a rational, risk averse, liquidity constrained agent can be written as the solution to the following problem:

$$\max_{\alpha} U_t = \max_{\alpha} E_t\left[W_{t+1}\right] - \frac{\gamma}{2} Var_t(W_{t+1}) \tag{2}$$

$$\begin{aligned} & s.t. \\ & W_{t+1} = W_t(1 + r_f + \alpha X_{t+1}) \\ & \alpha \in [-b_S, b_L] \end{aligned} \tag{3}$$

where W_t denotes the wealth at time t, γ represents the investor's level of risk aversion, r_f is the risk-free interest rate, α is the percentage of the initial wealth W_t invested in the risky asset, and X_{t+1} is the speculative return above the riskless interest rate obtained from investing in the risky asset. We implicitly assume that the expectations are conditional on the technical trading signal z_t, or that $E_t\left[X_{t+1}\right] = E\left[X_{t+1} \mid z_t\right]$. Assuming that the initial wealth is equal to one ($W_t = 1$), the optimal portfolio allocation can be expressed as:

$$\alpha^*_{RA} = \frac{E_t\left[X_{t+1}\right]}{\gamma Var_t(X_{t+1})}. \tag{4}$$

Since typical TTRs merely specify trading positions (long or short), only under very restrictive circumstances will these TTRs emerge as optimal trading rules. This is only possible, as mentioned before, in the case of risk neutrality combined with liquidity constraints. In this case, the optimal portfolio is a bang-bang solution and the investment position is determined based only on the sign of the expected return. More formally, for a risk-neutral investor ($\gamma = 0$), the optimal trading rule is given by:

$$\alpha^*_{RN} = \begin{cases} b_L \text{ if } E_t\left[X_{t+1}\right] > 0 \\ 0 \text{ if } E_t\left[X_{t+1}\right] = 0 \\ -b_S \text{ if } E_t\left[X_{t+1}\right] < 0. \end{cases} \tag{5}$$

Note, however, that this bang-bang solution under risk neutrality only collapses to the standard TTR if there is a positive relation between the sign of the chartist signal (z_t) and the expected excess return on the risky asset ($E_t\left[X_{t+1}\right]$). Comparing (1) and (5), one observes that only if the following equivalences hold will the optimal rule for a risk neutral agent be equivalent to the standard chartist trading rule ($\alpha^*_{RN} = \alpha_{CH}$):

$$\begin{aligned} & E_t[X_{t+1}] > 0 \Leftrightarrow z_t > 0, \\ & E_t[X_{t+1}] = 0 \Leftrightarrow z_t = 0, \\ & E_t[X_{t+1}] < 0 \Leftrightarrow z_t < 0. \end{aligned} \tag{6}$$

TTRs are, therefore, not necessarily irrational trading strategies. To the extent that there is a positive relation between the information variable z_t and the rationally expected direction of the return, the strategy is optimal for risk neutral agents. Moreover, given the continuity of the portfolio positions in γ, technical trading strategies could also function as near-rational trading rules for a specific subclass of risk averse agents, i.e. for $\gamma's$ close to zero. If the level of risk aversion is relatively high or the relation between the chartist signal and the rationally expected return is not positive, the total opportunity cost of following a TTR instead of the optimal strategy can become prohibitively high. Using eq. (2) and (3), this cost can be quantified as the difference between the utility obtained by a rational risk averse investor using his/her optimal trading strategy, $U_t^{RA}(\alpha_{RA}^*)$, and the utility achieved by the *same* investor now using the optimal trading strategy of a chartist trader, $U_t^{RA}(\alpha_{CH})$:

$$\begin{aligned} \Lambda_{TOT}(z_t) &= U_t^{RA}(\alpha_{RA}^*) - U_t^{RA}(\alpha_{CH}) && (7) \\ &= \left(\alpha_{RA}^* - \alpha_{CH}\right) E_t\left[X_{t+1}\right] - \frac{1}{2}\gamma\left(\alpha_{RA}^{*2} - \alpha_{CH}^2\right) Var_t\left(X_{t+1}\right). && (8) \end{aligned}$$

The above total cost of following the chartist trading rule instead of the optimal strategy can be decomposed into two effects. For this, we add and subtract from (7) the utility obtained by a rational risk averse investor making use of the optimal trading strategy of a risk neutral agent, $U_t^{RA}(\alpha_{RN}^*)$:

$$\begin{aligned} \Lambda_{TOT}(z_t) &= \left(U_t^{RA}(\alpha_{RA}^*) - U_t^{RA}(\alpha_{RN}^*)\right) + \left(U_t^{RA}(\alpha_{RN}^*) - U_t^{RA}(\alpha_{CH})\right) \\ &= \Lambda_{ALL}(z_t) + \Lambda_{EXP}(z_t). && (9) \end{aligned}$$

The first effect, denominated *allocational cost* (Λ_{ALL}), relates to the suboptmality of the portfolio allocation for a risk averse investor when using the optimal strategy of a risk neutral agent. Although both agents have the same expectation regarding future excess returns, $E_t\left[X_{t+1}\right]$, they adopt different trading strategies. Risk averse agents invest a percentage of their wealth in the risky asset taking into account the riskiness of the trading position. Risk neutral investors, on the other hand, always adopt a bang-bang strategy, i.e. they either take a full long or short position depending only on the sign of the expected excess return. Using again eq. (2) and (3), this cost can be expressed as:

$$\Lambda_{ALL}(z_t) = \left(\alpha_{RA}^* - \alpha_{RN}^*\right) E_t\left[X_{t+1}\right] - \frac{1}{2}\gamma\left(\alpha_{RA}^{*2} - \alpha_{RN}^2\right) Var_t\left(X_{t+1}\right). \quad (10)$$

The second effect, denominated *expectational cost* (Λ_{EXP}), refers to the difference in utility obtained by a risk averse investor making use of two different strategies: the one adopted by a risk neutral investor and the one used by a chartist trader. These two agents have in common the use of bang-bang strategies. Their trading positions differ, however, since the portfolio decision of a risk neutral agent is based on the sign of the rationally expected future excess return while the chartist trader only uses the sign of the trading signal for this purpose. In other words, this cost component expresses the loss due solely to the possible differences between the rational and the technical expectations[4] regarding the sign of future excess returns. If chartist beliefs are inconsistent with the rationally expected sign of the future excess return, following chartist trading rules results in a loss in expected terms. If, on the other hand, chartist beliefs are consistent with the rational ones, i.e. the equivalence in (6) holds, the trading strategies of both agents are the same, $\alpha^*_{RN} = \alpha_{CH}$, and this expectational cost drops out. The expression for the expectational cost is given by:

$$\Lambda_{EXP}(z_t) = (\alpha^*_{RN} - \alpha_{CH}) E_t [X_{t+1}] - \frac{1}{2}\gamma \left(\alpha^{*2}_{RN} - \alpha^2_{CH}\right) Var_t (X_{t+1}). \tag{11}$$

Note that only in two cases α^{*2}_{RN} can be different from α^2_{CH}, making the second term in (11) different from zero: (i) when the trading signal is equal to zero but the expected excess return on the risky asset is not, $z_t = 0$ and $E_t [X_{t+1}] \neq 0$. In this case, $\alpha_{CH} = 0$ but $\alpha^*_{RN} \neq 0$; and (ii) when the expected excess return is equal to zero for a trading signal different from zero, $z_t \neq 0$ and $E_t [X_{t+1}] = 0$. In this case $\alpha^*_{RN} = 0$ but $\alpha_{CH} \neq 0$. In all other cases, $\alpha^2_{RN} = \alpha^2_{CH}$ and the mentioned term vanishes. In practice, both cases have a zero probability of happening and the expectational cost can be considered as only a function of the difference in the optimal positions of the risk neutral investor and the chartist trader (i.e. the second term on the right-hand side of (11) drops out).

Observe also that both the allocational and the expectational costs are, by definition, nonnegative. For the allocational cost, the trading position α^*_{RA} is the one that maximizes the utility of a risk averse investor. Any other strategy will, therefore, result in a lower utility. Regading the expectational cost, the question is which suboptimal bang-bang strategy (α^*_{RN} or α_{CH}) results in a higher utility for the risk averse investor. Since this investor shares the same expectation regarding future excess returns as the risk neutral agent, the use of the latter's optimal strategy (α^*_{RN}) will result in a higher utility in comparison with the one from a chartist trader (α_{CH}), resulting in a positive expectational cost.

3 Empirical analysis

3.1 Data and TTRs

The empirical analysis is performed for a set of TTRs applied to the spot exchange rates of four currencies against the U.S. dollar: the German mark, the British pound, the Japanese yen, and the Swiss franc (DEM/USD, GBP/USD, JPY/USD, and CHF/USD, respectively). Exchange rates are expressed in the standard way as the price in the domestic currency of one U.S. dollar, considered here as the foreign currency. We use daily data for the above exchange rates obtained from Datastream for the period January 1, 1973 to March 25, 2003, yielding a total of 7887 observations.

The type of trading rule depends on the way the trading signal is computed. In this paper, we restrict the empirical analysis to the class of moving average trading rules. This is one of

[4] Note that the rational expectation is formed based on the same information set used by the chartist trader, i.e. the technical trading signal.

the most used trading rules since the early seventies and has been shown to generate excess profits through time in the foreign exchange market. Furthermore, it also seems to work well out-of-sample (see, for instance, Neely et al., 1997). Due to its widespread use, this choice also aims at reducing possible selection bias with respect to the chosen class of trading rules. The technical trading signal z_t for this class of rules is constructed based on a short and a long moving average window of past exchange rates:

$$z_t = \frac{\frac{1}{K}\sum_{j=0}^{K-1} S_{t-j}}{\frac{1}{L}\sum_{j=0}^{L-1} S_{t-j}} - 1, \tag{12}$$

where S_t denotes the exchange rate at time t, and $K(L)$ denotes the size or number of observations of the short (long) window of the moving average signal. The investment position of a chartist trader at each point in time is determined based on the sign of z_t, as expressed in eq. (1). We consider three types of trading signals (or rules) depending on the number of days incorporated in the short (K) and long (L) window of the moving average rule. The following trading rules are used. Rule 1: $K = 10$, $L = 50$; Rule 2: $K = 20$, $L = 100$; and Rule 3: $K = 40$, $L = 200$. Due to the backward-looking nature of the moving average signal, the effective number of data points in the sample of each trading rule depends on the size of the long window used to compute the trading signal. Finally, in the computation of the return from investing in the exchange rate market, we disregard the interest rate differential between the countries, as standard in the literature.[5]

Table 1 presents the mean excess return obtained by an agent using the technical trading signal to invest in the risky asset according to the bang-bang chartist strategy presented in (1). The analysis is done for the whole sample period assuming that the investor can only invest his/her own wealth ($b_S = b_L = 1$). One observes the standard result found in the literature that investing according to TTRs generates significant mean excess returns. In our case, both TTRs 1 and 2 generate significant excess returns at the 1% confidence level for the four currencies considered (third line of results for each rule). The TTR 3 only generates significant excess returns for the Japanese yen-U.S. dollar. Note also that the trading rule profits are not homogeneous across positive and negative signals (first and second line of results for each rule). Typically, trading rules are more profitable in one of the two investment positions.

Insert Table 1

3.2 Least squares prediction models

Central in the above analysis is the projection of expected future excess return moments on the technical trading signal z_t. Naturally, the type of model used to project these moments on

[5] The speculative return $X_{t+1} = \Delta e_{t+1} + r_f^* - r_f + \Delta e_{t+1} r_f^*$, where e denotes the natural logarithm of the exchange rate, r_f^* represents the risk-free return in the foreign country, is then simplified to $X_{t+1} \cong \Delta e_{t+1}$.

the set of chartist signals influences the results regarding the costs of using chartist strategies. In this paper, we try to strike a balance between the generality in the class of functions used to project moments and the computational cost of continuously updating these projections to take into account the "real-time" flow of information.[6] We, therefore, approximate the mapping between the first two excess return moments defined as

$$\begin{aligned} m_t\left(z_t\right) &= E_t[X_{t+1}] \\ v_t\left(z_t\right) &= Var_t(X_{t+1}) \end{aligned} \tag{13}$$

and the information set, or trading signal, at time t in terms of a Taylor expansion around the mean of z_t. We, furthermore, assume that agents follow the Least Squares (LS) learning principles, i.e. use Ordinary Least Squares (OLS) techniques to estimate and update the forecasting model according to the information set.[7] In practice, we use an expanding window regression framework to estimate the parameters of a Taylor expansion:

$$\begin{aligned} m_t(z) &= Z'\beta_t \\ v_t\left(z\right) &= Z'\delta_t \end{aligned} \tag{14}$$

where $Z \in R^{(P+1)\times 1}$ denotes the vector containing the independent variables: 1, z, z^2, ... , z^p. The parameter vectors β_t and δ_t are obtained from regressing observed excess returns on technical trading signals. We implicitly assume stationarity of the exchange rate changes and hence stationarity of the distribution of the trading signal z_t. We select the optimal order of the Taylor approximation based on the Akaike Information Criterion (AIC). We restrict the maximum order of the expansion to six and retain the one that minimizes the AIC. The orders of the Taylor approximations are also allowed to differ across the mean and variance equations and can be continuously updated based on the new flow of information. In the out-of-sample exercise presented below, this criterion is updated at a fixed number of new observations.

In order to save on space, the regression results are not presented here.[8] We illustrate the main results by presenting the expected mean and variance of excess returns conditioned on the chartist signal z_t for the CHF/USD exchange rate based on the TTRs 1 to 3 (see top panel of Figures 1 to 3, respectively). These results are computed using the whole sample period to estimate the coefficients in eq. (14) [9], equal to β_T and δ_T in this case. The orders of the

[6] Note that the most general technique available to optimize the trading position is the one proposed by Brandt (1999). He combines a nonparametric technique with the first order condition for the optimal portfolio composition in order to derive a mapping between the portfolio and the information variable. This approach could be used here as well. Nevertheless, this technique requires a significant computational effort in the continuous updating of the information set. In a previous version of this paper, we have analyzed the optimality of technical trading rules based on the Brandt technique (see Dewachter and Lyrio, 2002). The portfolio allocation obtained there corresponds closely to the ones reported in this paper.

[7] This LS learning coincides with Bayesian learning under a diffuse prior, in which the agents are both unsure of their models and parameter values (Sargent, 1993). We thank the referee for pointing this out to us.

[8] All the results are, however, available upon request.

[9] For all exchange rates and trading rules, most of the computed coefficients are statistically significant at the 5% confidence level.

Taylor expansions used in this equation are, therefore, only computed once. The histogram of the trading signal is shown in the background. Two comments are to be made with respect to these results. First, for a certain range of the trading signal around zero, the regression results are approximately in line with the practice of technical traders, i.e. we observe a one-to-one relation between the trading signal and the expected future excess return (top-left panel of Figures 1 to 3). Considering the whole range of signals, however, one observes that this relation is both non-linear and non-monotonic. In fact, for more extreme signals, with a very low frequency of occurrence, this relation becomes inverted giving rise to a contrarian strategy.

Second, we find that the trading signal contains relevant information concerning future volatility. In line with the standard generalized autoregressive conditional heteroskedasticity (GARCH) literature, we see that the trading signals are correlated with future volatility, implying some predictability of z_t with respect to v_t. Also in line with the GARCH literature, we observe a generally symmetric relation between the trading signal and the variance.[10]

Insert Figures 1 to 3

In summary, the chartist signal contains significant information with respect to the future evolution of the exchange rate. Nevertheless, the information in the signal does not fully conform with the widely held beliefs of technical traders. Most importantly, the relation between the trading signal and the rationally expected future returns is non-linear and non-monotonic.

3.3 Optimal portfolio allocation

We now compute the optimal portfolio allocation based on eq. (4). This value depends on the sample used since the coefficients β_t and δ_t in eq. (14) vary through time. As in the previous section, we illustrate the main results based on the full sample period. The optimal portfolio strategy is given by the standard mean-variance optimal portfolio, conditioned on the signal z_t:

$$\alpha_{RA}^{*}(z_t) = \frac{m_T(z_t)}{\gamma v_T(z_t)} = \frac{Z_t'\beta_T}{\gamma Z_t'\delta_T}. \tag{15}$$

The four bottom panels of Figures 1 to 3 depict the optimal portfolio composition for the CHF/USD exchange rate in function of the observed trading signal z_t. The dashed line shows the unrestricted optimal portfolio, while the full line shows the optimal portfolio for an investor constrained to invest his/her own wealth ($b_S = b_L = 1$). Note that the optimal portfolio deviates significantly from the bang-bang strategy adopted by a risk neutral agent. For weak absolute trading signals, the optimal portfolio is typically less aggressive than the risk neutral solution ($|\alpha_{RA}^{*}| < 1$). The optimal rule is also clearly non-linearly related to the trading signal. While for weak signals there seems to be a trend following strategy, i.e. go

[10] For a detailed study concerning the relation between a trading signal and both the mean and volatility of stock returns, see Brock et al. (1992).

long when $z > 0$ and go short when $z < 0$, for strong positive or negative signals the optimal trading strategy tends to become contrarian, i.e. go short if $z >> 0$ and long if $z << 0$.

3.4 The costs of using TTRs

In this section, we assess the opportunity cost for rational risk averse investors of using the above mentioned TTRs . We consider four types of investors according to their level of risk aversion ($\gamma = 1, 5, 10, 20$), ranging from a relatively aggressive investor to a very risk averse one. As mentioned before, we adopt symmetric liquidity constraints by stipulating that an agent can only invest his/her own wealth (equal to one at the beginning of each period, $W_t = 1$), implying that $b_S = b_L = 1$.

In this paper, we restrict ourselves to an out-of-sample analysis of the opportunity cost of using TTRs. In-sample measures would only represent an accurate value of this opportunity cost if the parameter estimates in (14) would not vary through time. We verify, however, that these estimates are significantly different depending on the size of the sample used.[11] The out-of-sample analysis starts in November 01, 1976 (point 1001 out of a total of 7887 observations) and ends at the end of the sample (March 25, 2003).

The out-of-sample opportunity cost of using TTRs can be computed in two ways. One is based on the agent's expectations regarding the exchange rate mean and variance. It is then the direct computation of eq. (10) and (11) based on expanding window regressions of (14) up to time t and on the resulting time-varying optimal portfolio expressed in (**??**). Denoting T_0 and T respectively as the beginning and end points of the out-of-sample regressions, the average opportunity cost components can be computed as:

$$\overline{\Lambda}_{ALL} = \frac{1}{(T - T_0 + 1)} \sum_{t=T_0}^{T} \left[\left(\alpha^*_{RA}(z_t) - \alpha^*_{RN}(z_t) \right) m_t(z_t) - \frac{1}{2}\gamma \left(\alpha^{*2}_{RA}(z_t) - \alpha^2_{RN}(z_t) \right) v_t(z_t) \right]$$

$$\overline{\Lambda}_{EXP} = \frac{1}{(T - T_0 + 1)} \sum_{t=T_0}^{T} \left[\left(\alpha^*_{RN}(z_t) - \alpha_{CH}(z_t) \right) m_t(z_t) - \frac{1}{2}\gamma \left(\alpha^{*2}_{RN}(z_t) - \alpha^2_{CH}(z_t) \right) v_t(z_t) \right].$$

The above computations give us the out-of-sample average opportunity cost of using a TTR in terms of the *expected* mean and variance of the excess returns. As mentioned before, these costs are, by definition, nonnegative. They do not reflect, however, the realized opportunity cost faced by the investor. For this reason and in order to save on space, the results from these computations are not presented here.[12]

We focus here on the second variant of the out-of-sample computations of the opportunity cost components, which is based on the sample realizations of the mean and variance of the excess exchange rate returns.[13] As in the previous case, we first compute expanding window regressions for the expected mean and variance of the excess exchange rate return based on

[11] Although we omit the regression results here they are available upon request.
[12] These results are, however, available upon request.
[13] We thank the referee for suggesting this approach.

the technical trading signal at time t, as given in (14).[14] This allows us to compute the time t optimal portfolio given by (**??**). We then use the realized time $t+1$ exchange rate return to compute the agent's realized return from investing in the risky asset. For the risk averse agent, this is equal to $r_t^{RA} = \alpha_{RA}^*(z_t)\Delta e_{t+1}$, where e_t denotes the natural logarithm of the exchange rate, with a sample mean equal to

$$\overline{r}^{RA} = \frac{1}{(T-T_0)}\sum_{t=T_0}^{T-1} r_t^{RA}.$$

From eq. (2) and (3) and the above definitions, we can compute the out-of-sample average mean-variance utility of a risk averse investor using his/her optimal trading strategy:

$$\begin{aligned} \overline{U}^{RA}(\alpha_{RA}^*) &= 1 + r_f + \overline{r}^{RA} - \frac{1}{2}\gamma\frac{1}{(T-T_0-1)}\sum_{t=T_0}^{T-1}(r_t^{RA} - \overline{r}^{RA})^2 \\ &= 1 + r_f + \widetilde{U}^{RA}(\alpha_{RA}^*), \end{aligned}$$

where we denote $\widetilde{U}^{RA}(\alpha_{RA}^*)$ as the respective scaled out-of-sample mean-variance utility. The out-of-sample mean-variance utility of the *same* risk averse investor now using the optimal trading strategy of a risk neutral investor (α_{RN}^*) and of a chartist trader (α_{CH}) can be computed in the same way and are denoted, respectively, as $\overline{U}^{RA}(\alpha_{RN}^*)$ and $\overline{U}^{RA}(\alpha_{CH})$. The average allocational and expectational costs incurred by the risk averse agent in using the TTR can be computed as a function of these utilities:

$$\overline{\Lambda}_{ALL} = \overline{U}^{RA}(\alpha_{RA}^*) - \overline{U}^{RA}(\alpha_{RN}^*) = \widetilde{U}^{RA}(\alpha_{RA}^*) - \widetilde{U}^{RA}(\alpha_{RN}^*)$$

$$\overline{\Lambda}_{EXP} = \overline{U}^{RA}(\alpha_{RN}^*) - \overline{U}^{RA}(\alpha_{CH}) = \widetilde{U}^{RA}(\alpha_{RN}^*) - \widetilde{U}^{RA}(\alpha_{CH})$$

with the total average opportunity cost being equal to $\overline{\Lambda}_{TOT} = \overline{\Lambda}_{ALL} + \overline{\Lambda}_{EXP}$. $\widetilde{U}^{RA}(\alpha_{RN}^*)$ and $\widetilde{U}^{RA}(\alpha_{CH})$ denote the scaled out-of sample mean-variance utility obtained by a risk averse investor when using the strategy of a risk neutral agent and of a chartist trader, respectively. Note that in this out-of-sample case one would not necessarily expect a non-negative average utility for the optimal trading strategy ($\widetilde{U}^{RA}(\alpha_{RA}^*)$). Only if the implied model predicts in a accurate way the out-of-sample expected excess return and variance, one would expect this value to be positive.

Tables 2 to 4 present the scaled out-of-sample mean-variance utility of a risk averse investor making use of each of the three trading strategies under consideration (α_{RA}^*, α_{RN}^*, α_{CH}) and for the TTRs 1 to 3, respectively. We also present in these tables the resulting opportunity cost components computed based on these utilities. We can derive a number of observations from these results. First, the chartist signal contains ex post a significant value for rational risk averse agents (see positive values for $\widetilde{U}^{RA}(\alpha_{RA}^*)$). This is the case since these signals are used to compute the rationally expected mean and variance of the excess returns. The obtained utility is higher for more aggressive agents (low γ). This does not imply, however,

[14] The orders of the Taylor approximations in (14) are computed at every 10 new observations.

that rational risk averse agents should turn to technical trading *strategies*. In fact, in most cases the use of a technical trading strategy generates a negative scaled utility for the rational risk averse agent (see negative values for $\widetilde{U}^{RA}(\alpha_{CH})$). This value is in a few times positive for agents with a very low level of risk aversion. As a result, one observes a rather significant total opportunity cost of using TTRs, ranging from approximately 0 to 16% per year.

Insert Tables 2 to 4

Second, one observes that in most cases for levels of risk aversion up to 10, expectational costs constitute the main component of the total opportunity cost of using a certain TTR. For higher levels of risk aversion, allocational inefficiencies dominate this total opportunity cost, becoming prohibitively large. Overall, expectational costs range from 1% to 10% per year, which seem quite substantial. Since these costs do not depend on the agent's level of risk aversion, they constitute a lower bound on the total opportunity cost of using technical rules. Even if the above costs were considered as reasonable, the allocation costs tend to increase significantly with the increase in the level of risk aversion. Figure 4 illustrates this cost decomposition for the GBP/USA exchange rate applying the TTR 2 (see Table 3). Aggressive investors (low γ) using chartist rules are mainly concerned about possible expectational errors, or errors in their judgement regarding the sign of the expected exchange rate return. Most part of the opportunity cost for conservative investors (high γ) making use of technical rules derive from the allocation of wealth in a non-optimized way, i.e. from using a bang-bang strategy.

Insert Figure 4

In summary, although chartist signals contain ex post relevant and valuable information to rational risk averse investors, TTRs are not an efficient or near-efficient way to incorporate this information. TTRs fail to map accurately signals into trading positions both due to expectational and allocational inefficiencies. The total cost for a risk averse investor of using such rules is, therefore, prohibitively high. Note that in this paper we restrict the information used by the rational investor to the same information set used by the chartist trader. To answer the main question of this paper, we conclude that TTRs are not near-rational equivalents to optimal trading rules.

4 Conclusions

The main goal of this paper was to answer the basic question whether or not TTRs could be interpreted as near rational investment strategies for a class of risk averse agents. Based on the above analysis, we conclude that they cannot be interpreted in that way. We find that allocational and expectational costs generate prohibitively high welfare costs to rational agents.

The relative importance of these two cost components varies clearly with the level of risk aversion. For low levels of risk aversion, allocational costs tend to be very low and expectational costs impose, therefore, most of the cost of using TTRs. Since expectational costs do not depend on the level of risk aversion, they can be seen as a lower bound on the opportunity cost for risk averse agents of using chartist rules. This type of cost alone should prevent investors from using the technical trading signal in order to apply chartist trading strategies. Allocational costs increase with the level of risk aversion and tend to equate with the level of expectational costs for levels of risk aversion around 10. For higher levels of risk aversion, allocational costs are clearly the dominating effect within the total opportunity cost of using such rules. The results hold in general for the three moving average rules and for each of the exchange rates analyzed in this paper.

References

[1] Brandt, M. (1999), "Estimating Portfolio and Consumption Choice: A Conditional Euler Equations Approach", *Journal of Finance* 54 (5), 1609-1645.

[2] Brock, W., J. Lakonishok, and B. LeBaron (1992), "Simple Technical Trading Rules and the Stochastic Properties of Stock Returns", *Journal of Finance* 47 (5), 1731-1764.

[3] Dewachter, H. (2001), "Can Markov Switching Models Replicate Chartist Profits in the Foreign Exchange Market", *Journal of International Money and Finance* 20, 25-41.

[4] Dewachter, H. and M. Lyrio (2002), "The Economic Value of Technical Trading Rules: A Nonparametric Utility-based Approach", CES Discussion Paper 02.03, KULeuven.

[5] Gençay, R. (1999), "Linear, Non-linear and Essential Foreign Exchange Rate Prediction with Simple Technical Trading Rules", *Journal of International Economics* 47, 91-107.

[6] LeBaron, B. (1992),"Do Moving Average Trading Rule Results Imply Nonlinearities in Foreign Exchange Markets?", Working Paper, Department of Economics, University of Wisconsin-Madison.

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

[8] LeBaron, B. (2002), "Technical Trading Profitability in Foreign Exchange Markets in the 1990's", Working Paper, Graduate School of International Economics and Finance, Brandeis University.

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

[10] Sargent, T. (1993), *Bounded Rationality in Macroeconomics*, Oxford: Clarendon Press.

[11] Skouras, S. (2001a), "Financial Returns and Efficiency as Seen by an Artificial Technical Analyst", *Journal of Economic Dynamics and Control* 25, 213-244.

[12] Skouras, S. (2001b), "Learning to Profit with Discrete Investment Rules", *Quantitative Finance* 1, 284-288.

[13] Taylor, S. (1980), "Conjectured Models for Trends in Financial Prices, Tests and Forecasts", *Journal of the Royal Statistical Society* A, 143, 338-362.

Table 1: Annual mean excess return for the selected TTRs

	DEM/USD	GBP/USD	JPY/USD	CHF/USD
	TTR 1: K=10, L=50			
Return (p.a.) $z < 0$	0.07559***	0.03253	0.10317***	0.09025***
Return (p.a.) $z > 0$	0.05004*	0.06104**	0.05070*	0.03713
Return (p.a.)	0.06308***	0.04702***	0.07661***	0.06446***
	TTR 2: K=20, L=100			
Return (p.a.) $z < 0$	0.07200***	0.01666	0.08865***	0.07448**
Return (p.a.) $z > 0$	0.05441*	0.04615*	0.03685	0.02659
Return (p.a.)	0.06386***	0.03223*	0.06292***	0.05201**
	TTR 3: K=40, L=200			
Return (p.a.) $z < 0$	0.02319	0.00982	0.05915**	0.05357*
Return (p.a.) $z > 0$	0.00650	0.03860	0.01131	0.00657
Return (p.a.)	0.01538	0.02430	0.03738*	0.03245

The analysis is done for the whole sample period. The agent invests according to eq. (1) and is constrained to invest his/her own wealth ($b_S = b_L = 1$).
Returns are presented in per annum terms by multiplying the daily returns by the number of trading days, taken here to be 262. The entry *Return (p.a.)* $z < 0$ measures the annual mean excess return obtained when the signal z was negative, implying a short position. Analogously, *Return (p.a.)* $z > 0$ measures the annual mean excess return obtained when the signal z was positive, implying a long position. The entry *Return (p.a.)* expresses the annual mean excess return from trading according to both positive and negative signals.
***,** and * indicate that the averages are statistically different from zero at the 1%, 5% and 10% confidence level, respectively.

Table 2: Average out-of-sample mean-variance utility and cost decomposition of TTRs (TTR 1: K=10, L=50)

	Level of risk aversion (γ)			
	1	5	10	20
		DEM/USD		
$\widetilde{U}^{RA}(\alpha^*_{RA})$	0.04609	0.03309	0.02327	0.01356
$\widetilde{U}^{RA}(\alpha^*_{RN})$	0.04834	0.02561	-0.00281	-0.05966
$\widetilde{U}^{RA}(\alpha_{CH})$	0.00023	-0.02253	-0.05098	-0.10787
$\overline{\Lambda}_{ALL}$	-0.00225	0.00749	0.02608	0.07322
$\overline{\Lambda}_{EXP}$	0.04811	0.04811	0.04811	0.04811
$\overline{\Lambda}_{TOT}$	0.04586	0.05560	0.07419	0.12133
		GBP/USD		
$\widetilde{U}^{RA}(\alpha^*_{RA})$	0.04869	0.03358	0.01860	0.00630
$\widetilde{U}^{RA}(\alpha^*_{RN})$	0.04693	0.02752	0.00327	-0.04525
$\widetilde{U}^{RA}(\alpha_{CH})$	0.00996	-0.00946	-0.03374	-0.08230
$\overline{\Lambda}_{ALL}$	0.00176	0.00606	0.01533	0.05154
$\overline{\Lambda}_{EXP}$	0.03696	0.03696	0.03696	0.03696
$\overline{\Lambda}_{TOT}$	0.03872	0.04302	0.05230	0.08850
		JPY/USD		
$\widetilde{U}^{RA}(\alpha^*_{RA})$	0.08906	0.06976	0.04841	0.03069
$\widetilde{U}^{RA}(\alpha^*_{RN})$	0.08749	0.06326	0.03297	-0.02760
$\widetilde{U}^{RA}(\alpha_{CH})$	0.00663	-0.01766	-0.04803	-0.10876
$\overline{\Lambda}_{ALL}$	0.00157	0.00650	0.01544	0.05829
$\overline{\Lambda}_{EXP}$	0.08084	0.08084	0.08084	0.08084
$\overline{\Lambda}_{TOT}$	0.08241	0.08734	0.09628	0.13913
		CHF/USD		
$\widetilde{U}^{RA}(\alpha^*_{RA})$	0.07143	0.05379	0.03011	0.01390
$\widetilde{U}^{RA}(\alpha^*_{RN})$	0.06912	0.04007	0.00376	-0.06886
$\widetilde{U}^{RA}(\alpha_{CH})$	0.00574	-0.02335	-0.05972	-0.13245
$\overline{\Lambda}_{ALL}$	0.00231	0.01372	0.02635	0.08276
$\overline{\Lambda}_{EXP}$	0.06337	0.06337	0.06337	0.06337
$\overline{\Lambda}_{TOT}$	0.06568	0.07709	0.08972	0.14613

$\widetilde{U}^{RA}(\alpha^*_{RA})$ denotes the scaled out-of-sample mean-variance utility of a risk averse investor using his/her optimal trading strategy (α^*_{RA}). $\widetilde{U}^{RA}(\alpha^*_{RN})$ and $\widetilde{U}^{RA}(\alpha_{CH})$ denote the scaled out-of-sample mean-variance utility of the *same* risk averse investor now using the optimal trading strategy of a risk neutral investor (α^*_{RN}) and of a chartist trader (α_{CH}), respectively. $\overline{\Lambda}_{ALL}$, $\overline{\Lambda}_{EXP}$, and $\overline{\Lambda}_{TOT}$ denote the average allocational cost, expectational cost, and total opportunity cost, respectively, in terms of utility, incurred by a risk averse investor using the optimal trading strategy of a chartist trader (α_{CH}). All entries are in per annum terms.

Table 3: Average out-of-sample mean-variance utility and cost decomposition of TTRs (TTR 2: K=20, L=100)

	Level of risk aversion (γ)			
	1	5	10	20
	DEM/USD			
$\widetilde{U}^{RA}(\alpha^*_{RA})$	0.06297	0.04436	0.03174	0.01745
$\widetilde{U}^{RA}(\alpha^*_{RN})$	0.06067	0.03794	0.00953	-0.04728
$\widetilde{U}^{RA}(\alpha_{CH})$	-0.00289	-0.02565	-0.05410	-0.11100
$\overline{\Lambda}_{ALL}$	0.00230	0.00642	0.02221	0.06473
$\overline{\Lambda}_{EXP}$	0.06355	0.06355	0.06355	0.06355
$\overline{\Lambda}_{TOT}$	0.06585	0.06997	0.08576	0.12828
	GBP/USD			
$\widetilde{U}^{RA}(\alpha^*_{RA})$	0.01900	0.01205	0.00513	-0.00091
$\widetilde{U}^{RA}(\alpha^*_{RN})$	0.01763	-0.00179	-0.02607	-0.07463
$\widetilde{U}^{RA}(\alpha_{CH})$	-0.00246	-0.02188	-0.04616	-0.09473
$\overline{\Lambda}_{ALL}$	0.00137	0.01384	0.03120	0.07372
$\overline{\Lambda}_{EXP}$	0.02008	0.02008	0.02008	0.02008
$\overline{\Lambda}_{TOT}$	0.02145	0.03393	0.05128	0.09380
	JPY/USD			
$\widetilde{U}^{RA}(\alpha^*_{RA})$	0.08191	0.06891	0.04905	0.03375
$\widetilde{U}^{RA}(\alpha^*_{RN})$	0.08240	0.05817	0.02787	-0.03271
$\widetilde{U}^{RA}(\alpha_{CH})$	-0.01922	-0.04351	-0.07388	-0.13461
$\overline{\Lambda}_{ALL}$	-0.00050	0.01074	0.02118	0.06646
$\overline{\Lambda}_{EXP}$	0.10161	0.10161	0.10161	0.10161
$\overline{\Lambda}_{TOT}$	0.10111	0.11235	0.12279	0.16807
	CHF/USD			
$\widetilde{U}^{RA}(\alpha^*_{RA})$	0.06800	0.04115	0.02623	0.01528
$\widetilde{U}^{RA}(\alpha^*_{RN})$	0.06423	0.03517	-0.00114	-0.07378
$\widetilde{U}^{RA}(\alpha_{CH})$	0.00735	-0.02174	-0.05810	-0.13083
$\overline{\Lambda}_{ALL}$	0.00377	0.00597	0.02737	0.08906
$\overline{\Lambda}_{EXP}$	0.05687	0.05687	0.05687	0.05687
$\overline{\Lambda}_{TOT}$	0.06064	0.06284	0.08424	0.14593

$\widetilde{U}^{RA}(\alpha^*_{RA})$ denotes the scaled out-of-sample mean-variance utility of a risk averse investor using his/her optimal trading strategy (α^*_{RA}). $\widetilde{U}^{RA}(\alpha^*_{RN})$ and $\widetilde{U}^{RA}(\alpha_{CH})$ denote the scaled out-of-sample mean-variance utility of the *same* risk averse investor now using the optimal trading strategy of a risk neutral investor (α^*_{RN}) and of a chartist trader (α_{CH}), respectively. $\overline{\Lambda}_{ALL}$, $\overline{\Lambda}_{EXP}$, and $\overline{\Lambda}_{TOT}$ denote the average allocational cost, expectational cost, and total opportunity cost, respectively, in terms of utility, incurred by a risk averse investor using the optimal trading strategy of a chartist trader (α_{CH}). All entries are in per annum terms.

Table 4: Average out-of-sample mean-variance utility and cost decomposition of TTRs (TTR 3: K=40, L=200)

	Level of risk aversion (γ)			
	1	5	10	20
	DEM/USD			
$\widetilde{U}^{RA}(\alpha^*_{RA})$	0.01507	0.00664	0.00069	0.00021
$\widetilde{U}^{RA}(\alpha^*_{RN})$	0.02796	0.00521	-0.02322	-0.08010
$\widetilde{U}^{RA}(\alpha_{CH})$	0.01605	-0.00670	-0.03515	-0.09204
$\overline{\Lambda}_{ALL}$	-0.01290	0.00142	0.02391	0.08031
$\overline{\Lambda}_{EXP}$	0.01191	0.01191	0.01191	0.01191
$\overline{\Lambda}_{TOT}$	-0.00099	0.01333	0.03582	0.09222
	GBP/USD			
$\widetilde{U}^{RA}(\alpha^*_{RA})$	0.02484	0.01750	0.01365	0.00936
$\widetilde{U}^{RA}(\alpha^*_{RN})$	0.02640	0.00698	-0.01730	-0.06584
$\widetilde{U}^{RA}(\alpha_{CH})$	-0.01357	-0.03300	-0.05728	-0.10584
$\overline{\Lambda}_{ALL}$	-0.00156	0.01052	0.03095	0.07520
$\overline{\Lambda}_{EXP}$	0.03997	0.03997	0.03997	0.03997
$\overline{\Lambda}_{TOT}$	0.03841	0.05049	0.07092	0.11517
	JPY/USD			
$\widetilde{U}^{RA}(\alpha^*_{RA})$	0.04824	0.02970	0.01588	0.00858
$\widetilde{U}^{RA}(\alpha^*_{RN})$	0.04341	0.01913	-0.01121	-0.07190
$\widetilde{U}^{RA}(\alpha_{CH})$	-0.00716	-0.03145	-0.06182	-0.12256
$\overline{\Lambda}_{ALL}$	0.00483	0.01057	0.02709	0.08049
$\overline{\Lambda}_{EXP}$	0.05056	0.05056	0.05056	0.05056
$\overline{\Lambda}_{TOT}$	0.05539	0.06113	0.07765	0.13105
	CHF/USD			
$\widetilde{U}^{RA}(\alpha^*_{RA})$	0.05853	0.03803	0.02287	0.01129
$\widetilde{U}^{RA}(\alpha^*_{RN})$	0.05543	0.02636	-0.00997	-0.08262
$\widetilde{U}^{RA}(\alpha_{CH})$	-0.01097	-0.04007	-0.07643	-0.14916
$\overline{\Lambda}_{ALL}$	0.00311	0.01167	0.03284	0.09391
$\overline{\Lambda}_{EXP}$	0.06639	0.06639	0.06639	0.06639
$\overline{\Lambda}_{TOT}$	0.06950	0.07806	0.09923	0.16030

$\widetilde{U}^{RA}(\alpha^*_{RA})$ denotes the scaled out-of-sample mean-variance utility of a risk averse investor using his/her optimal trading strategy (α^*_{RA}). $\widetilde{U}^{RA}(\alpha^*_{RN})$ and $\widetilde{U}^{RA}(\alpha_{CH})$ denote the scaled out-of-sample mean-variance utility of the *same* risk averse investor now using the optimal trading strategy of a risk neutral investor (α^*_{RN}) and of a chartist trader (α_{CH}), respectively. $\overline{\Lambda}_{ALL}$, $\overline{\Lambda}_{EXP}$, and $\overline{\Lambda}_{TOT}$ denote the average allocational cost, expectational cost, and total opportunity cost, respectively, in terms of utility, incurred by a risk averse investor using the optimal trading strategy of a chartist trader (α_{CH}). All entries are in per annum terms.

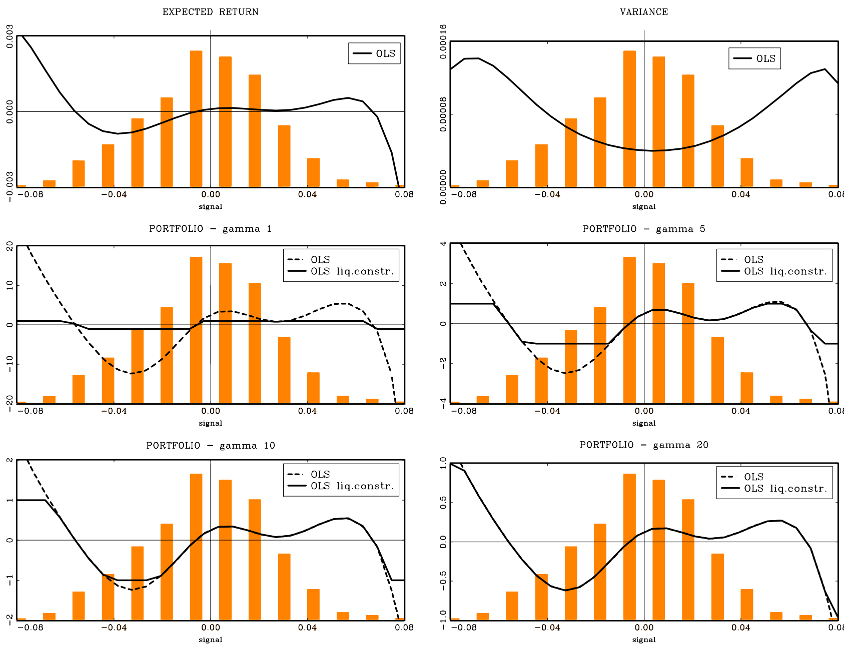

Figure 1: Daily expected return, variance and portfolio choice for the CHF/USD using the full sample (TTR 1: K=10, L=50)

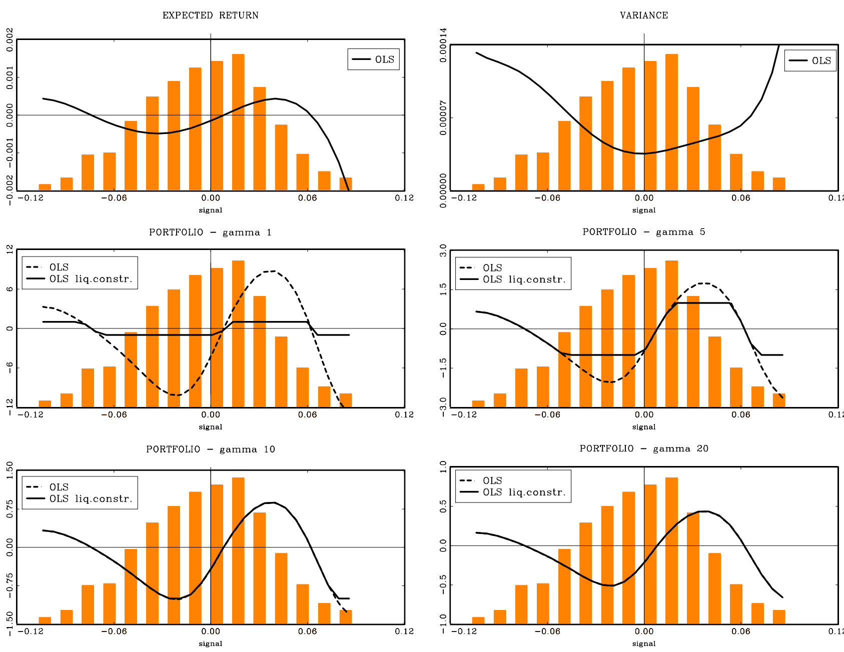

Figure 2: Daily expected return, variance and portfolio choice for the CHF/USD using the full sample (TTR 2: K=20, L=100)

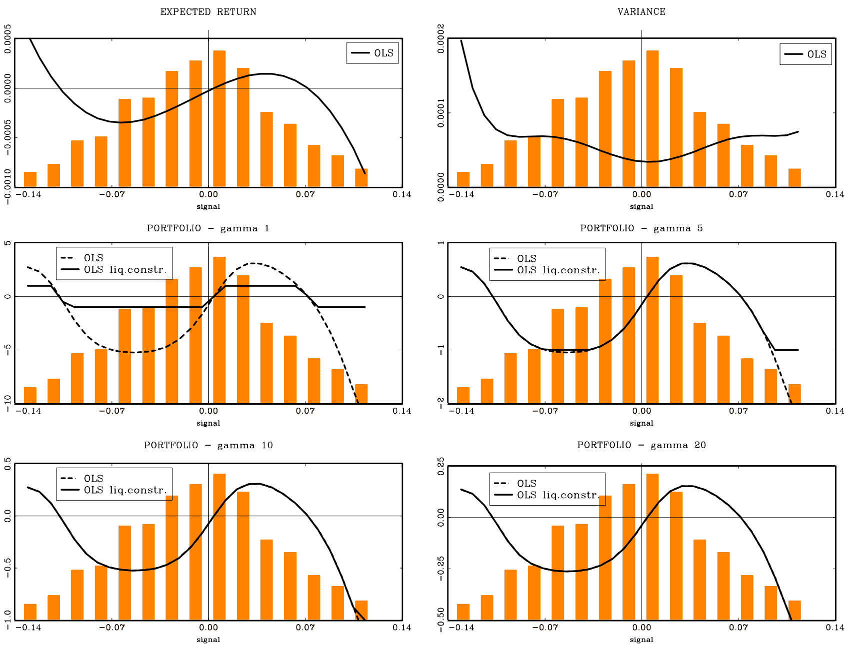

Figure 3: Daily expected return, variance and portfolio choice for the CHF/USD using the full sample (TTR 3: K=40, L=200)

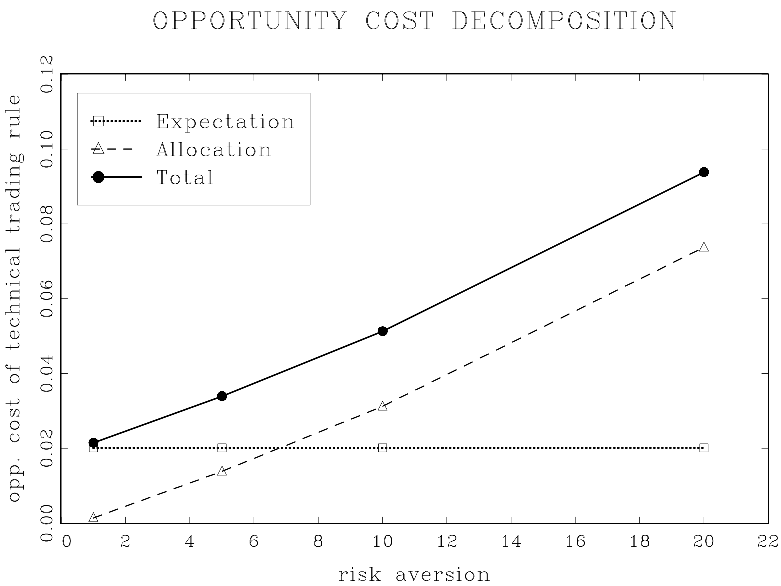

Figure 4: Annual opportunity cost decomposition for the GBP/USA for the out-of-sample period (TTR 2: K=20, L=100)