

Behavioral Forecasting

Harshit Singh, Puneet Chhabra, Rachit Prasad, Sudeep Tandon

INTRODUCTION

Since the 1970s, the Rational Expectations Efficient Market model (REEM) has dominated thinking about movements of asset prices, in particular, exchange rates. It is based on two propositions - first, price changes can only occur because of unexpected movements (news) in underlying fundamental variables (inflation, output growth, interest rates, etc.); and second, the link between prices and fundamentals is a stable one. These models also assume that each agent who trades in the market has access to all the information that can affect an asset's price, and that he/she will trade in exactly the manner which maximizes his/her utility.

Although REEM models like those in [1] continue to remain popular, they have frequently failed to explain empirical observations. A large body of research [2], [3], [4] has found a disconnection between changes in exchange rates and changes in fundamentals. Various empirical anomalies like 'excess volatility' have also been uncovered ([5], [6]), which again reject the idea of a stable link between rates and fundamental economic indicators.

The approach employed in this model is based on the above body of work, and is in particular an extension of [11]. The model proposed in [11] was used to explain bubbles and crashes in historical foreign exchange data. Our model seeks to tap the forecasting power of the same model by forecasting price movements in stock and foreign exchange markets. In any freely traded market, movements of price will be determined by the way all investors have responded (traded) in the previous period. The problem of prediction then reduces to having a model that can capture dynamics of trading at any time.

The model assumes *bounded rationality*, whereby an agent uses simple heuristics to process information and determine his/her set of actions. This view is supported by a growing literature of behavioral finance theory. The 'fitness' of these heuristics is then controlled *ex post* by checking their

Behavioral Forecasting

profitability, and by switching to the more profitable rules. The switching mechanism we use is the same as proposed in [11].

BACKGROUND

Observations from exchange rate data suggest two major departures from the perfect rationality model. One, it is not practical for an agent to have access to all information and at the same time, and two, be able to efficiently process it. There also exist differences in how investors view the same information, and in how they expect that to affect the asset's price. In the aforementioned case, agents do not play their respective equilibriums, and deviations from perfect rationality start to significantly affect the dynamics of the market.

Empirical failures of REEM models have led to new attempts to model exchange rates. One such approach recognizes that among heterogeneous market agents, there exist differences in beliefs about the behavior of exchange rates. These different beliefs introduce nonlinear features in the dynamics of exchange rates. This approach was initiated in [7], which studied the anomalous behavior of dollar prices between June 1984 and February 1985. It proposed a theory of existence of traders with conflicting forecasting techniques, namely 'chartists' and 'fundamentalists', whose interplay causes the 'excess volatility' in markets. If there exist chartist-like traders who tend to forecast by extrapolating recent trends, then their actions can exacerbate swings in the exchange rate. The work also presents empirical evidence of traders changing their strategies depending on whether they were making profits or losses. A survey of forecasting firms between 1978 and 1985 [7] revealed that most firms with a fundamentalist strategy shifted to a chartist strategy when they started making losses during 'excess volatility' from the period 1984 onwards. Dynamic shifts over time in the weight given to different forecasting techniques are a source of change in demand for asset, and hence of a change in their prices. [8], [9], [10] further the aforementioned work to develop a model in the context of stock prices.

Behavioral Forecasting

MODEL DESCRIPTION

The model we use is an extension of the work by Grauwe et. al. [12]. Investors are divided into two broad classes – fundamentalists, who trade based on their notion of “fair value” of the asset, and chartists, who trade based on historical price movements. The dynamic interaction of these heterogeneous agents causes observable price changes. However, investors are rigid in the strategies that they employ but learn continuously depending on relative profitability of their rules.

The fundamentalist expects the change in price to be towards his notion of fair value.

$$E^f(\Delta_{t+1}S) = -\alpha (S_t - S^*) \dots \dots \dots (1)$$

where

f corresponds to fundamentalists,

S_t is the asset price at time t ,

S^* denotes the fundamental or fair value of the asset,

α is the speed of mean-reversion

In the above formula, the fundamentalist expects next period’s price to move towards the fundamental, or in other words, mean-revert. The left hand side is the symbolic representation for the change from today to tomorrow (when using daily forecasting) that the fundamentalist expects.

A model is also necessary to describe what the chartist expects the change in price to be. We implement a trend-and-wave model as shown below. This closely resembles many aspects of ‘technical analysis’ used by traders in the market.

$$E^c(\Delta_{t+1}S) = a_0 + b_0t + a_1\sin(b_1t + c_1) + a_2\sin(b_2t + c_2) \dots \dots \dots (2)$$

The parameters of the model are (a_i, b_i) ($i = 0, 1, 2$). Thus, the chartist projects the trend and cyclicity forward to forecast the price. This is a significant departure from the model proposed for the chartist in [12].

Behavioral Forecasting

After making the prediction, both classes of investors are assumed to place bets in the market independently. The change in prices caused by investors is expected to be

$$E(\Delta_{t+1}S) = w_f * E^f(\Delta_{t+1}S) + (1 - w_f) * E^c(\Delta_{t+1}S) \dots \dots \dots (3)$$

This gives us the forecast of the model. The expected change in the price is represented as a linear combination of changes expected by both kinds of agents. w_f is the weight given to the fundamentalist's expectation. Equation (3) is linear but the model is non-linear because of the inherent switching model used to calculate w_f .

$$w_f = \exp(\gamma P_f) / [\exp(\gamma P_f) + \exp(\gamma P_c)] \dots \dots \dots (4)$$

where

γ is the learning parameter,

P_f is the risk-adjusted cumulative profit made by the fundamentalists over a training period

P_c is the risk-adjusted cumulative profit made by the chartists over a training period.

Equation (4) can be better understood by the following graph:

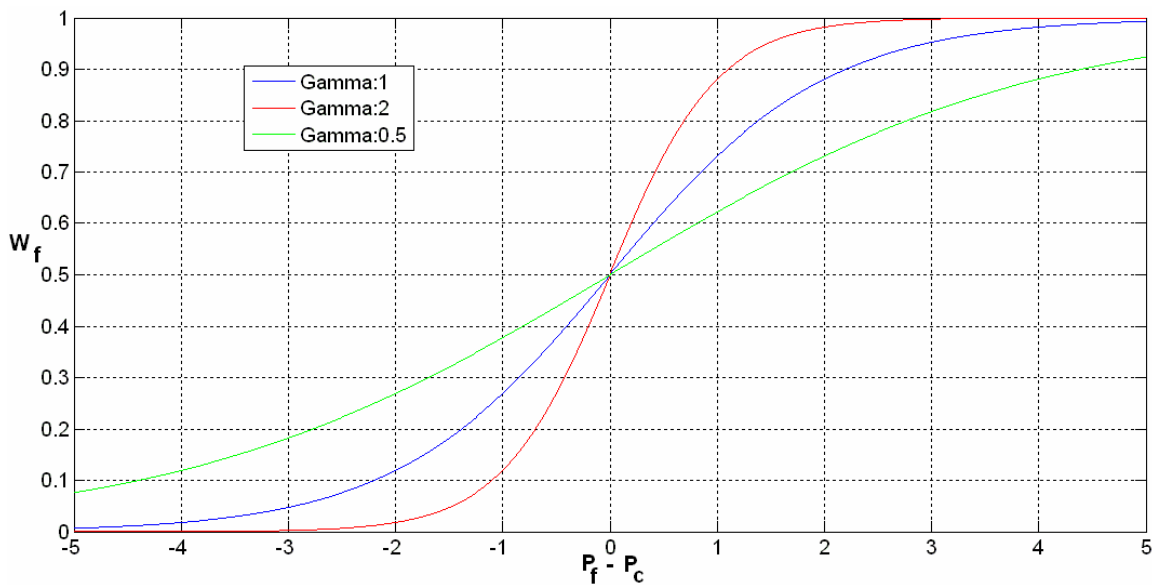


Fig 1. Weight dependence on relative profitability

X-axis: relative profitability of the fundamentalist with respect to the chartist.

Y-axis: weight given to the fundamentalist prediction

Behavioral Forecasting

We note that as the relative profitability of the fundamentalists increases, so does the weight of the fundamentalists in the prediction. When the profits made by the fundamentalists exceed those of the chartists by more than a certain amount, the model gives almost full weight ($w_f = 1$) to the fundamentalists. The same argument holds true from the chartists' perspective since $(1-w_f)$ is the weight given to the chartist's prediction.

The three curves illustrate the effect of the learning parameter γ . As γ increases, we observe that the more profitable strategy gets more emphasis. This γ reflects the rate of switching that occurs between the investor classes. The switching model therefore encapsulates the intended effects of profits and is symmetric with respect to the investor classes.

Risk-adjusted cumulative profits are calculated as follows:

$$P_f = \Sigma D_f - \mu\sigma_f \dots\dots\dots(5)$$

$$P_c = \Sigma D_c - \mu\sigma_c \dots\dots\dots(6)$$

where

subscripts f and c correspond to fundamentalists and chartists respectively,

ΣD denotes the sum of daily profits

σ denotes the standard deviation of profits made by the fundamentalists over a training period.

μ ($0 \leq \mu \leq 1$) is a risk aversion parameter

μ penalizes investors if the profits made are not consistent. Although a particular investor class may make more profits, switching to the more profitable strategy may not occur rapidly if the profits are very volatile.

Behavioral Forecasting

Figure 2 further illustrate the switching mechanism.

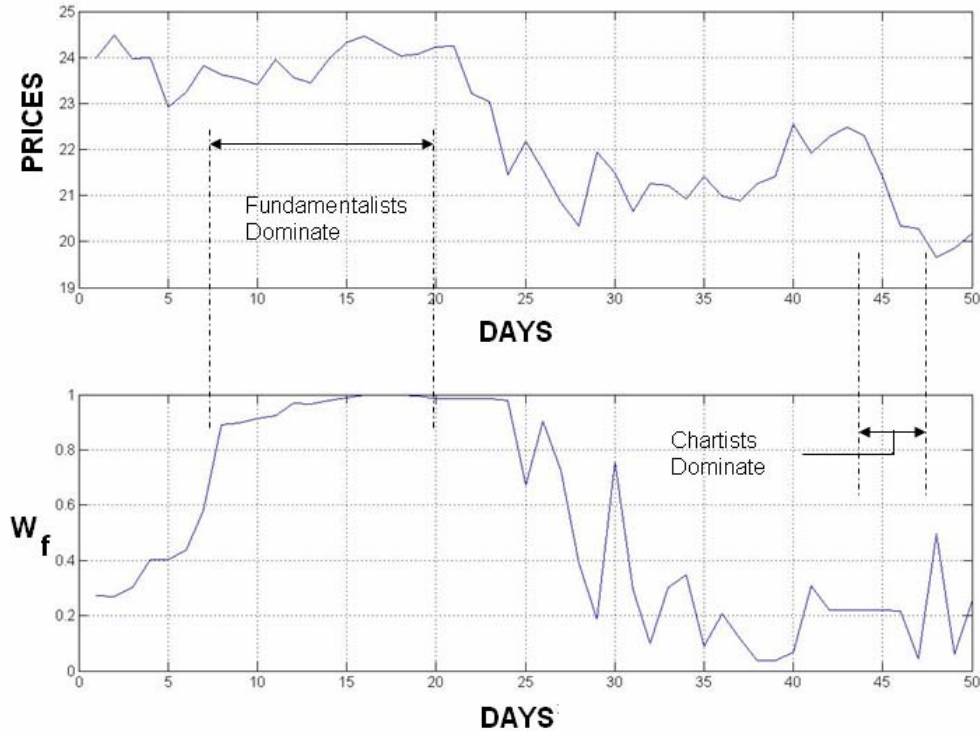


Fig 2: Illustration of switching between investor classes

The top plot is an arbitrary price series and the bottom plot is the way the weight of the fundamentalists changes in equation (3) to give the minimum error in prediction. We see that the fundamentalists dominate when the chartists' fit given in equation (2) is not accurate enough. This means that when the market is choppy, the chartists cannot get a good estimate of the parameters in equation (2) and so make less profit than a trader who trades based on intrinsic value. However, when the chartist model begins to make profits, he dominates the market. This is seen in the second half of the price series when the decreasing price trend resembles the negative of a sin-curve which the chartist catches on.

Training and Prediction

In this model α, γ, μ , and S^* are unknown parameters. We determine these parameters by calibrating the model over a training period of 15 days. We let the model run over an in-sample of 15 days on a known market realized price and search for the set of parameters that minimizes in-sample error of prediction of the model. For each day, the chartists looks back a period of 15 days to make a forecast, and chooses

Behavioral Forecasting

the parameters of equation (2) based on the 15 day look-back. Similarly the fundamentalists short the asset if the current price is above the ‘fair value’ and are long the asset if it is below. Thus, over a period of 15 days, we train the model by calculating the error in prediction.

Error of prediction for a single day during the training window is given by the difference between actual change and expected change.

$$\varepsilon = (\Delta_{t+1}S) - E(\Delta_{t+1}S) \dots \dots \dots (7)$$

During the training window, we have an error of prediction for each day, and we optimize the model to minimize the 2-norm of the error vector. The parameters to use for the forecast are determined from this training process. The forecasts are then made using equation (3).

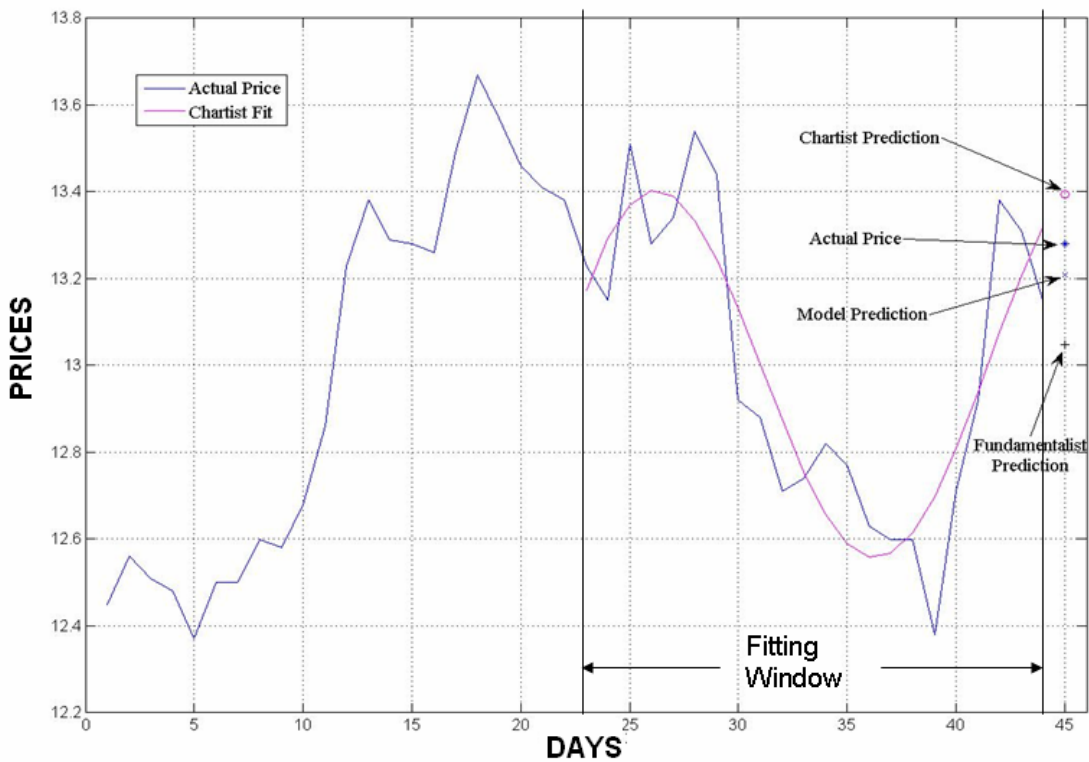


Fig 3: Illustration of training and forecasting

The training and forecasting process is explained by the Figure 3, where the training window used is of 22 days. The model here intends to make a prediction for the 45th day. The fundamentalists expect the price to move down because the price of the 44th day is higher than the fundamental value. The chartists

Behavioral Forecasting

looks back 22 days and fit equation (2). The way the chartists look at historical realized price is indicated by the sinusoidal curve. Using this, they forecast for the 45th day. The model predicted price derived from equation (3) is between the predictions of the chartists and the fundamentalists. This is observed to be close to the actual price itself. The difference between the actual price and model predicted price gives the error for the prediction on the 45th day. The optimization finds the optimal model parameters which minimize the 2-norm of the error vector. These parameters are then used in equation (3) to make the final prediction. The prediction for the 46th day is then made by moving the model forward by one day and following the above procedure.

Optimization

It is possible that the objective function may have multiple local optima. We therefore use a two step process:

- 1) The parameter space is divided into a grid of points
- 2) Each point on the grid is then chosen as a starting point and a quasi-Newton (BFGS) algorithm is used to find the local optima for this starting point.

The above procedure finds the local optimum that has the best fit over the training period. This enables us to be reasonably certain that the optimization procedure gives near-optimal results.

Behavioral Forecasting




Fig 4: Illustration of chartist and fundamentalist predictions

DATA

To explore the strengths and weaknesses of the model it was comprehensively tested over two different asset classes. Microsoft stock prices (NASDAQ: MSFT) from 04/28/1986 – 09/28/1989 were chosen to represent the stock asset class, while USDJPY exchange rates from 01/02/1975 – 11/15/1985 were chosen to represent the FOREX asset class. The choice of the specific tickers was driven primarily by availability by clean data over a sufficiently long period of time.

RESULTS AND DISCUSSION

The base model was tested first on the USDJPY price series, by using a training window length of 15 days, and with no transaction costs. The results shown in Figure 5, depicts the actual price series and the cumulative profits made over duration of 1200 days. Although, losses are seen for some days, the profits made for the remaining period are able to overcome those losses, resulting in a net gain. The model also comfortably beats the buy-and-hold strategy for the price series.

Behavioral Forecasting

Figure 6 shows the histogram for the daily returns of the model on USDJPY data. The returns are slightly biased towards the positive side, and on an average 6.32% of annualized returns are observed.

Figure 7 shows the plots for cumulative profits for both the investor classes – chartists and fundamentalists - on their own, and the overall model cumulative profits. Clearly, at any point in time, the model selects the better of the two strategies, and generates higher profits than either of them. This plot gives an insight into the working of the model, where through the mechanism of dynamic weight updates the agents switch to the more profitable strategy and the model combines the two strategies in the most profitable way.



Fig 5 Base Model – Cumulative model profits and price series for USDJPY for the period
01/02/1975 – 09/26/1979

Behavioral Forecasting




Fig 6 Base Model – Histogram of Daily Returns for USDJPY for the period 01/02/1975 – 09/26/1979

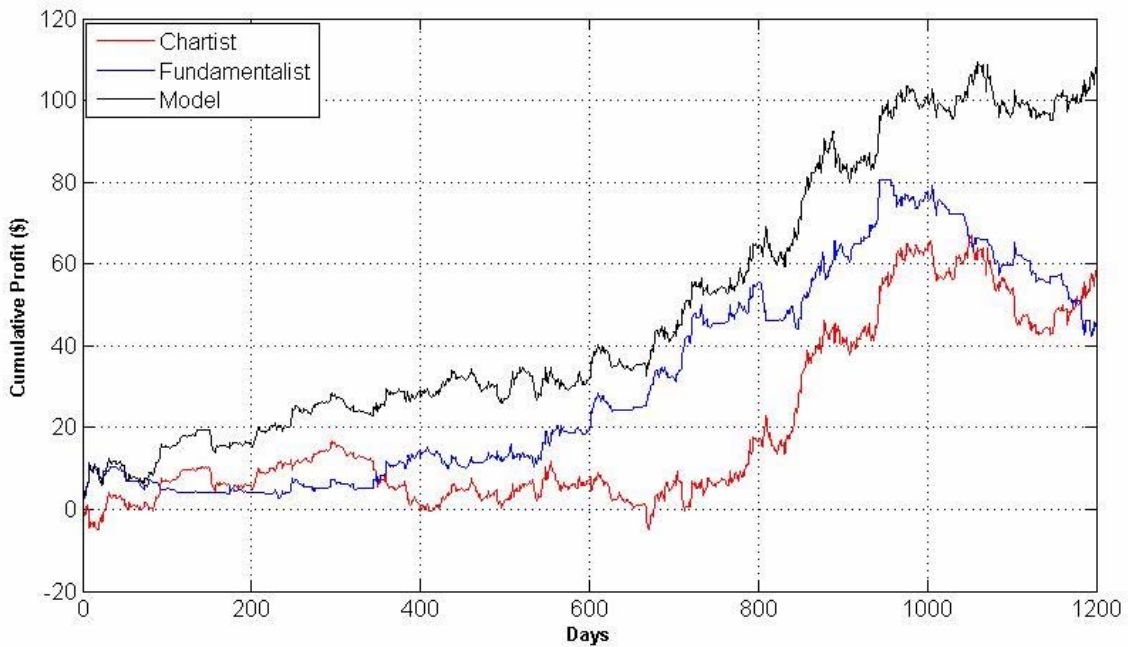


Fig 7 Base Model – Cumulative Profits of Investor Classes and combined model for USDJPY for the period 01/02/1975 – 09/26/1979

Behavioral Forecasting

Figure 8 shows results from the run on MSFT stock price data with no transaction costs and window length = 15. As can be seen, the profits made beat the buy-and-hold strategy for most part of the time-period under consideration.

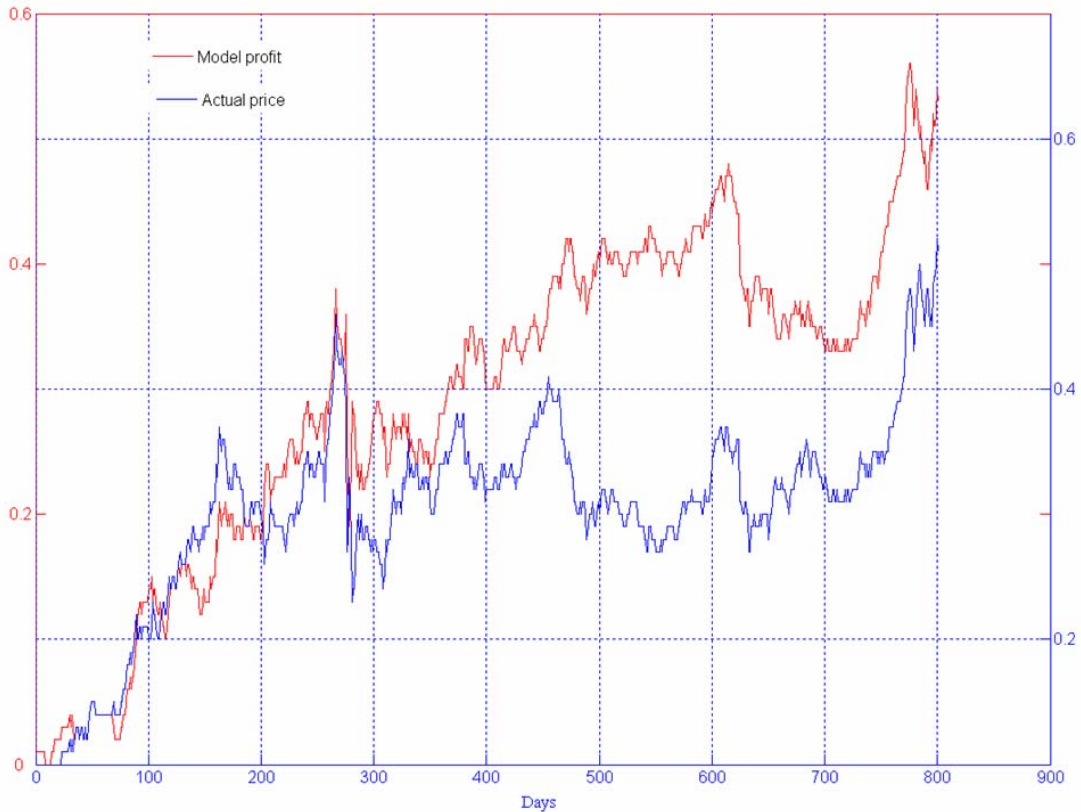


Fig 8 Base Model – Cumulative model profits and price series for MSFT for the period 04/28/1986 – 09/28/1989

EXTENSIONS

Various modifications were made to the base model in order to explore opportunities for performance improvement. Though the modifications did not always lead to significant improvements, but they helped in providing important insight into the working of the model. The major extensions are discussed below.

1. Binary Decision Model (BDM)

In contrast to the base model where the market agents endeavor to predict the exact asset price for the next period, in BDM the agents just predict the direction of the price movement, i.e. the forecasts comprise of whether prices in the next period will go up or down with respect to the current price. This approach is intuitively appealing as all that matters in making a buy or sell decision in the current period is the expected direction of price movements.

BDM aggregates the opinion of the two market agents through their respective weights and then decides whether to take a long or short position on the asset for the next period. The optimal parameters for the model are determined by minimizing number of wrong direction predictions rather than forecasting errors. The training procedure and computation of the weights is done in a way similar to the base model.



Fig 9 Binary Decision Model – Cumulative model profits and price series for USDJPY for the period 09/2000 – 06/2002

Figure 9 summarizes the cumulative profits earned by using BDM with USDJPY as the input price series. The model clearly beats the buy and hold strategy over the run period with an

Behavioral Forecasting

annualized return of over 16%. The hit rate for BDM (as defined earlier) across different price series was observed to be consistently higher than that from the base model.

2. Constant Parameter Model (CPM)

In this variant of the base model, the model parameters are fixed values based on empirical data. In particular, $\alpha = 0.1$, $\mu = 0.2$; $\gamma = 3.0$. The reason for this choice of parameters is the intuitive appeal behind these values. In particular, it is quite clear that α should have a low value since fundamentalist cannot expect the price to return to the fundamental price in 1 day. Similarly, value for γ higher than 3 might lead to rapid switching in a day, which is against intuition. We also take fundamental price at any day as the average of previous 100 days.

The CPM was tested on the similar USDJPY price series, as used for previous models. The resulting profits are shown in Figure 10. As can be seen, profits were made for the shown duration of around 700 days, the model beats the buy and hold strategy. Although, profits made were not as high as produced by the base model and binary model, which demonstrates the importance of optimization strategy that we have been using in the other models.




Fig 10 Constant Parameter Model – Cumulative profits and price series for USDJPY

Behavioral Forecasting

3. Moving Average Model

Another variation that was tried was to incorporate a moving average strategy over the base model. The strategy was to follow the action suggested by base model, when the moving average of cumulative profits for past m days was greater than the corresponding value for the past n days, where $n > m$. Otherwise, the base model's decision is reversed. The results from using this strategy were not better than those from using the base model.

4. Changing Window-Length

Comprehensive tests were performed on various price series of different asset classes, by changing the window length used for training the model to 22. For all the cases, we found that the results produced by using window length of 15 outperform the ones using that of 22. Hence, all the results reported here were generated by using window length = 15.

REFERENCES

- [1] Dornbusch, R., 1976. *Expectations and exchange rate dynamics*. Journal of Political Economy 84.
- [2] Meese, R., Rogoff, K., 1983. *Empirical exchange rate models of the seventies: Do they fit out of sample?* Journal of International Economics 14, 3–24.
- [3] Goodhart, C., Figliuoli, L., 1991. *Every minute counts in the foreign exchange markets*. Journal of International Money and Finance 10, 23–52.
- [4] Faust, J., Rogers, J., Swanson, E., Wright, J., 2002. *Identifying the effects of monetary policy shocks on exchange rates using high frequency data*. International Finance Discussion Papers, No. 739, Board of Governors of the Federal Reserve System, Washington, DC.
- [5] Baxter, M., Stockman, A., 1989. *Business cycles and the exchange rate regime - Some international evidence*. Journal of Monetary Economics 23 (May), 377–400.

Behavioral Forecasting

- [6] Flood, R., Rose, A., 1995. *Fixing the exchange rate regime: A virtual quest for fundamentals*. Journal of Monetary Economics 36 (August), 3–37.
- [7] Frankel, J. A., Froot K. A., 1990. *Chartists, Fundamentalists, and Trading in the Foreign Exchange Market*. American Economic Review 80, No. 2, 181–185.
- [8] Kirman, A., Teyssie`re, G., 2002. *Microeconomic models for long memory in the volatility of financial time series*. Studies in Nonlinear Dynamics and Econometrics 5, 281–302.
- [9] Brock, W., Hommes, C., 1998. *Heterogeneous beliefs and routes to chaos in a simple asset pricing model*. Journal of Economic Dynamics and Control 22, 1235–1274.
- [10] Gaunersdorfer, A., Hommes, C., Wagener, F., 2003. *Bifurcation routes to volatility clustering*. CeNDEF Working Paper 03-03, University of Amsterdam.
- [11] Grauwe, P., Grimaldi, M. *Exchange rate puzzles: A tale of switching attractors*. European Economic Review 50 (2006) 1–33.
- [12] Grauwe, P., Grimaldi, M. *The Exchange rate in a behavioral finance framework*. Princeton University Press (2006)