Dispersion of Beliefs in Foreign Exchange

Ron Jongen*

Willem F.C. Verschoor#

Christian C.P. Wolff§

Remco C.J. Zwinkels#

January, 2007

Abstract

The role of market microstructure and heterogeneity on behalf of market participants is gaining more attention in the foreign exchange literature. In this empirical paper we address three fundamental questions relating to the dispersed beliefs that market participants hold with respect to future exchange rates. First, we address the question of how we can measure dispersion of beliefs and show 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. Second, we question the source of the dispersion in beliefs and find that it seems to occur as a result of the combined effects of market participants holding individual information and attaching different weights to various elements from their information sets. Third, we examine the causal relation between the change in belief dispersion and the general condition of the market and 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 in Foreign Exchange

Abstract

The role of market microstructure and heterogeneity on behalf of market participants is gaining more attention in the foreign exchange literature. In this empirical paper we address three fundamental questions relating to the dispersed beliefs that market participants hold with respect to future exchange rates. First, we address the question of how we can measure dispersion of beliefs and show 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. Second, we question the source of the dispersion in beliefs and find that it seems to occur as a result of the combined effects of market participants holding individual information and attaching different weights to various elements from their information sets. Third, we examine the causal relation between the change in belief dispersion and the general condition of the market and find that market volatility Granger-causes trader heterogeneity.

1. Introduction

In the previous decades there has been 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 dispersion of beliefs 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 dispersed beliefs comprise of theoretically and how one can measure and quantify dispersion of beliefs, specifically in survey datasets of individual forecasts. We focus on the question what is the source of the possible dispersion. Next we employ these insights to a survey dataset of individual expectations.

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 behavior from the panelists, 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. (2006), 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.

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 dispersed beliefs. In section 4 we introduce the data used for this analysis. Section 4 examines whether expectations are dispersed and in section 5 we try to give an answer as to what the sources of dispersion of beliefs is. Section 6 links dispersion of beliefs to market uncertainty. Finally, section 7 concludes.

2. Dispersion of Beliefs in Foreign Exchange

In this article we are first and fore mostly interested in the expectations, or beliefs, that market participants, or agents, hold with respect to future levels of an exchange rate and in particular to what extent and by which causes these differ from each other. When the expectations of a group of agents differ (for instance, one agent expects a higher appreciation of a currency than another agent, or one agent expects an appreciation while the other expects a depreciation), we state that their beliefs are dispersed. ‘Dispersion of beliefs’ therefore explicitly refers to the difference in *expectations* that the various market participants have with respect to the future status of the market. In related literature, the term ‘heterogeneity’ is often used to express similar concepts. Yet, this term is used in a much broader setting, for instance to refer to the market participants themselves (heterogeneous agents) or to the type of information they posses (heterogeneous information). In order to avoid confusion about the terminology, we prefer the term dispersion of beliefs instead of the term heterogeneity, whenever we talk about the difference in agents’ expectations.

There are broadly three commonly used explanations for the arising of dispersion of beliefs in financial markets. First of all, one strand of the literature argues that dispersion of beliefs arises because of information asymmetries. 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. Even when market participants use the same techniques in forming expectations and try to make the best forecast conditional on the information they have, their expectations will be inherently different due to the fact that all have a different information set. This asymmetry in information may arise because of several reasons. First, when the concept of asymmetric information was introduced in the New Classical Theory of the macro economy, the key reason why agents were assumed to be unable to obtain information that is public in other parts of the economy was that the transmission of public information was rigid.

Another source of asymmetric information is the natural informational advantage that some players have over others. For example, Peiers (1997) reports on how central banks regularly execute their interventionary transactions through selected large commercial, often domestic banks instead of going through the brokers’ market. The central bank can thus benefit from the strong market presence and liquidity from its key intermediary, and the intermediary commercial bank has privileged information regarding the central bank’s activities. Information asymmetry is therefore caused by the fact that central bank activity is revealed in multiple stages. Covrig and Melvin (2002) furthermore attribute part of the private information in the foreign exchange market to the customer order flow: large domestic banks will receive the foreign exchange business of large domestic corporations and this confers a temporary informational advantage on large domestic banks. The costs of asymmetric information to market participants can be substantial even today. By analyzing trades and quotes in the Deutschemark - U.S. Dollar market, Payne (2003) shows that asymmetric information can account for around 60 percent of the average bid-ask spread.

There are, on the other hand, several arguments and empirical considerations against the claim that all of the 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 are all competitive in the search for information, one would expect that each of them should possess essentially the same information. Second, even when a particular market participant possesses private, inside 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.

A second strand of the literature therefore 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 how they should interpret this information. A popular approach in the discussion about why agents interpret information differently is the rational beliefs theory due to Kurz (1994). In this theory, the disagreement is caused by the fact that economic agents do not know the structural relations of the economy, such as demand and supply functions. 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. They form their probability beliefs about the future by using the empirical distribution that is derived from the relative frequency of events in the past. Their own experience and success in forming accurate forecasts will likely play a role in selecting and valuing information. Given a set of identical data, different agents will then add different weights to the various elements of the data, to express their relative importance which they think the various elements have when forming an expectation about the future.

The rational beliefs theory therefore stands in contrast to the more commonly used theory of rational expectations, where agents are supposed to know a lot about the structure of the economy. It was argued that it is unreasonable to assume that economic agents possess skills and structural knowledge that no human being can possess. Structural knowledge can be gradually acquired by agents though, for instance through learning, but the literature on learning has thus far not been successful in justifying a fully rational agent.

An important observation in the rational beliefs theory is that market risks are determined by the individual beliefs of the various market participants and that a change in the distribution of these beliefs therefore affects the volatility of the market. Kurz and Motolese (2001, p. 497) in fact claim that the "distribution of beliefs in the market is the most important propagation mechanism of economic volatility", in other words, much of the volatility in the market is generated by the individual beliefs of the market participants and the correlation between these beliefs.

A third 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 beliefs. On the one hand, fundamentalists think of the exchange rate as the output of an economic model, while on the other hand 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, and that the use of technical forecasting tools has increased over the past decades.

The realization that the presence of fundamentally different interacting market participants can shed new light on some of the anomalies in the financial economics literature has resulted in a strand of literature on how to model these heterogeneous agents (see Hommes, 2006, for an overview of this literature). Most heterogeneous agents models are based on the simple chartists – fundamentalist distinction as well, but revolutionary is the fact that agents can switch between trader types. Switching occurs as a result of differences in forecasting performance of the different strategies.

Simulations show that heterogeneous agents models are capable of replicating the stylized facts from financial markets, see e.g. De Grauwe and Grimaldi (2005). Periods in which (extrapolating) chartists are dominating the market are characterized by high volatility, while periods dominated by (mean-reverting) fundamentalists are characterized by low volatility. This combination yields an economic intuitive explanation for the existence of volatility clustering, heavy tailed return distributions, excess volatility, and slow mean reversion. The first estimation results indicate the existence of switching heterogeneous agents in different financial markets; see Boswijk et al. (2007) for the S&P500 and De Jong et al. (2006) for the EMS exchange rates.

The market microstructure approach uses dispersion of beliefs as well highlights the importance of the market microstructure theory in explaining the complex short-term behavior of the exchange rate (Evans and Lyons, 2002). Using order flow as a proximate determinant, they develop a model that is strikingly successful in accounting for realized exchange rate changes. In their analyses, they show that the order flow conveys dispersed information and that the distribution of information is an important determinant in short-run exchange rate movements. Moreover, the large volumes in the exchange rate market are explanatory for short-term movements and an indication that agents have different information or process information differently. In other words, a high dispersion of beliefs causes high volatility. As shown by Fan and Lyons (2001), informative trades are mixed with uninformative trades, indicating that market efficiency crucially depends on how markets accomplish the difficult task of aggregating dispersed information.

In the next sections we attempt to deal with the first two arguments of why dispersion of beliefs occurs, namely because of asymmetry in information and/or disagreement about how information should be interpreted. In a follow up paper we look at the relative role of fundamentally different types of market participants.

3. Data

To investigate the behavioral aspects of the beliefs of market participants, and in particular the dispersion thereof between participants, we use a unique database of survey forecasts. The individual forecasts are obtained from a survey conducted by Consensus Economics of London on a monthly basis among the 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 panelists companies are located worldwide, although they are all from developed economies (notably the U.S., Japan, U.K., France, Germany, Canada etc.). The forecasts are point forecasts for a large set of currencies against the U.S. Dollar or the Deutschemark (later the Euro) and are available for horizons of 1, 3, 12, and 24 months ahead. The names of the panelist companies are revealed.

Although survey participants have a few days time to return their forecasts, we learned that the vast majority send their responses by e-mail on the Friday before the publication day (usually the second Monday of the month). We consider this Friday to be the day on which their beliefs are formed and assume that their beliefs are translated one-to-one in a point forecast. There may be reasons for panelists not to reveal their true beliefs, though. 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. When the names of the forecasters are given in the survey publication, agents have an incentive to perform well in order to attract customers.

For the Friday on which the beliefs are thought to be formed, 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 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. Other data, such as interest rates, are also obtained through DataStream.

For our sample, we obtain the forecasts for the U.K. Pound, Japanese Yen and Euro against the U.S. Dollar 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 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

[1] Prior to January 1999 we use forecasts on the Deutschemark versus the U.S. Dollar. We transform these forecasts into Euro / U.S. Dollar forecasts using the official conversion rate.

1995 (Sample 1) and November 1995 through December 2004 (Sample 2). In this paper we focus on the 3- and 12-month-ahead forecasts since these have the highest response rates.

Insert table 1 here

Table 1 provides an overview of the survey dataset that is used and gives an indication of how many responses are obtained from the panelists per time period. We find that the median number of responding participants is over 100 per month for the Japanese Yen and Euro forecasts, while slightly less than 100 per month for U.K. Pound forecasts.

One way in which we may think of dispersed beliefs is that market participants disagree as to what the market will do in the future. In order to get an impression of how the dispersion of beliefs changes over time, we therefore need to define a measure of dispersion based on the cross-section of the forecasts. Frankel and Froot (1990) and Beine et al. (2002) for instance use functions of the standard deviation across respondents in a survey 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 beliefs of where the market is going become more dispersed. 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 number of methods to measure or quantify dispersion of beliefs is small because of the relatively scarcity of data on individual (survey) forecasts. We therefore follow the formal approach proposed by Shalen (1993), who describes the theoretical relationship between market volatility, volume, and the dispersion of beliefs. She develops a measure of the dispersion of beliefs that is related to the weighted average difference between the individual beliefs and the average beliefs about a particular asset. When $E_{i,t}(s)$ is agent *i*'s true belief about the future exchange rate and $E_t(s)$ the average belief, then Shalen defines the dispersion of beliefs as

$$H_t = E(\sum_i w_{i,t}(D_{i,t})^2) \equiv E(\sum_i w_{i,t}(E_{i,t}(s) - E_t(s))^2), \qquad (1)$$

where the weight $w_{i,t}$ is determined by the relative precision of agent *i*'s forecast of the most recent spot rate, while adjusting for her risk tolerance.

In a survey-based framework, we assume that the participants' forecasts are determined at the same moment when their beliefs about the future are generated, that their forecasts reflect their beliefs, and that the sample reflects the true population. The empirical, survey-based counterpart to the unknown $D_{i,t}$ is then given by

$$d_{i,t} \equiv s_{i,t}^{e} - \bar{s}_{t}^{e}, \tag{2}$$

where we define $s_{i,t}^{e}$ as the survey-based forecast of the future k-period ahead spot rate[2] and $\bar{s}_{t}^{e}$ the cross-sectional average. In case we make the assumption that the weight $w_{i,t}$ is equal for all market participants, then the resulting statistic is the variance of the individual survey forecasts. Taking the square root and dividing by the cross-sectional average $\bar{s}_{t}^{e}$ then presents the dispersion of belief statistic,

$$h_t = \frac{1}{\bar{s}_{t}^{e}} \left(\sum_{t} \frac{1}{n} (d_{i,t})^2 \right)^{1/2}, \tag{3}$$

as the coefficient of variation, defined as the cross-sectional standard deviation divided by the cross-sectional average, which is a dimensionless statistic that allows the comparison of the variability of populations that have a different mean, such as comparing two exchange rates.[3]

Insert figure 1 here

Figure 1 presents this coefficient of variation along two forecast horizons for the three currencies. In addition, high-low spreads are presented as a robustness check.[4] 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.

[2] We do not use a sub- or superscript for the forecast horizon k in this paper to keep the algebra tractable.

[3] We check to what extent our measure of dispersion is driven by either the cross-sectional standard deviation or the cross-sectional average. The correlation with the former is typically between 85 and 97 percent for all currency and forecast horizon combinations. The correlation with the latter is never more than 56 percent. Thus, the variation in our measure of dispersion is primarily driven by the volatility of the individual forecasts.

[4] Using the spread as a measure of dispersion 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) forecast directly influence the measure of dispersion. Therefore, we only use the spread as a robustness check.

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 agrees more about what will happen in the nearby future than in the distant future. It should be mentioned though that this is not evidence that market participants can better predict short term changes over long term changes.

Third, dispersion of beliefs is not constant over time, but seems to be high in some periods of considerable length, while low in others. In particular for the U.K. Pound and for the Euro dispersion of beliefs was obviously higher in the period of early 1990 to late 1993 than in the subsequent period. Similarly, for the Japanese Yen rate beliefs 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.

4. Are the Beliefs of Foreign Exchange Forecasters Dispersed?

One natural question is what we can learn from the distribution of the differences between the individual forecasts and the cross-sectional average, $d_{i,t}$. Under the null hypothesis that beliefs are not dispersed, these differences (or alternatively differences between contemporaneous individual forecasts) 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 the distribution, since under the null hypothesis of no dispersed beliefs 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.

We proceed by examining formally whether beliefs are dispersed. To do so we follow the methodology due to Ito (1990), who describes a way of how individual agents form their beliefs about future exchange rates. Suppose individual i generates a forecast[5] about the k-period ahead level of a particular exchange rate and that this forecast consists of a common structural part based on public information (that is common to all market participants) and an individual effect g_i. The individual forecast $s^e_{i,t}$ can then be described as

[5] Recall that we made the assumption that the survey-based forecast is as close as we can get to the true, yet unknown belief of the agent.

$$s_{i,t}^{e} = f(\Omega_t) + g_i + \varepsilon_{i,t}\,, \tag{4}$$

where $\varepsilon_{i,t}$ is a pure random disturbance term (with respect to *i* and *t*) that can, for instance, be the result of a measurement. In a similar way can the cross-sectional average of a set of market participants' forecasts be defined as

$$\bar{s}_{t}^{e} = f(\Omega_t) + \bar{g} + \bar{\varepsilon}_t\ . \tag{5}$$

Suppose that a normalization such that $\bar{g} = 0$ is possible (when $f(\Omega_t)$ contains a constant term). Subtracting the cross-sectional average from the individual forecast then gives

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

or

$$d_{i,t} = g_i + (\varepsilon_{i,t} - \bar{\varepsilon}_t\ )\,. \tag{6}$$

In this case it is not necessary to know the exact structure of $f(\Omega_t)$, as long as it is common to all agents. Table 2 presents summary statistics for the distribution of the individual expectations in excess of their cross-sectional average $d_{i,t}$, which should give us information about the individual effect g_i. The statistics are generated from the sample of pooled data.

Insert table 2 here

Several observations are noteworthy. Both the high-low spread (maximum minus minimum forecast) and standard deviation increase as the forecast horizon increases. This indicates an increase in dispersion of beliefs as the forecast horizon lengthens, as we have seen in figure 1. Furthermore, the deviations are significantly and positively auto-correlated and the pattern of autocorrelation is stronger on longer forecast horizons and for the second sample period.[6] Since the autocorrelation pattern does not extend beyond the first lag (not reported), we can conclude that the error terms in equation 6 follow either an AR(1) or MA process with an unknown number of lags. This correlation can be explained from the overlapping pattern of the survey forecasts: Hansen (1982) and Hansen and Hodrick (1980) demonstrate that, when the forecast horizon is longer than the observational frequency (which

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

is one month), the error term in equation 4 will be serially correlated up to a moving average process of order $k - 1$. The error component in equation 6 therefore inherits this structure. In the remainder of the article we will correct for this moving average structure.

Although the individual forecasts do not appear to be skewed in a particular direction, the distributions are consistently leptokurtic. While high kurtosis is not analogous to fat tail behavior, we may state that an excessive number of forecasts are at enough distant from the consensus forecast 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 of the distribution of individual deviations 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 the survey-based forecasts for a particular exchange rate are independently and identically distributed. We then define D as the set of all $d_{i,t}$ for all i and t. Define $d_{n-m,n}$ to be the $(n - m)$'th ascending order (or in other words, the m-th smallest expected change) from the set of individual expectations d. The Hill estimator can then be defined as

$$\hat{\alpha} = \left[\frac{1}{m} \sum_{M=0}^{m-1} \ln\left(\frac{d_{n-M,n}}{d_{n-m,n}} \right) \right]^{-1}, \tag{7}$$

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 D.

Insert table 3 here

Table 3 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^{th} moment in the distribution and indicates that extreme expectations occur more often than would be 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 significantly dispersed and the assumption of a representative agent can be rejected.

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 should be significantly bigger than the left tail. Using a two-sided *T*-test we formally test for equality of the left and right 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 Euro. 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. In the next section we go a step further and question what the source of this dispersion is.

5. Asymmetric Information with Individual and Idiosyncratic Effects

It remains a matter of debate whether dispersion of beliefs arise because market participants hold different information sets whereby part of the information is private, whether they attach different weights to some elements of a common information set, or a combination of both. To deal with this question we should ideally look at evidence of dispersed beliefs for each market participant separately. Reconsider equation 6

$$d_{i,t} = g_i + (\varepsilon_{i,t} - \bar{\varepsilon}_t).$$

Here, we assumed that the public information set Ω_t is common for all market participants. 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 forecasts in excess of the cross-sectional average on a constant. The significance of the g_i term indicates whether individual effects are present for agent *i*, or in other words that this agent bases her forecasts on an informational set that is unique for this agent.[7]

[7] 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: $s^e_{i,t} - s^e_{j,t} = g_i - g_j + (\varepsilon_{i,t} - \varepsilon_{j,t})$. However, the estimation results are virtually identical to those of equation 6. These results are available on request.

We therefore estimate equation 6 using a least squares regression for each market participant. Since the forecast horizon is longer than the observational frequency, the error component in equation 6 will be serially correlated up to a moving average process of order k – 1. While least squares estimates of equation 6 remain consistent in spite of the serially correlated residuals, the least squares standard errors are biased. Therefore, we adjust the standard errors nonparametrically using the procedure of Newey and West (1987, 1994). employing the Bartlett kernel. This correction will also be applied to all variations of this equation further on in the paper.

Thus far this reasoning assumes that dispersion of beliefs arises due to differences in information sets, i.e. ‘private’ information. In addition to holding private information, it is reasonable to assume that although the remainder information may be common to all market participants, they interpret this information differently and hence attach different weight to its various elements. We call the different weights idiosyncratic effects. Ito (1990) and MacDonald and Marsh (1996), for instance, 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.

We consider two models that may explain how market participants form their forecasts. First, we assume that agents follow a simple autoregressive forecasting rule, where they extrapolate the most recent trend in the exchange rate into the future. In essence this is a technical (or chartist) tool. Second, we assume that agents use an uncovered interest parity approach, where the expected change in the level of the exchange rate is related to the relative interest rate level of the domestic and the foreign country. This is on its turn a more fundamentals-based approach. If the fit of a certain model of expectation formation is better for one agent compared to the other and vice versa, this is not only a direct proof of dispersion of beliefs, but it also lays bare the source of the dispersion, namely on which variables the agents put different weights. Audretsch and Stadtmann (2005) determine the optimal model for each agent by using the R^2 of these types of regressions. These statistics are informative in that they reveal to what extent the forecast is based on observable information. The remaining variance is then caused by other variables, such as psychological, experience, etc.

If market participants form forecasts in an extrapolative way, then

$$s_{i,t}^{e} - s_{t} = h_{i} + \sum_{l=1}^{L} \beta_{l,i} (s_{t-l} - s_{t-l+1}) + \varsigma_{i,t} \ . \tag{8}$$

where L is the number of lags that is used in the forecast. This equation is a specification to equation 4, where we have subtracted s_t from both sides of the equality sign. In fact, we have impose a structure for the unknown $f(\Omega_t)$. Following the same procedure as in equations 4-6 we find that

$$d_{i,t} = h_i - \overline{h} + \sum_{l=1}^{L} (\beta_{l,i} - \overline{\beta}_l)(s_{t-l} - s_{t-l+1}) + \varsigma_{i,t} - \overline{\varsigma}_t \ . \tag{9}$$

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 10 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.

We continue the analysis by questioning whether uncovered interest parity (UIP) has appeal as an empirical fundamental model for the formation of beliefs and whether market participants have idiosyncratic beliefs as to the importance in the relative level of interest rates. Uncovered interest parity assumes that the change in the level of the exchange rate over a period k is related to the level of the k-period domestic interest rate relative to its foreign counterpart, such that

$$s_{i,t}^{e} - s_t = h_i + \gamma_i (i_t - i_t^{*}) + \varsigma_{i,t}\,, \tag{10}$$

where i_t is the k-period domestic interest rate and i_t^{*} the k-period foreign interest rate. Following the same line of reasoning as under equations (4)-(6), we can rewrite equation (10) as

$$d_{i,t} = h_i - \overline{h} + (\gamma_i - \overline{\gamma})(i_t - i_t^{*}) + \varsigma_{i,t} - \overline{\varsigma}_t \ . \tag{11}$$

We measure i_t and i_t^{*} as the domestic and foreign interbank rates, respectively, using the 3- and 12-month interbank rates for the two forecasting horizons.

The results for equations 6, 9 and 11 are summarized in table 4. In this table we present the percentages of market participants for which we find individual effects, idiosyncratic effects with respect to an extrapolating model, and idiosyncratic effects with respect to uncovered interest parity, all compared to the cross-sectional average. Results are for Wald tests on the significance of the parameters or combinations thereof as to discriminate between the various effects.

Insert table 4 here

For around 40 percent of the individual market participants we find evidence of individual effects, $h_i - \overline{h} \neq 0,$ on the short forecast horizon. When the forecast horizon extends to one year, this percentage increases to around 50 percent. In the second sample period the percentages decreases to around 30 percent for the short horizon, but remain at 50 percent for the long horizon.

Furthermore, in addition to these 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 15 (20) percent of the market participants show idiosyncratic effects with respect to extrapolation, so $\beta_{l,i} - \overline{\beta}_l \neq 0,$ in the short run for the first (second) sample period.[8] This effect is more pronounced at the longer forecast horizon, for both sample periods and all exchange rates. The percentage of respondents with significant individual effects remains comparable after introducing the idiosyncratic extrapolation. The percentage of respondents that exhibit both individual and idiosyncratic extrapolation effects, so the ones for which $h_i - \overline{h} \neq 0$ and $\beta_{l,i} - \overline{\beta}_l \neq 0,$ is in general considerably lower than the total effects, suggesting that respondents show either individual or idiosyncratic effects. The combination of both effects occurs more frequently for the longer horizon beliefs with up to 32 percent of the agents subject to both effects.

The percentage of respondents with significant idiosyncratic UIP effects is larger than extrapolation effects, implying that respondents' beliefs on the effect of the interest rate differential on the exchange rate are highly dispersed. Specifically, for the first period, we find that around 25 (40) percent of the respondents exhibit significant UIP effects for the short (long) horizon. For the second period, these percentages are both around 35 percent. The percentage of respondents with individual effects after including the idiosyncratic UIP effects is again comparable to the results without UIP effects. Contrary to the case of the idiosyncratic extrapolation effects, we find that the majority of respondents that have

[8] Beware that we used two lags in the specification of the extrapolation model. For the remainder of the article, we therefore mean with $\beta_{l,i} - \overline{\beta}_l$ a vector of both $\beta_{1,i} - \overline{\beta}_1$ and $\beta_{2,i} - \overline{\beta}_2$.

idiosyncratic UIP effects also have individual effects, that is the percentage of agents for which $h_i - \overline{h} \neq 0$ and $\gamma_i - \overline{\gamma}_l \neq 0$ is between 15 to 25 percent, also for the shorter forecast horizon.

A claim that is occasionally made with these types of expectation formation models is that when one particular model is examined in a single equation, the results may be affected by a missing variable bias. It is likely that agents use various models when forming a belief about the future instead of just one technique and therefore the omission of a set of variables in an equation with just a single explanatory variable may lead to this bias. In order to obviate this missing variable bias, we provide as a robustness test a combination of equations (6), (9) and (11) that contains individual effects and both technical (lags) and fundamental (interbank) information:

$$d_{i,t} = h_i - \overline{h} + \sum_{l=1}^{L} (\beta_{l,i} - \overline{\beta}_l)(s_{t-l} - s_{t-l+1}) + (\gamma_i - \overline{\gamma})(i_t - i_t^*) + \varsigma_{i,t} - \overline{\varsigma}_t \ . \tag{12}$$

Again, we use two lags for L and the standard errors are robust to the overlapping data problem. Estimation results for this equation are displayed in table 5.

Insert table 5 here

The results for the combined test indicate that the results from the individual regressions in table 4 are not driven by missing variable bias: the percentages of dispersion of beliefs remain at a comparably high level. For the first period, the fraction of respondents with individual effects decreases somewhat compared to the results in table 4, from 40 to around 35 percent for the 3-months horizon and from 50 to somewhat over 40 percent for the 12-months horizon. The dispersion vis-à-vis the extrapolation, on the other hand, increases, from 15 (20) to 20 (30) percent for the short (long) horizon. The same conclusion holds for the idiosyncratic UIP effects: from 25 (40) to 30 (over 40) percent for the short (long) horizon. For the second period, we cannot draw any conclusions per effect, but we can per currency. For the U.K. Pound and the euro, we observe that all percentages decrease; for the Japanese Yen there is a positive tendency.

Insert table 6 here

We look at various combinations of the hypothesis with respect to the three effects. Table 6 reports percentages of market participants for which at least one of the three effects has been identified and splits the sample up into subsets for which the indicated equality or inequality holds.[9] The results corroborate what we found earlier. This is concluded from the fact that the highest percentages in table 6 can be found in the top three rows of each panel, i.e. the percentages of respondents with a single significant effect. This conclusion holds particularly for the first period, but can also be distilled for the second. Specifically, respondents that do combine two effects, combine the individual with the idiosyncratic UIP effect, or, to a lesser extent combine all three effects. Finally, in all three tables (4, 5, and 6) it comes forward that the individual effect is the most important source of dispersion, followed by the idiosyncratic UIP effect and, third, the idiosyncratic extrapolative effect.

We should mention two drawbacks of the above approach. First is the matter of variable selection. The exact set of variables that is commonly used to form beliefs is not known. Second, some of the variables used in forming beliefs are inherently unobservable, such as the state of mind of the agent 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 beliefs are dispersed 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 lead to biased results, since it assumes a permanent difference in weight through time. Therefore, the question of whether beliefs are dispersed and if so, what the source of this dispersion is, should ideally be approached from various angles.

In this section we have examined the question of what the source of dispersion of beliefs is. Using various techniques we showed that the dispersion of beliefs arises 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 a relationship between the dispersion of beliefs are certain market characteristics.

[9] For instance, $c_i - \overline{c} \neq 0 \wedge \beta_{l,i} - \overline{\beta}_l = 0 \wedge \gamma_i - \overline{\gamma} \neq 0$ indicates the percentage of agents for which we find individual effects, idiosyncratic effects due to different weights for UIP, but no idiosyncratic effects with respect to different weights for the most recent change in the exchange rate (extrapolation).

6. Causality and Market Volatility

In the previous section we have shown that the dispersion in beliefs is caused by the fact that market participants hold private information and that they attach different weights to the various elements from the set of common information. The next step is to see whether the degree of dispersion 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 dispersion of beliefs to market volatility.[10] 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 dispersion and market volatility because the expectations of the different groups are opposite to each other.[11] 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 dispersion 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 dispersion of beliefs affects the volatility in the market. 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, from surveys conducted by Money Market Services International and The Economist, that dispersion of opinion affects the volume of trade and thereby also market volatility.

[10] See Hommes (2006) for an extensive survey.
[11] 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.

We study the causal relation between market volatility and dispersion by means of Granger causality tests. Our approach is particularly useful in that we consider a long time period, we use different measures of dispersion and volatility, and we focus on the causality question instead of correlations. In the previous sections we argued that dispersion can be measured in various ways. Similarly, volatility can be quantified by several different methods. Altogether, we examine twelve different dispersion–volatility measure combinations in order to check the robustness of the results.

In a standard setup, the Granger-causality test regresses the *lagged* values of two variables X and Y on X and Y itself such that the test gives an indication concerning the causality between X and Y instead of the correlation. The setup of the standard Granger test is slightly altered here. In our analysis, in the case of dispersion as the dependent variable, we do not use the lagged but the contemporaneous value of market volatility[12]. The reason for this is that dispersion 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 U.K. Pound we only use the first set (1989 – 1995) since the second set contains two-monthly observations. The three and twelve month forecasts 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 dispersion measures are the coefficient of variation and the high-minus-low range. The volatility measures are the monthly absolute returns, the monthly squared returns and a GARCH - measure.[13] 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.

Insert table 7 here

---

[12] Plus lagged values of volatility (if necessary) and dispersion itself.

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

Table 7 presents the results for the three exchange rates. For the Euro, looking at the absolute and squared returns for the three-month horizon, we find a significantly positive causal relation running from the variance to dispersion for both dispersion measures. The twelve months horizon gives a more mixed image. For the coefficient of variation we find similar results as for the three months expectations, but concerning the range 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 dispersion 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.

For the Japanese Yen we see a significant positive causal link from the market variance to dispersion for both dispersion variables and horizons, at least when variance is measured by absolute and squared returns. In the case of the GARCH measure we find a two-way causal relation for both dispersion 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.K. Pound, finally, reflect a relatively homogeneous conclusion. For all variance and dispersion 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 dispersion 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 dispersion 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.

7. Conclusion

In this article we examined whether expectations of future exchange rates are dispersed in that agents have different beliefs about the future path of the exchange rates. We approach this problem using a panel of survey forecasts 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 beliefs are dispersed 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 dispersion of beliefs 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 dispersion of beliefs 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

Adler, M., and B. Dumas, (1984). Exposure to Currency Risk: Definition and Measurement, *Financial Management*, 13(2): 41-50.

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

Ang, A, Bekaert, G., and Wei, M. (2006). Do Macro Variables, Asset Markets or Surveys Forecast Inflation Better?, *Journal of Monetary Economics*, forthcoming.

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.

Boswijk, H.P., Hommes, C.H., and Manzan, S. (2007). Behavioral Heterogeneity in Stock Prices, *Journal of Economic Dynamics and Control*, forthcoming.

Cavaglia, S., Verschoor, W. F. C., and Wolff, C. C. P. (1993). 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.

Covrig, V. and Melvin, M. (2002). Asymmetric Information and Price Discovery in the FX Market: Does Tokyo Know More About the Yen?, *Journal of Empirical Finance*, 9(3): 271-285.

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

De Jong, E., Verschoor, W.F.C., and Zwinkels, R. (2006). Behavioural Heterogeneity and Shift-Contagion: Evidence from the Asian Crisis. NiCE Working paper.

De Long, J.B., Shleifer, A., Summers, L.H., and Waldmann, R.J. (1990). Noise Trader Risk in Financial Markets, *Journal of Political Economy*, 98(4): 703-738.

Dokko, Y. and Edelstein, R.H. (1989). How Well do Economists Forecast Stock Market Prices? A Study of the Livingston Surveys, *American Economic Review*, 79(4): 865-871.

Elliott, G. and Ito, T. (1999). Heterogeneous Expectations and Tests of Efficiency in the Yen/Dollar Forward Exchange Rate Market, *Journal of Monetary Economics*, 43(2): 435-456.

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.

Fan, M. and Lyons, R.K. (2001). Customer Trades and Extreme Events in the Foreign Exchange, in: Mizen, P. (ed.) *Monetary History, Exchange Rates and Financial Markets*. Edward Elgar, Northampton.

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. *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.

Froot, K.A. (1989). New Hope for the Expectations Hypothesis of the Term Structure of Interest Rates, *Journal of Finance*, 44(2): 293-305.

Froot, K.A. and Thaler, R.H. (1990). Anomalies: Foreign Exchange, *Journal of Economic Perspective*, 4(3): 179-192.

Friedman, B.M. (1979). Interest Rate Expectations Versus Forward rates: Evidence from an Expectations Survey, *Journal of Finance*, 34(4): 965-973.

Friedman, B.M. (1980). Survey Evidence of the ‘Rationality’ of Interest Rate Expectations, *Journal of Monetary Economics*, 6(4): 453-465.

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.

Hansen, L.P., (1982). Large Sample Properties of Generalized Method of Moments Estimators, *Econometrica,* 50(4): 1029-1054.

Hansen, L.P. and Hodrick, R.J. (1980). Forward Exchange Rates as Optimal Predictors of Future Spot Rates: An Econometric Analysis, *Journal of Political Economy*, 88(5): 829-853.

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.

Jorion, P. (1990). The Exchange Rate Exposure of U.S. Multinationals, *Journal of Business*, 63(3): 331-345.

Keane, M.P. and Runkle, D.E. (1990). Testing the Rationality of Price Forecasts: New Evidence from Panel Data, *American Economic Review*, 80(4): 714-735.

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

Kurz, M. and Motolese, M. (2001). Endogeneous Uncertainty and Market Volatility, *Economic Theory*, 17(3): 497-544.

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

MacDonald, R. and Macmillan, P. (1994). On the Expectations View of the Term Structure, Term Premia and Survey-Based Expectations, *Economic Journal*, 104(425): 1070-1086.

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

MacDonald, R. and Torrance, T.S. (1988). On Risk, Rationality and Excessive Speculation in the Deutschmark-US Dollar Exchange Market: Some Evidence Using Survey Data, *Oxford Bulletin of Economics and Statistics*, 50(2): 107-123.

Marsh, I.W. and Power, D.M. (1996). A Note on the Performance of Foreign Exchange Forecasters in a Portfolio Framework, *Journal of Banking and Finance*, 20(3): 605-613.

Newey, W.K. and West, K.D. (1987). A Simple, Positive Semi-Definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix, *Econometrica*, 55(3): 703-708.

Newey, W.K. and West, K.D. (1994). Automatic Lag Selection in Covariance Matrix Estimation, *Review of Economic Studies*, 61(4): 631-653.

Payne, R. (2003). Informed Trade in Spot Foreign Exchange Markets: An Empirical Investigation, *Journal of International Economics*, 61(2): 307-329.

Peiers, B. (1997). Informed Traders, Intervention, and Price Leadership: A Deeper View of the Microstructure of the Foreign Exchange Market, *Journal of Finance*, 52(4):1589-1614.

Shalen, C.T. (1993). Volume, Volatility, and the Dispersion of Beliefs, *Review of Financial Studie*s, 6(2): 405-434.

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.

Table 1: Data Summary

	U.K. Pound	Japanese Yen	Euro (Deutschemark)
a) Period 1989:10 – 1995:2			
Survey frequency	Monthly	Monthly	Monthly
Min. # panelists	89	94	99
Max. # panelists	112	120	127
Median # panelists	99	108	112
b) Period 1995:11 – 2004:12			
Survey frequency	Bi-monthly	Monthly	Monthly
Min. # panelists	70	82	85
Max. # panelists	108	123	128
Median # panelists	89	103	102

Notes: The medians are rounded.

Figure 1: Coefficients of Variation and High-Low Spreads

Panel a: U.K. Pound

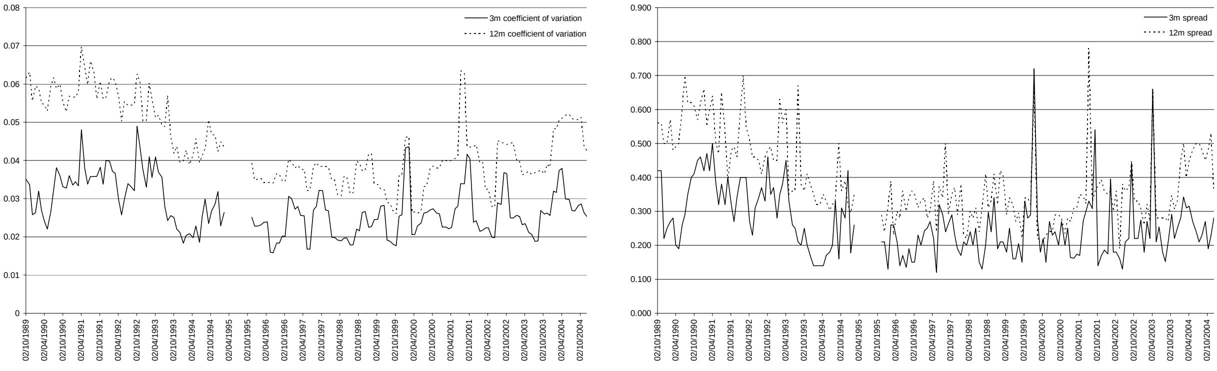

Panel b: Japanese Yen

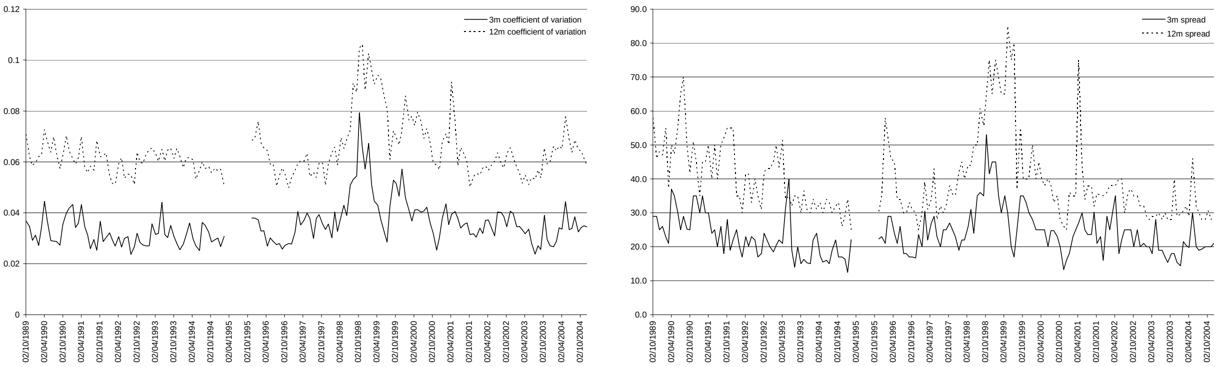

Panel c: Euro

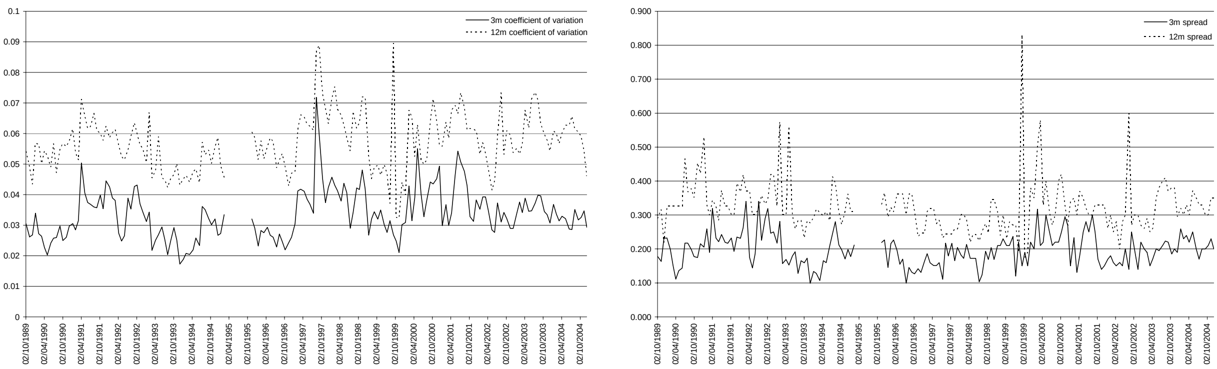

Table 2: Distributional Summary Statistics of $d_{i,t}$ (pooled sample)

	3 months forecast horizon			12 months forecast horizon		
	U.K. Pound	Japanese Yen	Euro	U.K. Pound	Japanese Yen	Euro
a) Period 1989:10 – 1995:2						
Maximum	0.1481	0.2127	0.1710	0.2563	0.2574	0.3239
Minimum	-0.1934	-0.1693	-0.1695	-0.3487	-0.3044	-0.2457
Standard deviation	0.0319	0.0317	0.0308	0.0536	0.0604	0.0536
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***
Autocorrelation (1st lag)	0.454***	0.467***	0.446***	0.610***	0.657***	0.613***
b) Period 1995:11 – 2004:12						
Maximum	0.1037	0.2287	0.1540	0.1399	0.3630	0.1371
Minimum	-0.1098	-0.1646	-0.1471	-0.1553	-0.2346	-0.3290
Standard deviation	0.0255	0.0378	0.0322	0.0400	0.0732	0.0554
Skewness	-0.0286	-0.0504	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***
Autocorrelation (1st lag)	0.303***	0.506***	0.541***	0.563***	0.738***	0.717***

Notes: The figures are in logarithms, so they can be seen as the percentage deviations from the consensus average expectations. A *, **, and *** indicates a rejection of the null hypothesis of normality at the 10, 5, and 1 percent significant level, respectively.

Table 3: Hill Estimates (pooled sample)

	3 months forecast horizon			12 months forecast horizon		
	U.K. Pound	Japanese Yen	Euro	U.K. Pound	Japanese Yen	Euro
a) Period 1989:10 – 1995:2						
	6364 obs.	6881 obs.	7150 obs.	6427 obs.	6955 obs.	7235 obs.
1th percentile	$m = 63$	$m = 68$	$m = 71$	$m = 64$	$m = 69$	$m = 71$
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***
5th percentile	$m = 318$	$m = 344$	$m = 357$	$m = 321$	$m = 347$	$m = 357$
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***
10th percentile	$m = 636$	$m = 688$	$m = 715$	$m = 642$	$m = 695$	$m = 715$
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**
b) Period 1995:11 – 2004:12						
	1106 obs.	2215 obs.	2207 obs.	1104 obs.	2212 obs.	2204 obs.
1th percentile	$m = 11$	$m = 22$	$m = 22$	$m = 11$	$m = 22$	$m = 22$
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*
5th percentile	$m = 55$	$m = 110$	$m = 110$	$m = 55$	$m = 110$	$m = 110$
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***
10th percentile	$m = 110$	$m = 221$	$m = 220$	$m = 110$	$m = 221$	$m = 220$
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 4: Percentages Market Participants with Individual or Idiosyncratic Effects

	3 months forecast horizon			12 months forecast horizon		
	U.K. Pound	Japanese Yen	Euro	U.K. Pound	Japanese Yen	Euro
a) Period 1989:10 – 1995:2						
Individual effect						
$h_i - \bar{h} \neq 0$	43.33	37.01	37.87	50.00	49.69	58.29
Idiosyncratic effect due to extrapolation						
$\beta_{l,i} - \bar{\beta}_l \neq 0$	22.00	9.74	14.20	32.24	13.21	29.71
$h_i - \bar{h} \neq 0$	44.00	35.71	39.64	52.63	42.14	57.14
$\gamma_i - \bar{\gamma} \neq 0 \wedge h_i - \bar{h} \neq 0$	11.33	5.19	7.10	18.42	6.92	21.71
Idiosyncratic effect due to UIP						
$\gamma_i - \bar{\gamma} \neq 0$	24.00	26.62	26.63	34.87	38.99	42.29
$h_i - \bar{h} \neq 0$	33.33	39.61	30.77	43.42	55.97	43.43
$\gamma_i - \bar{\gamma} \neq 0 \wedge h_i - \bar{h} \neq 0$	19.33	15.58	9.47	23.02	27.04	18.86
b) Period 1995:11 – 2004:12						
Individual effect						
$h_i - \bar{h} \neq 0$	31.03	29.03	32.26	44.83	61.29	48.39
Idiosyncratic effect due to extrapolation						
$\beta_{l,i} - \bar{\beta}_l \neq 0$	31.03	12.90	19.35	48.28	35.48	48.39
$h_i - \bar{h} \neq 0$	27.59	29.03	29.03	48.28	61.29	54.84
$\gamma_i - \bar{\gamma} \neq 0 \wedge h_i - \bar{h} \neq 0$	10.34	6.45	0.00	31.03	22.58	32.26
Idiosyncratic effect due to UIP						
$\gamma_i - \bar{\gamma} \neq 0$	37.93	29.03	41.94	34.48	38.71	38.71
$h_i - \bar{h} \neq 0$	31.03	35.48	35.48	44.83	38.71	48.39
$\gamma_i - \bar{\gamma} \neq 0 \wedge h_i - \bar{h} \neq 0$	17.24	19.35	25.81	20.69	25.81	25.81

Notes: This table presents the percentage of respondents for which we find significant individual and/or idiosyncratic effects, i.e. the percentage of respondents for which we find significant coefficients from estimating equations 6, 9, and 11.

Table 5: Percentages Market Participants with Individual or Idiosyncratic Effects (combined)

	3 months forecast horizon			12 months forecast horizon		
	U.K. Pound	Japanese Yen	Euro	U.K. Pound	Japanese Yen	Euro
a) Period 1989:10 – 1995:2						
$h_i - \overline{h} \neq 0$	33.33	39.61	34.91	43.42	44.03	40.00
$\beta_{l,i} - \overline{\beta}_l \neq 0$	25.33	14.29	20.12	36.84	22.64	32.57
$\gamma_i - \overline{\gamma}_l \neq 0$	27.33	29.22	29.59	38.16	40.88	45.71
b) Period 1995:11 – 2004:12						
$h_i - \overline{h} \neq 0$	17.24	29.03	35.48	31.03	41.94	51.61
$\beta_{l,i} - \overline{\beta}_l \neq 0$	27.59	9.68	12.90	37.93	38.71	32.26
$\gamma_i - \overline{\gamma}_l \neq 0$	34.48	29.03	38.71	27.59	48.39	32.26

Notes: This table presents the percentage of respondents for which we find significant individual and/or idiosyncratic effects, i.e. the percentage of respondents for which we find significant coefficients from estimating equation 12.

Table 6: Combinations of Effects

	3 months forecast horizon			12 months forecast horizon		
	U.K. Pound	Japanese Yen	Euro	U.K. Pound	Japanese Yen	Euro
a) Period 1989:10 – 1995:2						
$c_i - \bar{c} \neq 0$ $\beta_{l,i} - \overline{\beta}_l = 0$ $\gamma_i - \overline{\gamma} = 0$	9.33	22.08	15.38	11.84	15.09	14.29
$c_i - \bar{c} = 0$ $\beta_{l,i} - \overline{\beta}_l \neq 0$ $\gamma_i - \overline{\gamma} = 0$	12.67	8.44	5.92	16.45	4.40	7.43
$c_i - \bar{c} = 0$ $\beta_{l,i} - \overline{\beta}_l = 0$ $\gamma_i - \overline{\gamma} \neq 0$	4.00	11.04	13.02	7.89	15.72	18.86
$c_i - \bar{c} \neq 0$ $\beta_{l,i} - \overline{\beta}_l \neq 0$ $\gamma_i - \overline{\gamma} = 0$	2.00	1.30	5.33	5.92	5.66	5.14
$c_i - \bar{c} \neq 0$ $\beta_{l,i} - \overline{\beta}_l = 0$ $\gamma_i - \overline{\gamma} \neq 0$	12.67	13.64	7.69	15.79	12.58	6.86
$c_i - \bar{c} = 0$ $\beta_{l,i} - \overline{\beta}_l \neq 0$ $\gamma_i - \overline{\gamma} \neq 0$	1.33	1.95	2.37	4.61	1.89	6.29
$c_i - \bar{c} \neq 0$ $\beta_{l,i} - \overline{\beta}_l \neq 0$ $\gamma_i - \overline{\gamma} \neq 0$	9.33	2.60	6.51	9.87	10.69	13.71
b) Period 1995:11 – 2004:12						
$c_i - \bar{c} \neq 0$ $\beta_{l,i} - \overline{\beta}_l = 0$ $\gamma_i - \overline{\gamma} = 0$	6.90	6.45	9.68	10.34	6.45	32.26
$c_i - \bar{c} = 0$ $\beta_{l,i} - \overline{\beta}_l \neq 0$ $\gamma_i - \overline{\gamma} = 0$	10.34	6.45	9.68	17.24	12.90	12.90
$c_i - \bar{c} = 0$ $\beta_{l,i} - \overline{\beta}_l = 0$ $\gamma_i - \overline{\gamma} \neq 0$	20.69	9.68	9.68	10.34	6.45	9.68
$c_i - \bar{c} \neq 0$ $\beta_{l,i} - \overline{\beta}_l \neq 0$ $\gamma_i - \overline{\gamma} = 0$	3.45	3.23	0.00	10.34	0.00	0.00
$c_i - \bar{c} \neq 0$ $\beta_{l,i} - \overline{\beta}_l = 0$ $\gamma_i - \overline{\gamma} \neq 0$	0.00	19.35	25.81	6.90	16.13	3.23
$c_i - \bar{c} = 0$ $\beta_{l,i} - \overline{\beta}_l \neq 0$ $\gamma_i - \overline{\gamma} \neq 0$	6.90	0.00	3.23	6.90	6.45	3.23
$c_i - \bar{c} \neq 0$ $\beta_{l,i} - \overline{\beta}_l \neq 0$ $\gamma_i - \overline{\gamma} \neq 0$	6.90	0.00	0.00	3.45	19.35	16.13

Notes: This table presents the combination of effects in percentages of market participants.

Table 7: Granger Causality Tests

		Absolute returns		Squared returns		GARCH	
		Variance → Dispersion	Dispersion → Variance	Variance → Dispersion	Dispersion → Variance	Variance → Dispersion	Dispersion → Variance
a) U.K. Pound							
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** (+–)
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 (+–)
b) Japanese Yen							
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*** (++)
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** (+–)
c) Euro							
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 (++)
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 (++)

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 and dispersion of beliefs. Variance → Dispersion implies that the null hypothesis is that causality runs from variance to dispersion of beliefs. 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.