An Empirical Examination of Heterogeneity and Switching in Foreign Exchange Markets[a]

David Goldbaum[b]

Remco C.J. Zwinkels[c]

February 2013

Abstract

In order to study the expectation formation of financial institutions in the foreign exchange market we develop and apply a recursive selection and estimation algorithm to a dataset of surveyed foreign exchange market expectations. Responses are classified into two groups and forecasting models are endogenously determined within the groups. Estimation results reveal that a fundamentalist-chartist model is capable of explaining a large portion of foreign exchange market expectations. Allowing panelists to switch between models significantly improves the fit of the model, especially at the relatively shorter forecast horizons. We find that the fundamentalist model is increasingly used as the forecast horizon extends. Finally, results indicate that model choice is based on a combination of period-specific and individual-specific determinants.

Keywords: Heterogeneity, Discrete Choice, Foreign Exchange, Survey expectations

JEL Codes: F31, C35

---

[b] School of Finance and Economics, University of Technology Sydney; PO Box 123 Broadway; NSW 2007 Australia; email: david.goldbaum@uts.edu.au.
[c] Erasmus School of Economics, Erasmus University Rotterdam; PO Box 1738, 3000DR, Rotterdam, The Netherlands; email: zwinkels@ese.eur.nl.

1. Introduction

Economics abounds with scenarios in which agents make decisions based on predictions of the future value of endogenously determined variables. This is particularly true in financial markets. It is quite common for competing forecasting models to coexist, each with its own adherents, possibly with shifting popularity over time. The empirical objective of this investigation is to determine whether systematic heterogeneity in forecasts exists consistent with the use of multiple identifiable models, whether market participants do engage in model switching, and what determines the decision to switch to an alternative model. The exercise will be conducted using data on the forecasts of the foreign exchange spot price submitted by financial market institutions active in the markets.

Behavioral finance has documented numerous examples in which psychological factors influence financial decision making; see Barber and Odean (2013) for a recent overview of the literature. A distinction is typically made between biases in preferences (i.e., deviations from the traditional Von Neumann-Morgenstern Expected Utility Theory) and biases in beliefs, or expectations (i.e., deviations from rational expectations). Whereas prospect theory (Kahneman and Tversky, 1979) is a proper alternative for modeling preferences, there is no consensus on an alternative for modeling expectations. All we know is that individual expectations deviate from rationality (see, e.g., Cavaglia et al, 1994) and that there are a number of documented biases in expectation formation, such as overconfidence (Huisman et al., 2012) and wishful thinking (Ito, 1990).

Stepping away from the rationality approach introduces a large number of degrees of freedom. A substantial body of literature in economics and finance therefore models investors as heterogeneous and adaptive. The heterogeneity in expectations allows the interaction between traders behaving differently to impact the market. The heterogeneity can exist in a market at equilibrium or may keep the market out of equilibrium. The adaptation allows traders to select behavior appropriate for the perceived, possibly changing, market setting. The sensitivity of the market to the behavior of the traders can produce market destabilizing feedback loops.

Despite wide application, there is surprisingly little micro-level evidence on the empirical validity of adaptive heterogeneity. This paper contributes to the still emerging literature that empirically tests for the presence of adaptive heterogeneity and estimates adaptive heterogeneous agent models. The estimation employs the reported forecasts of

financial market institutions on major exchange rates over a variety of horizons and currencies. The analysis seeks evidence of heterogeneity in the models across financial institutions to generate their forecasts at a given point in time. Additionally, the analysis seeks evidence of model switching by individual institutions over time. Procedures are developed to address empirical challenges encountered in the analysis. Finally, we study which type of determinants trigger panelists to switch between models.

Grossman and Stiglitz (1980) establish a theoretical foundation supporting the sustainable co-existence of fundamental and market-based trading strategies. The market-based traders are fully rational and their presence is based on their ability to extract costly information from the price at a cost advantage. A market-based strategy can also survive based on superior performance relative to a fundamental strategy, as in Goldbaum and Panchenko (2010). Dynamics arise as traders switch between trading strategies, as is the case in Brock and Hommes (1997, 1998) where past relative performance determines popularity. The model employs the random element in the discrete choice model of Manski and McFadden (1981) to create heterogeneity in the individual-level choice among the available options. The environment highlights the potentially inherent instability of markets as the minority strategy performs better.

Heterogeneous adaptive agent models provide structure and insight to explanations for market phenomena. Simulations based on such models generate empirical phenomena replicating features of actual market data; examples for the foreign exchange markets are De Grauwe and Grimaldi (2005, 2006), De Grauwe and Markiewicz (2008), and Spronk et al. (2013). Less prevalent in the literature are direct empirical tests of the features that drive the agent-based models.

Only a handful of papers have sought to estimate adaptive heterogeneous agent models directly, including the Boswijk, Hommes, and Manzan (2007) determination that trader switching between trend following and mean reverting strategies contributes to swings in the S&P 5000. Evidence of heterogeneity can also be found in the MacDonald and Marsh (1996) survey of market participants documenting the heterogeneity of beliefs and the employment of different models. Branch (2004) finds evidence of adaptive heterogeneous behavior based on survey respondents' reported inflation forecasts.

Adaptive heterogeneity contributes to the empirical modeling of a number of financial market phenomena. Goldbaum and Mizrach (2008) use a model of adaptive heterogeneity to understand to allocation of new wealth into mutual funds. De Jong,

Verschoor and Zwinkels (2009) and (2010) find evidence of behavioral heterogeneity in setting equity prices across multiple markets and in foreign exchange rates respectively. Frijns, Lehnert, and Zwinkels (2010) find allowing multiple investor strategies simultaneously important for modeling the pricing of options. Markiewicz (2012) use model uncertainty among traders to explain shifts in volatility in foreign exchange markets.

Evidence in favor of switching has also been found at the individual level in experimental settings. Experiments involving market entry decisions often find a wide range of strategies have been employed by the participants that still combined to bring the market to the equilibrium number of entrants. Hommes et al (2005, 2007) identify four rule-of-thumb strategies employed by participants in a financial market setting rewarding conformity in expectation formation. Bloomfield and Hales (2002) show that even when forecasting a variable known to follow an exogenous random walk, participants switch between the simple rules of trend extrapolation and mean reversion.

A number of issues remains unresolved or are in need of empirical support. The current project also seeks to examine markets for evidence of adaptive heterogeneity and also to determine whether there is evidence in favor of the fundamentalist – chartist dichotomy in foreign exchange markets. As with Branch (2004), the current project seeks to model the reported forecast of survey participants and thus we use a direct measure of individual investor expectations.[1] This is in contrast to efforts to infer expectations from market realizations, such as Frankel and Froot (1990). Two features distinguish the current investigation from that of Branch (2004). First, currency markets offer an environment with strong direct positive feedback between market behavior and participant beliefs. This is especially true given the fact that the survey responses are from the large financial institutions that dominate the foreign exchange market. Second, Branch (2004) estimates the three alternative forecast models on the realized inflation data, imposing the resulting rules on the survey respondents. We estimate our alternative forecasting models on the reported forecasts themselves. This introduces endogeneity in the sorting of forecasts between the model alternatives and the estimation of the model which is accommodated for in the empirical analysis.

In our case, the data being employed is the exchange rate forecasts collected from participating international financial institutions. Each period includes forecasts for

---

[1] Assuming that the survey response is an unbiased proxy for the respondent's expectations.

the Euro/Dollar and Yen/Dollar exchange rates over the one, three, and twelve month forecast horizons for a number of individual institutions. Using the same data, Jongen et al. (2012) show that expectations are dispersed, and that panelists base expectations on fundamentalist/chartist types of considerations. Our results indicate that a combination of fundamentalism and chartism is indeed applied by the survey participants. Forcing panelists to be either fundamentalist or chartist over the full sample period does not improve the model fit. Allowing panelists to switch between the two models, however, does significantly improve the fit. The latter is especially true for the relatively shorter forecast horizons. Finally, we find that panelists use a combination of period specific and individual specific determinants for the switching decision.

The remainder of the paper is organized as follows. Section 2 presents the underlying model. Section 3 introduces the survey data used in the empirical section and Section 4 explains the empirical methodology applied. In Section 5 we present the results, and Section 6 concludes.

2. Forecasting Models and Model Choice

2.1 Fundamental model

A population of foreign exchange market participants, labeled N^F, consider the spot market exchange rate to be anchored by an underlying fundamental value. Holding beliefs consistent with the foreign exchange market described in Frankel and Froot (1990), these "fundamentalists" allow for realized deviation in the spot rate from the fundamental. Let e_t and f_t represent the date t logarithm of the spot and fundamental market exchange rates respectively. As modeled in Mark (1995), the fundamentalists presume the spot market rate tends to revert to the fundamental value such that over a k period horizon,

$$\Delta e_t^{t+k} = \beta_k (f_t - e_t) + v_{t+k,t}. \tag{1}$$

The notation Δe_t^{t+k} represents spot market innovation $e_{t+k} - e_t$ and β_k captures the perceived rate of reversion, with $\beta_k \in [0,1]$. Allow that f_t follows a driftless random walk so that

$$f_{t+1} = f_t + \eta_{t+1} \tag{2}$$

with η_t as the random fundamental innovation term.

Using notation $\beta = \beta_1$ and $v_{t+1} = v_{t,t+1}$, for $k = 1$,

$$\Delta e_t^{t+1} = \beta(f_t - e_t) + v_{t+1}. \tag{3}$$

Rearrange (3) to reveal that the value of e_{t+1} reflects a weighted combination of the previous spot price and the most recently realized fundamental,

$$e_{t+1} = (1-\beta)e_t + \beta f_t + v_{t+1}. \tag{4}$$

Using (2) and (3), we obtain a recursive expression of Δe_t^{t+1},

$$\Delta e_t^{t+1} = (1-\beta)\Delta e_{t-1}^t - v_t + \beta\eta_t + v_{t+1}, \tag{5}$$

revealing that the spot market exchange rate follows a predictable path of adjustment, shocked each period by two random processes. The innovation in the spot rate from t to $t+1$ includes a β weighted partial adjustment to the most recently observed innovation to the fundamental rate, η_t. Additionally, the future v_{t+1} causes e_{t+1} to deviate from the predictable adjustment towards f_t. The predictable element of (5), $(1-\beta)\Delta e_{t-1}^t - v_t$, is the mechanism through which prior shocks continue to influence current innovations. In the absence of shocks, the previously determined components would, over time, complete the previous partial adjustments to fundamental innovations and dissipate the influence of the transitory errors. [2]

Consistency between (1) and (5) implies $\beta_k = 1 - (1-\beta)^k$ and

$$v_{t+k,t} = \sum_{i=0}^{k-1}\{(1-\beta)^i v_{t+k-i} + [1-(1-\beta)^{i+1}]\eta_{t+k-1-i}\}. \tag{6}$$

As reflected in the coefficients in (6), the transitory uncertainty originating from the v_t induced deviation dissipate away while the persistent η_t innovations to the fundamental value accumulate over the forecast horizon.

As with the 1-period innovation, Δe_t^{t+k} can also be expressed recursively. The k-period equivalent of (5) is

$$\Delta e_t^{t+k} = (1-\beta)\Delta e_{t-1}^{t+k-1} - v_t + \beta\sum_{i=0}^{k-1}\eta_{t+i} + v_{t+k} \tag{7}$$

Expressed in terms of the observable spot and the underlying fundamental rates, using the notation $\hat{x}_t = E_{t-1}^F(x_t)$, (5) and (7) become, respectively,

---

[2] Without $-v_t$ in (5), the originally positive impact of v_t on e_t would be perpetuated into the current and future innovations through $(1-\beta)\Delta e_{t-1}^t$. The $-v_t$ reverses the sign of the influence and initiates the reversal of its influence on the future spot rates.

$$\Delta e_t^{t+1} = (1-\beta)\Delta e_{t-1}^{t} - (e_t - \hat{e}_t) + \beta(\Delta f_{t-1}^{t}) + (e_{t+1} - \hat{e}_{t+1}) \tag{8}$$

and

$$\Delta e_t^{t+k} = (1-\beta)\Delta e_{t-1}^{t+k-1} - (e_t - \hat{e}_t) + \beta(\Delta f_{t-1}^{t+k-1}) + (e_{t+k} - \hat{e}_{t+k}). \tag{9}$$

2.2 Chartist model

A second population of market participants, referred to as "chartists," employs market-based information to indicate future innovations in the spot market exchange rate. Let N^C represent the chartist population. The chartists employ a trend following model of exchange rate innovation believing that

$$\Delta e_t^{t+1} = \alpha(\Delta e_{t-1}^{t}) + \phi_{t+1} \tag{10}$$

with ϕ_t as the random innovation term. Over a k-period horizon,

$$\Delta e_t^{t+k} = \alpha_k(\Delta e_{t-1}^{t}) + \phi_{t,t+k} \tag{11}$$

is consistent with (10) for $\alpha_k = \alpha\left(\dfrac{1-\alpha^k}{1-\alpha}\right)$ and $\phi_{t,t+k} = \sum_{i=0}^{k-1}\dfrac{1-\alpha^{i+1}}{1-\alpha}\phi_{t+k-i}$.

2.3 forecasting

Let $E_{i,t}^{h}(\Delta e_t^{t+k})$ indicate the forecast of participant i where $h = F$ if $i \in N^F$ and $h = C$ if $i \in N^C$. The $i \in N^F$ participant develops a forecast model, $E_{i,t}^{F}(\Delta e_t^{t+k})$, based on (9), the product of which is

$$E_{i,t}^{F}(\Delta e_t^{t+k}) = (1-\beta)E_{i,t-1}(\Delta e_{t-1}^{t+k-1}) + \beta\Delta\hat{f}_{t-1}^{t} - (e_t - E_{i,t-1}(e_t)) + \varepsilon_{i,t,k}^{F}. \tag{12}$$

In (12), participant i's actual forecast from $t-1$, $E_{i,t-1}(\Delta e_{t-1}^{t+k-1})$, replaces the unobserved Δe_{t-1}^{t+k-1} in (9). Similarly, the participant's own forecast $E_{i,t-1}(e_t) = E_{i,t-1}(\Delta e_t) - e_{t-1}$ replaces $\hat{e}_t$. Also, the fundamentalist, while adherent to the notion of a fundamental exchange rate, does not get to observe its value. The market participant is thus forced to estimate it from other sources. Let $\hat{f}_t$ represent the fundamental market participants' belief about the value of f_t. Recognizing that $E_t(\Delta f_t^{t+k-1}) = E_t\left(\sum_{i=1}^{k-1}\eta_{t+i}\right) = 0$ and $E_t(e_{t+k} - \hat{e}_{t+k}) = E_t(v_{t+k}) = 0$ completes the transition from (9) to (12) with $\varepsilon_{i,t,k}^{F}$

capturing the idiosyncratic component of the k-period forecast horizon, $E(\varepsilon^F_{i,t,k}) = 0$ and $E(\varepsilon^F_{i,t,k}\varepsilon^F_{j,t,k})$, $j \neq i$.

There are two sources of heterogeneity among those employing the fundamental model. The idiosyncratic term, $\varepsilon^F_{i,t,k}$, captures trader specific differences between the forecasts of individual traders. These can be seen as the result of private information not otherwise captured by the model, deviation in the objective function from the presumed utility function, deviations resulting from heterogeneity in the estimate of the fundamental rate, or simply the result of randomness in the traders forecasting method. The presence of $e^F_{i,t,k}$ contributes to the second source of heterogeneity, the individual-specific choice patterns that perpetuate different individual forecasts histories appearing in the first and third terms on the RHS of (12).

Based on (11), participant $i \in N^C$ forecasts

$$E^C_{i,t}(\Delta e^{t+k}_t) = \alpha_k \Delta e^t_{t-1} + \varepsilon^C_{i,t,k} \tag{13}$$

with $\varepsilon^C_{i,t,k}$ capturing individual forecast deviation from the predicted model and the only source of heterogeneity in among the chartist population.

2.4 Model choice and estimation

Equations (12) and (13) identify two distinct methods for forecasting exchange rate innovation employing different foundational information. Variations in how these two forecasts models are combined and estimated indicate whether the data support different forms of heterogeneity and switching of methods among the financial institutions.

The benchmark model incorporates the two information sets, fundamental and chartist, into a single equation,

$$\hat{E}^B_{i,t}(\Delta e^{t+k}_t) = (1-\beta)\Delta e^t_{t-1} + \beta\Delta\hat{f}^t_{t-1} - (e_t - E_{i,t-1}(e_t)) + \alpha_k \Delta e^t_{t-1}. \tag{14}$$

Implicit in (14) is the notion that financial institutions incorporate both information sets in forming their forecasts. The model is misspecified in the presence of heterogeneity or switching between information sets.

The alternative empirical models incorporate heterogeneity. Forecasts are presumed to originate from the two distinct forecasting models; the fundamentalist model and chartist model. The models are evaluated in two different employment

settings. The forecast of the individual financial institution can be expressed as originating from

$$E_{i,t}(\Delta e_t^{t+k}) = \theta_{i,t}\hat{E}_t^F(\Delta e_t^{t+k}) + (1-\theta_{i,t})\hat{E}_t^C(\Delta e_t^{t+k}) + \varepsilon_{i,t} \quad (15)$$

where $\hat{E}_t^F(\Delta e_t^{t+k})$ and $\hat{E}_t^C(\Delta e_t^{t+k})$ are the fitted components of (12) and (13) respectively. That is,

$$\hat{E}_{i,t}^F(\Delta e_t^{t+k}) = (1-\beta)E_{i,t-1}(\Delta e_{t-1}^{t+k-1}) + \beta\Delta\hat{f}_{t-1}^t - (e_t - E_{i,t-1}(e_t)), \quad (16)$$

$$\hat{E}_{i,t}^C(\Delta e_t^{t+k}) = \alpha_k\Delta e_{t-1}^t. \quad (17)$$

The two different settings consider two different models of classification, each of which is capture by the process determining the value of $\theta_{i,t}$.

Classification 1 - Static model:

In this variation $\theta_{i,t} = \theta_i \; \forall t$ and $\theta_i \in \{0,1\}$. Thus, the forecast model used by financial institution i throughout the sample is either the fundamental model or the chartists model. This configuration allows for heterogeneity between institutions but no adaptation, i.e., no time-variation in the forecast method. The value of θ_i is determined by the average relative proximity of the forecast to the two models,

$$\theta_i = \begin{cases} 1 \text{ if } \sum_{t=1}^{T}\left(E_{i,t}(\Delta e_t^{t+k}) - \hat{E}_{i,t}^F(\Delta e_t^{t+k})\right)^2 < \sum_{t=1}^{T}\left(E_{i,t}(\Delta e_t^{t+k}) - \hat{E}_{i,t}^C(\Delta e_t^{t+k})\right)^2 \\ 0 \text{ if } \sum_{t=1}^{T}\left(E_{i,t}(\Delta e_t^{t+k}) - \hat{E}_{i,t}^C(\Delta e_t^{t+k})\right)^2 \leq \sum_{t=1}^{T}\left(E_{i,t}(\Delta e_t^{t+k}) - \hat{E}_{i,t}^F(\Delta e_t^{t+k})\right)^2. \end{cases} \quad (18)$$

Classification 2 –Dynamic model:

In this second variation, $\theta_{i,t} \in \{0,1\}$. The forecast model used by financial institution i in period t is either the fundamental model or the chartists model. Hence, we allow institutions to switch between models over time. The value of $\theta_{i,t}$ is determined by the relative proximity of the forecast to the two models,

$$\theta_{i,t} = \begin{cases} 1 \text{ if } \left(E_{i,t}(\Delta e_t^{t+k}) - \hat{E}_{i,t}^F(\Delta e_t^{t+k})\right)^2 < \left(E_{i,t}(\Delta e_t^{t+k}) - \hat{E}_{i,t}^C(\Delta e_t^{t+k})\right)^2 \\ 0 \text{ if } \left(E_{i,t}(\Delta e_t^{t+k}) - \hat{E}_{i,t}^C(\Delta e_t^{t+k})\right)^2 \leq \left(E_{i,t}(\Delta e_t^{t+k}) - \hat{E}_{i,t}^F(\Delta e_t^{t+k})\right)^2. \end{cases} \quad (19)$$

In the remainder of the paper, we estimate the benchmark (13), static (14-17), and dynamic (14-16, 18) models on a dataset of survey expectation, and determine which gives the best representation of the survey responses.

3. Data

3.1 Survey Expectations

To investigate the behavioral aspects of the forecasts of market participants, we use a unique database of survey-based exchange rate forecasts. The individual forecasts are obtained from a survey conducted by Consensus Economics of London on a monthly basis among leading market participants in the foreign exchange market, investment banks, and professional forecasting agencies. Examples of panelist companies are Morgan Stanley, Oxford Economic Forecasting, Deutsche Bank Research, and BNP Paribas. The full list of panelists is given in the Appendix. The panelists companies are located worldwide, although they are all from developed economies. The forecasts are point forecasts for a large set of currencies against the U.S. dollar and are available for horizons of 1, 3 and 12 months ahead. The names of the panelist companies are revealed.

Although survey participants have a few days to return their forecasts, we learned that the vast majority send their responses by e-mail on the Friday before the publication day, which is typically the second Monday of the month. We consider this Friday to be the day on which the forecasts are formed and assume that the beliefs are translated one-to-one in a point forecast. To verify that the information sets of market participants are not too diverse, all of the analyses throughout this study were re-estimated using spot data from various days surrounding this Friday, yet the overall results remain virtually unchanged.

There may be reasons for panelists not to reveal their true beliefs. One motive may be that agents do not want to expose their (private) information to other market participants. This effect is mitigated by the reputation effect that this survey can have. When the names of the forecasters are given in the survey publication (as is the case with our data), agents have an incentive to formulate a response that is close to their true informed belief. Therefore, in the remainder of the paper we will assume that the survey responses are an unbiased proxy of the institutions' actual expectations, and will use the terms interchangeably.

In this study we use the forecasts for the Japanese Yen and the Euro[3] against the U.S. dollar (i.e., in foreign currency per US Dollar) from 31 unique respondents for the period of November 1995 through December 2007, which are 146 monthly observations.[4] This period is of particular interest since it contains several financial crises, the introduction of a single monetary currency unit, and several large changes in the level of the exchange rates. The panel is unbalanced since the response rate of the individual market participants is less than 100 percent and due to market participants leaving the panel and subsequently replaced by others. Analyses are done on the 1, 3, and 12 months forecasting horizon. The 1 month forecasts are also used as a control variable for the models of the 3 and 12 months horizon (see Section 2).

< Insert Table 1 Here >

Table 1 presents the descriptive statistics of the survey data. Respondents are consistent across currencies and forecast horizons in answering the survey; i.e., if they answer one, they answer all six. The median response rate per period is 70%, which results in an average total number of observations per currency/horizon pair of approximately 2,900.

The descriptive statistics of the expected exchange rate returns in panel b) indicate there is a wide variety in answers, ranging from -30 to +45%. Median expected returns are all slightly negative, suggesting a median expected depreciation of the US Dollar. The kurtosis indicates that, as is the case for market returns, the distribution of expectations is heavy tailed. The expected returns are strongly auto correlated, which is inconsistent with actual FOREX returns. Partly, the autocorrelation in expectations is due to overlapping observations, i.e., the frequency of the data is higher than the forecast horizon. The autocorrelations for the one-month forecast horizon, however, which do not suffer from the overlapping observations issue, suggest that panelists also rely heavily on previous period's expectation in forming current expectations.

---

[3] The database also contains U.K. Pound expectations. However, for unknown reasons, the U.K. Pound expectations are only reported every other month. We do not use this currency because of the limited number of observations.

[4] Prior to January 1999 forecasts on the Deutschemark versus the U.S. Dollar are used. These forecasts are transformed into Euro / U.S. dollar forecasts using the official conversion rate.

3.2 Fundamental Value

The fundamentalist expectation, given by (12), includes a term representing the (partial) accommodation to perceived changes in the fundamental value, $\Delta \hat{f}_{t-1}^t$. Important subsequent question is the econometrician's choice for calculating the panelist's perceived fundamental exchange rate $\hat{f}_t$. The issue here is not necessarily to select an estimate that represents the true fundamental exchange rate, but a model that is appropriate for capturing the panelist's fundamental-based forecasts of the exchange rate. We propose a version of the monetary model introduced by Mark (1995) given by

$$\hat{f}_t = (m_t - m_t^*) + (y_t - y_t^*) \tag{20}$$

Here, m_t is the home money supply, m_t^* the foreign money supply, y_t the home income, and y_t^* the foreign income. The choice for this model is based on two arguments. First, the study by Mark (1995) is well known and one of the few persuasive studies illustrating the forecasting power of a fundamental model. Second, this fundamental value is relatively simple to implement; that is, it does not require any further estimation. This ensures that no additional choices or assumptions regarding the estimation process are required.

Data wise, all non-survey based data is retrieved from Datastream. M2 is used for the money supply m_t and industrial production for income y_t. Figure 1 displays the (log) exchange rates and fundamental exchange rates for both the Yen and the Euro.

< Insert Figure 1 Here >

4. Methodology

Unique to the current examination (to our best knowledge) is the fact that different models under consideration are endogenous to the traders employing them. Branch (2004), in contrasts, considers three exogenous models of inflation. Branch' (2004) naïve expectation model has $\pi_{t+1}^e = \pi_t$. The two more sophisticated models are a model of adaptive expectations and a VAR. In both cases, the parameters of the model are chosen to fit realized inflation. Thus, the models are optimized to minimize the error of the forecast of inflation rather than to capture the model employed by the forecaster.

In the presence of multiple forecasting models, each forecast should be employed to estimate only the model used to generate it. Unfortunately, the information on the model being employed to generate the forecast is not available to the researcher. The methodology outlined below includes a procedure for simultaneously estimating multiple models from the observed forecasts and classifying the forecasts according to the estimated models.

The approach in this examination is to have those traders employing the model indicate the parameters of the model[5]. By doing so, we do not impose (nor rule out) that forecasters minimize a certain forecast error; we rather give a direct description of expectation formation. This is accomplished by choosing the parameters to minimize the mean squared error of the forecast by those traders who employ the forecast. This involves some degree of simultaneity as the estimation of the model depends on how the individuals are sorted and the sorting depends on the model. Our solution is to estimate, sort, and then re-estimate over a number of iterations until the sorting and the model parameters settle.

Specifically, the estimation procedure for the discrete choice model is as follows. First, the fundamental and chartist models are estimated separately for each institution in a single equation using OLS. The initial distribution of agents over groups is done by estimating the two expectation formation models (12) and (13) individually per respondent. Based on best fit, given by either (18) or (19), each respondent or each observation is subsequently classified as either fundamentalist or chartist.[6] Next, the two rules are estimated using OLS in a single equation in a pooled setup, given by (15) – (17), using the initial distribution of respondents as values of $\theta_{i,t}$. The distribution of respondents across groups, i.e., $\theta_{i,t}$, is subsequently updated based on the new estimation results using (18) or (19), and the equation is re-estimated. This procedure is repeated until convergence, i.e. until respondents do not change groups anymore and coefficient estimates of the rules are constant. Generally, this occurs within ten iterations, conditional on the complexity of the model. As such, the classification of agents and the

---

[5] The finite mixture estimation methodology is closely related to ours (see Beard et al., 1991). Finite mixture, however, extracts the degree of heterogeneity in the constituents in a data set from the aggregated values. In our case, we have access to the individual expectations. Focusing solely on the aggregate expectation, therefore, would be inefficient use of information.

[6] There exists a certain path dependency conditional on the initial distribution of agents. Throughout, we perform a grid search over initial conditions to find the global optimum; we report the solution with the best fit.

actual expectation formation rules are being updated endogenously in the iteration process.

For the simultaneous estimation of the two models, the parameter constraints implied structurally by (12) are imposed in the empirical estimation of the two heterogeneous classification models (such as β and $(1-\beta)$). Given the endogenous allocation of observations between the two forecasting models, allowing the parameters to be determined freely runs the risk of enabling the parameters to adjust to alter the allocation of observations to the incorrect model to better fit the data. Once misallocated, the erroneously assigned forecast would contribute to a biased estimation of the final model parameters.

The autocorrelation in the residuals due to the overlapping data issue is captured by the construction of the models (12) and (13). Any remaining auto-correlation is accounted for by calculating Newey-West standard errors, as in MacDonald (2000).

5. Results

5.1 Estimation Results

Table 2 presents the results for the benchmark, static, and dynamic models estimated on the Euro/USD and Yen/USD exchange rate forecasts.

< Insert Table 2 Here >

At each forecast horizon and for both currencies, the benchmark model is found to provide some power to explain the forecasts, which is an indication that both the chartist and the fundamentalist derived information are relevant for the panelists. The exception is the 1-month forecasts for both currencies, for which we only find a significant α_k.[7] The fit of the benchmark model increases with the forecast horizon, indicating the increased use of the proposed models as the horizon extends.

The chartist coefficient α_k is negative and highly significant for all currency-horizon combinations in the benchmark model. A negative α_k implies that panelists expect a reversion of previous returns and therefore act like contrarians. The reversion increases on the longer horizons, which is explained by the fact that we include the 1-

---

[7] The slightly negative adjusted R^2 is caused by the term $(E_{i,t-1}(e_t) - e_t)$ in the fundamentalist model.

period lagged return for the *k*-period-ahead forecast. This finding is consistent with the literature on survey expectations; see MacDonald (2000) for an overview.[8] The estimated α_k's for the Yen are somewhat smaller than for the Euro, suggesting that contrarian behavior is stronger in the Yen/USD market than the Euro/USD market.

The fundamentalist coefficient β is, apart from the 1-month horizon, always positive and significant for the benchmark model. A positive β implies expected mean reversion towards the fundamental value. Note, that this finding serves two important purposes: 1) our fundamental value estimate is a reasonable approximation of the fundamental value perceived by panelists, and 2) the panelists expect mean-reversion towards this imposed fundamental value. When moving from the 3 to the 12-month forecast horizon, the estimated β decreases for both currencies. This suggests panelists' greater reliance on previous forecasts and reduced reliance on new fundamental information when forecasting over longer horizons than is consistent with the 1-period horizon forecast.

The second column, labeled static model, splits the sample of panelists in two groups, fundamentalists and chartists, based on best fit using (18) as selection mechanism. The assumption is that panelists remain of a certain type throughout the sample period. A number of differences arise when comparing the results of the benchmark models to those of the static models. First, with regards to the fundamentalists' reversion β, the effect size becomes notably stronger and significant for the 1-month horizon for both currencies. As for the chartist coefficient α_k, the effect size is also stronger compared to the benchmark model but the significance levels are equally high.

The percentage of panelists that uses the fundamentalist model in forming expectations increases as the forecast horizon increases. The increase is substantial for both currencies, with an increase from 3.1 to 95.5% for the Euro and from 7.7 to 92.6% for the Yen. Given that we have 31 unique panelists, this implies that only one panelist applies the chartist model for the 12-month horizon. Interestingly, this is the same panelist for both currencies: General Motors. This result suggests that panelists are consistent in their choice of model over different currencies. Hence, the choice of model

---

[8] On shorter horizons (less than one month), survey expectations are found to be of the momentum type. Apparently, the 1-month horizon is perceived to be the long run.

contains individual specific determinants, and not only currency or time specific determinants.

The image that arises when comparing the R^2 of the benchmark model to that of the static model is mixed, though. In three out of six cases, the model fit decreases. Given the increased flexibility of the model, this is an indication that the static model should not be seen as a proper description of the expectation formation process of panelists.

The final columns in Table 2, labeled dynamic, report the results of estimating the model captured by (16) and (17) in which panelists are allowed to update their forecasting strategy each period according to (19). Hence, instead of considering the average distance between the rule and the expectation over the full sample period, the selection procedure is applied per period and panelists are allowed to switch between the two models.

The increased flexibility changes the estimation results substantially. The chartist contrarian behavior becomes stronger for both currencies and all horizons. The effect size of the fundamental model β tends to decrease somewhat, apart from the 1-month Euro/USD expectations.

The percentage of periods in which the fundamental model is employed increases with the forecast horizon, as was the case with the static model, though the increase is in a narrower band ranging from 43.5 to 65.9% for the Euro and from 42.8 to 65.5% for the Yen. The use of the fundamental model is approximately equal across currencies. The autocorrelation in model choice increases as the forecast horizon increases, and ranges from 8 to 25%. This implies that panelist have less tendency to switch models for the longer horizons. In the 3- and 12-month horizons, the model choice in the Euro/USD expectations tends to be more persistent than the Yen.

The fit of the model is substantially higher in the dynamic case compared to the benchmark case. While the model fit did not give a clear improvement when moving from the benchmark model to the model with static classification, having dynamic classification does matter notably in the parameter estimates and the division of the population. At this point, however, we cannot conclude whether the dynamic model is *significantly* better than the benchmark model, because the additional number of degrees of freedom the dynamic model consumes is unclear. The next section will study this issue in more detail using a simulation setup.

The estimation results reported in Tables 2 suggest that the fundamentalist and the chartist forecasting rules are being used by the panelists in the survey. The fundamentalist-chartist dichotomy, as suggested in Frankel and Froot (1990) and subsequently often applied in the literature therefore appears a relevant classification, consistent with the findings of, among others, Allen and Taylor (1990, 1992), and Jongen et al. (2012). The flexibility of agents to change strategy is of great importance. There is a substantial improvement in the fit after introducing switching. This is direct evidence in favor of the heterogeneous agent models with switching, as introduced in Brock and Hommes (1997, 1998).

5.2 Selective power of the model

To ascertain the validity of the estimation results in Section 5.1 in general and the iterative selection procedure specifically, this section looks into the selective power of the procedure. As such, we proceed with two analyses. First, we look into the selective power of the methodology by checking the consistency of the selection mechanism. We divide the panel observations into two separate subsamples of the data set according to their identification as originating from either the fundamentalist or chartist forecasting model as determined by the estimated static and dynamic models. We proceed to then estimate anew both the fundamental model and the chartist model using each of the newly created subsamples in order to compare the fit of each model on each of the subsamples.

Second, we need to study whether the increase in R^2 of the dynamic model relative to the benchmark model is significant. The increase in model fit cannot be interpreted directly because it is unclear how many extra degrees of freedom the dynamic model consumes. For each observation, the selection procedure determines the optimal $\theta_{i,t}$. This, however, does not imply that the dynamic model consumes one additional degree of freedom per observation, because $\theta_{i,t}$ only takes the values zero or one. Furthermore, the $\theta_{i,t}$ are determined outside the regression equation. In other words, the benchmark model is not nested in the dynamic model.

To determine the significance of the dynamic model, we simulate the benchmark model for all six currency/horizon combinations. As such, we know that the data generating process of the simulated data is non-dynamic. Subsequently, we estimate both the benchmark and dynamic model on the simulated data and compare the increase in R^2 to the empirically observed increase in R^2. We run 1,000 simulations generating

the same number of observations as the empirical dataset (i.e., 31 respondents with 146 time-series observations). If the empirically observed increase in R^2 is exceeded by the simulations in less than 50 cases (5% of 1,000), we can conclude that the dynamic model gives a significantly better model fit than the benchmark model at the 5% confidence level.

Table 3 presents the results of the first method, the consistency checks, checking the discriminative power of the static model (i.e., selection based on 17). Table 4 similarly presents the results checking the discriminative power of the switching model (i.e., selection based on 18).

< Insert Table 3 Here >

< Insert Table 4 Here >

Given the qualitative similarity between Tables 3 and 4, we will discuss the results simultaneously. Focusing on the R^2 of the different Group/Model combinations, we observe a clear pattern. When estimated on the fundamental identified observations, the fit of the newly estimated fundamental model is considerably better than the fit of the chartist model. Likewise, when estimated on the chartist identified observations, the fit of the newly estimated chartist model is considerably better than the fit of the fundamental model.[9] These results are indications that the selection mechanism correctly discriminates between two subsamples of survey observations, consistent with the two models.

The R^2 of the benchmark model is generally equal to the R^2 or not significantly higher than the R^2 of the fundamental model for the subsample of fundamentalists. Hence, the chart model does not contain additional explanatory power for the fundamentalists.

The coefficients β and α_k are typically significant for both models per subsample. The difference between the coefficients estimated for the chartist subsample and fundamentalist subsample, however, are typically sizeable. The estimated β for the chartist subsamples is often larger than unity. This is inconsistent with the fundamental

---

[9] The negative R^2's for some of the estimations of the fundamental model are caused by the $(E_{i,t-1}(e_t) - e_t)$ term. Because this term does not have a coefficient, it introduces the possibility that the left hand side variable in the regression has a lower variance than the right hand side variables.

model.[10] These observations are another indication that the iterative selection procedure did a good job in discriminating between two subsamples of observations.

Having established that the empirical strategy contains selective power, i.e., it correctly classifies panelists and individual expectations of panelists as being of the fundamentalist or the chartist type, the subsequent question is whether the selection significantly increases the explanatory power of the model. Table 5 presents the results of the simulation study explained above.

< Insert Table 5 Here >

The top panel of Table 5 presents the empirical difference in R^2 between the benchmark model and the switching model, as observed in Table 2, as well as the distribution of differences in R^2 retrieved from the Monte Carlo simulation. As suggested above, if the empirically observed difference is exceeded in less than 50 out of 1000 simulations, we can conclude that the dynamic model adds significantly to the explanatory power of the model.

For both currencies for the 1 and 3- month horizons, the null-hypothesis of no difference in explanatory power between the benchmark and dynamic models is rejected at the 0.1% level. For the 12-month horizon, the null is rejected at the 5% level for the Euro, while it is rejected only at a 20.0% level for the Yen. Hence, for all but one currency/horizon pair, the results suggest that the panelists apply a dynamic model in forming expectations. Dynamics appear to play a larger role for the relatively short horizons. This is consistent with the observation from Table 2 that panelists use the fundamental model in the majority of cases with a relatively large persistence (AC); this already suggests that panelists do not often switch between models at the longest horizon.

Panels B and C of Table 5 examine the difference between the R^2 of the fundamental model and the R^2 of the chartist model for the fundamentalist subsample and chartist subsample, respectively. As in Panel A, we compare the empirically observed difference from Table 4 with the distribution of simulated differences.

Panel B indicates that the observed difference in R^2 is never matched by the simulations for the 3-month horizon. For the other horizons, the difference is not

---

[10] A $1 < \beta < 2$ implies fundamentalist believe in an oscillating, rather than smooth, convergence process in the spot rate towards the fundamental value.

significant. For the chartist group in Panel C, however, we again observe that the null of equal fit for both models is rejected at the 0.1% level.. Hence, the empirical consistency results for the dynamic model, as presented in Table 4, could not have been obtained from the benchmark data generating process. In other words, the true data generating process is significantly better captured by the dynamic model containing both fundamentalist and chartist elements.

5.3 Model choice

A natural subsequent question to ask is what are the determinants of model choice. We distinguish between two types of determinants, namely those specific to the individual panelist and those specific to the current market circumstance. The former causes a single panelist to choose a certain model across all currency/horizon combinations whereas the latter causes different individuals to choose a certain model for a single currency/horizon combination.

Within the heterogeneous agent literature, the agents switch based on past profitability of the models; see Brock and Hommes (1998) or De Grauwe and Grimaldi (2006) for the foreign exchange market. Because the profitability measure is equal for all agents, this is an example of a period specific determinant of model choice. Because the performance measure is conditional on the developments in a certain market, this causes model choice across individuals to be correlated for a specific currency/horizon combination. Also, if model choice is consistent in different forecast horizons, model choice based on period specific determinants induces correlation between the model choices within a currency for the different horizons.

< Insert Figure 2 Here >

To illustrate the period specific determinants of model choice, Figure 2 presents the evolution of the average percentage fundamentalists (i.e., the average percentage of panelists that chooses the fundamental model per period) for all six currency / horizon pairs. If model choice is not based on period specific determinants, this average per period should be roughly constant and changes in this average should be random. First of all, the figures illustrate that there is ample time variation in the choice of models. The percentages range from zero – all panelists apply the chartist model – to one – all panelists apply the fundamental model. Apart from short term patterns, we can also

discern longer term patterns, such as the decrease in the choice of fundamental model for the 12-month Euro forecasts over the first 25 periods. Interestingly, the same pattern can be seen for the 3-month Euro forecasts. More generally there are similar patterns in the averages, especially within a currency. In addition, it can be seen that the choice of fundamental model increases as the horizon increases, as was already concluded from the results in Table 2. All in all, the figures suggest that there are period specific factors triggering panelists to choose either model.

On the other hand, agents might have an unconditional preference for a certain model, regardless of its (time-varying) performance. Status quo bias (Kahneman and Tversky, 1974) induces people to prefer the current state of affairs, as any deviation from the status quo is perceived as a loss. This induces the choice of model to be independent of the performance of the models, even though the outcome might be detrimental to overall forecasting performance. Because the status quo is specific to the panelist, this is an individual specific determinant of model choice and introduces a correlation between model choices for one individual across different currency/horizon combinations.

< Insert Figure 3 Here >

Figure 3 illustrates the percentage of periods in which each individual panelist chooses the fundamental model. If there are no individual specific determinants of model choice, the averages per individual should be roughly equal and differences between individuals should be random. In Figure 3, there are clear differences between the average tendencies of panelists to choose the fundamental model, ranging from zero – always chartist – to one – always fundamentalist. Again we observe the increase in the use of fundamental model as the forecast horizon increases. Interestingly, we observe some similarities among all six figures. The most striking similarity across all figures is the continuously low score for panelist number six, which is Barclays Bank. Apart from the 12-month Euro forecast, Barclays Bank consistently scores as (one of the) panelist making least use of the fundamental model.

For a somewhat more formal analysis of the relation between model choices across time and across individuals, Table 6 presents the pooled, within period, and within individual correlations across currency and horizons between model choices.

< Insert Table 6 Here >

Panel A displays the pooled correlation between model choices; it thus combines the period- and individual- specific elements. The correlations between different horizons within a certain currency, in upper left and lower right quadrants, show highly significant positive correlations. For both the Euro/USD and Yen/USD forecast, there is a strong relation between the model choices for different forecast horizons. The correlations between the one and three month forecasts are highest (around 40%), followed by the correlations between the three and twelve month forecasts (around 26%). This is an indication that period-specific determinants matter for model choice, as choice is expected to be consistent across horizons.

The correlations between currencies, in the lower left quadrant, are also positive and oftentimes significant as well. The between currency correlations are typically close to 5%. The choices for the 1-month Euro forecasts are significantly correlated to the choices for all three Yen horizons. The twelve month Euro forecast is only related to the twelve month Yen forecasts. This is an indication that individual specific determinants matter for model choice.

Panel B displays the within period correlations; i.e., the correlations between the 146 average uses of the fundamental model per period as displayed in Figure 2. This indicates whether average model choice per period is correlated across the six currency/horizon combinations and is therefore an illustration of the importance of period specific determinants of model choice. Panel B yields a rather different image than Panel A. In Panel B, only the within currency correlations in the lower right and upper left quadrants are positive and significant. The between currency correlations are all insignificant. This implies that the period specific determinants that cause panelists to choose a certain model do not correlate strongly between the two currencies. The within currency correlations, though, are all substantially higher than we observed in Panel A. This is an indication that the within currency correlations between model choices are for an important part driven by period specific determinants. Hence, the external factor(s) causing panelists to choose a certain model at one horizon also causes panelists to choose that model at other horizons.

Panel C gives the within individual correlations, i.e., the correlations between the 31 time-series average choices per individual as displayed in Figure 3. This indicates whether an individual has the tendency to make the same choice across the six

22

currency/horizon combinations. Panel C generally gives the opposite result of Panel B: The within individual correlations in Panel C are typically high and significant for the between currency correlations in the lower left quadrant and less so for the within currency correlations (the latter especially for the Yen) in the upper left and lower right quadrants. The between currency correlations in the lower left quadrant reach up to 67% for the model choice between the 1-month Euro and 1-month Yen forecasts. Given that the between currency correlations are not significant in Panel B whereas they are in Panel A and in Panel C, we can conclude that the between currency correlations are driven by individual specific determinants.

The results show that panelists have certain idiosyncratic preferences for a certain model next to period specific determinants such as lagged profitability, and as a result have a higher tendency to apply that particular model for other forecasts they make, regardless the currency or horizon that is considered and regardless the current market conditions. If the panelists in our sample trade on their expectations, this might have important consequences for the market in general. Specifically, consider the chartist model. The chartist model always forecasts a contrarian pattern in the exchange rate. If agents decide to apply the chartist model across different currencies based on individual preferences, regardless the market conditions, this might induce excess cross-currency correlations.

6. Conclusion

A model has been developed to examine the behavior of financial institutions when forming forecasts of future exchange rate innovations over a variety of currencies and time horizons. The model allows for market participants to switch between different strategies for forming expectations. A model based on two strategies is developed, a fundamental strategy by which predictions concerning future exchange rates are based on exchange rate fundamentals and a chartist strategy by which market based information serves as a predictor of future exchange rates.

The empirical analysis suggests that the switching model is useful for explaining the heterogeneity in the forecasts of the different banking institutions that took part in the survey. Allowing the panelists to switch strategies during the sample period improved the fit of the model significantly, especially at the relatively short forecast horizons. In addition, we find that model choice is based on a combination of period specific and individual specific determinants. The former means that certain market

circumstances induce panelists to choose a particular model. The latter implies that if a panelist chooses a model for one currency/horizon combination, she has the tendency to make the same choice for the other currency/horizon combinations. It provides an attractive narrative of market behavior that is consistent with stylized facts.

Allen and Taylor (1990) document the use of chartist techniques among foreign exchange traders. Individual traders explain that it is not necessarily that they believe that charting captures fundamentals, but that the market can be driven by chartists since they are so plentiful in the foreign exchange markets. For this reason, it is important to include chartist tools when considering trades. Presumably, the same is true when forming predictions. The fact that forecasts of financial institutions appear to be driven, at times, by a chartist model may be a reflection of the fact that institutions believe that the market based information is informative about market innovations away from fundamentals. The results could also be considered supportive of the notion that market based information is useful for predicting fundamental innovations supported by private information not available to the modeler. The latter interpretation is consistent with Grossman and Stiglitz (1980) and other papers that argue in favor of the use of chartists techniques to extract information from the market.

Bibliography

Allen, H. and M.P. Taylor (1990). Charts, Noise and Fundamentals in the London Foreign Exchange Market, Economic Journal 100(400): 49-59.

Barber, B. and T. Odean (2013). The Behavior of Individual Investors, in: Constantinides, G. (eds), Handbook of the Economics of Finance, Elsevier.

Beard, T.R., S.B. Caudill, and D.M. Gropper (1991). Finite Mixture Estimation of Multiproduct Cost Functions, Review of Economics and Statistics 73(4): 654-664.

Bloomfield, R. and J. Hales (2002). Predicting the Next Step of a Random Walk: Experimental Evidence of Regime-Shifting Beliefs, Journal of Financial Economics 65: 397-414.

Boswijk, H.P., C.H. Hommes and S. Manzan (2007). Behavioral Heterogeneity in Stock Prices, Journal of Economic Dynamics and control 31: 1938-1970.

Branch, W.A (2004). The Theory of Rational Heterogeneous Expectations: Evidence from Survey Data on Inflation and Expectations. The Economic Journal 114: 592–621

Brock, W. and C.H. Hommes (1997). A Rational Route to Randomness, Econometrica 69: 1059-1095.

Brock, W. and C.H. Hommes (1998). Heterogeneous Beliefs and Routes to Chaos in a Simple Asset Pricing Model, Journal of Economic Dynamics and Control 22: 123-1274.

Cavaglia, S., W.F.C. Verschoor, and C.C.P. Wolff, (1994). On the Biasedness of Forward Foreign Exchange Rates: Irrationality or Risk Premia?, Journal of Business, 1994, 67: 321-43.

Frankel, J.A. and Froot, K.A., (1990). Chartists, Fundamentalists, and Trading in the Foreign Exchange Market. American Economic Review 80(2):181-185.

Frijns, B., T. Lehnert, and R.C.J. Zwinkels (2010). Behavioral Heterogeneity in the Option Market, Journal of Economic Dynamics and Control 34(11): 2273-2287.

Goldbaum, D. and B. Mizrach (2008). Estimating the Intensity of Choice in a Dynamic Mutual Fund Allocation Decision, Journal of Economic Dynamics and Control 32(12): 3866-3876.

Goldbaum, D. and V. Panchenko (2010). Learning and adaptation's impact on market Efficiency, Journal of Economic Behavior and Organization 76: 635-653.

De Grauwe, P. and M. Grimaldi (2005). Heterogeneity of Agents, Transaction Costs and the Exchange Rate, Journal of Economic Dynamics and Control 29: 691-719.

De Grauwe, P. and M. Grimaldi (2006). Exchange Rate Puzzles: A Tale of Switching Attractors, European Economic Review 50(1): 1-33.

De Grauwe, P. and A. Markiewicz (2008). Learning to Forecast the Exchange Rate: Two Competing Approaches, CESifo Working Paper 1747.

Grossman, S.J., Stiglitz, J.E., (1980). On the impossibility of informationally efficient markets. The American Economic Review 70(3), 393-408.

Huisman, R., N. Van der Sar, and R.C.J. Zwinkels (2012). A New Measurement Method of Investor Overconfidence, Economics Letters 114: 69-71.

Hommes, C.H., Sonnemans, J., Tuinstra, J. and Velden, H. van de, (2005), Coordination of expectations in asset pricing experiments, Review of Financial Studies 18, 955-980.

Hommes, C.H., Sonnemans, J., Tuinstra, J. and Velden, H. van de, (2007), Learning in cobweb experiments, Macroeconomic Dynamics 11 (Supplement 1), 8-33.

Ito, T. (1990), Foreign exchange rate expectations: Micro survey data, American Economic Review, 80(3): 434-449.

De Jong, E., W.F.C. Verschoor, and R.C.J. Zwinkels (2009). Behavioral Heterogeneity and Shift-Contagion: Evidence from the Asia Crisis, Journal of Economic Dynamics and Control 33(11): 1929 – 1943.

De Jong, E., W.F.C. Verschoor, and R.C.J. Zwinkels (2010). Heterogeneity of Agents and Exchange Rate Dynamics: Evidence from the EMS, Journal of International Money and Finance 29(8): 1652-1669.

Jongen, R., C.C.P. Wolff, W.F.C. Verschoor, and R.C.J. Zwinkels (2012). Explaining dispersion in foreign exchange expectations: A heterogeneous agent approach, Journal of Economic Dynamics and Control 36(5): 719 – 735.

Kahneman, D. and A. Tversky (1979). Prospect theory: An analysis of decision under risk, Econometrica 47: 263-291.

MacDonald, R. and I.W. Marsh (1996). Currency Forecasters are Heterogeneous: Confirmation and Consequences, Journal of International Money and Finance 15(5): 665-685.

MacDonald, R. (2000). Expectation Formation Risk in Three Financial Markets: Surveying what the Surveys Say, Journal of Economic Surveys 14(1): 69-100.

Manski, C.F., McFadden, D., (1981). Structural Analysis of Discrete Data with Econometric Applications (MIT Press, Cambridge, MA).

Mark, N.C. (1995). Exchange Rates and Fundamentals: Evidence on Long-Horizon Predictability, American Economic Review 85(1): 201-218.

Markiewicz, A. (2012). Model uncertainty and exchange rate volatility, International Economic Review 53, 815-844.

Spronk, R., W.F.C. Verschoor, and R.C.J. Zwinkels (2013). Carry trade and foreign exchange rate puzzles, European Economic Review, forthcoming.

Taylor, M.P. and H. Allen (1992). The Use of Technical Analysis in the Foreign Exchange Market, Journal of International Money and Finance 11(3): 304-314.

Appendix

Table A1: Panelists

1.	ABN AMRO Bank
2.	Bank of America
3.	Bank of Tokyo Mitsubishi
4.	Bankers Trust Company
5.	BNP Paribas
6.	Barclays Bank
7.	Barclay's Capital
8.	Chase Manhattan
9.	Citigroup
10.	Commerzbank
11.	Credit Suisse First Boston
12.	Deutsche Bank Research
13.	Dresdner Kleinwort Wasserstein
14.	General Motors
15.	Global Insight
16.	HSBC MIDLAND
17.	Imperial Chemical Inds
18.	Industrial Bank of Japan
19.	ING Barings
20.	JP Morgan Chase
21.	Merrill Lynch
22.	Morgan Stanley
23.	NatWest Group
24.	Nomura Research Institute
25.	Oxford Economic Forecasting
26.	Royal Bank of Canada
27.	SBC Warburg
28.	Societe General
29.	Standard Chartered Bank
30.	UBS Warburg
31.	Westdeutsche Lbank

Tables and figures

Table 1: Data

	Euro			Yen		
	1 Month	3 Months	12 Months	1 Month	3 Months	12 Months
a) # Observations						
Min. # panelists / period	12	14	14	13	15	15
Max. # panelists / period	24	24	24	24	24	24
Median # panelists / period	19	20	20	19	20	20
Min. # periods / panelist	3	12	12	3	12	12
Max. # periods / panelist	143	144	144	143	144	144
Median # periods / panelist	101	102	102	102	103	103
Total # observations	2825	2932	2930	2835	2941	2940
b) Descriptive statistics						
Median	-0.0036	-0.0118	-0.0397	-0.0030	-0.0060	-0.0343
Maximum	0.1136	0.1366	0.2110	0.1985	0.3331	0.4512
Minimum	-0.1700	-0.2036	-0.2883	-0.1667	-0.1667	-0.3050
Standard deviation	0.0246	0.0370	0.0686	0.0301	0.0473	0.0857
Skewness	-0.3430	-0.2080	0.0565	0.3216	0.3332	0.5532
Kurtosis	5.7290	3.9952	2.8560	6.1241	4.7625	4.3544
Autocorrelation (1^{st} lag)	0.3900	0.5090	0.7550	0.4170	0.5710	0.7410

Notes: Table presents the number of observations per period and per respondent (Panel a) as well as the descriptive statistics of the expected log-changes in the exchange rate, i.e. $\ln(E_{i,t}(e_{t+k})) - \ln(e_t)$ over all panelists and periods (Panel b).

Table 2: Estimation Results

	1 month			3 months			12 months		
	Benchmark	Static	Dynamic	Benchmark	Static	Dynamic	Benchmark	Static	Dynamic
					Panel A: EURO				
c^f	-0.0005*	-0.0024	0.0007	-0.0070***	-0.0060***	0.0038***	-0.0060***	-0.0055***	0.0022**
	(-1.6896)	(-0.8814)	(1.4442)	(-8.8630)	(-5.8113)	(6.2747)	(-5.2846)	(-5.6806)	(2.3790)
β	0.0497	0.7571***	0.7701***	0.8398***	0.7658***	0.5495***	0.2715***	0.2574***	0.1641***
	(1.0463)	(7.4665)	(38.821)	(35.804)	(32.011)	(35.095)	(17.329)	(21.005)	(14.285)
c^c		-0.0056***	-0.0074***		-0.0150***	-0.0235***		-0.0127***	-0.0629***
		(-11.680)	(-15.095)		(-17.0242)	(-35.0380)		(-3.3024)	(-48.655)
α_k	-0.2027***	-0.2860***	-0.4389***	-0.2548***	-0.4162***	-0.4295***	-0.4983***	-0.3356**	-0.5884***
	(-9.3149)	(-16.999)	(-26.886)	(-8.9567)	(-13.330)	(-20.086)	(-14.968)	(-2.4209)	(-17.436)
$Adj. R^2$	-0.0395	0.1035	0.5469	0.1571	0.1966	0.6322	0.6568	0.6292	0.7934
% fun		0.0305	0.4354		0.4366	0.5062		0.9546	0.6592
AC			0.0840			0.1711			0.3036
					Panel B: YEN				
c^f	-0.0003	0.0037*	-0.0019***	-0.0017**	-0.0010	-0.0002	-0.0065***	-0.0065***	-0.0004
	(-0.8836)	(1.8707)	(-3.5330)	(-2.0742)	(-0.8246)	(-0.2466)	(-4.9575)	(-5.8260)	(-0.3611)
β	0.0607	1.2845***	1.2097***	0.8362***	0.8012***	0.6315***	0.3129***	0.3032***	0.2033***
	(1.3734)	(21.559)	(61.071)	(41.772)	(35.656)	(33.852)	(17.6968)	(24.128)	(14.811)
c^c		-0.0025***	-0.0023***		-0.0064***	-0.0090***		-0.0492***	-0.0352***
		(-4.3826)	(-3.7285)		(-5.8980)	(-9.3190)		(-13.096)	(-21.863)
α_k	-0.3648***	-0.3504***	-0.4231***	-0.5127***	-0.4867***	-0.7956***	-0.6970***	-0.6311***	-0.8137***
	(-15.615)	(-20.597)	(-22.156)	(-17.563)	(-14.982)	(-30.179)	(-22.932)	(-5.4470)	(-21.582)
$Adj. R^2$	0.1316	0.1713	0.6232	0.3329	0.2530	0.6526	0.6782	0.6145	0.7820
% fun		0.0769	0.4279		0.4533	0.5313		0.9260	0.6545

Table 3: Estimation Results Comparison Static Model

	EURO						YEN					
Currency:												
Sample:	Fundamentalist			Chartist			Fundamentalist			Chartist		
Model:	Fundamental	Chart	Bench	Fundamental	Chart	Bench	Fundamental	Chart	Bench	Fundamental	Chart	Bench
1 MONTH												
c	-0.002	0.001	-0.001	-0.001**	-0.006**	0.000	0.004*	-0.012**	0.000	0.001***	-0.002**	0.000
	(-1.458)	(0.277)	(-0.583)	(-2.548)	(-7.254)	(-1.609)	(1.840)	(-2.521)	(0.087)	(2.993)	(-2.528)	(-0.952)
β	0.757***		0.033	0.998***		0.043	1.285***		-0.105	1.239***		0.076*
	(10.020)		(0.131)	(44.823)		(0.883)	(13.784)		(-0.692)	(49.252)		(1.664)
α_k		0.025	0.163**		-0.286***	-0.212***		-0.439***	-0.455***		-0.350***	-0.359***
		(0.289)	(2.134)		(-14.676)	(-9.727)		(-5.548)	(-5.419)		(-14.529)	(-14.770)
R^2	-0.096	-0.012	0.141	-0.305	0.105	-0.039	0.173	0.116	0.407	-0.228	0.164	0.080
3 MONTHS												
c	-0.006***	-0.012***	-0.006***	-0.008***	-0.015***	-0.008***	-0.001	-0.007**	0.000	-0.003**	-0.006***	-0.003**
	(-5.801)	(-5.379)	(-5.702)	(-6.970)	(-9.617)	(-7.333)	(-0.800)	(-2.554)	(-0.366)	(-2.642)	(-3.179)	(-2.622)
β	0.766***		0.772***	0.902***		0.910***	0.801***		0.811***	0.864***		0.867***
	(30.226)		(29.574)	(25.064)		(25.211)	(32.215)		(31.473)	(25.981)		(26.182)
α_k		-0.324***	-0.210***		-0.416***	-0.290***		-0.425***	-0.444***		-0.487***	-0.569***
		(-7.617)	(-5.226)		(-12.436)	(-7.469)		(-7.220)	(-9.320)		(-13.798)	(-16.130)
R^2	0.275	0.053	0.297	-0.033	0.114	0.022	0.336	0.074	0.417	0.029	0.149	0.233
12 MONTHS												
c	-0.006***	-0.039***	-0.006***	-0.010***	-0.013***	-0.010***	-0.007***	-0.030***	-0.006***	-0.024***	-0.049***	-0.023***
	(-4.642)	(-12.175)	(-5.050)	(-2.727)	(-2.433)	(-2.748)	(-5.009)	(-7.378)	(-4.500)	(-3.697)	(-8.167)	(-3.734)
β	0.257***		0.259***	1.131***		1.135***	0.303***		0.301***	0.626***		0.629***
	(18.435)		(18.780)	(5.863)		(6.149)	(17.842)		(17.516)	(5.288)		(5.028)
α_k		-0.639***	-0.522***		-0.336**	-0.080		-0.626***	-0.699***		-0.631***	-0.675***
		(-14.130)	(-17.270)		(-2.032)	(-0.351)		(-12.704)	(-22.090)		(-5.796)	(-6.267)
R^2	0.636	0.070	0.683	-0.197	0.046	-0.205	0.634	0.060	0.709	-0.148	0.111	-0.020

Notes: Table presents results for the consistency estimations; i.e., we estimated the two models on the two separate subsamples in order to judge the selective power of the selection mechanism. R^2 is adjusted R^2; Newey-West standard errors in parenthesis; *, **, *** represents significance at the 10, 5, and 1% level, respectively.

Table 4: Estimation Results Comparison Switching Model

Currency:	EURO						YEN					
Sample:	Fundamentalist			Chartist			Fundamentalist			Chartist		
Model:	Fundamental	Chart	Bench	Fundamental	Chart	Bench	Fundamental	Chart	Bench	Fundamental	Chart	Bench
1 MONTH												
c	0.001	-0.004***	0.001*	-0.003***	-0.007***	-0.003***	-0.002***	-0.005**	-0.002***	0.004***	-0.002***	0.000
	(1.418)	(-2.769)	(1.926)	(-4.224)	(-15.234)	(-4.576)	(-3.553)	(-2.647)	(-3.893)	(4.794)	(-3.785)	(0.518)
β	0.770***		0.563***	1.204***		-0.179***	1.210***		0.951***	1.261***		-0.377***
	(39.503)		(12.044)	(38.822)		(-2.815)	(60.837)		(21.278)	(31.701)		(-6.304)
α_k		-0.071**	0.172***		-0.439***	-0.445***		-0.261***	-0.237***		-0.423***	-0.430***
		(-2.271)	(8.278)		(-26.883)	(-16.701)		(-6.106)	(-11.511)		(-22.091)	(-13.522)
R^2	0.672	0.004	0.681	-1.559	0.357	-0.764	0.770	0.054	0.777	-1.828	0.356	-0.775
3 MONTHS												
c	0.004***	-0.004*	0.004***	-0.024***	-0.024***	-0.025***	0.000	-0.006**	0.000	-0.005***	-0.009***	-0.007***
	(6.281)	(-1.833)	(6.454)	(-21.882)	(-35.553)	(-24.224)	(-0.247)	(-2.174)	(0.147)	(-3.540)	(-9.617)	(-5.601)
β	0.550***		0.553***	1.327***		1.344***	0.631***		0.624***	1.135***		1.252***
	(34.904)		(35.371)	(36.338)		(36.064)	(33.871)		(33.706)	(28.868)		(38.009)
α_k		-0.348***	-0.122***		-0.429***	-0.403***		-0.094	-0.140***		-0.796***	-0.960***
		(-6.873)	(-4.005)		(-20.073)	(-11.869)		(-1.608)	(-4.750)		(-29.830)	(-27.460)
R^2	0.719	0.044	0.724	-1.087	0.261	-0.858	0.733	0.002	0.740	-0.653	0.496	0.054
12 MONTHS												
c	0.002**	-0.026***	0.002**	-0.053***	-0.063***	-0.053***	0.000	-0.029***	0.000	-0.031***	-0.035***	-0.031***
	(2.401)	(-6.493)	(2.329)	(-19.619)	(-50.595)	(-20.591)	(-0.366)	(-5.667)	(0.032)	(-11.665)	(-22.436)	(-13.809)
β	0.164***		0.169***	0.907***		0.885***	0.203***		0.198***	0.867***		0.897***
	(14.347)		(15.373)	(24.251)		(23.444)	(14.957)		(14.760)	(20.295)		(23.724)
α_k		-0.671***	-0.372***		-0.588***	-0.639***		-0.485***	-0.495***		-0.814***	-1.021***
		(-9.569)	(-11.269)		(-17.560)	(-11.711)		(-6.080)	(-15.902)		(-21.485)	(-20.303)
R^2	0.836	0.057	0.853	-0.755	0.260	-0.450	0.845	0.024	0.871	-0.543	0.352	0.011

Notes: Table presents results for the consistency estimations; i.e., we estimated the two models on the two separate subsamples in order to judge the selective power of the selection mechanism. R^2 is adjusted R^2; Newey-West standard errors in parenthesis; *, **, *** represents significance at the 10, 5, and 1% level, respectively.

Table 5 Simulation Results

	Euro			Yen		
	1 month	3 months	12 months	1 month	3 months	12 month
R^2 Switching Model - R^2 Benchmark Model						
Observed	0.586	0.475	0.137	0.492	0.320	0.10
Simulation:						
Max	0.181	0.358	0.146	0.145	0.273	0.13
3rd perc	0.154	0.310	0.113	0.116	0.217	0.10
2nd perc	0.147	0.292	0.105	0.111	0.204	0.09
1st perc	0.142	0.254	0.097	0.106	0.193	0.08
Min	0.123	0.197	0.069	0.086	0.149	0.05
Prob.	<0.000	<0.000	0.050	<0.000	<0.000	0.19
Fundamental Sample: R^2 Fundamental Model - R^2 Chartist Model						
Observed	0.668	0.675	0.779	0.716	0.730	0.82
Simulation:						
Max	0.687	0.610	0.845	0.738	0.662	0.82
3rd perc	0.663	0.563	0.829	0.713	0.639	0.80
2nd perc	0.653	0.461	0.824	0.706	0.631	0.80
1st perc	0.642	0.441	0.820	0.698	0.622	0.79
Min	0.528	0.381	0.799	0.566	0.398	0.76
Prob.	0.120	0.000	1.000	0.160	0.000	0.05
Chartist Sample: R^2 Fundamental Model - R^2 Chartist Model						
Observed	-1.916	-1.347	-1.015	-2.184	-1.149	-0.89
Simulation:						
Max	-0.270	0.287	-0.453	-0.434	0.164	-0.51
3rd perc	-0.691	0.175	-0.548	-0.618	-0.643	-0.59
2nd perc	-0.735	0.131	-0.569	-0.657	-0.659	-0.60
1st perc	-0.777	-0.638	-0.588	-0.698	-0.678	-0.62
Min	-0.949	-0.757	-0.657	-0.918	-0.739	-0.68
Prob.	<0.000	<0.000	<0.000	<0.000	<0.000	<0.00

Notes: This table presents the empirically observed differences in R^2 from Tables 2 and 4
distribution of differences in R^2 as generated by the simulations. Max represents the max
percentile; Min is the minimum. Prob. is the percentage of simulations (out of 1,000 runs)
observe a larger change in R^2 than empirically observed; this represents the probability
model is not the actual law of motion.

Table 6 Correlations between model choices

|		Euro			Yen		
		1 Month	3 Months	12 Months	1 Month	3 Months	12 Months
		Panel A: Pooled Correlations					
Euro	1 Month	1.000					
	3 Months	0.381***	1.000				
		(20.529)					
	12 Months	0.137***	0.260***	1.000			
		(6.870)	(13.383)				
Yen	1 Month	0.050**	0.020	0.003	1.000		
		(2.494)	(0.980)	(0.163)			
	3 Months	0.052**	0.044**	0.015	0.403***	1.000	
		(2.604)	(2.216)	(0.768)	(21.94)		
	12 Months	0.050**	0.053**	0.051**	0.134***	0.256***	1.000
		(2.477)	(2.637)	(2.530)	(6.746)	(13.21)	
		Panel B: Within Period Correlations					
Euro	1 Month	1.000					
	3 Months	0.548***	1.000				
		(7.843)					
	12 Months	0.336***	0.506***	1.000			
		(4.260)	(7.012)				
Yen	1 Month	0.030	-0.126	-0.057	1.000		
		(0.356)	(-1.513)	(-0.685)			
	3 Months	0.087	-0.069	-0.105	0.500***	1.000	
		(1.041)	(-0.823)	(-1.261)	(6.898)		
	12 Months	0.029	-0.010	0.128	0.169**	0.425***	1.000
		(0.345)	(-0.114)	(1.543)	(2.056)	(5.620)	
		Panel C: Within Individual Correlations					
Euro	1 Month	1.000					
	3 Months	0.514***	1.000				
		(3.230)					
	12 Months	0.693***	0.489***	1.000			
		(5.178)	(3.021)				
Yen	1 Month	0.672***	0.207	0.362**	1.000		
		(4.882)	(1.142)	(2.092)			
	3 Months	-0.047	-0.116	-0.263	0.256	1.000	
		(-0.251)	(-0.627)	(-1.466)	(1.427)		
	12 Months	0.632***	0.299*	0.291*	0.562***	0.033	1.000
		(4.393)	(1.686)	(1.637)	(3.657)	(0.175)	

Notes: This table presents the correlations between the model choices of panelists for the six currency / horizon combinations. Panel A gives the pooled correlations; Panel B gives the within period correlations (using the 146 averages over individuals); Panel C gives the within individual correlations (using the 31 averages over periods). T-values in parentheses; *, **, *** represents significance at the 10, 5, and 1% level, respectively.

Figure 1: Fundamental Exchange Rates
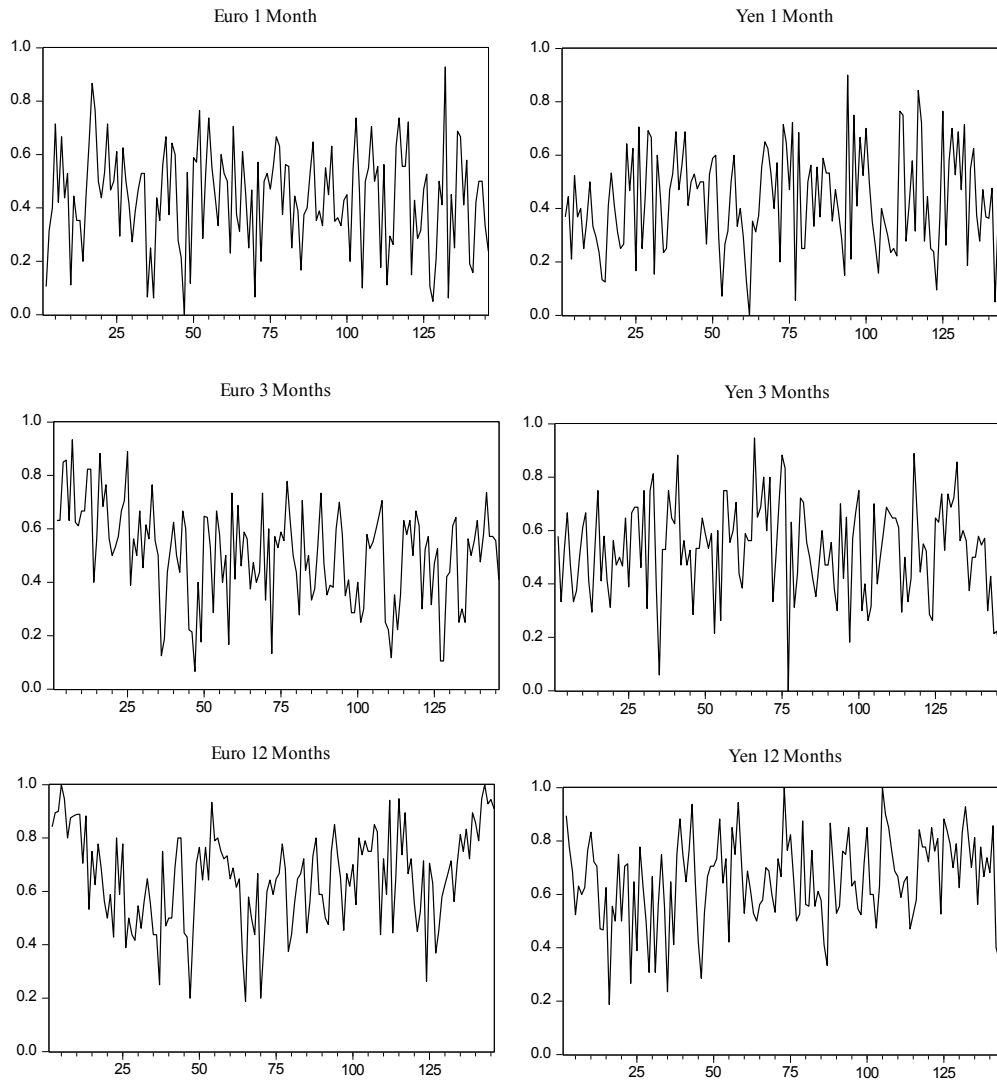
Euro   Japanese Yen
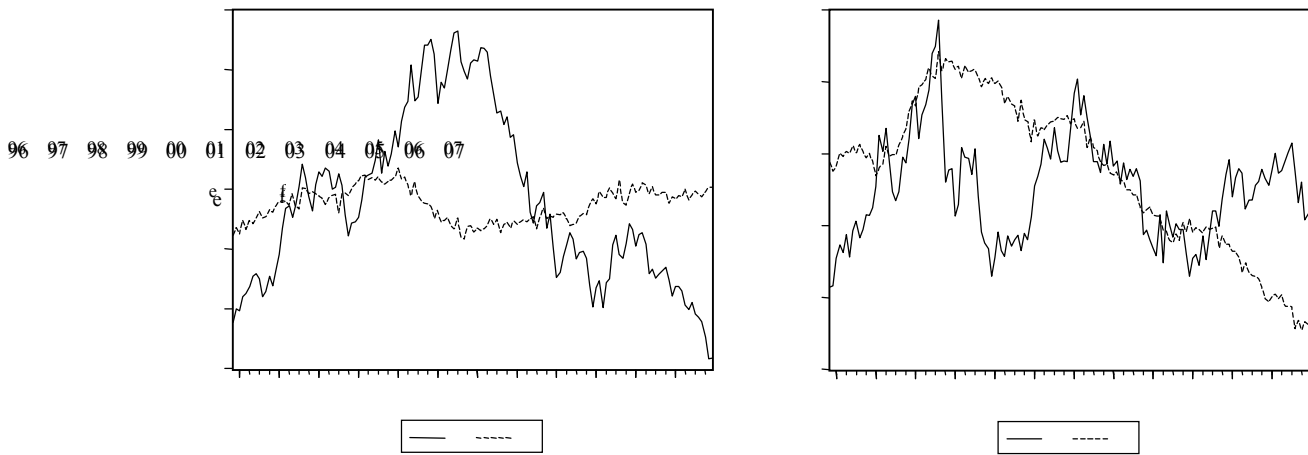
Notes: Figure depicts the log fundamental and market exchange rates. Fundamental calculated using (19).

Figure 2: Average Percentage of Fundamentalists over Time

Notes: Figure presents the cross-sectional average choice for the fundamental value over time, retrieved from the dynamic model.

Figure 3: Average Percentage of Fundamentalist Expectations per panelists
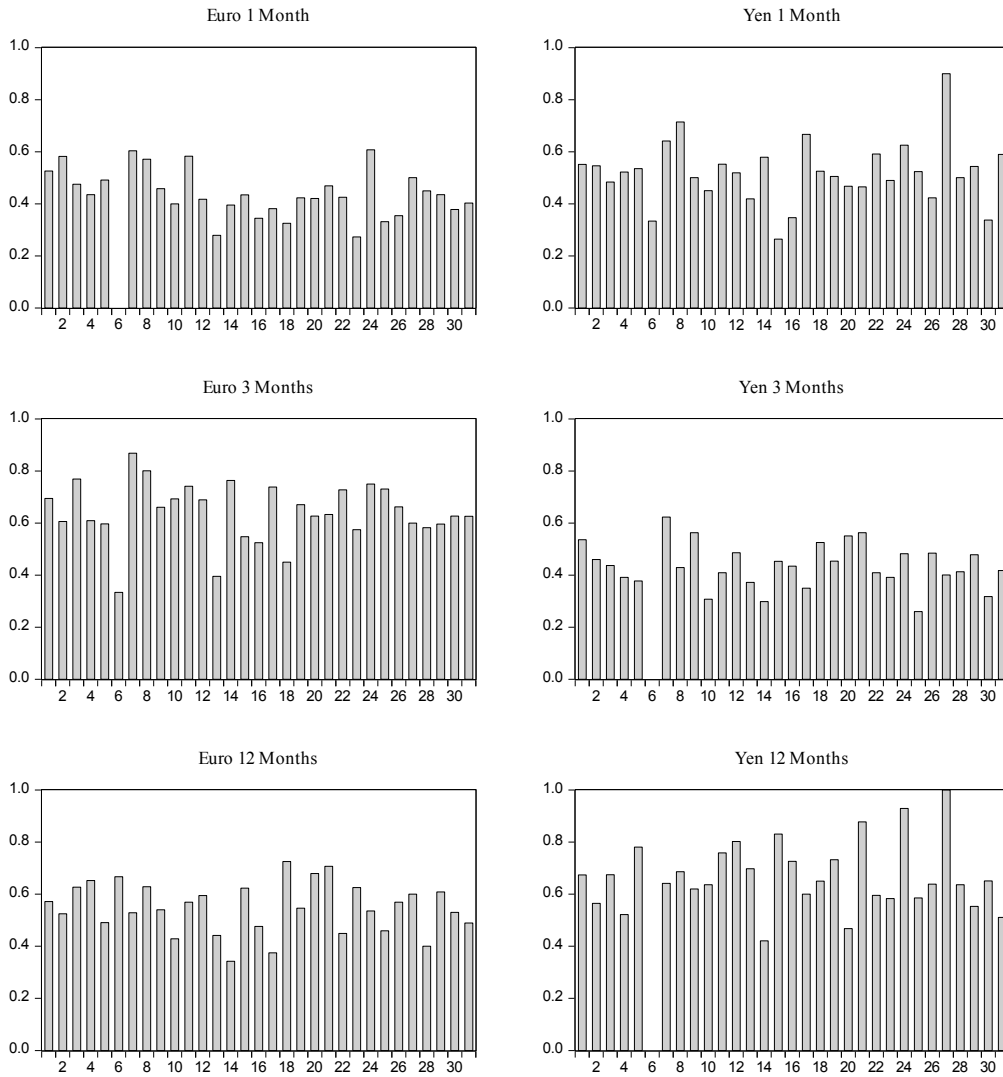
Notes: Figure presents the time-series average choice for the fundamental value per panelist, retrieved from the dynamic model.