Dispersion of Beliefs and Market Volatility in the Foreign Exchange Market

Ron Jongen*

Willem F.C. Verschoor#

Christian C.P. Wolff§

Remco C.J. Zwinkels#

March, 2006

Abstract

This paper examines the existence of heterogeneous expectations among market participants in the foreign exchange market by using a data set of individual market expectations for the major currencies, and approaches the formation of expectations from a bounded-rationality approach. We find that that there are distinct periods of high and low dispersion in which market participants disagree as to what will happen to the future level of the exchange rates. Furthermore, we document that the frequency at which extremist differences in expectations among market participants occurs, is higher than that what would occur under normality. Dispersion of beliefs seems to occur as a result of the combined effects of market participants holding individual expectations and attaching different weights on various elements from their information sets. Finally, we find that market volatility Granger-causes trader heterogeneity.

Keywords: Exchange rate expectations, heterogeneity, dispersion of beliefs, bounded rationality, tail behaviour, survey data.

JEL classification: F31

---

* Maastricht University, The Netherlands.

§ Maastricht University, The Netherlands, and CEPR, United Kingdom.

Nijmegen School of Management, Radboud University Nijmegen, The Netherlands.

Correspondence address: Department of Finance, Maastricht University, P.O. Box 616, 6200MD Maastricht, The Netherlands. Phone: +31 (43) 3883838; e-mail: c.wolff@finance.unimaas.nl (C.C.P. Wolff).

Dispersion of Beliefs and Market Volatility in the Foreign Exchange Market

Abstract

This paper examines the existence of heterogeneous expectations among market participants in the foreign exchange market by using a data set of individual market expectations for the major currencies, and approaches the formation of expectations from a bounded-rationality approach. We find that that there are distinct periods of high and low dispersion in which market participants disagree as to what will happen to the future level of the exchange rates. Furthermore, we document that the frequency at which extremist differences in expectations among market participants occurs, is higher than that what would occur under normality. Dispersion of beliefs seems to occur as a result of the combined effects of market participants holding individual expectations and attaching different weights on various elements from their information sets. Finally, we find that market volatility Granger-causes trader heterogeneity.

1. Introduction

In the previous decades we have seen an increase in the number of studies that attempt to explain various aspects of the foreign exchange market. The interest in this area does not come as a surprise, since the large amount of foreign exchange that is traded worldwide is far in excess of what is required for trade in goods and services. It therefore seems that the foreign exchange market is a market 'on its own' and that this market, because of its large volume, is highly liquid and efficient (Froot and Thaler, 1990). For this reason, market participants are said to have equal access to information and form their expectations about future events in a uniform, rational manner.

In the academic literature, there has therefore been a tendency to consider a representative agent approach whenever one needed to form an idea about exchange rate expectations. That is, expectations in the foreign exchange market are assumed to be formed in a rational way where all market participants optimize all information as efficiently as possible and all market participants have similar access to all information. The assumption of homogeneous, rational expectations has consequently been used in a variety of fields within the finance literature.

An example of a strand of the literature that has relied on the assumption of rational expectations is the forward premium puzzle discussion (see Hodrick, 1987, or Engel, 1996, for an overview). Forward exchange rates have often been found to be biased predictors of future spot rates. Based on the assumption of rational, homogeneous expectations, most studies have attributed this bias to the existence of time-varying risk premia. However, Frankel and Froot (1987) and Cavaglia et al. (1993, 1994) have questioned to what extent these interpretations are biased by the possible false assumption of rationality; Ito (1990) examines whether the homogeneity assumption is valid and Frankel and Froot (1990) condition the use of information on the forecasting horizon.

Another example is the literature on foreign exchange rate exposure, popularized by the influential work of Adler and Dumas (1985) and Jorion (1990). This literature assumes that unexpected exchange rate changes affect the returns of companies, whether or not these companies have foreign operations. Based on the assumption of rationality, the unexpected changes are commonly approximated by realized changes. Yet after a few decades of research, the empirical evidence of whether domestic firms are exposed to currency risk still remains inconclusive and puzzling. It is for this reason that Gao (2000) questions the rationality assumption and explicitly tries to model the expected change in exchange rates using macroeconomic variables.

From these examples and the existence of other anomalies, such as the excess volume of trade in the foreign exchange market, it becomes clear that the notion of rational expectations is losing more and more ground. Instead, the focus is shifting in the direction of bounded rationality, and the accompanying heterogeneity of agents' expectations. New insights into how market participants form their expectations are therefore warranted.

Among the reasons of the popularity of rational expectations is the relative ease of (mathematical) modeling and, especially in empirical work, the lack of alternatives. The latter issue has been resolved by the introduction of survey datasets. Ever since, the use of survey data is not uncommon in the finance literature and an increase of the use of surveys in various areas of the finance literature is observed. For example, Friedman (1979, 1980), Froot (1989) and MacDonald and Macmillan (1994), have used interest survey data in tests for identifying term premiums and examining the rationality of expectations of future interest rates and concluded that predictions were biased and respondents did not efficiently exploit the information contained in past interest movements. Similarly, Dokko and Edelstein (1989) review the usefulness of the Livingston forecasts of stock market rates of return and find evidence of adaptive behavior in the forecasts. Keane and Runkle (1990) use survey forecasts of the GNP deflator and find that expectations are rational, and MacDonald and Torrance (1988) use survey data on expected changes in money aggregates with U.K. data and find that these survey measures of expectations are extremely useful, for, unlike statistical methods for generating estimates, they are truly exogenous.

But particularly in the foreign exchange rate market literature there has been a considerable amount of interest in the exploration of survey-based expectations for years in order to understand the behavior of foreign exchange market participants. Frankel and Froot (1987) and Cavaglia et al. (1993, 1994) use survey data on foreign exchange expectations to examine whether the failure of the forward premium puzzle can be attributed to irrational behavior on behalf of market participants or due to the existence of time-varying risk premia and Marsh and Power (1996) and Elliott and Ito (1999) examine the forecast performance of survey-based exchange rate forecasts.

In this paper we add to the evidence of heterogeneity in expectations using a survey dataset. Our contribution lies in the fact that we approach the issues from a more fundamental viewpoint compared to the existing literature. Furthermore, we use a broader dataset and a richer arsenal of tests. We start with a fundamental discussion on what heterogeneous expectations comprise of theoretically and how one can measure and quantify heterogeneity in expectations, specifically in survey datasets of individual expectations. We focus on the question what is the source of the possible heterogeneity. Next we employ these insights to a survey dataset of individual expectations.

We find that that there are distinct periods of high and low dispersion where market participants disagree as to what will happen to the future level of the exchange rates. Furthermore, we document that the frequency at which extremist differences in expectations among market participants occurs is higher than that what would occur under normality. Dispersion of beliefs seems to occur as a result of the combined effects of market participants holding individual expectations and attaching different weights on various elements of the set of public information. Finally, we find that market volatility Granger-causes trader heterogeneity.

The remainder of the article is presented as follows. In Section 2 we examine the rationale behind heterogeneous expectations. In Section 3 we introduce the data used for this analysis. Section 4 examines whether expectations are heterogeneous and in Section 5 we link dispersion of beliefs to market uncertainty. Section 6 concludes.

2. Heterogeneity in Foreign Exchange Markets

Heterogeneity is a concept that is used in a variety of ways in the foreign exchange literature. There are broadly two commonly used explanations for the existence of dispersion of beliefs in financial markets. One strand of the literature argues that dispersion of beliefs arises because of the asymmetry in information. Different market participants are assumed to hold different sets of information, whereby part of the information is common for all participants and part is private. The concept of asymmetric information was first introduced in the New Classical Theory of the macro economy, where agents were assumed to be unable to obtain information that is public in other parts of the economy, and where agents are rational in the Muth (1961) sense in that they use that information that is available to them in the best possible way to form their expectations of a particular variable. The asymmetry in information may be caused by the rigidity in the transmission of public information, such that the heterogeneity in agents' beliefs is caused by an artificial informational assumption (Kurz and Motolese, 2001).

There are several arguments and empirical considerations against the claim that dispersion of beliefs is caused by the existence of private information. First, since financial markets and foreign exchange markets in particular are dominated by large institutions that all have huge resources and are competitive in the search for information, all should possess essentially the same information. Second, even when some market participant possesses inside private information, it is still illegal to trade on such information. Since developed financial markets are heavily monitored by a community of regulators, brokers, and governmental institutions, one would expect that trading on such private information is scarce, or at minimum it is unrealistic to assume that a significant part of the volatility in foreign exchange markets is caused by illegal trade.

Another strand of the literature assumes that all market participants hold different beliefs about economic variables even when there is no difference in the information that is available to them. The difference in beliefs arises because agents disagree about the interpretation of this information. To argue why the difference in interpretation occurs we can follow the rational belief theory due to Kurz (1994), which assumes that heterogeneity of beliefs is caused by the fact that economic agents do not know the structural relations of the economy. Agents only have 'information' or 'empirical knowledge', which is readily observable from the economy, usually in the form of a large amount of data concerning the past performance of an asset or economy in general. Agents form their opinions about the future by using the empirical distribution that is derived from the occurrence of events in the past. Their own experience and success in forming accurate forecasts will likely play a role in selecting and valuing information.

An alternative explanation for why market participants hold different beliefs about the future may be related to the existence of fundamentally different types of market participants. In the financial economics literature there have been several attempts to present models with different types of investors who in essence all have the same information. For instance, De Long et al. (1990) present a model where noise traders, with no access to inside information, act irrationally on noise as if it were information that would give them an advantage over rational arbitrageurs. These noise traders can earn higher expected returns, because of their own destabilizing influence and not because they bear more of the fundamental risk.

Frankel and Froot (1986, 1990) popularize the view that the foreign exchange market is dominated by two types of market participants that differ in which information they use for forming their expectations. Fundamentalists think of the exchange rate as an economic model, while chartists predominantly use the exchange rate's own history as input in their expectations formation process. They show that the increased use of technical analysis (chartist behavior) in the late 1980s could well explain why the U.S. Dollar was so far above its long-run equilibrium at that time, or why the volume of foreign exchange trading worldwide has grown so tremendously. Allen and Taylor (1990) and Taylor and Allen (1992) were among the first to show by means of surveys that at least 90 percent of market participants place some weight on technical advise, although this is used predominantly for shorter-term forecasts.

The realization that the presence of heterogeneous interacting market participants can shed new light on some of the anomalies in the financial economics literature has resulted in a strand of literature, which is mostly simulation-based, on how to model these heterogeneous agents (see Hommes, 2006, for an overview of this literature). For instance, heterogeneous agents models can easily generate large trading volume consistent with empirical observations.

3. Data

Every second Monday of each calendar month Consensus Economics of London publishes results from a survey among up to 150 professional market participants and forecasting agencies for their subjective expectations of a large number of exchange rates. Forecasts are made for a variety of horizons ranging from 1 to 24 months ahead.

For our panel, we obtain expectations on the U.S. Dollar / U.K. Pound (USA/UKI), Japanese Yen / U.S. Dollar (JAP/USA), and U.S. Dollar / Euro (USA/EUR) exchange rates for the period of October 1989 through December 2004.[1] 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 some of the exchange rates. The panel is unbalanced since the response rate of the individual market participants is less than 100 percent and since market participants left the panel and were replaced by others. Due to data unavailability we have to split the sample in two periods: October 1989 through February 1995 (Sample 1) and November 1995 through December 2004 (Sample 2). In this paper we focus on the 3- and 12-month-ahead forecasts.

Although survey participants have a few days time to return their expectations, we learned that the vast majority send their responses by e-mail on the Friday before the publication day (usually second Monday of the month). We consider this Friday to be the day on which the expectations are formed. On this Friday, we obtain spot and forward exchange rate series with different maturities to match with the survey data. All spot rate series are obtained through Datastream. To verify that the information sets of market participants are not too diverse, all

[1] Prior to January 1999 we use expectations on the Deutschemark / U.S. Dollar. We transform these expectations into U.S. Dollar / Euro expectations.

of the analyses throughout this study were re-estimated using spot data from various days surrounding this Friday, yet the overall results remained virtually unchanged. Other data, such as interest rates, are also obtained through DataStream.

Typical concerns when using survey data in any setting are whether this data reflects the true market's expectations, whether the expectations are biased because of strategic behaviour from the panellists, or whether forecasts from surveys are of any good in an out-of-sample forecast setting—a criteria that has often been put forward to evaluate the quality of survey expectations. It should be noted that for survey data in the present setting it is most important that the survey expectations reflect the market's sentiment at the time they are formed, that is, the survey data should reflect expectations, nothing more than that. While it is not the primary concern that the expectations outperform other forecasting techniques, there is a consensus view that expectations from surveys in general perform no worse than any other forecast technique. We can learn much about the usefulness of survey-measures of expectations from related fields. Ang et al. (2005), for instance, provide recent evidence that expectations from various surveys on inflation consistently deliver better forecasts than time-series models, models based on the yield curve, and forecasts based on the Phillips curve, which highlights the usefulness of survey measures of expectations. Elliott and Ito (1999) find that in the foreign exchange market portfolios based on survey expectations produce small, but positive, profits.

Another point of criticism is that there may be reasons for forecasters not to give their true expectation. One motive may be that agents do not want to expose their (private) information to other market participants. This effect can be mitigated by the reputation effect that surveys can have. If the names of the forecasters are given in the survey publication, agents have an incentive to perform well in order to attract customers.

Reputation effects are not only relevant for customers of financial institutions, but are also present between financial institutions. Beine et al. (2003) show that there is a clear leader-follower pattern in the expectation formation process. If forecasters notice that a certain forecaster performs well, his or her forecast becomes a focal point and other forecasters will follow up. Note that if this 'herding' occurs on a large enough scale, or if a single respondent is large, and forecasters actually act upon the given expectations, the expectations can become self fulfilling. Therefore, a better understanding of the mechanisms that drive individual expectations is warranted.

Finally one can question the relation between foreign exchange expectations and foreign exchange returns. Heterogeneous expectations literature (see Hommes 2000) assumes that the realized exchange rate return is equal to the market expectation of the exchange rate return and the market expectation is equal to the weighted average of individual expectations. Given that we assume to be in a boundedly rational environment, this direct relation, especially with survey forecasts, is questionable for a number of reasons. First, behaviour of market participants is conditional on their expectations, but also on their endowment, risk aversion and incentives. Therefore, the relation between expectations and realizations might not be one-to-one. Second, the market microstructure literature (see Lyons, 2001) stresses the importance of the organization and setup of the market; furthermore, traders use the exchange rate expectation, often formed by another division of the company, as a source of information next to others like market sentiments. However, the relation between expectations and realizations is of less importance for this paper as we mainly focus on (the source of) heterogeneity in expectations.

4. Are Foreign Exchange Forecasts Heterogeneous?

The number of methods to measure or quantify heterogeneity is small because of the relatively scarcity of data on individual (survey) expectations. The available measures can be roughly classified into three groups. The first group simply quantifies the dispersion of beliefs across agents at one point in time by using a volatility-based approach. One way in which we may think of heterogeneity is that market participants disagree as to what the market will do in the future. Frankel and Froot (1990) use the standard deviation across respondents in a survey and Beine et al. (2002) use the cross-Section coefficient of variation as proxies for this disagreement. Alternatively, when the maximum and minimum forecasts among a group of participants go further apart from one period to the other, or when the expected changes go in opposite directions, one may also think that the belief of where the market is going becomes more diverse. Although the magnitudes of the expected changes of various market participants may naturally vary, it makes a great difference whether the expected changes are in the same direction or in the opposite.

The drawback of these measures is that they do not provide an absolute measure of heterogeneity, but only a relative measure. It can only be used to compare the degree of heterogeneity between two time periods or markets, but cannot conclude that a certain group of forecasters is significantly heterogeneous or not. On the other hand, the standard deviation (or any function thereof) has considerable appeal when one wants to find answers to questions like whether dispersion in beliefs is higher in some periods, or whether belief dispersion is causing market volatility or vice versa. Furthermore, it is a really intuitive measure and relatively easy to compute.

The second type of measure is more advanced in the sense that it is able to conclude whether agents are heterogeneous or not. Under the null hypothesis of homogeneity, the differences between contemporaneous individual forecasts (or the difference between the consensus forecast and the individual forecast) are independently and identically distributed (i.i.d.). This assumption is directly testable by examining the distributional properties of these differences. Of particular interest is the tail behavior of this distribution, since under the null hypothesis of homogeneity extreme differences should not occur at a rate higher than normal. Furthermore, the average difference between forecasters should be zero, and there should not be any form of autocorrelation. Yet, even though this second measure can differentiate between homogeneity and heterogeneity, it does not say anything about the source of the heterogeneity.

In a third approach one is able to determine the source of heterogeneity. This method estimates directly the underlying process that is forming the expectations. If the fit of a certain model of expectation formation is better for one agent compared to the other and vice versa, this is a direct proof of heterogeneity, and it lays bare the source of the heterogeneity (i.e., on which variables the agents put different weights). Audretsch and Stadtmann (2005) determine the optimal model for each agent. The R^2 of these type of regressions is informative in that it reveals to what extent the forecast is based on observable information. The remaining variance is then caused by other variables, such as psychological, experience, etc. Ito (1990) and MacDonald and Marsh (1996) regress the difference between the forecasts of agents on a set of variables that is possibly used in forming expectations, such that the difference in weight put on a certain variable in the expectation formation model is estimated directly.

Although this method determines whether agents are heterogeneous or not and on what grounds, there are still some drawbacks. First is the matter of variable selection. The exact set of variables used to form an expectation is not known. Second, some of the variables used in forming expectations are inherently unobservable, such as the state of mind of the forecaster during the expectation formation process, or even the weather condition at that particular time. Furthermore, the weight agents put on a certain variable might change through time (see, for instance, the scapegoat models in Bachetta and Wincoop, 2004). If agents are heterogeneous by definition, the adjustment of the weight given to a certain variable in the expectation formation process does not change simultaneously or equally among agents. Especially in these regression-based approaches this effect could result in biased results, since it assumes a constant difference in weight through time. Therefore, the question of whether expectations are heterogeneous and if so, what the source of this heterogeneity is, should ideally be approached from various angles.

We start our analysis of heterogeneity in beliefs by a visual inspection of volatility-based measures. Since the standard deviation is unit-dependent we take instead the coefficient of variation, defined as the cross-sectional standard deviation divided by the cross-sectional average. This allows us to compare the measure of heterogeneity across currencies. Figure 1 presents the coefficient of variation along two forecast horizons for the three currencies. In addition, high-low spreads are presented as a robustness check.[2] Several findings are noteworthy. First of all, the spread and coefficient of variation reveal similar patterns. This ensures us that both measures seem to be capable of capturing dispersion in beliefs, in other words, the variation in the spreads is not caused by the extremist expectations of outliers.

Second, the dispersion in beliefs increases with the forecast horizon. When one would see dispersion of beliefs as a sign of uncertainty, this would indicate that the market feels more certain about what will happen in the nearby future than in the distant future.

Third, dispersion of beliefs is certainly not constant over time, but seems to be high in some periods of considerable length, while low in others. In particular for the U.S. Dollar / U.K. Pound and for the U.S. Dollar / Euro dispersion of beliefs was obviously higher from early 1990 to late 1993 than in the subsequent period. Similarly, for the Japanese Yen / U.S. Dollar rate expectations became increasingly dispersed from early 1998 onwards until the end of 2001. This period is in the aftermath of the Asian crisis starting in mid-1997, which would lead us to believe that market instability is a driving factor of dispersion in beliefs. This would be an interesting argument, for if would suggest that a dramatic event such as the Asian crisis may have an effect on the expectations of market participants even 2 year after date.

We proceed by examining formally whether expectations are heterogeneous and, if so, whether the heterogeneity arises because of private information or that individuals have different opinions about the importance of some variables. To do so we follow the procedure due to Ito (1990). Suppose that individual i forms an expectation about the k-period ahead level of a particular exchange rate and that this forecast consists of a structural part based on

[2] Using the spread as a measure of heterogeneity may be subject to one severe problem. Since the spread only depends on two contemporaneous observations, non-representative market participant who often produce extreme (outlier) expectations directly influence the measure of heterogeneity. Therefore, we only use the spread as a robustness check.

public information (that is common to all market participants) and an individual effect g_i. Define this expectation $s_{i,t}^{e(k)}$ as

$$s_{i,t}^{e(k)} = f(\Omega_t) + g_i + \varepsilon_{i,t}^{(k)}, \tag{1}$$

where and $\varepsilon_{i,t}^{(k)}$ is a random disturbance term. The cross-sectional average of a set of market participants' forecasts is then defined as

$$\bar{s}_t^{e(k)} = f(\Omega_t) + \bar{g} + \bar{\varepsilon}_t^{(k)}. \tag{2}$$

Suppose that a normalization such that $\bar{g} = 0$ is possible, then subtracting the cross-sectional average from the individual forecast results gives that

$$s_{i,t}^{e(k)} - \bar{s}_t^{e(k)} = g_i + (\varepsilon_{i,t}^{(k)} - \bar{\varepsilon}_t^{(k)}). \tag{3}$$

Table 1 presents summary statistics for the distribution of the individual expectations in excess of their cross-sectional average, $s_{i,t}^{e(k)} - \bar{s}_t^{e(k)}$. The samples all consist of pooled data. Both the high-low spread (maximum minus minimum forecast) and standard deviation increase as the forecast horizon increases. This indicates an increase in heterogeneity as the forecast horizon increases, as we have seen in the Figure 1.

The highly significant positive autocorrelation in the deviations from the average furthermore indicate that individuals are persistent in their opinion relative to the average opinion. If a certain forecaster expects a lower (higher) than average change in period t, he or she also expects a lower (higher) than average change in period t+1. This cannot be the case when individuals are homogeneous, since deviations from the average should then be identically and individually distributed. This effect is stronger on longer forecast horizons and in the second sample[3].

Although the individual expectations do not appear to be skewed in a particular direction, the distributions are consistently leptokurtic. Furthermore, while high kurtosis is not analogous to fat tail behavior, we may state that an excessive number of expectations are at

[3] Except for the U.S. Dollar / U.K. Pound, but this is due to the fact that forecasts are given two-monthly in the second set.

enough distant from the consensus that they render the distribution non-normal. Jarque-Bera statistics for the normality of the individual expectations are included and corroborate the rejection of the normality assumptions under high levels of significance.

A formal analysis of the tail behavior can be given by examining a measure for the tail index. Hill (1975) defines a tail index estimate by looking at the order statistics of a series. Suppose that expectations on a particular exchange rate are independently and identically distributed. We can then pool all individual expectations (in excess of the cross-sectional mean) in a vector *x*. Define $X_{n-m,n}$ to be the (n – m)-th ascending order statistic (or in other words, the *m*-th smallest expected change) from the set of individual expectations *x*. The Hill estimator can then be defined as

$$\hat{\alpha} = \left[\frac{1}{m} \sum_{M=0}^{m-1} \ln\left(\frac{X_{n-M,n}}{X_{n-m,n}} \right) \right]^{-1} \tag{4}$$

where *m* is the number of tail observations used for estimating the tail index. The estimate $\hat{\alpha}$ can be interpreted as the highest moment that exists from the sample. So when the tail estimates are sufficiently low, we can question the existence of higher moments of the distribution and tails appear fat.

We choose the cutoff point *m* by using a certain percentage of the lowest observations in *x*. Table 2 provides tail index estimates for the 1, 5, and 10 percent tails of the distribution. Estimates for both upper (right) and lower (left) tails are included separately to detect any tail asymmetries. The number of tail observations *m* that is used to estimate the tail indices is reported above the estimates. It can be seen that the distribution is fat-tailed, with indices either just above or below 4, depending on the percentile. This would even question the existence of the 4$^{\text{th}}$ moment in the distribution and indicates that extreme expectations occur more common than normal. In other words, the number of expected changes that is extremely deviating from the mean expectation is such that we can state that market participants' beliefs are heterogeneous.

A question that appears naturally is whether for a particular currency the market expects more upward trends than downwards trends. For instance, when a sufficiently large number of extreme expected upward changes occur relative to expected downward changes, then the right tail of the distribution is bigger than the left. Using a two-sided *T*-test we formally test for equality of the upper and lower tail index estimates and find that we can only consistently reject the null hypothesis of equal tails for the 12-month-ahead forecasts of the US Dollar / Euro exchange rate. For most other series do the tail indices look similar and hence may we assume that tails are symmetric here. The tail indices are in general also larger for the longer-term forecast horizons.

Finally, the question remains whether heterogeneous beliefs arise because market participants hold different information, whether they attach different weights to some elements of that information, or both. Furthermore, if all market participants would make a few excessively large number of extreme expectations in the sample period under investigation, but for the rest would make estimates that are close to the consensus (or in general close to all other market participants), then we would hardly consider these expectations to be heterogeneous. Yet, the Hill estimator could indicate in such a case that there would be evidence of tail behavior (and hence heterogeneous beliefs) since the few number of extreme expectations of each market participant would add up to a considerable number when all market participants are pooled.

To deal with these issues we should ideally look at evidence of heterogeneous beliefs for each market participant separately. Reconsider Equation (3):

$$s_{i,t}^{e(k)} - \bar{s}_{t}^{e(k)} = g_i + (\varepsilon_{i,t}^{(k)} - \bar{\varepsilon}_{t}^{(k)})$$

Here, we assumed that the public information $f(\Omega_t)$ is common to all market participants at each sample period. For this reason, it is unnecessary to know the exact structure of $f(\Omega_t)$. Since the random disturbance terms both have expected value equal to zero, the individual effect g_i can be estimated from a regression of the individual expectation in excess of the cross-sectional average on a constant. We therefore estimated Equation (3) for each market participant. In case that the residuals are serially correlated, an AR(1) component is added to account for this. The significance of the g_i term indicates whether the forecaster shows individual effects. In addition, if the difference in the individual effects of two individuals (say i and j) is to be estimated, a similar method can be used:

$$s_{i,t}^{e(k)} - s_{j,t}^{e(k)} = g_i - g_j + (\varepsilon_{i,t}^{(k)} - \varepsilon_{j,t}^{(k)})\,. \tag{5}$$

Table 3 gives the percentages of market participants that show individual effects, where individuals are compared versus the cross-sectional average (Eq. 3) and versus each other (Eq.

5). For around 35 percent of the individual market participants we find evidence of individual effects on the short forecast horizon. When the forecast horizon extends to one year, this percentage increases to around 40 percent. In the second sample period the percentages are a little less, but still around 30 percent.

Thus far this reasoning assumes that heterogeneity arises due to differences in information sets, i.e. ‘private’ information. It is reasonable to assume that although information may be common to multiple market participants, they interpret this information differently and hence attach different weight to its various elements. These different weights we call idiosyncratic effects. We consider two models that may explain how market participants form their forecasts. Suppose that market participants form forecasts in an extrapolative way, such that

$$s_{i,t}^{e(k)} - s_t = h_i + \sum_{l=1}^{L} \beta_{l,i} (s_{t-l} - s_{t-l+1}) + \varsigma_{i,t}^{(k)} \ . \tag{6}$$

where L is the number of lags that is used in the forecast. Following the same procedure as in (3) we find that

$$s_{i,t}^{e(k)} - \overline{s}_t^{e(k)} = h_i - \overline{h} + \sum_{l=1}^{L} (\beta_{l,i} - \overline{\beta}_l)(s_{t-l} - s_{t-l+1}) + \varsigma_{i,t}^{(k)} - \overline{\varsigma}_t \ . \tag{7}$$

In other words, the structural part of the information set and the values of the regressors are the same for all market participants, yet the coefficients (or weights) differ. We estimate Equation (6) for all market participants. We use only two lags to preserve model parsimony, while still capturing most of the dynamics. Since we find that serial correlation usually does not extend beyond the second lag, we feel that this specification represents an extrapolative expectations model sufficiently. Table 4 presents the percentages of market participants for which we find individual and/or idiosyncratic effects. Results are for Wald tests on the significance of combinations of the parameters as to discriminate between individual and idiosyncratic effects.

We see that in addition to the earlier-mentioned individual effects, there is proof that a large group of market participants attach different weights on the information that is in the most recent history of the exchange rates. On average about 20 percent of the market participants show these idiosyncratic effects, although the effect is more pronounced at the shorter forecast horizon, for both sample periods and all exchange rates.

We continue the analysis by questioning whether the uncovered interest parity (UIP) has appeal as an empirical model for the formation of expectations and whether market participants have idiosyncratic beliefs as to the importance in the relative level of interest rates. Consider

$$s_{i,t}^{e(k)} - s_t = d_i + \gamma_i (i_t^{(k)} - i_t^{*(k)}) + \vartheta_{i,t}^{(k)}, \tag{8}$$

where $i_t^{(k)}$ is the k-period domestic interest rate and $i_t^{*(k)}$ the k-period foreign interest rate. Following the same line of reasoning as under (3), we can rewrite (8) as

$$s_{i,t}^{e(k)} - \bar{s}_t^{e(k)} = d_i - \bar{d} + (\gamma_i - \bar{\gamma})(i_t^{(k)} - i_t^{*(k)}) + \vartheta_{i,t}^{(k)} - \bar{\vartheta}_t^{(k)}, \tag{9}$$

We measure $i_t^{(k)}$ and $i_t^{*(k)}$ as the domestic and foreign interbank rates, respectively. Estimates are reported in Table 5. The parameter $(\gamma_i - \bar{\gamma})$ is significant in about 20 percent of the cases, where the effect is more pronounced on the longer forecast horizon.

Our findings that the weight put on extrapolation is more heterogeneous in the short forecast horizon and the weight put on the UIP is more heterogeneous in the long forecast horizon corroborates the results of Frankel and Froot (1990). They state that short horizon forecasts are more based in technical analyses while long term forecasts are more based in (macro-) economic information.

In order to use the panel-character of our disaggregated dataset to the fullest, we also estimate the abovementioned equations (5), (7) and (9) in pooled format. To determine whether respondents are heterogeneous or not, we first estimate the Equations with a single coefficient for all respondents, second estimate the same Equations with respondent-specific coefficients, and third compare the log likelihoods of the two regressions with a likelihood ratio test[4]. In the case of equation (5) this boils down to estimating the panel first with a common intercept and second as a fixed effect model. Equations (7) and (9) are first estimated with a common intercept, β and γ. In the second estimation the common intercept is kept, but β and γ are made respondent-specific. As a result, we estimate N - 1 (with N the number of respondents) coefficients more in the second round.

---

[4] The likelihood ratio is Chi-square distributed with 188 degrees of freedom for the first period and 30 for the second. This results from the fact that there are 189 and 31 respondents with sufficient observations for the first and second period respectively.

In all regressions we add an AR(1) term in order to correct for the autocorrelation in the deviation of individual expectations from the consensus. Furthermore, we estimate the models with a White correction. In the case of equation (7), we now only add one lag ($L = 1$) because including two lags would result in estimating two coefficients per respondent, which is too much. Table 6 presents the results.

For all but one cases we find a significant increase in the log-likelihood after introducing respondent-specific coefficients. We do not find a significant improvement for the USD/UKP 3-month horizon when using extrapolative techniques. This is a strong indication that respondents are indeed heterogeneous. To be more specific, they put significantly different weights on a number of variables when forming expectations. The individual coefficients (not reported) indicate that roughly 15 to 20 percent of the respondents have a significant positive coefficient; the same holds for significant negative coefficients. These results corroborate our findings from the separate regressions in the previous Tables.

In this Section we have examined the question of whether foreign exchange forecasters are heterogeneous in that they have dispersed beliefs. Using various techniques we showed that indeed expectations are heterogeneous and that the dispersion of beliefs occurs both due to the existence of individual effects and idiosyncratic effects, where market participants attach different weights to variables that are observed by all market participants contemporaneously. In the next Section we attempt to find an answer to the question whether market characteristics affect trader heterogeneity.

5. Dispersion of Beliefs and Market Volatility

In the previous section we have shown that there is indeed heterogeneity among foreign exchange expectations. The next step is to see whether the degree of heterogeneity is in any way linked to market characteristics. So, contrary to the previous Sections where we compared forecasts vis-à-vis each other and the consensus, we will now study the degree of dispersion vis-à-vis the market itself. To be more specific, we will examine the relation between dispersion and market uncertainty.

The heterogeneous expectations literature indicates that there is a direct causal relation running from trader heterogeneity to market volatility[5]. However, it also provides us with contrasting hypotheses concerning the sign of the effect. As the market price is moving away from the fundamental price, there is a negative relation between trader heterogeneity and

[5] See Hommes (2006) for an extensive survey.

market volatility because the expectations of the different groups are opposite to each other[6]. As the market price moves more and more away from the fundamental price, the fundamentalists get driven out of the market. The technical analysts remain active, continue to push the market price away from the fundamental price and increase market volatility as the number of technical analysts increases and their expectations become self-fulfilling.

If the market price is moving towards the fundamental price, however, there is a positive relation between trader heterogeneity and volatility as the different groups expect similar directions of change. In this situation volatility is rising as all traders active on the market push the price in the same direction; there is no counter acting force which is active when the price is moving away from the fundamental price and fundamentalists are still active. Both groups remain active as both strategies are profitable.

The market microstructure literature provides another mechanism by which trader heterogeneity affects markets. Evans and Lyons (2002) argue that high turnover in markets reflects great dispersion of opinions among traders. A number of authors have looked into this relation empirically. MacDonald and Marsh (1996) examine the relation between trader heterogeneity and trading volume, with the same survey dataset as us, and find a significant positive relation. Frankel and Froot (1990) conclude, for survey datasets from MMS International and The Economist, that dispersion of opinion affects the volume of trade and thereby also market volatility.

We study the causal relation between market volatility and trader heterogeneity by means of Granger causality tests. The added value of our analysis is that we consider a longer period compared to MacDonald and Marsh (1996), we use different measures of both trader heterogeneity and volatility, and we focus on the causality question instead of correlations. In the previous sections we argued that heterogeneity can be measured in various ways. Similarly, volatility can be quantified by several different methods. Therefore, we examine twelve different heterogeneity–volatility measure combinations in order to check the robustness of the results.

In order to shed light on the direction of causality, we use Granger-causality tests. The setup of the standard Granger test is slightly altered. In a standard setup, the Granger-causality test regresses the *lagged* values of both X and Y on X and Y such that the test gives an indication concerning the causality between X and Y instead of the correlation. In our analysis, in the case of trader heterogeneity as the dependent variable, we do not use the lagged but the

[6] Fundamentalists' expect the price to return to the fundamental rate while technical analysts expect the price to continue to move in the direction it has been moving in the previous periods.

contemporaneous value of market volatility[7]. The reason for this is that trader heterogeneity at time t is formed *at the end* of the month, when forecasts are submitted, and concerns forecasts for the *coming* month(s). Volatility at time t, on the other hand, is formed *during* that past month (given that t is measured in months). Therefore, including the contemporaneous value of volatility is informative in the causality question; excluding it would imply dismissing one month of information.

Since this analysis does not require a time-series in the strict sense, it is possible to use the total sample, from October 1989 to December 2004, with a gap of six months in 1995, which are 175 monthly observations. For the UK we only use the first set (1989 – 1995) since the second set contains two-monthly observations. The three and twelve months horizons are available for the total sample. We focus on the first and second lag of the Granger causality test as we assume that there is no relation between the variables over periods longer than two months. The heterogeneity measures are the cross-sectional standard deviation, coefficient of variation, range and range percentage. The volatility measures are monthly absolute returns, monthly squared returns and a GARCH - measure[8]. Since all variables are not normally distributed, especially the absolute and squared returns, we apply a Box-Cox adjustment to both the left and right hand side variables.

Table 7 presents the results for the U.S. Dollar / Euro exchange rate. Looking at the absolute and squared returns for the three-month horizon, we find a significantly positive causal relation running from the variance to heterogeneity for all heterogeneity measures. The twelve months horizon gives a more mixed image. For the coefficient of variation and the standard deviation we find similar results as for the three months expectations, but concerning the range and range percentage measures we find positive significant results running from the heterogeneity to the variance.

The GARCH volatility measure renders a different image altogether. For the three months horizon we find, for all heterogeneity measures, a significant causal relation running from trader heterogeneity to market uncertainty; positive for the first lag, negative for the second. For the twelve months horizon we find a similar pattern, but in the opposite causal direction.

Table 8 presents the result for the Japanese Yen / U.S. Dollar exchange rate. For the absolute and squared returns we find a significant positive causal link from the market variance to trader heterogeneity for all heterogeneity measures and horizons. In the case of the

[7] Plus lagged values of volatility (if necessary) and heterogeneity itself.

[8] The GARCH measure was constructed by estimating a GARCH(1,1) Equation for daily data and taking monthly averages of the daily conditional volatilities.

GARCH – measure we find a two-way causal relation for all heterogeneity measures and horizons. The signs are in general positive for the first lag, negative for the second for the three months horizon and both positive for the twelve months horizon.

The results for the U.S. Dollar / U.K. Pound exchange rate in Table 9, finally, reflect a relatively homogeneous conclusion. For all variance and heterogeneity measures we find a significantly positive causal link from market uncertainty to dispersion of opinion. For a small number of combinations we find a two-way relation, but there does not seem to be a pattern.

In general we can conclude that the causal relation between market volatility and trader heterogeneity tends to be significant and positive for different measures of both trader heterogeneity and market volatility. This result corroborates the findings of MacDonald and Marsh (1996) but is opposite to the results of Frankel and Froot (1990), who find causality running in the opposite direction. Furthermore, the results are not in line with the theoretical predictions that heterogeneity affects volatility. This might be caused by the fact that we do not have data for the total market, but only a subset of traders. Since our traders are relatively homogeneous in nature (most of them are financial institutions based in London), our observations of trader heterogeneity might be biased downward. Given that we observe a subset of the total set of traders who cannot influence the market significantly, current uncertainty can cause uncertainty concerning the future, hence relatively diverse expectations.

6. Conclusion

In this article we examined whether expectations of future exchange rates are heterogeneous in that agents have different beliefs about the future path of the exchange rates. We approach this problem using a panel of survey expectations for the major three exchange rates along several forecast horizons. Using several measures for dispersion of beliefs we find that there are distinct periods of high and low dispersion where market participants disagree as to what will happen to the future level of the exchange rates. For the Japanese Yen versus the U.S. Dollar we even find that the Asian crisis that began in mid-1997 preceded an almost two-year period of highly increased dispersion of beliefs for this rate.

We test formally whether expectations are heterogeneous using an extreme value approach by examining the tail index estimates and conclude that the frequency at which extremist differences in expectations among market participants occurs, is higher than that what would occur under normality. We furthermore attempt to answer whether heterogeneity in expectations occurs because market participants hold different information sets or whether they attach different weights to commonly-held elements from their information sets. We find evidence for both. The extent of individual expectations seems to increase as the forecast horizon lengthens. The dispersion based on extrapolation decreases and dispersion based in interest rate differences increases as the forecast horizon lengthens. These results corroborate the fundamentalist/chartist literature.

Finally, we examine whether heterogeneity in expectations is influenced by volatility in the market, as suggested earlier by the increased dispersion of beliefs in the aftermath of the Asian crisis, or that heterogeneity in expectations is causal to the volatility in the market. We find that a causal relation between market volatility and trader heterogeneity tends to be significant and positive for different measures of both trader heterogeneity and market volatility.

Several issues remain unanswered. First of all, we question what the role of different 'types' of market participants is in the above analysis. We find different weights in different forecast horizons, but when one group of market participants uses an extrapolative, or any other chartist way of forecasting, and another group uses a fundamentals-based approach, we would like to see whether the number of 'chartists' versus 'fundamentalists' is related to market uncertainty or dispersion of beliefs. Second, we question whether market participants switch between various forecasting techniques (for instance, chartist or fundamentalist approaches) when the market becomes more volatile, and hence dispersion of beliefs increases. We feel that further investigation of these issues is warranted.

References

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

Audretsch, D.B. and G. Stadtmann (2005). Biases in FX-forecasts: Evidence from Panel Data, *Global Finance Journal* 16: 99-111.

Bachetta, P. and Van Wincoop, E. (2005). Rational Inattention: A Solution to the Forward Discount Puzzle, NBER Working Paper # 11633.

Beine, M., Bénassy-Quéré, A. and Colas, H. (2003). Imitation Amongst Exchange-Rate Forecasters: Evidence from Survey Data, *CEPII Working paper* 2003-08.

Cavaglia, S., Verschoor, W. F. C., and Wolff, C. C. P. (1993b). Further evidence on exchange rate expectations. *Journal of International Money and Finance*, 12(1):78–98.

Cavaglia, S. M. F. G., Verschoor, W. F. C., and Wolff, C. C. P. (1994). On the biasedness of forward foreign exchange rates: Irrationality or risk premia. *Journal of Business*, 67(3):321–343.

Engel, C. (1996). The forward discount anomaly and the risk premium: A survey of recent evidence. *Journal of Empirical Finance*, 3(2):123–192.

Evans, M.D.D. and Lyons, R.K. (2002). Order Flow and Exchange Rate Dynamics, *Journal of Political Economy*, 110(1): 170-180.

Frankel, J.A. and Froot, K.A. (1986). Understanding the US Dollar in the Eighties: The Expectations of Chartists and Fundamentalists, *Economic Record*, 62(supplement): 24-38.

Frankel, J. A. and Froot, K. A. (1987). Using survey data to test standard propositions regarding exchange rate expectations. *The American Economic Review*, 77(1):133–153.

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

Gao, T. (1990). Exchange Rate Movements and the profitability of U.S. Multinationals, *Journal of International Money and Finance*, 19(1): 117-134.

Hill, B.M. (1975). A Simple General Approach to Inference About the Tail of a Distribution, *Annals of Statistics*, 3: 1163-1173.

Hodrick, R. J. (1987). *The Empirical Evidence on the Efficiency of Forward and Futures Foreign Exchange Markets*. Harwood Academic Publishers, Chur, Switzerland.

Hommes, C.H. (2006). Heterogeneous Agents Models in Economics and Finance, In: Tesfatsion, L. and Judd, K.L. (eds.), *Handbook of Computational Economics*, Volume 2: Agent-Based Computational Economics, Elsevier Science.

Ito, T. (1990). Foreign Exchange Rate Expectations, *American Economic Review* 80: 434-449.

Kurz, M. (1994). On the Structure and Diversity of Rational Beliefs, *Economic Theory*, 4(6): 877-900.

Kurz, M. and Motolese, M. (2000). Endogeneous Uncertainty and Market Volatility, mimeo Stanford University.

Lyons (2001). *The Microstructure Approach to Exchange Rates*, MIT-Press.

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

Muth, J.F. (1961). Rational Expectations and the Theory of Price Movements, *Econometrica*, 29(3): 315-335.

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

Figure 1: Coefficients of Variation and High-Low Spreads

Panel a: US Dollar / British Pound

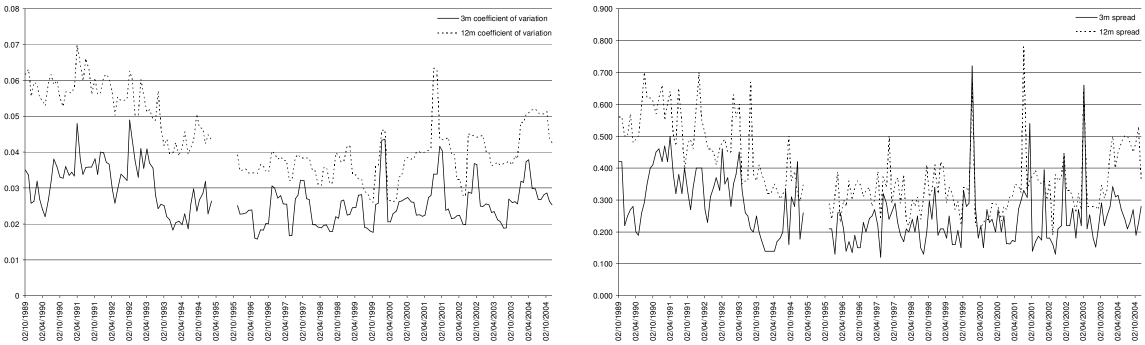

Panel b: Japanese Yen / US Dollar

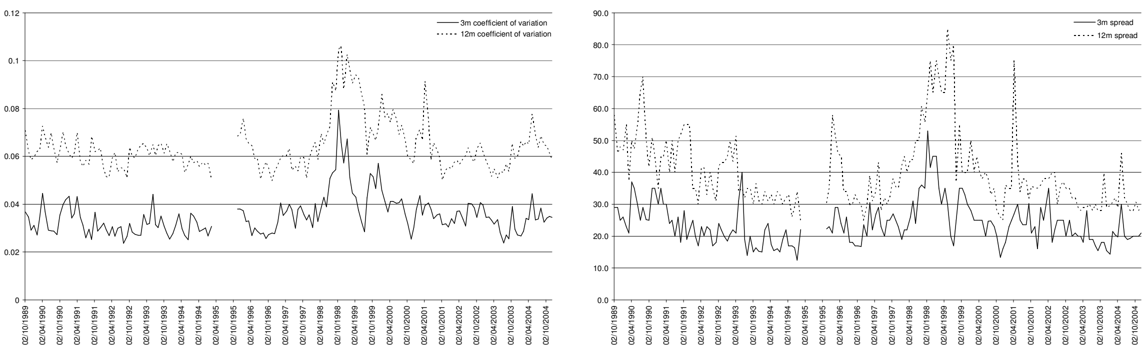

Panel c: US Dollar / Euro

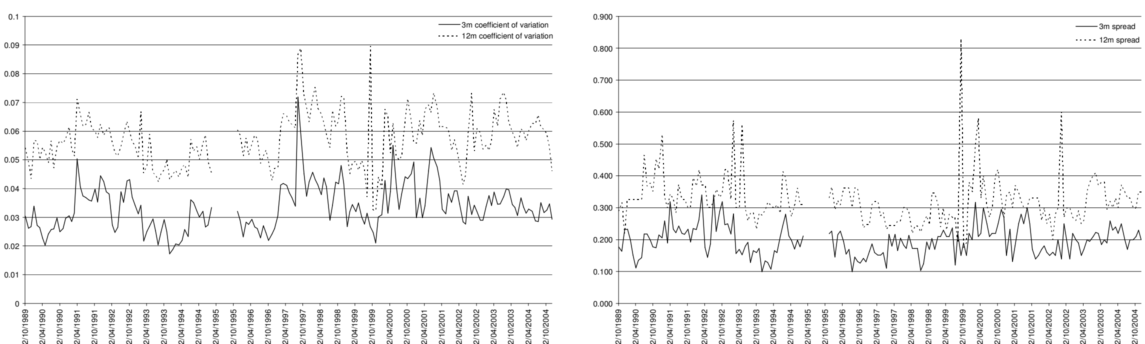

Table 1: Distributional Summary Statistics of the Individual Expectations (pooled)

	3 months horizon (k = 3)			12 months horizon (k = 12)		
	USA/UKI	JAP/USA	USA/EUR	USA/UKI	JAP/USA	USA/EUR
a) Period 1989:10 – 1995:2						
Maximum	0.0643	0.2127	0.0743	0.1113	0.2574	0.1407
Minimum	-0.0840	-0.1693	-0.0736	-0.1515	-0.3044	-0.1067
Standard deviation	0.0134	0.0317	0.0134	0.0232	0.0604	0.0233
Skewness	-0.2143	0.1208	0.1573	-0.0533	0.0927	0.2159
Kurtosis	5.6855	4.9181	4.5204	4.2819	3.5607	3.9239
Bera-Jarque	1961.02***	1071.56***	718.14***	443.08***	101.06***	313.52***
Auto-correlation (1 lag)	0.454***	0.458***	0.446***	0.610***	0.657***	0.613***
b) Period 1995:11 – 2004:12						
Maximum	0.0451	0.0993	0.0669	0.0607	0.1577	0.0595
Minimum	-0.0477	-0.0715	-0.0639	-0.0674	-0.1019	-0.1429
Standard deviation	0.0111	0.0164	0.0140	0.0174	0.0318	0.0240
Skewness	-0.0286	-0.0503	0.0072	-0.2962	0.4178	-0.5962
Kurtosis	4.5418	4.3235	4.4025	4.1308	3.7645	3.8750
Bera-Jarque	109.69***	162.60***	180.89***	74.96***	118.22***	200.87***
Auto-correlation (1 lag)	0.303***	0.506***	0.541***	0.563***	0.738***	0.717***

Notes: A *, **, and *** indicates a rejection of the null hypothesis of normality at the 10, 5, and 1 percent significant level, respectively.

Table 2: Hill Estimates (pooled sample)

	3 months horizon (k = 3)			12 months horizon (k = 12)		
	USA/UKI	JAP/USA	USA/EUR	USA/UKI	JAP/USA	USA/EUR
Sample 1	6364 obs.	6881 obs.	7150 obs.	6427 obs.	6955 obs.	7235 obs.
1 percent	$m = 63$	$m = 68$	$m = 71$	$m = 64$	$m = 69$	$m = 71$
α_l	3.682	5.981	4.917	4.846	6.383	5.460
α_r	5.430	4.521	5.085	4.981	5.557	16.147
T: $\alpha_l = \alpha_r$	2.115**	-1.605	0.208	0.156	-0.810	5.283***
5 percent	$m = 318$	$m = 344$	$m = 357$	$m = 321$	$m = 347$	$m = 357$
α_l	3.173	3.280	3.429	3.846	4.197	3.806
α_r	3.372	3.368	3.274	4.138	3.823	5.168
T: $\alpha_l = \alpha_r$	0.768	0.348	-0.618	0.925	-1.226	4.009***
10 percent	$m = 636$	$m = 688$	$m = 715$	$m = 642$	$m = 695$	$m = 715$
α_l	2.403	2.837	2.719	2.908	3.393	3.136
α_r	2.436	2.648	2.865	2.913	2.984	3.486
T: $\alpha_l = \alpha_r$	0.244	-1.277	0.989	0.033	-2.388**	1.994**
Sample 2	1106 obs.	2215 obs.	2207 obs.	1104 obs.	2212 obs.	2204 obs.
1 percent	$m = 11$	$m = 22$	$m = 22$	$m = 11$	$m = 22$	$m = 22$
α_l	4.325	6.728	4.300	4.932	8.789	4.843
α_r	3.858	5.989	6.271	6.557	3.048	8.794
T: $\alpha_l = \alpha_r$	-0.268	-0.385	1.216	0.657	-2.894***	1.846*
5 percent	$m = 55$	$m = 110$	$m = 110$	$m = 55$	$m = 110$	$m = 110$
α_l	3.428	3.487	3.435	2.958	5.390	3.992
α_r	3.497	3.173	3.474	4.240	4.186	5.996
T: $\alpha_l = \alpha_r$	0.106	-0.696	0.085	1.839*	-1.850*	2.917***
10 percent	$m = 110$	$m = 221$	$m = 220$	$m = 110$	$m = 221$	$m = 220$
α_l	2.568	2.744	2.949	2.658	3.270	3.041
α_r	2.606	2.856	2.612	2.888	3.495	4.002
T: $\alpha_l = \alpha_r$	0.119	0.420	-1.267	0.615	0.698	2.835***

Note: This Table gives the upper, α_r, and lower α_l, tail index estimates for the 1, 5, and 10th percentile of the tails of the distribution of the individual expectations in excess of the cross-sectional average. The T-statistic is for the equality of tails. A *, **, or *** refers to rejection of the null hypothesis that the tail index estimates are equal at a 10, 5, or 1 percent significance level, respectively. m refers to the number of tail observations that are included for the Hill estimates.

Table 3: Percentage Market Participants Showing Heterogeneous Effects

	3 months horizon (k = 3)			12 months horizon (k = 12)		
	USA/UKI	JAP/USA	USA/EUR	USA/UKI	JAP/USA	USA/EUR
a) Period 1989:10 – 1995:2						
versus consensus	37.33	31.61	35.50	38.16	41.51	45.71
versus each other	31.76	33.64	34.39	40.09	42.26	41.25
b) Period 1995:11 – 2004:12						
versus consensus	27.59	19.35	19.35	27.59	32.26	41.94
versus each other	22.87	23.63	25.54	29.01	29.12	33.17

Notes: Table 3 presents the percentage of respondents for which we find significant individual expectations (i.e. $g_i \neq 0$ in equation (3) and $g_i - g_j \neq 0$ in equation (5).

Table 4: Percentage Market Participants Showing Idiosyncratic Effects When Forecasts are Extrapolative

	3 months horizon (k = 3)			12 months horizon (k = 12)		
	USA/UKI	JAP/USA	USA/EUR	USA/UKI	JAP/USA	USA/EUR
a) Period 1989:10 – 1995:2						
versus consensus						
H_0: $\beta_1 = \beta_2 = 0$	17.33	24.52	27.80	17.76	12.58	16.00
H_0: $\beta_1 = \beta_2 = 0$, $h_i = 0$	37.33	41.29	40.82	45.39	40.88	40.00
versus each other						
H_0: $\beta_1 = \beta_2 = 0$	18.73	18.73	23.54	16.42	11.56	13.44
H_0: $\beta_1 = \beta_2 = 0$, $h_i = 0$	35.88	35.88	39.42	40.18	39.33	38.54
b) Period 1995:11 – 2004:12						
versus consensus						
H_0: $\beta_1 = \beta_2 = 0$	41.38	32.26	45.16	13.79	9.68	19.35
H_0: $\beta_1 = \beta_2 = 0$, $h_i = 0$	55.17	41.94	41.94	27.59	32.26	28.57
versus each other						
H_0: $\beta_1 = \beta_2 = 0$	36.69	33.89	37.02	12.15	12.92	19.86
H_0: $\beta_1 = \beta_2 = 0$, $h_i = 0$	43.59	38.11	43.00	28.24	25.92	30.89

Notes: Table 4 presents the percentage of respondents for which we find significant differences in weight put on extrapolation (i.e. $(\beta_{l,i} - \overline{\beta}_l) \neq 0$ and $(\beta_{l,i} - \overline{\beta}_l) \neq 0, h_i - \overline{h} = 0$ in Equation (7)).

Table 5: Percentage Market Participants Showing Idiosyncratic Effects When Forecasters use Uncovered Interest Parity

	3 months horizon (k = 3)			12 months horizon (k = 12)		
	USA/UKI	JAP/USA	USA/EUR	USA/UKI	JAP/USA	USA/EUR
a) Period 1989:10 – 1995:2						
versus consensus						
H_0: $\gamma_1 = 0$	17.33	18.71	19.53	25.00	24.53	29.14
H_0: $\gamma_1 = 0$, $d_i = 0$	34.67	39.35	39.05	50.00	52.20	54.29
versus each other						
H_0: $\gamma_1 = 0$	17.82	21.58	22.00	26.49	22.89	29.07
H_0: $\gamma_1 = 0$, $d_i = 0$	35.60	39.48	40.14	49.86	49.55	53.35
b) Period 1995:11 – 2004:12						
versus consensus						
H_0: $\gamma_1 = 0$	31.03	16.13	19.35	27.59	19.35	25.81
H_0: $\gamma_1 = 0$, $d_i = 0$	51.72	24.00	28.57	31.03	36.67	48.15
versus each other						
H_0: $\gamma_1 = 0$	31.13	24.24	22.91	23.48	18.62	18.38
H_0: $\gamma_1 = 0$, $d_i = 0$	36.99	30.65	33.58	36.09	28.50	32.33

Notes: Table 5 presents the percentage of respondents for which we find significant differences in weight put on interest differences (i.e. $(\gamma_{l,i} - \bar{\gamma}_l) \neq 0$ and $(\gamma_{l,i} - \bar{\gamma}_l) \neq 0, d_i - \bar{d} \neq 0$ in Equation (9)).

Table 6: Pooled Heterogeneity Tests

	3 months horizon (k = 3)			12 months horizon (k = 12)		
	USA/UKI	JAP/USA	USA/EUR	USA/UKI	JAP/USA	USA/EUR
a) Period 1989:10 – 1995:2						
Constant	388.12***	443.46***	397.74***	487.26***	445.50***	520.70***
Extrapolative	135.42	525.58***	577.58***	301.26***	360.64***	318.92***
Uncovered Interest Parity	408.70***	264.92***	416.32***	494.96***	215.04*	424.16***
b) Period 1995:11 – 2004:12						
Constant	64.69***	57.828***	81.01***	56.87***	46.40**	80.67***
Extrapolative	143.262***	109.452***	158.03***	54.03***	60.60***	93.42***
Uncovered Interest Parity	58.902***	61.894***	69.34***	45.71**	60.85***	54.07***

Notes: Table 6 presents the likelihood ratios between the restricted models (coefficient is constant across respondents) and unrestricted models (coefficients are respondent-specific). Chi-square cut-off values for the first period are 213.24; 220.99 and 236.03 for the 10, 5 and 1 percent level respectively (188 d.f.); for the second period these values are 40.26; 43.77 and 50.89 (30 d.f.).

Table 7: Granger Causality USD/EUR

		Absolute returns		Squared returns		GARCH	
		Var → Hetero	Hetero → Var	Var → Hetero	Hetero → Var	Var → Hetero	Hetero → Var
Coefficient of Variation	3M	0.001^{***} (++)	0.587 (++)	0.001^{***} (++)	0.606 (++)	0.790 (+-)	0.000^{***} (+-)
	12M	0.009^{***} (++)	0.172 (-+)	0.009^{***} (++)	0.180 (-+)	0.039^{**} (+-)	0.043^{**} (+-)
Standard Deviation	3M	0.000^{***} (++)	0.599 (+-)	0.000^{***} (++)	0.617 (++)	0.529 (--)	0.000^{***} (+-)
	12M	0.005^{***} (++)	0.243 (-+)	0.005^{***} (++)	0.241 (-+)	0.037^{**} (+-)	0.043^{**} (+-)
Range	3M	0.004^{***} (++)	0.561 (++)	0.002^{***} (++)	0.496 (++)	0.016^{**} (+-)	0.000^{***} (+-)
	12M	0.168 (++)	0.087* (-+)	0.139 (++)	0.092* (-+)	0.017^{**} (+-)	0.398 (+-)
Range Percentage	3M	0.013^{**} (++)	0.666 (++)	0.008^{***} (++)	0.598 (++)	0.062^{*} (+-)	0.001^{***} (+-)
	12M	0.178 (++)	0.080* (-+)	0.140 (++)	0.089* (-+)	0.004^{***} (+-)	0.347 (+-)

Notes: This Table presents *p*-values of the Granger causality tests for two lags for the three and twelve month horizons (3M and 12M) between market volatility (Var) and trader heterogeneity (Hetero). Var → Hetero implies that the nul hypothesis is that causality runs from variance to heterogeneity. The sign of the coefficient for the independent variables on the first and second lag are given in parentheses. *, **, *** denotes significance at the 10, 5 and 1% level respectively.

Table 8: Granger Causality JAP/USD

		Absolute returns		Squared returns		GARCH	
		Var → Hetero	Hetero → Var	Var → Hetero	Hetero → Var	Var → Hetero	Hetero → Var
Coefficient of Variation	3M	0.000*** (++)	0.774 (++)	0.000*** (++)	0.795 (++)	0.001*** (+-)	0.000*** (+-)
	12M	0.010** (++)	0.598 (++)	0.014** (++)	0.652 (+-)	0.007*** (++)	0.008*** (++)
Standard Deviation	3M	0.000*** (++)	0.565 (++)	0.000*** (++)	0.654 (++)	0.050** (+-)	0.000*** (+-)
	12M	0.026** (++)	0.548 (++)	0.035** (++)	0.660 (+-)	0.072* (++)	0.006*** (++)
Range	3M	0.000*** (++)	0.311 (++)	0.000*** (++)	0.358 (-+)	0.002*** (+-)	0.000*** (+-)
	12M	0.156 (++)	0.509 (++)	0.180 (++)	0.600 (++)	0.021** (++)	0.018** (+-)
Range Percentage	3M	0.000*** (++)	0.398 (-+)	0.000*** (++)	0.411 (-+)	0.000*** (+-)	0.000*** (+-)
	12M	0.118 (++)	0.442 (++)	0.132 (++)	0.501 (++)	0.009*** (++)	0.044** (+-)

Notes: This Table presents *p*-values of the Granger causality tests for two lags for the three and twelve month horizons (3M and 12M) between market volatility (Var) and trader heterogeneity (Hetero). Var → Hetero implies that the nul hypothesis is that causality runs from variance to heterogeneity. The sign of the coefficient for the independent variables on the first and second lag are given in parentheses. *, **, *** denotes significance at the 10, 5 and 1% level respectively.

Table 9: Granger Causality USD/UKI

		Absolute returns		Squared returns		GARCH	
		Var → Hetero	Hetero → Var	Var → Hetero	Hetero → Var	Var → Hetero	Hetero → Var
Coefficient of Variation	3M	0.000*** (++)	0.099* (+-)	0.000*** (++)	0.081* (+-)	0.010** (++)	0.025** (++)
	12M	0.000*** (++)	0.047** (++)	0.000*** (++)	0.046** (++)	0.054* (++)	0.172 (++)
Standard Deviation	3M	0.000*** (++)	0.018** (+-)	0.000*** (++)	0.014** (+-)	0.141 (+-)	0.014** (++)
	12M	0.005*** (++)	0.020** (++)	0.007*** (++)	0.022** (++)	0.157 (+-)	0.054* (++)
Range	3M	0.000*** (++)	0.102 (+-)	0.000*** (++)	0.097* (+-)	0.045** (+-)	0.077* (++)
	12M	0.012** (++)	0.095* (++)	0.016** (++)	0.095* (++)	0.245 (++)	0.533 (++)
Range Percentage	3M	0.000*** (++)	0.315 (+-)	0.000*** (++)	0.309 (+-)	0.011** (+-)	0.171 (++)
	12M	0.013** (++)	0.294 (++)	0.014** (++)	0.285 (-+)	0.044** (++)	0.775 (--)

Notes: This Table presents *p*-values of the Granger causality tests for two lags for the three and twelve month horizons (3M and 12M) between market volatility (Var) and trader heterogeneity (Hetero). Var → Hetero implies that the nul hypothesis is that causality runs from variance to heterogeneity. The sign of the coefficient for the independent variables on the first and second lag are given in parentheses. *, **, *** denotes significance at the 10, 5 and 1% level respectively.