Open access • Journal Article • DOI:10.1111/J.1538-4616.2009.00209.X

Disagreement and Biases in Inflation Expectations — Source link

Carlos Capistrán, Allan Timmermann

Institutions: Bank of Mexico, University of California, San Diego

Published on: 01 Jun 2006 - Journal of Money, Credit and Banking (Blackwell Publishing Inc)

Topics: Survey of Professional Forecasters, Real interest rate, Inflation and Rational expectations

Disagreement and Biases in Inflation Expectations*

Carlos Capistrán
Banco de Mexico

Allan Timmermann
UCSD and Creates

February 18, 2008

Abstract

Empirical work documents substantial disagreement in inflation expectations obtained from survey data. Furthermore, the extent of such disagreement varies systematically over time in a way that reflects the level and variance of current inflation. This paper offers a simple explanation for these facts based on asymmetries in the forecasters' costs of over- and under-predicting inflation. Our model implies biased forecasts with positive serial correlation in forecast errors and a cross-sectional dispersion that rises with the level and the variance of the inflation rate. It also implies that forecast errors at different horizons can be predicted through the spread between the short- and long-term variance of inflation. We find empirically that these patterns are present in inflation forecasts from the Survey of Professional Forecasters. A constant bias component, not explained by asymmetric loss and rational expectations, is required to explain the shift in the sign of the bias observed for a substantial portion of forecasters around 1982.

*We are grateful to two anonymous referees as well as the editor, Pok-sang Lam, for many valuable comments on the paper. We have also benefited from discussions with seminar participants at Banco de México, UCSD, the Latin American Meetings of the Econometric Society 2007, and the 2006 Meetings of the Society of Computational Economics. Gabriel López-Moctezuma provided excellent research assistance. The opinions in this paper are those of the authors and do not necessarily reflect the point of view of Banco de México.

1 Introduction

Differences in agents' beliefs and their importance to economic analysis has been emphasized by economists as early as Pigou (1927) and Keynes (1936).[1] A better understanding of the source of such disagreement has become increasingly important in view of the recent widespread use of macroeconomic models with heterogenous agents. Noting that disagreement about inflation is correlated with a host of macroeconomic variables, Mankiw, Reis and Wolfers (2003, p.2) go as far as suggesting that ".. disagreement may be a key to macroeconomic dynamics." This view is consistent with the theoretical models of Lucas (1973) and Townsend (1983) where heterogeneity in agents' beliefs play a key role.

Inflation forecasting is an area where disagreements appear to be particularly significant. Strong differences in inflation forecasts are found even at short forecast horizons and among professional forecasters with access to many common sources of information. Among 29 forecasters that participated in the Survey of Professional Forecasters in the third quarter of 2004, one-quarter-ahead forecasts of the annualized inflation rate ranged from 0.88% to 3.94% per annum. In the event, this range of three percentage points was one and a half times greater than the actual inflation rate of 1.98%.[2] Similar disagreements about future inflation have been found among different types of economic forecasters (professional economists versus lay consumers, Carroll (2003) and Mankiw, Reis and Wolfers (2003)), for different commodity groups (Mankiw, Reis and Wolfers (2003)), and across different sample periods (Zarnowitz and Braun (1992)).

A variety of explanations have been offered to explain these findings. Central to these is an assumption that agents have heterogenous information so that dispersion in beliefs reflects differences in information sets. Alternatively, differences may reflect heterogeneity in the rate at which agents update their beliefs. Mankiw and Reis (2002), and Carroll (2003) propose an elegant staggered updating model for expectations in which only a fraction of agents update their beliefs every period. Using this model, Mankiw, Reis and Wolfers (2003) are able to account for a number of features of inflation, including the extent of the observed disagreement and a variety of properties

[1] Pesaran and Weale (2006) and Hommes (2006) review the literature on expectations and the role of heterogeneity.

[2] Forecasts reported by the Survey of Professional Forecasters use the output deflator to measure inflation. The numbers reported here are for the annualized quarterly change between the third and fourth quarter of 2004. The Survey of Professional Forecasters is maintained by the Federal Reserve Bank of Philadelphia. Croushore (1993) provides detailed information on the construction of this data.

of the median forecast error. Cukierman and Wachtel (1979) suggest that both differences in expectations about the future rate of inflation and most of the changes over time in the variance of inflation are driven by the variance of aggregate demand shocks, but their empirical results only refer to the relation between periods with large variances in the rate of inflation and periods with large variances in inflation expectations, without measuring the demand shocks. There is also a related literature on model uncertainty and heterogeneity. Brock and Hommes (1997) and, more recently, Branch (2007) study agents who choose between different forecasting models each period. Under a variety of model selection rules they are able to account for disagreement among agents and generate considerable time-variation in dispersion across agents' beliefs.

None of these explanations, however, can provide an entirely satisfactory explanation for the biases observed in inflation expectations and the positive relationships between the cross-sectional dispersion in inflation beliefs and the level of the inflation rate. This goes to the heart of how we model inflation expectations and the reason why they differ among agents. Modelling heterogeneity without addressing these empirical relations could introduce dynamics in economic models that are at odds with reality. The impact, in particular in macroeconomics and finance where inflation expectations play a key role, is potentially large, for instance in models of the determination of the Phillips curve trade-off between unemployment and inflation (Mankiw and Reis (2002)), the determination of aggregate demand through the effect on consumption and investment (Clarida, Galí and Gertler (1999)) and the determination of stock prices (Fama (1991)).

This paper proposes a different explanation for how dispersion in inflation beliefs evolves over time and why it is correlated with both the level and volatility of inflation.[3] Our explanation relies on three mechanisms, namely asymmetric loss, heterogeneity in agents' loss functions and a constant loss component. Asymmetric loss captures the idea that the cost of over- and underpredicting inflation may be very different. Suppose that, for a particular agent, the cost of under-predicting inflation is higher than the cost of over-predicting it. Then it is optimal for this agent to bias the forecast so that on average he over-predicts inflation, thereby reducing the probability of costly under-predictions. Furthermore, if costs are increasingly large, the larger in absolute value the

[3]See Carlson and Valev (2003), Carroll (2003), Cukierman and Wachtel (1979), Mankiw, Reis and Wolfers (2003) and Souleles (2004) for evidence of a positive link between dispersion in inflation beliefs and the level of the inflation rate. A similar link between the variance of measured inflation and the cross-sectional dispersion has been established by Cukierman and Wachtel (1979) and Mankiw, Reis and Wolfers (2003).

forecast error (i.e., assuming loss is convex), then the optimal bias under rational expectations will be greater the higher the variance of the predicted variable. Finally, if the variance of the predicted variable is time-varying, then the optimal bias also becomes time-varying.

Turning to the second mechanism, heterogeneity in agents' loss functions means that periods with high degrees of macroeconomic uncertainty about the price level also coincide with periods where dispersions and biases in beliefs should be greater. Provided that there is heterogeneity across agents in their degree of loss asymmetry, such biases can drive dispersion across forecasters, giving rise to a positive relation between the variance of inflation and dispersion in beliefs. Moreover, if the variance of the inflation rate increases as the level of inflation goes up, then the dispersion in beliefs will also rise with the inflation rate. Both effects occur even if (a) agents are fully rational and their beliefs are updated every period and formed as conditional expectations (no belief distortions); and (b) agents have access to identical information and have identical beliefs about the mean and variance of future inflation.

The effects of asymmetric loss and cross-sectional heterogeneity are explored under the assumption of rational expectations. However, these two mechanisms, on their own, fall short of explaining an important feature of the survey data, namely the shift in the sign of the bias observed for a substantial portion of forecasters around 1982. We observe that many forecasters went from systematically under-predicting inflation prior to 1982 to over-predicting it in the period that followed. We show that the third mechanism, namely a constant bias component, capturing agents' tendency to over-predict inflation, can help explain this. This constant tendency to overpredict inflation is held against the time-varying tendency of many agents to under-predict inflation which is induced by asymmetry in their loss function and which gets stronger, the higher the level of inflation volatility. Prior to 1982, inflation was very volatile and so the asymmetry effect dominated the constant bias component and the overall effect was for agents to under-predict inflation. After 1982, inflation volatility came down and so the constant bias component dominated and forecasters tended to under-predict inflation.

The rest of the paper proceeds as follows. After a brief discussion of sources of asymmetric loss, Section 2 presents a theory of forecasts under asymmetric loss and explores its implications for the cross-sectional distribution of beliefs. Empirical evidence of asymmetries in forecasters' loss functions and evidence in support of our theory on the relation between inflation forecasts and inflation uncertainty is presented in Section 3. After a discussion of alternative explanations for dispersion in inflation beliefs, we conclude the paper in Section 4.

2 Inflation Forecasting under Asymmetric Loss

As pointed out by Mankiw and Reis (2002), an understanding of the microfoundations for agents' expectations is important to a theory of heterogeneity in expectation formation. For this reason we first briefly review three possible reasons for asymmetric loss, namely a utility cost explanation, a psychological explanation and a strategic explanation−see Elliott, Komunjer and Timmermann (forthcoming) for a more detailed discussion. We then propose a simple model that captures asymmetric loss and explore its implications for the cross-section of inflation beliefs.

2.1 Why Asymmetric Loss?

The most obvious explanation of asymmetric loss comes from the underlying economic 'primitives' of the decision problem that the inflation forecast is supposed to inform. Inflation forecasts matter for decisions on portfolio allocations, production levels, wage negotiations etc., so asymmetries in the costs of these factors due to over- or under-predicting the inflation rate should also affect the properties of the optimal forecast. Elliott, Komunjer and Timmermann (forthcoming) show how to derive asymmetric loss functions based on constant absolute or relative risk aversion utility functions combined with relations linking the forecast to the decision maker's actions.

Turning to the second explanation of asymmetric loss, a large literature in psychology has studied how peoples' judgments are affected in situations with different consequences of over- as opposed to under-assessment of a random event. In a comprehensive survey of this literature, Weber (1994) argues that the "asymmetric-loss-function interpretation provides a psychological explanation for observed judgments and decisions under uncertainty and links them to other judgment tasks" (p. 228). This literature also finds that the direction of a 'misestimate'–and thus the shape of the loss function–depends on the perspective of the forecaster in a way that reflects the consequences of a forecast error. For example, in experiments where individuals were asked to estimate the price of a car, when subjects took a buyer's perspective–a case where overestimates of the car's true price were more costly than underestimates–they tended to underestimate the price. Conversely, when subjects took the seller's perspective, the reverse happened and overestimates were more common (Birnbaum and Stegner (1979)).

While some forecasters may over-predict and others under-predict a particular outcome, a given individual's perspective, as reflected in the tendency to overweight or underweight the outcome, appears to be quite stable over time. Weber and Kirsner (1997) conclude that "... individuals differ in the relative emphasis they put on outcomes at the low (security) end of the distribution or at the high (potential) end of the distribution, and that this tendency is a stable, dispositional, *individual-difference* characteristic. Security-minded individuals are assumed to overweight outcomes at the low end of the distribution whereas potential-minded individuals do the opposite." (p. 42). Psychological factors may thus explain why some forecasters overweight high inflation outcomes relative to low inflation outcomes when forming their beliefs.

Finally, strategic explanations (e.g., Ehrbeck and Waldmann (1996); Laster, Bennette and Geoum (1999); Ottaviani and Sørensen (2006)) argue that asymmetries in the information available to forecasters versus their clients can be responsible for biases. This class of models views professional forecasters as agents of end users (the principals) so the role of the better-informed agents is to generate signals that are used to inform the principals' actions. Forecasters are remunerated based on their clients' assessments of their skills and so the aim of the forecast is to affect the clients' views. Models of strategic behavior can give rise to biases and asymmetric costs as the forecaster takes into consideration how a forecast will affect her future career path.

Empirical evidence on asymmetric loss has been found in inflation forecasts and in forecasts of other economic variables. For example, Capistrán (2006) finds that the Federal Reserve over-predicts inflation during the Volcker-Greenspan era as a result of the large costs associated with high inflation. Similarly, Ito (1990) finds that exporters and importers in Japan have different expectations over the dollar/yen exchange rate, with exporters expecting yen depreciation and importers expecting yen appreciation in his data. Elliott, Komunjer, and Timmermann (forthcoming) find that rejections of forecast rationality in survey data on real output growth and inflation may largely have been driven by the assumption of symmetric loss, whereas Patton and Timmermann (2007b) find that over-predictions of output growth are costlier for the Federal Reserve than under-predictions, particularly during periods with low economic growth.

2.2 Representation of Asymmetric Loss

Suppose that an economic agent is interested in predicting inflation h steps ahead, π_{t+h}, by means of information available at time t. We denote the associated h-step-ahead forecast by $f_{t+h,t}$ and the forecast error by $e_{t+h,t} = \pi_{t+h} - f_{t+h,t}$. What constitutes a good forecast depends on the forecaster's objectives as reflected in the loss function, $L(\cdot)$, that weights the costs of over- and under-predictions of different sizes. While it is commonplace to assume mean squared error (MSE) loss, i.e. $L(e) = e^2$, this loss function assumes that positive and negative forecast errors of equal magnitude lead to identical losses. Hence if inflation is forecast to be 2%, outcomes of 0% and 4% produce the same cost. This may well not be a reasonable assumption to make on economic grounds. Granger and Newbold (1986, p. 125) argue that "An assumption of symmetry for the cost function is much less acceptable" (than an assumption of a symmetric forecast error distribution).

We model asymmetric loss through the Linex loss function proposed by Varian (1974) and later adopted by Zellner (1986), Christoffersen and Diebold (1996, 1997) and Patton and Timmermann (2007a). This loss function captures asymmetries through a single parameter, ϕ, and takes the form:

$$L\left(e_{t+h,t};\phi\right) = \frac{1}{\phi^2}\left[\exp\left(\phi e_{t+h,t}\right) - \phi e_{t+h,t} - 1\right]. \tag{1}$$

If $\phi > 0$, the loss function is almost linear for $e_{t+h,t} < 0$ and becomes increasingly steep for positive values of $e_{t+h,t}$. Conversely, if $\phi < 0$, the loss function becomes increasingly steep for large, negative values of $e_{t+h,t}$. As $\phi \to 0$, the loss approaches symmetric, MSE loss.

To characterize an optimal forecast under this loss function, suppose that, conditional on information available to the forecaster at time t, Ω_t, inflation has a Normal distribution with conditional mean and variance $\mu_{t+h,t} = E\left[\pi_{t+h}|\Omega_t\right]$ and $\sigma^2_{t+h,t} = var\left[\pi_{t+h}|\Omega_t\right]$:[4]

$$\pi_{t+h|t} \sim N(\mu_{t+h,t}, \sigma^2_{t+h,t}). \tag{2}$$

Under assumptions (1) and (2) and rational expectations it is easy to show that the optimal forecast

[4] Conditional normality of the inflation rate is not nearly as strong an assumption as that of unconditional normality and allows the mean and variance of inflation to change over time–properties that we shall later see are crucial in capturing stylized facts of inflation forecasts.

that minimizes expected loss satisfies (see, e.g., Zellner (1986)):

$$\tilde{f}^*_{t+h,t} = \mu_{t+h,t} + \frac{\phi\sigma^2_{t+h,t}}{2}. \tag{3}$$

We extend this model by allowing for a constant bias, π_b^h, which is separately identifiable from the bias component that is induced by asymmetric loss provided that inflation volatility is time-varying:

$$f^*_{t+h,t} = \mu_{t+h,t} + \frac{\phi\sigma^2_{t+h,t}}{2} - \pi_b^h, \tag{4}$$

The negative sign in front of π_b^h turns out to be more convenient. Positive values of π_b^h indicate that agents tend to underpredict inflation while negative values suggest that they overpredict it. The constant bias component is not part of agents' first order condition under rational expectations and is indeed equal to zero if agents form rational expectations.

This simple model has some surprising implications which we next explore. We first analyze the simple benchmark case where forecasters have identical information so it is easy to track the effects of loss asymmetry on the properties of the average forecast and on dispersion in forecasts. Subsequently we relax the assumption of identical information and allow for heterogeneous information among forecasters.

2.3 Dispersion in Forecasts under Homogeneous Information

Let $\Omega_{t,i}$ be the information set of forecaster i at time t and consider the special case where forecasters' information sets are identical, $\Omega_{t,i} = \Omega_{t,j} = \Omega_t$ for all $i, j = 1, ..., N$, where N is the number of forecasters. Since forecasters share the same information, it follows from (4) that forecaster i's $h-$step-ahead forecast at time t, $f^*_{t+h,t,i}$, is:

$$f^*_{t+h,t,i} = \mu_{t+h,t} + \frac{\phi_i}{2}\sigma^2_{t+h,t} - \pi^h_{b,i}, \tag{5}$$

where ϕ_i is the asymmetry parameter and $\pi^h_{b,i}$ is the constant bias component of forecaster i. Even when forecasters have identical beliefs about the distribution of the predicted variable, their optimal forecasts will still differ provided that they have different degrees of loss asymmetry, as reflected in the parameter ϕ_i, or differ with regard to the constant bias component, $\pi^h_{b,i}$.

The bias in the forecast, defined as $bias_{t+h,t,i} = E[\pi_{t+h} - f^*_{t+h,t,i}|\Omega_t]$, follows directly from (2) and (5):

$$bias_{t+h,t,i} = \pi^h_{b,i} - \frac{\phi_i}{2}\sigma^2_{t+h,t}. \tag{6}$$

Even under rational expectations $(\pi^h_{b,i} = 0)$, this bias will generally be non-zero, provided that loss is asymmetric ($\phi_i \neq 0$).

Aggregating across forecasters, we get the mean forecast $\overline{f}^*_{t+h,t} = \frac{1}{N}\sum_{i=1}^N f^*_{t+h,t,i}$:

$$\begin{aligned} \overline{f}^*_{t+h,t} &= \frac{1}{N}\sum_{i=1}^N \mu_{t+h,t} + \frac{1}{N}\sum_{i=1}^N \left[\frac{\phi_i}{2}\sigma^2_{t+h,t}\right] - \frac{1}{N}\sum_{i=1}^N \pi^h_{b,i} \\ &= \mu_{t+h,t} + \frac{\bar{\phi}}{2}\sigma^2_{t+h,t} - \bar{\pi}^h_b, \end{aligned} \tag{7}$$

where $\bar{\phi} = N^{-1}\sum_{i=1}^N \phi_i$ is the average asymmetry parameter and $\bar{\pi}^h_b = N^{-1}\sum_{i=1}^N \pi^h_{b,i}$ is the cross-sectional average of the constant bias component. Hence the mean or "consensus" forecast varies with the conditionally expected inflation rate ($\mu_{t+h,t}$) and with the variance of inflation ($\sigma^2_{t+h,t}$)–i.e. it follows an ARCH-in-mean process–and the coefficient associated with the variance term is proportional to the average asymmetry parameter, $\bar{\phi}$. If all forecasters have symmetric loss ($\phi_i = 0, i = 1, ..., N$), then this average is zero[5] and the mean forecast is unbiased under rational expectations, although it may be biased if the constant bias component is not zero in the aggregate. A testable implication of (7) is that the mean forecast is correlated with the variance of inflation if the forecasters have asymmetric loss and is otherwise unrelated to it.

A generally overlooked implication of asymmetric loss is that it can induce serial correlation in the errors associated with an optimal forecast, even at the one-period horizon. The reason is that the optimal bias component, $-(\phi_i/2)\sigma^2_{t+h,t}$, is time-varying, so any persistence in the conditional variance translates into persistence in the forecast error. To see this, note that the first-order autocovariance between the one-step-ahead forecast error at time t and $t+1$ is given by $E[(\pi_{t+1} - f^*_{t+1,t,i})(\pi_t - f^*_{t,t-1,i})|\Omega_t]$. Suppose that time-variations in the conditional variance in (2) follow a standard GARCH(1,1) process with the current estimate of next period's variance ($\sigma^2_{t+1,t}$) reflecting current squared inflation innovations (ε^2_t) and last period's volatility estimate ($\sigma^2_{t,t-1}$): $\sigma^2_{t+1,t} = \alpha_0 + \alpha_1\varepsilon^2_t + \beta_1\sigma^2_{t,t-1}$. It can then be shown that the first-order autocovariance in the optimal

[5]This component is also zero if all forecasters have asymmetric loss but the asymmetries cancel out, i.e. $\bar{\phi} = 0$.

forecast errors is given by:[6]

$$Cov\left(e^*_{t+1,t,i}, e^*_{t,t-1,i}\right) = \frac{\phi_i^2\alpha_0^2\alpha_1^2(\alpha_1+\beta_1)}{2(1-(\alpha_1+\beta_1)^2-2\alpha_1^2)(1-\alpha_1-\beta_1)^2}. \quad (8)$$

Serial dependence in the error associated with an optimal forecast reflects the persistence in the conditional variance, as measured by $\alpha_1 + \beta_1$, and also depends on the degree of asymmetry in the loss, ϕ_i. The higher the persistence of the inflation process and the greater the loss asymmetry ($|\phi_i|$), the higher the persistence in the forecast error. Furthermore, since only ϕ_i^2 enters in (8), our theory predicts a positive first-order autocovariance irrespective of the sign of the asymmetry parameter. In contrast, as shown by Granger and Newbold (1986), the one-step forecast error should not be serially correlated under MSE loss.

Turning next to the cross-sectional dispersion in inflation forecasts, $\bar{s}_{t+h,t}$, this can be derived as follows:

$$\begin{aligned}
\bar{s}_{t+h,t} &\equiv \left[\frac{1}{N}\sum_{i=1}^{N}\left(f^*_{t+h,t,i}-\bar{f}^*_{t+h,t}\right)^2\right]^{1/2} \\
&= \left[\frac{1}{N}\sum_{i=1}^{N}\left(\mu_{t+h,t}+\frac{\phi_i}{2}\sigma^2_{t+h,t}-\pi^h_{b,i}-\left(\mu_{t+h,t}+\frac{\bar{\phi}}{2}\sigma^2_{t+h,t}-\bar{\pi}^h_b\right)\right)^2\right]^{1/2} \\
&= \left[\frac{1}{N}\sum_{i=1}^{N}\left(\left(\frac{\sigma^2_{t+h,t}}{2}\left(\phi_i-\bar{\phi}\right)\right)^2+\left(\pi^h_{b,i}-\bar{\pi}^h_b\right)^2-\sigma^2_{t+h,t}\left(\phi_i-\bar{\phi}\right)\left(\pi^h_{b,i}-\bar{\pi}^h_b\right)\right)\right]^{1/2}. \quad (9)
\end{aligned}$$

Under rational expectations, this expression simplifies to

$$\bar{s}_{t+h,t} = \frac{\sigma^2_{t+h,t}}{2}\left[\frac{1}{N}\sum_{i=1}^{N}\left(\phi_i-\bar{\phi}\right)^2\right]^{\frac{1}{2}}, \quad (10)$$

and the dispersion in inflation forecasts varies with the conditional variance of inflation, $\sigma^2_{t+h,t}$, multiplied by a term reflecting heterogeneity in the asymmetry parameters, ϕ_i. This result is able to account for the stylized fact that the dispersion in beliefs across forecasters and the conditional variance of inflation are positively correlated. Conversely, in the special case where forecasters have identical loss ($\phi_i = \bar{\phi}$ for all i), the dispersion is *not* correlated with the variance of in-

[6]This equation requires that $(\alpha_1+\beta_1)^2 + 2\alpha_1^2 < 1$ and $\alpha_1 + \beta_1 < 1$ so the variance exists.

flation even if loss is asymmetric ($\bar{\phi} \neq 0$). In the general case, heterogeneity in the constant bias component shifts the dispersion upwards and also introduces a time-varying interaction term, $-\sigma^2_{t+h,t}\left(\phi_i - \overline{\phi}\right)\left(\pi^h_{b,i} - \bar{\pi}^h_b\right)$, whose sign will depend on the cross-sectional relation between the constant bias component and agents' loss asymmetry parameters.

This simple model can explain why the dispersion in inflation forecasts is positively correlated with the level of inflation provided that the conditional variance of the inflation rate depends positively on the level of inflation–as found empirically by Ball and Cecchetti (1990) and Grier and Perry (1998). Higher levels of inflation appear to translate into higher time-series variability in inflation and this gives rise to a greater cross-sectional dispersion since differences in inflation forecasts are driven by the variance in our model (see equation (9)). The inter-quartile range provides a particular measure of the belief dispersion across forecasters and has the advantage of being robust to extreme inflation forecasts. Suppose for simplicity that there is no heterogeneity in the constant bias component, i.e. $\pi^h_{b,i} = \pi^h_{b,j}$ for all i, j. Then the inter-quartile range simplifies to:

$$f^*_{t+h,t,0.75} - f^*_{t+h,t,0.25} = \left(\frac{\phi_{0.75} - \phi_{0.25}}{2}\right)\sigma^2_{t+h,t}, \tag{11}$$

where $\phi_{0.25}$ and $\phi_{0.75}$ refer to the 25 and 75 percentiles of the cross-sectional distribution of $\phi-$values. Clearly our model implies a positive relation between the dispersion in inflation forecasts and the variance of inflation provided that the degree of asymmetry varies across forecasters. Heterogeneity in the constant bias component will not overturn this result unless there is a systematic positive correlation across forecasters between ϕ_i and $\pi_{b,i}$.

Furthermore, the difference in inflation forecasts across two forecasters, i and j, is given by:

$$f^*_{t+h,t,i} - f^*_{t+h,t,j} = \left(\frac{\phi_i - \phi_j}{2}\right)\sigma^2_{t+h,t} - (\pi^h_{b,i} - \pi^h_{b,j}). \tag{12}$$

Whether the ranking of the predictions (or prediction errors) produced by two forecasters changes over time therefore depends on the size of both bias components. Suppose the constant bias component is identical for the two forecasters, i.e. $\pi^h_{b,i} = \pi^h_{b,j}$. Then their ranking should not change provided that the degree of loss asymmetry is constant: If forecaster i dislikes positive forecast errors (underpredictions) more than forecaster j, so $\phi_i > \phi_j$, then forecaster i's predictions should generally be higher than those of forecaster j. Rankings continue to be time-invariant provided that the sign of $(\phi_i - \phi_j)$ is the opposite of the sign of $(\pi^h_{b,i} - \pi^h_{b,j})$, or if differences across agents in one bias component are 'small' relative to differences in the other component.

Another implication of our theory that, to the best of our knowledge, has not previously been considered is the effect of asymmetric loss on biases in the term structure of forecast errors. Surveys such as the Survey of Professional Forecasters (SPF) ask participants to forecast inflation at multiple horizons and this can be exploited to test our theory. Under our model the expected value of the differential between, say, forecaster i's four-step and one-step forecast error is given by

$$E_t[e_{t+4,t,i} - e_{t+1,t,i}] = (\pi^4_{b,i} - \pi^1_{b,i}) - \frac{\phi_i}{2}(\sigma^2_{t+4,t} - \sigma^2_{t+1,t}). \tag{13}$$

Differences between the forecast errors at different points of the term structure of inflation rates should thus be predictable by means of the spread between the long-run and short-run conditional variance. There is no particular sign implication, however, as this depends on the relative size of the fixed bias components $(\pi^4_{b,i} - \pi^1_{b,i})$ as well as the sign of ϕ_i and on whether $\sigma^2_{t+4,t}$ exceeds or falls below $\sigma^2_{t+1,t}$.

The sign of the bias in individual agents' forecasts can also change. To see this, recall that the individual biases are given by

$$bias_{t+h,t,i} = \pi^h_{b,i} - \frac{\phi_i}{2}\sigma^2_{t+h,t}.$$

In the absence of a constant bias component, the sign of the bias is always the same, i.e. negative if $\phi_i > 0$ and positive otherwise. However, when the signs of ϕ_i and $\pi^h_{b,i}$ are the same, fluctuations in $\sigma^2_{t+h,t}$ may cause the sign of the overall bias to shift. For example, if both ϕ_i and $\pi^h_{b,i}$ are negative, then the overall bias will be negative if $\sigma^2_{t+h,t}$ is low and positive otherwise. Conversely, if both ϕ_i and $\pi^h_{b,i}$ are positive, then the overall bias will be positive under low inflation volatility and negative if inflation volatility is sufficiently high.

2.4 Belief Dispersion under Heterogeneous Information

We next relax the assumption that all forecasters have the same information set, so that $\Omega_{t,i} \neq \Omega_{t,j}$ for at least one pair $i \neq j$, $i, j = 1, ..., N$. It then follows that, conditional on $\Omega_{t,i}$, $\pi_{t+h} \sim$

$N(\mu_{t+h,t,i}, \sigma^2_{t+h,t,i})$ and so forecaster i's forecast is:

$$f^*_{t+h,t,i} = \mu_{t+h,t,i} + \frac{\phi_i}{2}\sigma^2_{t+h,t,i} - \pi^h_{b,i}, \tag{14}$$

while the bias conditional on $\Omega_{t,i}$ equals $bias_{t+h,t,i} = \pi^h_{b,i} - (1/2)\phi_i\sigma^2_{t+h,t,i}$.

Aggregating across forecasters, the average inflation forecast now becomes:

$$\begin{aligned} \overline{f}^*_{t+h,t} &= \frac{1}{N}\sum_{i=1}^{N}\mu_{t+h,t,i} + \frac{1}{N}\sum_{i=1}^{N}\frac{\phi_i}{2}\sigma^2_{t+h,t,i} - \frac{1}{N}\sum_{i=1}^{N}\pi^h_{b,i} \\ &= \overline{\mu}_{t+h,t} + \frac{1}{2}\tilde{\sigma}^2_{t+h,t} - \overline{\pi}^h_b, \end{aligned} \tag{15}$$

where $\bar{\mu}_{t+h,t} = N^{-1}\sum_{i=1}^{N}\mu_{t+h,t,i}$ and $\tilde{\sigma}^2_{t+h,t} = N^{-1}\sum_{i=1}^{N}\left[\phi_i\sigma^2_{t+h,t,i}\right]$. The mean forecast now depends on (i) the average expected inflation, (ii) a weighted average of the variance of inflation $\tilde{\sigma}^2_{t+h,t}$ where the weights are the scaled loss asymmetry parameters, $\phi_i/(2N)$, and (iii) the average constant bias. When the underlying values of $\mu_{t+h,t,i}$ and $\sigma^2_{t+h,t,i}$ are unobserved, equation (15) has broadly similar effects as the result obtained under homogenous information. In both cases the testable implication is that the conditional variance of inflation affects the (cross-sectional) mean inflation forecast if the forecasters have asymmetric loss and fails to do so if they have symmetric loss.[7]

Unsurprisingly, when information is heterogenous the dispersion in inflation forecasts is also

[7] This result holds more broadly for general asymmetric loss functions and distributions of inflation, see Granger (1969) and Patton and Timmermann (2007a).

affected by differences in beliefs about first and second moments of inflation:

$$\begin{aligned}
\bar{s}^2_{t+h,t} &= \frac{1}{N}\sum_{i=1}^{N}\left(\mu_{t+h,t,i}+\frac{\phi_i}{2}\sigma^2_{t+h,t,i}-\pi^h_{b,i}-\left(\overline{\mu}_{t+h,t}+\frac{1}{2}\tilde{\sigma}^2_{t+h,t}-\bar{\pi}^h_b\right)\right)^2 \\
&= \frac{1}{N}\sum_{i=1}^{N}\left[\left(\mu_{t+h,t,i}-\overline{\mu}_{t+h,t}\right)-(\pi^h_{b,i}-\bar{\pi}^h_b)+\frac{1}{2}\left(\phi_i\sigma^2_{t+h,t,i}-\tilde{\sigma}^2_{t+h,t}\right)\right]^2 \\
&= \frac{1}{N}\sum_{i=1}^{N}\left[\left(\mu_{t+h,t,i}-\overline{\mu}_{t+h,t}\right)^2+(\pi^h_{b,i}-\bar{\pi}^h_b)^2+\frac{1}{4}\left(\phi_i\sigma^2_{t+h,t,i}-\tilde{\sigma}^2_{t+h,t}\right)^2\right] \qquad (16) \\
&\quad +\frac{1}{N}\sum_{i=1}^{N}\left(\mu_{t+h,t,i}-\overline{\mu}_{t+h,t}\right)\left(\phi_i\sigma^2_{t+h,t,i}-\tilde{\sigma}^2_{t+h,t}\right)-\frac{1}{N}\sum_{i=1}^{N}(\pi^h_{b,i}-\bar{\pi}^h_b)(\phi_i\sigma^2_{t+h,t,i}-\tilde{\sigma}^2_{t+h,t}) \\
&\quad -\frac{2}{N}\sum_{i=1}^{N}(\pi^h_{b,i}-\bar{\pi}^h_b)\left(\mu_{t+h,t,i}-\overline{\mu}_{t+h,t}\right).
\end{aligned}$$

Here $\bar{\sigma}^2_{t+h,t}=(1/N)\sum_{i=1}^{N}\sigma^2_{t+h,t,i}$. Dispersion in inflation forecasts now depends on differences in beliefs about the mean of future inflation, as captured by (i) a time-varying and (ii) a constant term; (iii) differences in beliefs about the variance of the inflation rate, weighted by the loss asymmetry parameters; (iv) a term reflecting systematic correlations between mean and variance forecasts; (v) the cross-sectional covariance between the constant bias component and the weighted variance forecast; and (vi) the cross-sectional covariance between the constant bias component and the conditional mean forecast..

The literature has so far focused on the first term, i.e., the variance in inflation expectations across forecasters with heterogeneous information (Mankiw, Reis and Wolfers (2003)). In our model this is only one of six factors driving the dispersion in beliefs across forecasters. Even if agents agree on the conditional mean of the inflation rate, we should still expect to see cross-sectional dispersion in inflation forecasts provided that agents disagree about the variance of the future inflation rate and have asymmetric loss. Furthermore, as shown earlier, if agents agree on both the mean and variance of the inflation rate, we should still observe cross-sectional dispersion provided that the loss asymmetry differs across forecasters. Moreover, the dispersion across forecasters continues to be driven by a combination of differences in information sets, differences in the constant bias component and differences in the costs associated with forecast errors.

2.5 Summary of Theoretical Implications

We summarize the implications of our model in the following claims for inflation forecasts:

Claim 1 *Inflation forecasts are generally biased with a bias that can be positive or negative depending on whether over- or under-predictions are costliest;*

Claim 2 *The magnitude (absolute value) of the bias in the inflation forecast increases in the conditional variance of the inflation rate;*

Claim 3 *Forecast errors should be positively serially correlated if the conditional variance of the inflation rate is persistent;*

Claim 4 *The mean differential between forecast errors at long and short horizons is predictable by means of the spread between the long-term and short-term conditional variance;*

Claim 5 *The ranking of inflation forecasts across forecasters remains constant over time provided that agents have homogenous beliefs about first and second moments of future inflation and either (i) the constant and time-varying bias components pull in the same direction or (ii) one of these components always dominates;*

Claim 6 *The cross-sectional dispersion in inflation forecasts increases as a function of the conditional variance of the inflation rate;*

Claim 7 *If the conditional variance rises with the level of the inflation rate, then the cross-sectional dispersion in inflation forecasts will also increase with the inflation rate.*

Claim 8 *The sign of the bias in the forecast error may shift across periods with high and low inflation volatility.*

Asymmetric loss coupled with rational expectations is the key mechanism underpinning the first four claims. The properties are robust although in rare cases the time-varying and constant bias components may cancel out. Claims 5-7 rely on heterogeneity in the degree of asymmetry in agents' loss functions. Claim 5, in particular, does depend on the relative significance of the constant and time-varying bias components. Both a constant bias component and asymmetric loss are important in generating the shift in the sign of the overall bias referred to in the last claim.

3 Empirical Evidence

In this section we use data from the Survey of Professional Forecasters (SPF) to investigate if our claims for the dispersion in inflation beliefs stand up to empirical scrutiny. The SPF data has a relatively extensive cross-sectional and time-series coverage. Interestingly, the questionnaire that participants in the SPF fill out simply asks for a forecast of a variety of variables without being explicit about the objective of the prediction. Thus, it is not specified if the forecaster should report the mean, median, mode or some other weighted average of possible inflation outcomes. This suggests the possibility that individual forecasters use different weighting functions (reflecting their different perspectives or use of inflation forecasts) to compute their forecasts.

3.1 Data

We use one and four-step-ahead inflation forecasts from the SPF with inflation measured as the annualized quarterly change in the output deflator. The sample starts with forecasts made in the third quarter of 1968 and ends with forecasts made in the third quarter of 2004. There are between 9 and 75 forecasters at each point in time, with a median of 34 participants. Our analysis includes forecasters with at least 30 non-zero forecasts. This leaves us with data on 62 forecasters. Figure 1 shows the number of forecasters included each quarter. Consistent with Croushore (1993), there was a steady decline in the number of forecasters from the early 1970s to the end of the 1980s. In 1990, the Philadelphia Fed took over the survey from the American Statistical Association and the NBER and added new forecasters to the survey. From then on, the number of forecasters shows no apparent trend.

For actual values of the output deflator we use both real-time and fully revised data obtained from the real-time database available at the Federal Reserve Bank of Philadelphia's web site. Fully revised data is the last revision as of April 2005, whereas real-time data corresponds to the second revision available for a given data vintage.[8]

[8]See Croushore and Stark (2000) for a description of the real-time data set.

3.2 Evidence from Individual Forecasters

3.2.1 Biases in Forecasts

To test our model's implication for forecast biases, Figure 2 shows the histogram of the individual mean forecast errors computed from the SPF using real-time data for actual values. The histogram is skewed to the right which implies a preponderance of positive mean forecasts errors (so under-prediction occurs more frequently). Indeed, the mean forecast error, averaged across time and across forecasters, is 0.27 and 0.22 percentage points per annum using revised and real-time data, respectively. Hence, on average the forecasters under-predicted next quarter's inflation rate by about a quarter of a percent per year. For comparison, the average standard deviation of the forecast error is 1.65 (revised data) and 1.86 (real-time data).

For each forecaster a test of unbiasedness can be carried out by regressing the forecast errors on a constant and applying a t-test. When applied to our data, we find evidence of a significant bias for 37 (60%) and 31 (50%) out of 62 forecasters using revised and real-time data for the actual values, respectively. Hence, consistent with Claim 1, at least half of the forecasters in our sample had biases sufficiently large to be significant at the 5% critical level. Interestingly, the distribution of the positive mean forecast errors is far more spread out than the distribution of negative mean forecast errors, indicating that large under-predictions are more common than large over-predictions of inflation.[9]

Under rational expectations, our model links forecast biases to the conditional variance, $\sigma^2_{t+h,t}$, and the degree of loss asymmetry, ϕ_i. Given any two of these components (bias, conditional variance and loss asymmetry), the third can be imputed from equation (6). To get a sense of the magnitude of the asymmetry parameters, we need to obtain an estimate of the one-step-ahead conditional variance of inflation. To this end, we estimate a GARCH model to the revised inflation rate series, allowing for autoregressive dynamics in the inflation rate. More specifically, we fit a GARCH(1,1)

[9] Since the number of forecasts as well as the standard deviation of the forecast errors differ across the individual forecasters, within a given bin in Figure 2 some survey participants produce a significant mean error, while others do not.

model of the form:

$$\begin{aligned}
\pi_{t+1,t} &= \lambda_0 + \lambda_1 \pi_{t,t-1} + \lambda_2 \pi_{t-3,t-4} + \varepsilon_{t+1}, \\
\varepsilon_{t+1} &\sim N(0, \sigma^2_{t+1,t}), \\
\sigma^2_{t+1,t} &= \omega + \alpha_1 \varepsilon_t^2 + \beta_1 \sigma^2_{t,t-1}.
\end{aligned} \tag{17}$$

Starting from Engle (1982), such ARCH models have been used extensively to estimate time-variations in the conditional variance of inflation. Engle and Kraft (1983) and Bollerslev (1986) apply the methodology to US inflation as measured by the change in the output deflator. Our model builds on Bollerslev's finding that an AR(4) model for the mean of inflation and a GARCH(1,1) model for the conditional variance provide a good description of US inflation.

Estimates of this model are reported in the first column of Table 1.[10] Both the α_1 and β_1 estimates are highly significant and add up to 0.98, suggesting that inflation volatility is time-varying and highly persistent. This is an important observation in view of our earlier remarks that inflation volatility acts as a transmission mechanism for dispersion in our model.

Armed with estimates of the conditional variance of inflation, an estimate of the asymmetry parameter for each forecaster is obtained, under rational expectations, as $\widehat{\phi}_i = \frac{-2}{T}\sum_{t=1}^{T} -e_{t+h,t,i}/\widehat{\sigma}^2_{t+h,t}$, where $\widehat{\sigma}^2_{t+h,t}$ is our estimate of the conditional variance of inflation.[11] When $\pi_b^h \neq 0$, instead we run a time-series regression of $e_{t+h,t,i}$ on a constant and $\sigma^2_{t+h,t}$ and obtain our estimator of ϕ_i from the slope coefficient on the variance term. Histograms of the asymmetry parameters under both scenarios are presented in Figure 3. Some forecasters have positive asymmetry parameters (under-prediction is more costly than over-prediction) whereas others have negative asymmetry parameters (over-prediction is more costly). Furthermore, we applied a simple t-test to find out if the asymmetry parameters are different from zero. Under rational expectations, using revised inflation data we reject the null hypothesis for 33 forecasters (53%) at the 5% level. With real-time data, we reject the null for 31 forecasters (50%).[12] When $\pi_b^h \neq 0$, using revised inflation data we

[10] We included four lags of inflation in the mean equation, but found that only lags one and four were significantly different from zero and thus eliminated the second and third lags. The results are robust to including all four lags.

[11] Ideally, one would use the conditional variance estimated from data on the individual forecast errors, $\widehat{\sigma}^2_{t+h,t,i}$. See below for an exercise along these lines.

[12] This test should be interpreted with caution since the null, $\phi = 0$, lies on the boundary of the parameter space. For this reason we also applied the Elliott, Komunjer, and Timmermann (2005) procedure to estimate the asymmetry parameter under rational expectations and an asymmetric quadratic loss function. Because this class naturally nests

reject the null hypothesis for 17 forecasters (27%) at the 5% level, whereas with real-time data, we reject the null for 16 forecasters (26%). We find that the constant bias component is significantly different from zero for 22 forecasters (35%) using revised data and for 20 forecasters (32%) using real-time data.

Asymmetric loss is not a very attractive explanation of dispersion in inflation beliefs if the required degrees of asymmetry are very large. To facilitate economic interpretation of the results, Table 2 presents summary statistics of the distribution of the estimated asymmetry parameters. Similarly, Figure 4 presents a plot of a traditional quadratic (MSE) loss and Linex loss functions evaluated at two different values of the asymmetry parameter, -1 and 1.5, both of which are close to the interquartile values reported in Table 2. For the forecaster with the positive asymmetry parameter, under-predicting inflation by one standard deviation is a little more than twice as costly as over-predicting it by the same amount, whereas for the forecaster with the negative asymmetry parameter over-prediction is a little less than twice as costly as under-prediction. These values do not appear overly large and so our explanation is consistent with economically modest degrees of differences in loss associated with over- and under-predictions.

Next we carry out an exercise where we allow $\sigma^2_{t+h,t,i}$ to vary across forecasters. This is important because we have to address the possibility that different information sets (as suggested, among others, by Mankiw and Reis (2002)) or different forecasting models (as suggested, among others, by Branch (2007)) could be causing the differences among forecasters that we are observing. The SPF includes a variable (PRPGDP) that gives the probability that the annual–average inflation falls in a particular range (i.e., the survey asks the forecasters to assign probabilities to each of a number of pre–determined intervals in which inflation might fall).[13] Although so far we have used the forecasts for next-quarter inflation, the histograms are only available for the annual–average so we switch variables only for this exercise. Each quarter, respondents report their probabilities of inflation falling in certain bins, so during the first quarter of each year they provide a 4-step-ahead forecast, whereas in the last quarter they report a 1-step-ahead forecast. In order to increase the number of observations available, for this exercise we assume that the asymmetry parameter does not change

MSE loss, a test of symmetry is readily available. We reject symmetry for 63% of the forecasters at the 5% level when revised data is used, whereas the figure is 55% for real-time data.

[13] See Rich and Tracy (2003) for a description of the properties of these histograms.

with the forecast horizon and simply pool the data for all the quarters available for each forecaster.[14] We have 54 forecasters that provided at least 15 quarters of histograms. We calculate the mean of the individual histograms as $\widehat{\mu}_{t+h,i} = \sum_{k=1}^{n} P_{t+h,t,i}(k) Mid_{t+h,t}(k)$, where n is the number of histograms available for forecaster i, $P_{t+h,t,i}(k)$ denotes the probability forecaster i attaches to interval k at time t for the h-step-ahead forecast and $Mid_{t+h,t}(k)$ denotes the mid-point of each interval.[15] Then we calculate the variance as $\widehat{\sigma}_{t+h,t,i} = \sum_{k=1}^{n} P_{t+h,t,i}(k) \left[Mid_{t+h,t}(k) - \widehat{\mu}_{t+h,i}\right]^2 - \frac{\omega_t^2}{12}$, where $\frac{\omega_t^2}{12}$ is Sheppard's correction.[16] Using real-time data for the actual values and assuming rational expectations, we find that the estimated asymmetry parameters are between -6 and 3, with 78% of the forecasters' asymmetry parameters falling between -2 and 2. Furthermore, we computed a simple t-test to find out if the asymmetry parameters are different from zero. We could reject the null hypothesis of symmetry for 16 forecasters (30%). These results further support our argument, and suggest that at least part of the heterogeneity does not come from either differences in the information sets or in the forecasting models.[17]

3.2.2 Mean Forecast Errors Across Forecasters

Equations (7) and (15) imply that the mean inflation forecast depends on the conditional variance of inflation. Assuming homogenous information, the mean bias is:

$$\begin{aligned} E\left[\pi_{t+h} - \overline{f}^{*}_{t+h,t}|\Omega_t\right] &= E\left[\pi_{t+h} + \bar{\pi}_b^h - \mu_{t+h,t} - \frac{\overline{\phi}}{2}\sigma^2_{t+h,t}|\Omega_t\right] \\ &= \bar{\pi}_b^h - \frac{\overline{\phi}}{2}\sigma^2_{t+h,t}. \end{aligned} \tag{18}$$

Under our model, the mean forecast error is predicted by variations in the conditional variance. Conversely, under MSE loss, the time-varying bias should be zero. This implication can be tested

[14] Our sample starts in the third quarter of 1981 and ends in the third quarter of 2004.

[15] To deal with the fact that the lowest and highest intervals are open-ended, we treat them as closed intervals and set their width equal to that of the interior intervals.

[16] This is a correction for the variance of grouped distributions which assumes that the frequencies are concentrated at the mid-points of intervals with width ω, see Kendall and Stuart (1963).

[17] Apart from the corrections explained in Rich and Tracy (2003), we also had to re-scale the forecasts of previous surveys when there was a change in base year (six changes in our sample), because the forecasters reported the forecasts using the old base. See the documentation provided by the Federal Reserve Board of Philadelphia.

through a regression of the form:

$$\left(\pi_{t+h} - \overline{f}^*_{t+h,t}\right) = \delta_0 + \delta_1 \sigma^2_{t+h,t} + \varepsilon_{t+h},$$

where, under asymmetric loss $\delta_1 \neq 0$. In this regression δ_0 captures the cross-sectional average bias.

To test this implication, we regress the mean forecast errors on a constant and the conditional variance obtained from the GARCH(1,1) model. The results are presented in columns one (revised data) and three (real-time data) of Table 3. In both cases, and consistent with Claim 2, the coefficient of the conditional variance is significant and positive, suggesting a negative value of $\overline{\phi}$.[18]

3.2.3 Serial Correlation

To test the implication of our theory (Claim 3) that the forecast errors are positively serially correlated, we used the SPF data to calculate Ljung-Box tests for the null hypothesis of zero first-order autocorrelation. We could reject the null of no autocorrelation in 56% and 43% of the cases with revised and real-time data, respectively. Figure 5 presents the histogram of the autoregressive coefficients estimated using real-time data to compute forecast errors. The histogram is skewed to the right with very few negative values and, as predicted by our theory, all the significant autocorrelation coefficients are positive.

3.2.4 Term Structure of Inflation Forecasts

To test the term structure implications of our model (equation (13) or Claim 4), we first generated one-step and four-step-ahead forecasts of the conditional variance from the GARCH(1,1) model using the parameter estimates in the first column of Table 1. Next, we regressed for each forecaster the error differential, $e_{t+4,t,i} - e_{t+1,t,i}$, on the variance differential, $\sigma^2_{t+4,t} - \sigma^2_{t+1,t}$. With the revised inflation data we could reject the null hypothesis that the coefficient on the variance differential is zero for 67% of all forecasters while with real time data we rejected the null 69% of the time. This suggests that the term structure implications of our model for the forecast errors at the annual and quarterly horizons are consistent with our findings.

[18]The biggest (in absolute value) biases occur during the period where inflation volatility was at its peak, i.e. in the late 1960s and early 1970s.

3.2.5 Rank Preservation

Assuming homogenous information our model (Claim 5) suggests that the ranking of pairs of individual forecasters, i and j, remains constant through time provided that the conditional variance of inflation, $\sigma^2_{t+h,t}$, does not cross the value $\left|2(\pi^h_{b,i} - \pi^h_{b,j})/(\phi_i - \phi_j)\right|$. This condition is always satisfied in the absence of a constant bias component or if there is no heterogeneity in the constant bias component, i.e. $\pi^h_{b,i} = \pi^h_{b,j}$. It also holds if $\sigma^2_{t+h,t}$ is either large or small relative to the cross-sectional heterogeneity in the loss asymmetry parameter versus the heterogeneity in the fixed bias component. Alternatively, the ranking of individual forecasters will remain constant provided that the sign of $(\phi_i - \phi_j)$ is the same as the sign of $(\pi^h_{b,j} - \pi^h_{b,i})$, irrespective of the variance of inflation. To test these ranking implications, we looked at forecasters with more than 15 forecasts during both the 1980s and the 1990s. Since these were very different periods with high and low inflation volatility, respectively, this provides a difficult test for our theory.[19] Using real-time data, we ranked the forecasters by their mean forecast errors and calculated Spearman's rank correlation between the two sub-samples to test if the ranking is preserved. The rank correlation is 0.71. This correlation is greater than zero at the 5% level (the p-value is 0.04), providing evidence that the ranks are indeed preserved, consistent with our theory.[20]

3.3 Evidence from the Cross-Section of Forecasters

The theory presented in Section 2 relates the conditional variance of inflation to the conditional mean of the inflation rate and the cross-sectional dispersion in inflation beliefs. As a first step towards investigating if these relations hold empirically, Figure 6 presents scatter plots of inflation, its conditional variance (fitted from a GARCH(1,1) model), and the mean and standard deviation across forecasters who participated in the SPF. Consistent with the literature cited in the introduction, there is a positive relation between the variables.

[19] The variance of inflation is even higher during the 70s, but unfortunately there are only two forecasters that had at least 15 forecasts for each decade.

[20] Olds (1938) provides a table of small sample critical values for this type of test.

3.3.1 Dispersion Across Forecasters

Our sixth claim, and an implication of equations (9) and (16), was that the dispersion across forecasters depends on the conditional variance of inflation. To test this we project the cross-sectional dispersion on a constant (resulting from heterogeneity in agents' constant bias component) and the conditional variance in inflation obtained from the GARCH(1,1) model:

$$\overline{s}_{t+h,t} = \gamma_0 + \gamma_1 \sigma^2_{t+h,t} + \varepsilon_{t+h}. \tag{19}$$

If the forecasters have asymmetric loss we expect to find $\gamma_1 > 0$, while under symmetric loss $\gamma_1 = 0$. Furthermore, if the conditional variance affects the dispersion across forecasters, but the explanation is different from the one given in this paper, then it could be the case that $\gamma_1 < 0$. The outcome of this regression is reported in column five of Table 3. The coefficient associated with the conditional variance is significant and positive as predicted.

Our theory is also consistent with the inter-quartile range across forecasters being positively correlated with the variance of inflation. Under homogeneous information an implication of equation (11) is that this range equals the variance of inflation times half the inter-quartile range of the asymmetry parameters. From Table 2, this range is close to two (it is 1.50 and 1.83 with real-time and revised data, respectively). Therefore we would expect the inter-quartile range across forecasters to be roughly equal to the variance of inflation. This implication of our theory is quite strong since it suggests a one-to-one relation between the dispersion across forecasters, measured by the inter-quartile range, and the conditional variance of inflation.

To see if this is borne out by our data, we plot time series for the interquartile range and the conditional variance in Figure 7, while Figure 8 presents a scatter plot. During the 25 years covered by the Volcker-Greenspan period (1979:3 to 2004:4), the two series are very close and resemble the 45 degree line in Figure 8. The correlation between the two series is 0.72. This implication of our model holds remarkably well, given our strong assumptions of homogeneity of information and conditional normality of inflation.

Interestingly, our model is also consistent with the finding in Mankiw, Reis and Wolfers (2003) that inflation expectations fan out in the period following a large shock to the inflation rate. This corresponds in our model to a large value of ε_t in (17) which will increase the conditional volatility in the next period and hence increases the cross-sectional spread.

To obtain the positive relationship between the level of inflation and the dispersion across forecasters contained in Claim 7, we need a positive relation between the level and the variance of inflation. Figure 9 indicates a strong relationship between the level of the inflation rate and disagreement across inflation forecasters. Using data from the SPF, this figure plots time-series of the interquartile range of inflation forecasts against the level of the inflation rate. While this figure is suggestive of a positive relationship, a more formal analysis is still required. To this end we use once again the ARCH methodology to explore any effects from the level of inflation on the variance of inflation. The conditional variance model that best describes our data is presented in the second column of Table 1 and takes the form:

$$\sigma^2_{t+1,t} = \omega + \delta\pi_t. \tag{20}$$

As before, the mean inflation rate is best described by an AR(4) model. The lagged level of the inflation rate now affects its conditional variance with a positive coefficient. No other ARCH terms were significant in the volatility specification and are hence omitted. The model indicates that when the level of inflation is high, the conditional variance of inflation also tends to be high.[21] This result, together with our previous findings of a positive effect of the conditional variance on the dispersion in beliefs across forecasters, explains the stylized fact that there is a positive relation between the dispersion across forecasters and the level of inflation. Furthermore, as can be seen from the new estimates presented in columns 2, 4 and 6 in Table 3, the earlier results based on the GARCH(1,1) specification continue to hold when the conditional variance is estimated instead from the model with the lagged inflation rate.

3.4 Changes in the Sign of the Bias in 1982

The final claim establishes that, under appropriate conditions, our model is consistent with a shift in the sign of the bias of the individual agents' inflation forecasts. This is important since the bias of a substantial portion of the survey forecasters appears to have changed sign around 1982.[22]

[21] A positive relation between the level of inflation and inflation uncertainty has also been documented empirically for data on the US and G7 countries by Ball and Cecchetti (1990), Grier and Perry (1998), and Logue and Willett (1976).

[22] As reported by Croushore (1998), the bias of the mean forecast also changes sign around 1982.

Suppose we split the sample in 1982q1 and consider the biases of forecasters with at least 30 non-zero forecasts and at least one forecast before and after the split. This yields 29 forecasters. Of those, 24 have a positive bias before 1982 and a negative bias after 1982. Moreover, of the 18 forecasters with a significant positive pre-1982 bias, 13 generate a significant negative bias after 1982.

Recall that the bias in our model is given by $\pi_{b,i}^h - \frac{\phi_i}{2}\sigma_{t+h,t}^2$. Hence we cannot explain this sign shift with our model if the asymmetry parameter, ϕ_i, is time-invariant and the constant bias component equals zero. To match the finding that the bias in the pre-1982 sample is positive while the post-1982 bias is negative, we must therefore have:

$$\begin{aligned}\pi_{b,i}^h - \frac{\phi_i}{2}\bar{\sigma}_1^2 &> 0 \\ \pi_{b,i}^h - \frac{\phi_i}{2}\bar{\sigma}_2^2 &< 0,\end{aligned} \tag{21}$$

where $\bar{\sigma}_1^2$ and $\bar{\sigma}_2^2$ denote the average conditional variances before and after 1982, respectively. This in turn requires that:

$$\begin{aligned}\pi_{b,i}^h &> \frac{\phi_i}{2}\bar{\sigma}_1^2 \\ \pi_{b,i}^h &< \frac{\phi_i}{2}\bar{\sigma}_2^2.\end{aligned} \tag{22}$$

To test if these conditions hold simultaneously, we first computed the average conditional variances pre- and post-1982q1 using the GARCH(1,1) model reported in Table 1. The resulting averages are $\widehat{\bar{\sigma}}_1^2 = 2.28$ and $\widehat{\bar{\sigma}}_2^2 = 0.75$, so our empirical analysis suggests that $\bar{\sigma}_1^2 > \bar{\sigma}_2^2$. This means that the inequalities (22) can only hold provided that $\pi_{b,i}^h < 0$ and $\phi_i < 0$ for the forecasters for whom we observe a change in the sign of their bias from positive to negative. This is a testable proposition. To see if this holds in our data, for the 13 forecasters that significantly change sign (ids: 15, 31, 39, 40, 43, 51, 66, 69, 70, 72, 82, 86, 89), only forecaster 66 has a positive asymmetry parameter. Among them, the median estimated asymmetry parameter is $\widehat{\phi} = -0.5$, while the smallest one is $\widehat{\phi} = -2.4$. Given these magnitudes, a $\pi_{b,i}^h$ of -0.2% (for the median asymmetry parameter) or a $\pi_{b,i}^h$ of -0.9% (for the smallest one) is sufficient to induce a change in the sign of the bias in 1982 as a result of the change in the inflation variance. Larger values of the constant bias are required, however, to match the biases in the post-82 period which must be consistent with:

$$\pi_{b,i}^{h} = bias_{i,t+h,t} + \frac{\phi_i}{2}\sigma_{i,t+h,t}^{2}.$$

Since both terms are negative for most forecasters in our sample, the absolute value of the constant bias must exceed the absolute value of the post-82 bias estimate which on average is -1.04%, and ranges between -1.7% and -0.5%.

Hence the change in the sign of the bias of this group of 13 forecasters is consistent with a model that includes a constant bias component as well as a variance-dependent bias component. Prior to 1982 many forecasters underpredicted inflation while after 1982 they have generally overpredicted it. This is consistent with a tendency for this group of forecasters to overpredict inflation which−due to asymmetric loss aversion−is held against an effort to underpredict it, which gets stronger the higher the variance. The latter effect dominated prior to 1982 and so these forecasters underpredicted inflation up to this point in time. Conversely, the variance effect diminished after 1982, a period when inflation was overpredicted by the 13 forecasters.

It is worth pointing out that if we further split the post-1982q1 period into two sub-samples, pre and post-1993q3, and repeat the exercise of looking for a possible change in sign, we find that a clear majority of forecasters do not generate any change in the sign of their mean errors. This evidence suggests that the apparent change in bias for a subset of our forecasters may be a one-off event. In particular, changes around 1982 could be related to several changes that affected inflation in the early 1980s, such as those that occurred in monetary policy as reported for example by Clarida, Gali and Gertler (1999) and the changes in the NAIRU in the early 1980s reported by, for example, Staiger, Stock and Watson (1997).

3.5 Alternative Explanations

Our analysis started from the observation that the traditional model of symmetric loss and symmetric information cannot explain why inflation forecast errors appear to be biased and serially correlated with a dispersion that varies both with the level of inflation and its variance. Furthermore, as argued by Mankiw, Reis and Wolfers (2003), simple models such as adaptive expectations also fall short of explaining the patterns observed in survey data.

Some of the observed properties of inflation forecasts can be explained in the context of the staggered updating model proposed by Carroll (2003) and Mankiw and Reis (2002). If only a fraction of agents update their forecasts each period, aggregate forecast errors can be serially correlated and dispersion in inflation forecasts across agents can be quite large at any point in time. Both implications are consistent with the data. This theory also has some limitations, however, and does not appear to be able to explain why forecast errors are biased on average. Nor has this theory been used to explain time-variations in inflation dispersion and its link to the conditional variance and the level of the inflation rate. Mankiw, Reis and Wolfers (2003 p. 41) note that "The sticky-information model gives no reason to find a systematic relation between the level of inflation and the extent of disagreement." In their model, changes over time in macroeconomic information drive inflation disagreements.

The analysis provided by Carroll (2003) suggests that a constant fraction of agents read news articles about inflation. If a high level of inflation is associated with high volatility and a high news coverage and hence more frequent updating, we would expect a smaller bias in such periods. In contrast our model predicts a larger bias (of either sign depending on whether ϕ_i is positive or negative) which appears to be consistent with the data.

Interestingly, Carroll identifies rational forecasts with the SPF forecasts based on findings reported by Croushore (1998) that these forecasts are unbiased. Using data from 1968:4-1996:4, third revision data for the realized values and a standard forecast efficiency regression, Croushore reports a p-value of 0.867 for the null hypothesis of efficient forecasts. We get very similar results for the mean forecast with a p-value of 0.55 using data from 1969:2-2004:3. However, in line with Croushore, we also find that the average forecast is biased in the first part of the data sample. Furthermore, as we have shown, the unbiasedness results hold only for the aggregate data and is generally rejected for the individual forecasters. As we found in our analysis, some forecasters have negative biases while others have positive biases in their forecasts, so it may be difficult to detect biases in the mean or median forecast.[23]

Another possible explanation of our findings is agents' learning, see Evans and Honkapohja

[23] Mankiw, Reis and Wolfers (2003) report strong evidence that rationality can be rejected for the median inflation forecasts in the SPF data. Bonham and Cohen (2001) also find that unbiasedness regressions do not share the same coefficients across forecasters when analyzing the SPF forecasts of the GDP deflator, with the implication that unbiasedness cannot be tested using the consensus, but has to be tested at the level of the individual forecaster.

(2001) for a recent survey. As shown by Timmermann (1993) in the context of a model for stock prices, learning can induce serial correlation in forecast errors even when agents use fully optimal estimators and update these using standard recursive algorithms. One implication of the learning explanation that appears at odds with the data is, however, that in a stationary environment, learning effects should taper off –something we do not seem to find in our data as seen by the rise in the interquartile range after 2000. Furthermore, agents' learning should proceed faster in a volatile environment where information flows faster, leading to smaller subsequent heterogeneity in beliefs, a point that again seems difficult to square with the observation that disagreement increases in more volatile environments. While learning undoubtedly plays an important role in understanding how agents form expectations, without a more structured model for how agents differ in their priors, what types of models underlie their expectations and an understanding of how new forecasters (with a short learning experience) enter and old forecasters (with a longer learning experience) leave the survey, it is difficult to formally test this alternative explanation.

An alternative to explaining the shift in 1982 in the sign of some of the forecasters' bias is to let the asymmetry coefficients, ϕ_i, be time-varying and subject even to a change in sign. Gordon and St-Amour (2000) propose a regime switching model for preferences which could be adopted to the asymmetry parameter here. Another possibility is regime switches in belief distortions as modeled by Cecchetti, Lam and Mark (2000). The fact that we do not find similar evidence of changes in the signs of the forecasters' biases after 1982 suggests that if such changes in the asymmetry parameter or in belief distortions occurred, they are likely to be rare.

4 Conclusion

Little is known about how agents arrive at the beliefs reported in survey data. As pointed out by Carroll (2003, page 270) "there appears to have been essentially no work proposing and testing positive alternative models for how empirical expectations are formed." We explored an alternative explanation in this paper, namely that agents weight the consequences of over- and under-predictions very differently and as a result calculate their forecasts under asymmetric loss with a shape of the loss function that differs across agents.

Under asymmetric loss, the stylized finding that the conditional variance of inflation varies over time translates into optimal inflation forecasts that have a time-varying bias even under rational expectations. This simple observation provides the basis for an explanation of many of the stylized facts reported in the literature on survey measures of inflation forecasts in addition to suggesting some new stylized facts that were confirmed to hold empirically. In particular, asymmetric loss in conjunction with time-varying volatility is able to explain why inflation uncertainty drives the disagreement among inflation forecasters. The combination of asymmetric loss and rational expectations falls short of explaining why a substantial portion of individual forecasters change from underpredicting inflation to overpredicting it around 1982. To explain this a constant bias component is required which also shows the limitations of a pure rational expectations story for survey participants' behavior. While our explanation is undoubtedly only part of the story and other explanations such as inertia in how agents update their expectations (Mankiw and Reis (2002), Carroll (2003)) as well as differences in information also play a role, it generates a variety of new testable propositions that seem to be borne out when tested on survey data.

References

[1] Ball, L., Cecchetti, S.G., 1990. Inflation and uncertainty at long and short horizons. Brookings Papers on Economic Activity 1, 215-245.

[2] Birnbaum, M.H., Stegner, S.E., 1979. Source credibility in social judgment: bias, expertise and the judge's point of view. Journal of Personality and Social Psychology 37, 48-74.

[3] Bollerslev, T., 1986. Generalized autoregressive conditional heteroskedasticity. Journal of Econometrics 31, 307-27.

[4] Bonham, C.S., Cohen, R., 2001. To aggregate, pool, or neither: testing the rational expectations hypothesis using survey data. Journal of Business and Economic Statistics 19, 278-291.

[5] Branch, W.A., 2007. Sticky information and model uncertainty in survey data on inflation expectations. Journal of Economic Dynamics and Control 31, 245-276.

[6] Brock, W.A., Hommes, C.H., 1997. A rational route to randomness. Econometrica 65, 1059-1160.

[7] Capistrán, C., 2006. Bias in federal reserve inflation forecasts: is the Federal Reserve irrational or just cautious? Banco de México Working Paper Series: #2006-14.

[8] Carlson, J.A., Valev, N.T., 2003. Sources of dispersion in consumer inflation forecasts. Applied Economic Letters 10, 77-81.

[9] Carroll, C.D., 2003. Macroeconomic expectations of households and professional forecasters. Quarterly Journal of Economics 118, 269-298.

[10] Cecchetti, S.G., P-S Lam and N.C. Mark, 2000. Asset pricing with distorted beliefs: are equity returns too good to be true? American Economic Review 90, 787-805.

[11] Christoffersen, P.F., Diebold F.X., 1996. Further results on forecasting and model selection under asymmetric loss. Journal of Applied Econometrics 11, 561-572.

[12] Christoffersen, P.F., Diebold F.X., 1997. Optimal prediction under asymmetric loss. Econometric Theory 13. 808-817.

[13] Clarida, R., Galí, J., Gertler, M., 1999. The science of monetary policy: a new Keynesian perspective. Journal of Economic Literature 36, 1661-1707.

[14] Croushore, D., 1993. Introducing: the Survey of Professional Forecasters. Federal Reserve Bank of Philadelphia Business Review November/December, 3-13.

[15] Croushore, D., 1998. Evaluating inflation forecasts. Federal Reserve Bank of Philadelphia Working Paper Series: #98-14.

[16] Croushore, D., Stark, T., 2000. A funny thing happened on the way to the data bank: A real-time data set for macroeconomists. Federal Reserve Bank of Philadelphia Business Review September/October, 15-27.

[17] Cukierman, A., Wachtel, P., 1979. Differential inflationary expectations and the variability of the rate of inflation: theory and evidence. American Economic Review 69, 595-609.

[18] Elliott, G., Komunjer, I., Timmermann, A., 2005. Estimation and testing of forecast rationality under flexible loss. Review of Economic Studies 72, 1107-1125.

[19] Elliott, G., Komunjer, I., Timmermann, A., forthcoming. Biases in macroeconomic forecasts: irrationality or asymmetric loss?. Journal of the European Economic Association.

[20] Ehrbeck, T., Waldmann, R., 1996. Why are professional forecasters biased? agency versus behavioral explanations. Quarterly Journal of Economics 111, 21-40.

[21] Engle, R.F., 1982. Autoregressive conditional heteroskedasticity with estimates of the variance of U.K. inflation. Econometrica 50, 987-1008.

[22] Engle, R.F., Kraft, D., 1983. Multiperiod forecast error variances of inflation estimated from ARCH models, in: Zellner, A. (Ed.), Applied Time Series Analysis of Economic Data, Bureau of the Census, Washington, D.C., pp. 293-302.

[23] Evans, G.W., Honkapohja, S., 2001. Learning and expectations in macroeconomics. Princeton University Press, Princeton, NJ.

[24] Fama, E., 1991. Efficient markets: II. Fiftieth Anniversary Invited Paper. Journal of Finance 46, 1575-1617.

[25] Gordon, S. and St-Amour, P., 2000. A preference regime model of bull and bear markets. American Economic Review 90, 1019-1033.

[26] Granger, C.W.J., 1969, Prediction with a generalized cost function. Operational Research 20, 199-207.

[27] Granger, C.W.J., Newbold, P., 1986. Forecasting Economic Time Series. 2nd edition. Academic Press, Orlando, FL.

[28] Grier, K., Perry, M.J., 1998. On inflation and inflation uncertainty in the G-7 countries. Journal of International Money and Finance 17, 671-689.

[29] Hommes, C.H., 2006. Heterogeneous agent models in economics and finance, in: Handbook of Computational Economics II: Agent-Based Computational Economics. North-Holland, Amsterdam.

[30] Ito, T., 1990. Foreign exchange rate expectations: micro survey data. American Economic Review 80, 434-449.

[31] Kendall, M., Stuart, A., 1963. The Advanced Theory of Statistics. Vol. 1. Griffin, London.

[32] Keynes, J.M., 1936. The General Theory of Unemployment, Interest and Money. Harcourt, Brace and World, New York.

[33] Kiefer, N.M., Vogelsang, T.J., 2002. Heteroskedasticity-Autocorrelation robust standard errors using the Bartlett kernel without truncation. Econometrica 70, 2093-2095.

[34] Laster, D., Bennett, P., Geoum, I.S., 1999. Rational bias in macroeconomic forecasts. Quarterly Journal of Economics 114, 293-318.

[35] Logue, D.E., Willett, T.D., 1976. A note on the relation between the rate and the variability of inflation. Economica 43, 151-158.

[36] Lucas, Robert E., Jr., 1973. Some international evidence on output-inflation tradeoffs. American Economic Review 63, 326-334.

[37] Mankiw, G.N., Reis, R., 2002. Sticky information versus sticky prices: a proposal to replace the new Keynesian Phillips curve. Quarterly Journal of Economics 117, 1295-1328.

[38] Mankiw, G.N., Reis, R., Wolfers, J., 2003. Disagreement about inflation expectations. NBER Macroeconomics Annual.

[39] Olds, E.G., 1938. Distribution of sums of squares of rank differences for small samples. Annals of Mathematical Statistics 9,133-148.

[40] Ottaviani, M., Sørensen, P.N., 2006. The strategy of professional forecasting. Journal of Financial Economics 81, 441-466.

[41] Patton, A., Timmermann, A., 2007a. Properties of optimal forecasts under asymmetric loss and nonlinearity. Journal of Econometrics 140, 884-918.

[42] Patton, A., Timmermann, A., 2007b. Testing forecast optimality under unknown loss. Journal of the American Statistical Association 102, 1172-1184.

[43] Pesaran, M.H., Weale, M., 2006. Survey expectations. Handbook of Economic Forecasting. North-Holland, Amsterdam.

[44] Pigou, A.C., 1927. Industrial Fluctuations. Macmillan, London.

[45] Rich, R., Tracy, J., 2003. Modeling uncertainty: predictive accuracy as a proxy for predictive confidence. Manuscript, Federal Reserve Bank of New York.

[46] Souleles, N.S., 2004. Expectations, heterogeneous forecast errors, and consumption: micro evidence from the Michigan Consumer Sentiment Surveys. Journal of Money, Credit and Banking 36, 39-72.

[47] Staiger, D., Stock, J.H., and Watson, M.W. 1997. The NAIRU, unemployment and monetary policy. The Journal of Economic Perspectives 11, 33-49.

[48] Timmermann, A., 1993. How learning in financial markets generates excess volatility and predictability in stock prices. Quarterly Journal of Economics 108, 1135-1145.

[49] Townsend, R.M., 1983. Forecasting the forecasts of others. The Journal of Political Economy 91, 546-588.

[50] Varian, H., 1974. A Bayesian approach to real estate assessment, in: Fienberg S.E., Zellner, A. (Eds.) Studies in Bayesian Econometrics and Statistics in Honor of L.F. Savage. North-Holland, Amsterdam pp. 195-208.

[51] Weber, E.U., 1994. From subjective probabilities to decision weights: the effect of asymmetric loss functions on the evaluation of uncertain outcomes and events. Psychological Bulletin 115, 228-242.

[52] Weber, E.U., Kirsner, B., 1997. Reasons for rank-dependent utility evaluation. Journal of Risk and Uncertainty 14, 41-61.

[53] Zarnowitz, V., Braun, P., 1992. Twenty-two years of the NBER-ASA quarterly economic outlook surveys: aspects and comparisons of forecasting performance. NBER, Working Paper Series: 3965.

[54] Zellner, A., 1986. Biased predictors, rationality and the evaluation of forecasts. Economics Letters 21, 45-48.

Table 1: Conditional Variance of Inflation. Sample 1968:4 - 2004:4.

	Mean	
λ_0	2.45***	3.03***
	(0.79)	(0.51)
λ_1	0.60***	0.58***
	(0.07)	(0.06)
λ_2	0.32***	0.34***
	(0.06)	(0.05)
	Variance	
ω	0.02	-1.53***
	(0.02)	(0.28)
α_1	0.12**	
	(0.06)	
β_1	0.86***	
	(0.06)	
δ		0.35***
		(0.06)
T	145	144
Log-likelihood	-224.15	-214.37

Notes: Standard errors in parenthesis.

** $p < 0.05$. ***$p < 0.01$.

Table 2: Summary Statistics of the Distribution of the Linex Asymmetry Parameter Estimates.

	No Constant Bias		With Constant Bias	
Statistic	Revised Data	Real-time Data	Revised Data	Real-time Data
Quantile 0.25	-0.85	-0.74	-0.78	-0.66
Quantile 0.50	-0.20	-0.17	0.01	0.17
Quantile 0.75	0.98	0.76	0.53	0.61
Inter-quartile Range	1.83	1.50	1.31	1.27
Maximum	3.25	3.18	3.11	2.95
Minimum	-2.15	-1.80	-4.25	2.95
Mean	0.06	0.08	-0.04	0.01

Table 3: Inflation Forecasts and Uncertainty.

Dependent variable	$\pi_{t+1} - \overline{f}_{t+1,t}$ (Revised)	$\pi_{t+1} - \overline{f}_{t+1,t}$ (Revised)	$\pi_{t+1} - \overline{f}_{t+1,t}$ (Real-time)	$\pi_{t+1} - \overline{f}_{t+1,t}$ (Real-time)	$\overline{s}_{t+1,t}$	$\overline{s}_{t+1,t}$
Intercept	-0.44	-0.27	-0.34	-0.25	0.62*	0.93**
	(-2.80)	(-1.30)	(-2.38)	(-1.65)	(4.49)	(5.30)
Variance	0.34**		0.28*		0.35**	
	(5.45)		(4.17)		(5.25)	
Variance (Level)		0.18***		0.18***		0.10
		(9.24)		(7.31)		(3.15)
T	144	144	144	144	144	144
R^2	0.07	0.10	0.04	0.07	0.28	0.13

Notes: t-statistics with standard errors calculated using a Bartlett kernel without truncation in parenthesis. Critical values for the asymptotic distribution of the t-statistics are reported in Kiefer and Vogelsang (2002).

* $p < 0.10$. ** $p < 0.05$. ***$p < 0.01$.

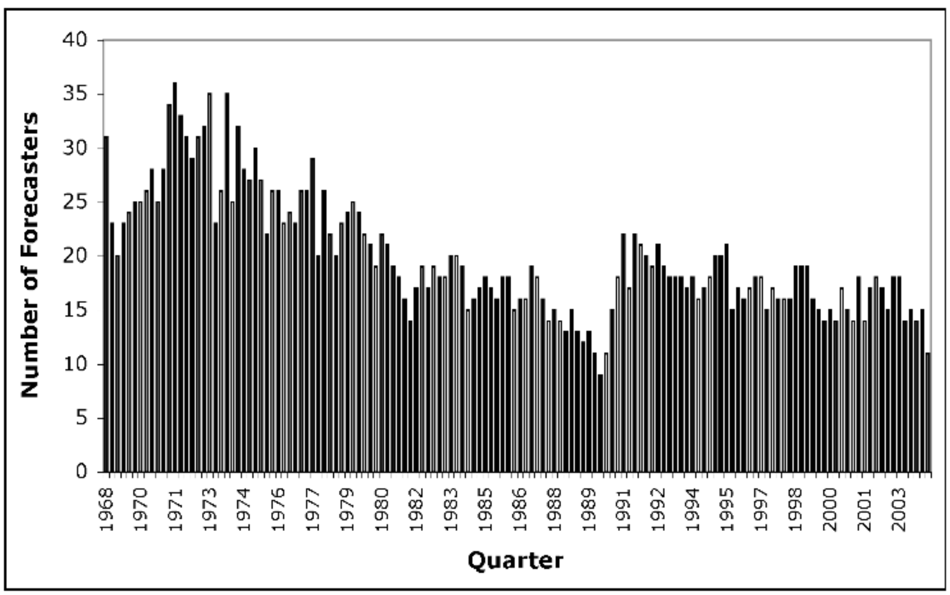

Figure 1: Number of forecasters, only forecasters with at least 30 non-zero forecasts, Survey of Professional Forecasters (One-step-ahead forecasts).

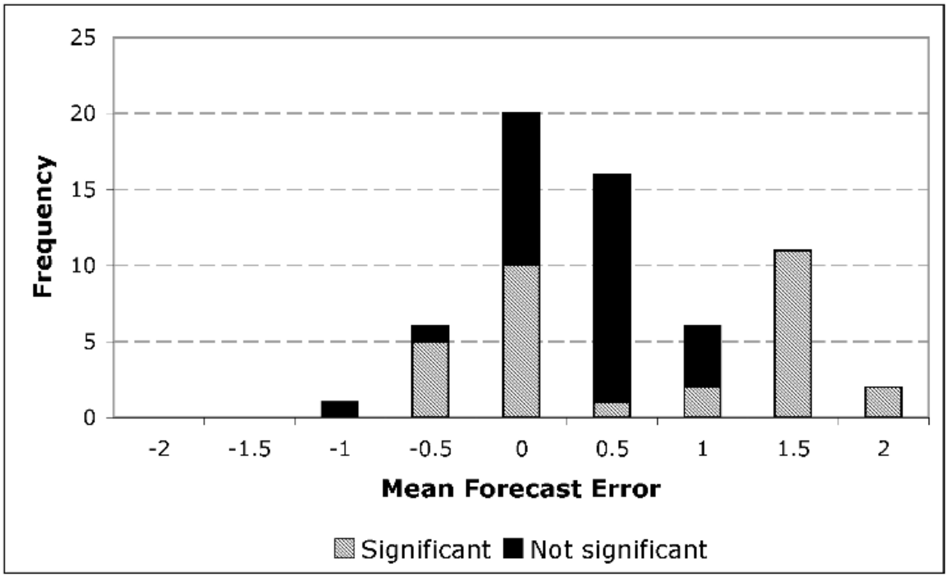

Figure 2: Histogram of mean forecast errors, Survey of Professional Forecasters (One-step-ahead forecasts, real-time data).

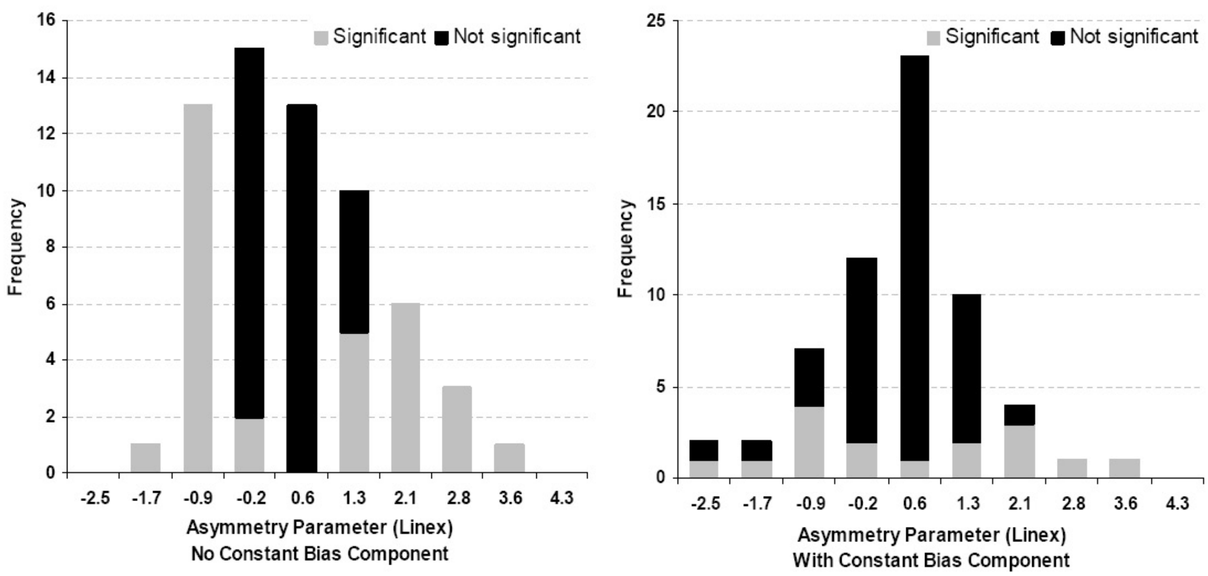

Figure 3: Histogram of linex asymmetry parameter estimates, Survey of Professional Forecasters (One-step-ahead forecasts, real-time data).

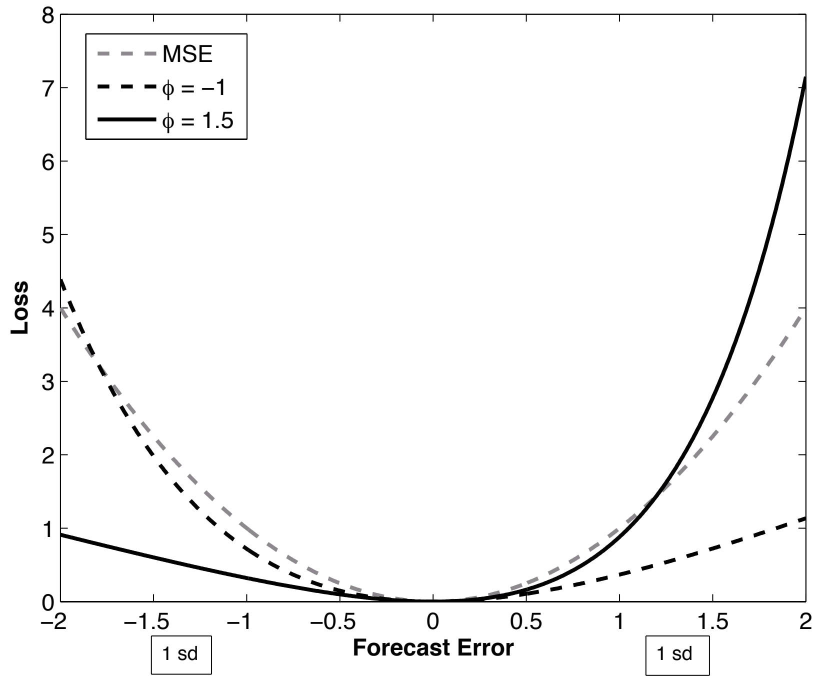

Figure 4: Degree of asymmetry.

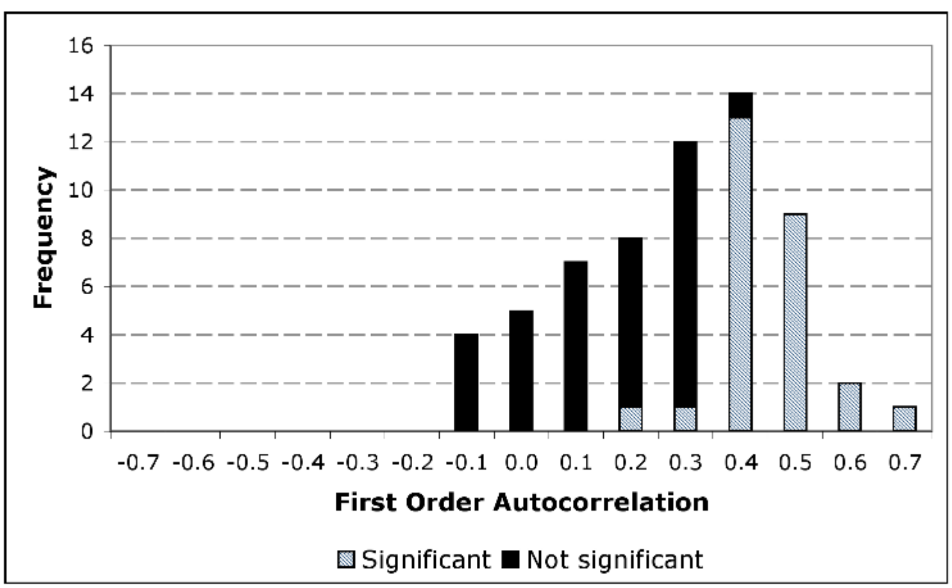

Figure 5: Histogram of first order autocorrelations, Survey of Professional Forecasters (One-step-ahead forecasts, real-time data).

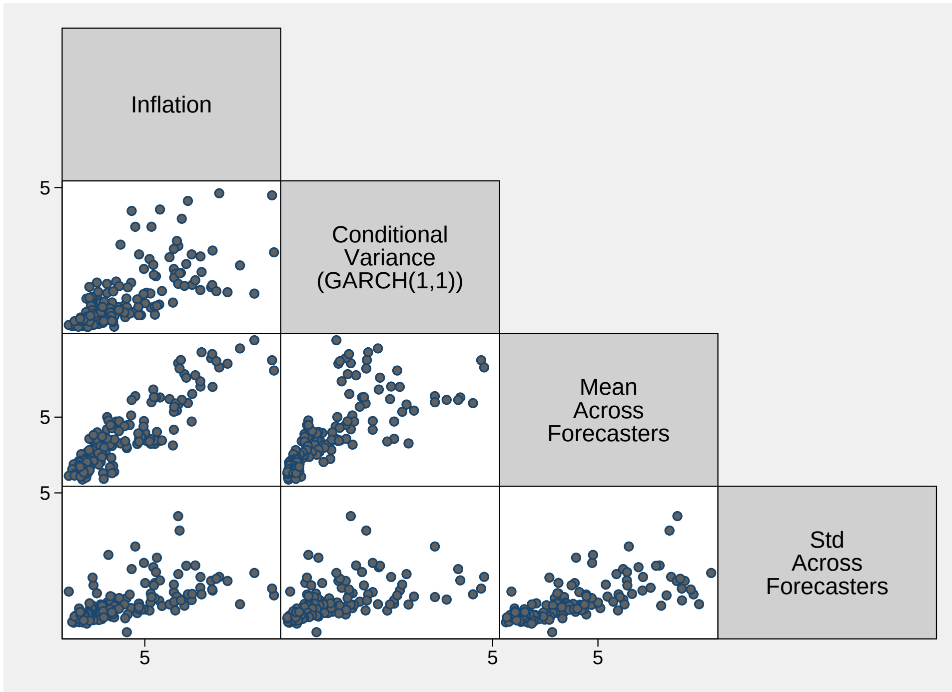

Figure 6: Scatter plot of location and scale of inflation and one-step-ahead inflation forecasts, Survey of Professional Forecasters.

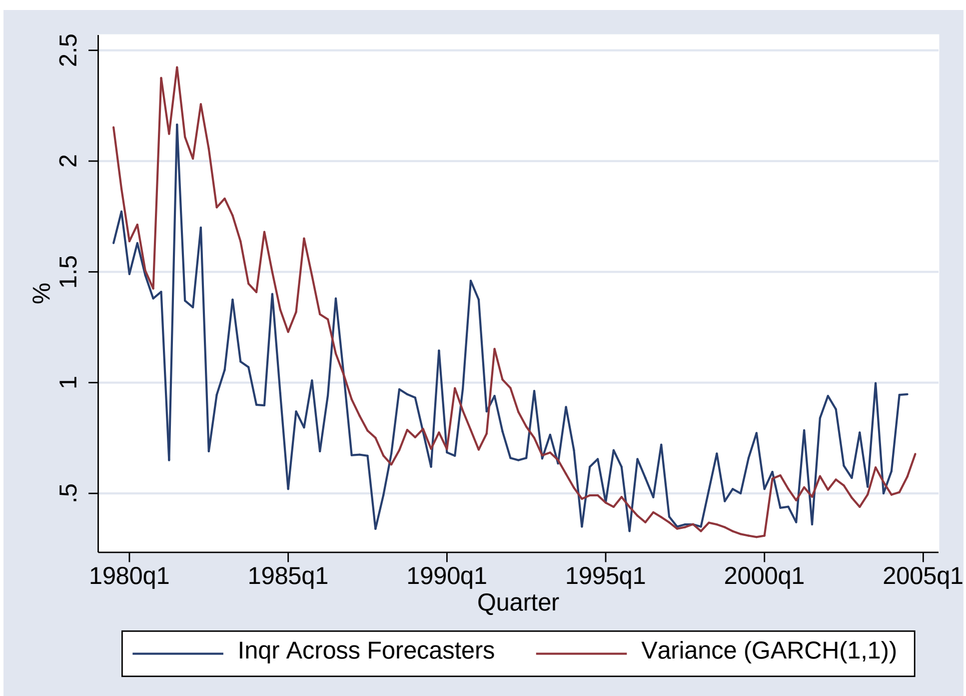

Figure 7: The inter-quartile range across one-step-ahead forecasts of inflation from the Survey of Professional Forecasters and the conditional variance of inflation calculated using a GARCH(1,1) model.

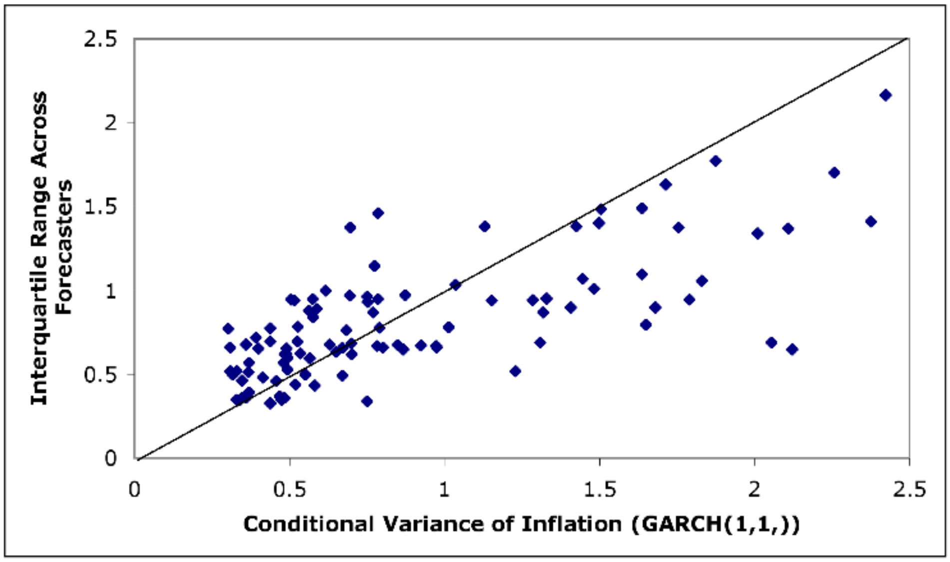

Figure 8: The inter-quartile range across one-step-ahead forecast of inflation from the Survey of Professional Forecasters and the conditional variance of inflation calculated using a GARCH(1,1) model.

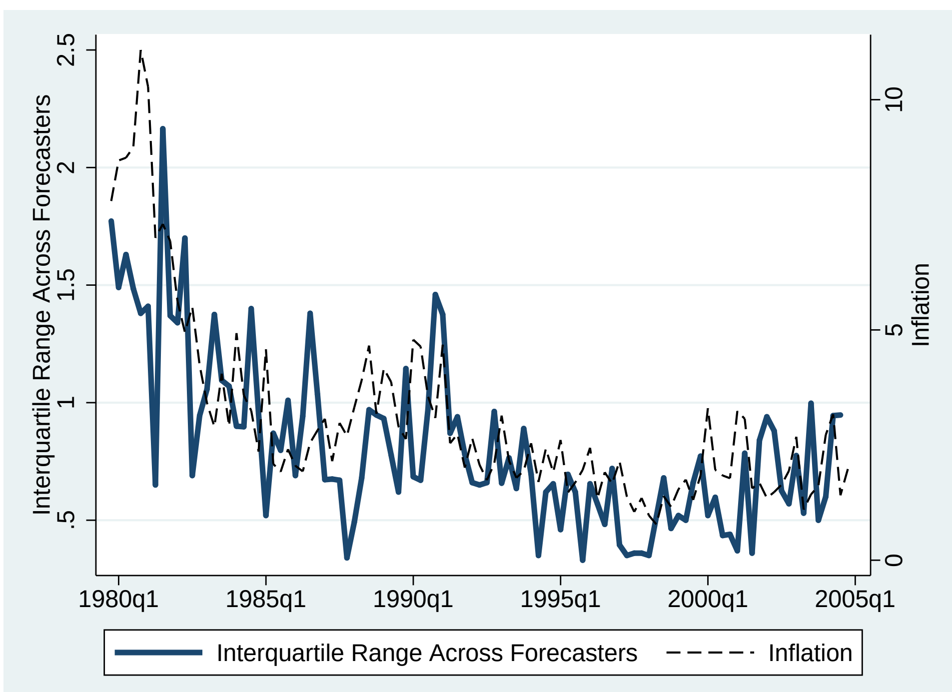

Figure 9: Disagreement across one-step-ahead inflation forecasters as measured by inter-quartile range, Survey of Professional Forecasters.