Disagreement and Biases in Inflation Expectations[*]

Carlos Capistrán[†]   Allan Timmermann[‡]

Banco de México   UCSD

December 5, 2005

Abstract

Recent 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 biases in forecasters' ranks should be preserved over time and 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.

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 (2005) and Hommes (2005) 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.

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 our 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 is based on differences in forecasters' gains and losses from mistakes in their inflation forecasts. Simple intuition explains our results. 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 it over-predicts inflation, thereby reducing the probability of costly under-predictions. Furthermore, if costs are increasingly large, the larger in absolute value the forecast error (i.e., assuming loss is convex), then the optimal bias 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: Periods with high degrees of macroeconomic uncertainty about the price level 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,

---

[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 dispersion has been established by Cukierman and Wachtel (1979) and Mankiw, Reis and Wolfers (2003).

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.

A natural question that arises from this story is why forecasters should have asymmetric costs associated with over- and under-predicting inflation. Most obviously, asymmetric costs may simply reflect differences in the cost or benefits of the economic actions that are based on the inflation forecasts. Another reason for asymmetric costs comes from the literature that views professional forecasters as agents of end users (the principals) such that the role of the better-informed agents is to generate signals that are used to inform the principals' actions. In such situations it is commonly found that it is optimal for the agent to behave strategically and bias the forecasting signal (Ehrbeck and Waldmann (1996); Laster, Bennette and Geoum (1999); Ottaviani and Sørensen (2005)). A third and final reason is related to psychological factors that have the effect of biasing the weights on the low and high ends of the distribution of possible outcomes. Empirical evidence on the implications of asymmetric costs has been found in inflation forecasts and in forecasts of other important economic variables. For example, Capistrán (2005) 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.[4]

The rest of the paper proceeds as follows. After a brief discussion of sources of asymmetric loss, Section 2 presents a theory of optimal 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.

---

[4]See Leitch and Tanner (1991), a financial application to interest rate forecasting, for another example of the use of an objective function other than mean squared error (MSE).

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 discuss three possible reasons for asymmetric loss, namely a utility cost explanation, a psychological explanation and a strategic explanation. 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 (2005) 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 overpredict and others underpredict 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 (see, e.g., Ehrbeck and Waldmann (1996); Ottaviani and Sørensen (2005)) argue that asymmetries in the information available to forecasters versus their clients can be responsible for biases. These models assume that forecasters are remunerated based on their clients' assessments of their skills and that 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.

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 (2004). 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 $\sigma^2_{t+h,t} = var\left[\pi_{t+1}|\Omega_t\right]$: [5]

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

Under assumptions (1) and (2) it is easy to show that the optimal forecast that minimizes expected loss satisfies (see, e.g., Zellner (1986)):

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

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 (3) that forecaster $i's$ optimal $h-$step-ahead forecast at time t, $f^*_{t+h,t,i}$, is:

$$f^*_{t+h,t,i} = \mu_{t+h,t} + \frac{\phi_i}{2}\sigma^2_{t+h,t}, \tag{4}$$

---

[5]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.

where ϕ_i is the asymmetry parameter 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.

The bias in the optimal forecast, defined as $bias_{t+h,t,i} = E[\pi_{t+h} - f^*_{t+h,t,i}|\Omega_t]$, follows directly from (2) and (4):

$$bias_{t+h,t,i} = -\frac{\phi_i}{2}\sigma^2_{t+h,t}. \tag{5}$$

In general this will be non-zero, provided that loss is asymmetric ($\phi_i \neq 0$).

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

$$
\begin{aligned}
\overline{f}^*_{t+h,t} &= \frac{1}{N}\sum_{i=1}^{N}\mu_{t+h,t} + \frac{1}{N}\sum_{i=1}^{N}\left[\frac{\phi_i}{2}\sigma^2_{t+h,t}\right] \\
&= \mu_{t+h,t} + \frac{\overline{\phi}}{2}\sigma^2_{t+h,t},
\end{aligned}
\tag{6}
$$

where $\overline{\phi} = N^{-1}\sum_{i=1}^{N}\phi_i$ is the average asymmetry parameter. 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 scaled by the average asymmetry parameter, $\overline{\phi}$. If all forecasters have symmetric loss ($\phi_i = 0, i = 1, ..., N$), then this average is zero and the mean forecast is unbiased.[6] A testable implication of (6) 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, $-(\phi_i/2)\sigma^2_{t+h,t}$, 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

---

[6]The average is also zero if all forecasters have asymmetric loss but the asymmetries cancel out, $\overline{\phi} = 0$.

then be shown that the first-order autocovariance in the optimal forecast errors is given by:[7]

$$Cov\left(e_{t+1,t,i}^{*}, e_{t,t-1,i}^{*}\right) = \frac{\phi_i^2 \alpha_0^2 \alpha_1^2 (\alpha_1 + \beta_1)}{2(1 - (\alpha_1 + \beta_1)^2 - 2\alpha_1^2)(1 - \alpha_1 - \beta_1)^2}. \tag{7}$$

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 (7), 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, this is given by:

$$\begin{aligned}
\bar{s}_{t+h,t} &\equiv \left[\frac{1}{N}\sum_{i=1}^{N}\left(f_{t+h,t,i}^{*} - \bar{f}_{t+h,t}^{*}\right)^2\right]^{\frac{1}{2}} \\
&= \left[\frac{1}{N}\sum_{i=1}^{N}\left(\mu_{t+h,t} + \frac{\phi_i}{2}\sigma_{t+h,t}^2 - \left(\mu_{t+h,t} + \frac{\bar{\phi}}{2}\sigma_{t+h,t}^2\right)\right)^2\right]^{\frac{1}{2}} \\
&= \frac{\sigma_{t+h,t}^2}{2}\left[\frac{1}{N}\sum_{i=1}^{N}\left(\phi_i - \bar{\phi}\right)^2\right]^{\frac{1}{2}}. \tag{8}
\end{aligned}$$

The dispersion in inflation forecasts varies with the conditional variance of inflation, $\sigma_{t+h,t}^2$, scaled by a term reflecting cross-sectional variation or 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 inflation even if loss is asymmetric ($\bar{\phi} \neq 0$).

Our simple model can also 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

---

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

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 (8)). 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. Under assumptions (1) and (2), the inter-quartile range is:

$$f^*_{t+h,t,0.75} - f^*_{t+h,t,0.25} = \left(\frac{\phi_{0.75} - \phi_{0.25}}{2} \right) \sigma^2_{t+h,t}, \tag{9}$$

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.

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

$$f^*_{t+h,t,i} - f^*_{t+h,t,j} = \left(\frac{\phi_i - \phi_j}{2} \right) \sigma^2_{t+h,t}. \tag{10}$$

The ranking of the predictions (or prediction errors) produced by two forecaster should therefore not change over time provided that the degree of loss asymmetry is constant: If forecaster i dislikes positive forecast errors more than forecaster j, so $\phi_i > \phi_j$, then forecaster i's predictions should generally be higher than those of forecaster j. This implication is easy to test empirically.

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 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

$$E_t[e_{t+4,t,i} - e_{t+1,t,i}] = -(1/2)\phi_i(\sigma^2_{t+4,t} - \sigma^2_{t+1,t}). \tag{11}$$

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 sign of ϕ_i as well as on whether $\sigma^2_{t+4,t}$ exceeds or falls below $\sigma^2_{t+1,t}$, which in turn will depend on whether the initial volatility level is above or below its long-run average.

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 optimal forecast is:

$$f^*_{t+h,t,i} = \mu_{t+h,t,i} + \frac{\phi_i}{2}\sigma^2_{t+h,t,i}, \tag{12}$$

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

Aggregating across forecasters, the average inflation forecast now becomes:

$$
\begin{aligned}
\overline{f}^*_{t+h,t} &= \frac{1}{N}\sum_{i=1}^{N}\mu_{t+h,t,i} + \frac{1}{N}\frac{\sum_{i=1}^{N}\phi_i\sigma^2_{t+h,t,i}}{2} \\
&= \overline{\mu}_{t+h,t} + \frac{1}{2}\tilde{\sigma}^2_{t+h,t},
\end{aligned} \tag{13}
$$

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 and (ii) a weighted average of the variance of inflation $\tilde{\sigma}^2_{t+h,t}$ where the weights are the scaled loss asymmetry parameters, $\phi_i/(2N)$. When the underlying values of $\mu_{t+h,t,i}$ and $\sigma^2_{t+h,t,i}$ are unobserved, equation (13) 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.[8]

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

---

[8]This result holds more broadly for general asymmetric loss functions and distributions of inflation, see Granger (1969) and Patton and Timmermann (2004).

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

$$
\begin{aligned}
\bar{s}^2_{t+h,t} &= \frac{1}{N}\sum_{i=1}^{N}\left(\mu_{t+h,t,i} + \frac{\phi_i}{2}\sigma^2_{t+h,t,i} - \left(\overline{\mu}_{t+h,t} + \frac{1}{2}\tilde{\sigma}^2_{t+h,t}\right)\right)^2 \\
&= \frac{1}{N}\sum_{i=1}^{N}\left[\left(\mu_{t+h,t,i} - \overline{\mu}_{t+h,t}\right) + \frac{\left(\phi_i\sigma^2_{t+h,t,i} - \overline{\phi}\bar{\sigma}^2_{t+h,t}\right)}{2} - \frac{1}{N}\sum_{i=1}^{N}\frac{(\phi_i - \overline{\phi})}{2}\sigma^2_{t+h,t,i}\right]^2 \\
&= \frac{1}{N}\sum_{i=1}^{N}\left[\left(\mu_{t+h,t,i} - \overline{\mu}_{t+h,t}\right)^2 + \left(\frac{\phi_i\sigma^2_{t+h,t,i} - \overline{\phi}\bar{\sigma}^2_{t+h,t}}{2}\right)^2\right] \qquad (14) \\
&\quad + \left(\frac{1}{N}\sum_{i=1}^{N}\frac{(\phi_i - \overline{\phi})}{2}\sigma^2_{t+h,t,i}\right)^2 + \frac{1}{N}\sum_{i=1}^{N}\left(\mu_{t+h,t,i} - \overline{\mu}_{t+h,t}\right)\left(\phi_i\sigma^2_{t+h,t,i} - \tilde{\sigma}^2_{t+h,t}\right).
\end{aligned}
$$

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 (i) differences in beliefs about the mean of future inflation; (ii) differences in beliefs about the variance of the inflation rate, weighted by the loss asymmetry parameters; (iii) a term reflecting the cross-sectional covariance between the loss asymmetry parameters and variance forecasts; and (iv) a term reflecting systematic correlations between mean and variance forecasts: If those forecasters who anticipate a relatively high mean of the inflation rate also anticipate a relatively high variance, then the cross-sectional dispersion will be greater.

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 four 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.

Letting $\left(\phi_i\sigma^2_{t+h,t,i}\right)_q$ be the qth quantile of the cross-sectional distribution of $\left(\phi_i\sigma^2_{t+h,t,i}\right)$, the inter-quartile range across forecasters under heterogeneous information is now:

$$
f^*_{t+h,t,0.75} - f^*_{t+h,t,0.25} = \left(\mu_{t+h,t,0.75} - \mu_{t+h,t,0.25}\right) + \frac{1}{2}\left(\left(\phi_i\sigma^2_{t+h,t,i}\right)_{0.75} - \left(\phi_i\sigma^2_{t+h,t,i}\right)_{0.25}\right). \qquad (15)
$$

As before, the dispersion across forecasters is driven by a combination of differences in the information sets 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 *Under homogenous beliefs about first and second moments of future inflation, the ranking of inflation forecasts across forecasters remains constant over time;*

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.*

3 Empirical Evidence

In this section we use data from the Survey of Professional Forecasters (SPF) and from the Livingston Survey (LS) 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 but does not identify the type of forecaster. In contrast, the shorter and smaller LS data identifies the affiliation of the forecaster which allows us to relate forecasting biases to the type of forecaster.

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-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. We include forecasters with at least 30 non-zero forecasts. This leaves us with data on 62 forecasters.

The LS data consists of six-month-ahead inflation forecasts, with inflation measured as the annualized biannual change in the consumer price index (CPI). The sample begins during the second semester of 1992 (which is the time when current quarter figures started to get included in the Survey) and runs through the first semester of 2004. This data contains information on the affiliation of the forecasters. Most affiliations have very few observations, so only those corresponding to Industry, Academic, and Banking will be considered. Given the short sample, we keep forecasters with at least 10 non-zero forecasts, resulting in 7 academic forecasters, 13 forecasters from industry, and 16 from banking institutions.

For actual values of the output deflator and CPI 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.[9]

---

[9]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 1 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.

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 (3). 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) model of the form:

$$
\begin{aligned}
\pi_{t+1,t} &= \lambda_0 + \lambda_1 \pi_{t,t-1} + \lambda_2 \pi_{t-3,t-4} + \varepsilon_{t+1}, \\
\varepsilon_{t+1} &\sim N(0, \sigma^2_{t+1,t}), \\
\sigma^2_{t+1,t} &= \omega + \alpha_1 \varepsilon_t^2 + \beta_1 \sigma^2_{t,t-1}.
\end{aligned}
\tag{16}
$$

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. 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 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.[10] A histogram of the asymmetry parameters estimated with real-time data is presented in Figure 2. As with the mean forecast errors, the histogram is skewed to the right and some forecasters have positive asymmetry parameter (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. With 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%).[11]

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 3 presents a plot of a traditional quadratic (Mean Squared Error) 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

---

[10]Ideally, one would use the conditional variance estimated from the series of individual forecast errors, $\widehat{\sigma}^2_{t+h,t,i}$. Two reasons prevent us from doing this. First, for some forecasters we only have 30 observations; second, the time-series of individual forecast errors are not always contiguous (i.e., a forecaster may drop out from the survey for a while and then return).

[11]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 Elliot, Komunjer, and Timmermann (2005) procedure to estimate the asymmetry parameter under an asymmetric quadratic loss function. Because this class naturally nests Mean Squared Error 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.

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.

3.2.2 Mean Forecast Errors Across Forecasters

Equations (6) and (13) imply that the mean inflation forecast depends on the conditional variance of inflation. Under homogenous information, the mean bias is:

$$
\begin{aligned}
E\left[\pi_{t+h} - \overline{f}^*_{t+h,t}|\Omega_t\right] &= E\left[\pi_{t+h} - \overline{\mu}_{t+h,t} - \frac{\overline{\phi}}{2}\sigma^2_{t+h,t}|\Omega_t\right] \\
&= -\frac{\overline{\phi}}{2}\sigma^2_{t+h,t}.
\end{aligned}
\tag{17}
$$

Under our model, the mean forecast error is predicted by variations in the conditional variance. Conversely, under MSE loss, the bias should be zero. This implication can be tested through a regression of the form:

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

where, under asymmetric loss $\delta_1 \neq 0$.

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 $\overline{\phi}$.

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 4 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

16

autocorrelation coefficients are positive.

3.2.4 Term Structure of Inflation Forecasts

To test the term structure implications of our model (equation (11) 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.

3.2.5 Rank Preservation

Under homogenous information our theory implies (Claim 5) that the ranking of mean forecasts and forecast errors is preserved over time, despite changes in the variance of inflation. To test such 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.[12] 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 0.714. 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.[13]

3.2.6 Bias and Type of Forecaster

We also calculated mean forecast errors for each forecaster using the data from the Livingston Survey (LS) which provides the forecaster's affiliation. Figure 5 presents a histogram of the bias estimates. Forecasters with an academic affiliation produce mean forecast errors closer to zero than forecasters in industry, which in turn have less dispersed mean forecast errors than forecasters

---

[12]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.

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

from the banking sector. Tests of unbiasedness cannot reject the null of a zero bias for any of the academic forecasters, but reject it at the 5% critical level for 13% of the forecasters in banks, and for 15% of the forecasters in industry using real-time data. If belief dispersion were driven entirely by differences in information sets, then this evidence would suggest that (a) academic forecasters have more homogenous information, and (b) academics' information is better (in the sense that it produces a smaller bias) than the information of the forecasters in banks and industry. These suggestions are difficult to defend. Alternatively, under asymmetric loss (even with homogeneous information sets) the evidence suggests that forecasters in academia have more symmetric losses from over- or under-predicting inflation than forecasters in industry or in the banking sector.

It is difficult to come up with good explanations for why academic forecasters should have strongly asymmetric loss and thus bias their forecasts. In contrast, a simple example serves to illustrate why forecasters in the banking sector may have asymmetric loss. Consider a bank with a large net asset position in fixed-interest bonds. If inflation rises by more than the predicted amount, the interest rate is likely to rise and so bond values will decline, translating into greater losses for the bank. Conversely, if the inflation and interest rate come out lower than anticipated, then the bank stands to gain although it clearly suffers opportunity costs relative to the potential gains it could have earned had it correctly predicted the lower interest rate (in which case it might have held a bond portfolio with a longer duration). In addition, increases in the interest rate typically have different effects on bond prices than decreases of similar size due to the convex relationship between interest rates and bond prices. Clearly the effect on the bank's utility or loss of greater-than-expected and smaller-than-expected inflation are not symmetric and the bank may prefer negative forecast errors (inflation comes out lower than expected) rather than negative ones (inflation comes out higher than expected). The converse argument can be worked out for a bank with a net liability exposure to interest rate risk.

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 devia-

tion across forecasters who participated in the SPF. Consistent with the literature cited in the introduction, there is a positive relation between the variables.

3.3.1 Dispersion Across Forecasters

Our sixth claim, and a direct implication of equations (8) and (14), was that the dispersion across forecasters depends on the conditional variance of inflation. To test this we estimate the following model:

$$\bar{s}_{t+h,t} = \gamma \sigma^2_{t+h,t} + \varepsilon_{t+h}. \tag{18}$$

The dependent variable is now the cross-sectional standard deviation of inflation forecasts and the regressors are simply a constant and the conditional variance from the GARCH(1,1) model. If the forecasters have asymmetric loss we expect to find $\gamma > 0$, while under symmetric loss $\gamma = 0$. 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 $\gamma < 0$. 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 also predicts that the inter-quartile range across forecasters is positively correlated with the variance of inflation. Under homogeneous information an implication of equation (9) 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 implication 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 (16) 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:

$$\sigma^2_{t+1,t} = \omega + \delta\pi_t. \tag{19}$$

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.[14] 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.

---

[14]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).

3.4 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 elegant 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 1969:2-2003:4. However, in line with Croushore, we also find that the average forecast is biased in the first part of the data sample that finishes in 1983. 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.[15]

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

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 proposed 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. This simple observa-

---

[15]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.

tion 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. 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., and S.G. Cecchetti, 1990, Inflation and Uncertainty at Long and Short Horizons, *Brookings Papers on Economic Activity*, Brookings Institution Press.

[2] Birnbaum, M.H., and S.E. Stegner, 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., and R. Cohen, 2001, To Aggregate, Pool, or Neither: Testing the Rational Expectations Hypothesis Using Survey Data, *Journal of Business and Economic Statistics,* 19, 278–291.

[5] Capistrán, C., 2005, Bias in Federal Reserve Inflation Forecasts: Is the Federal Reserve Irrational or Just Cautious, *UCSD discussion paper 2005-05*.

[6] Carlson, J.A., and N.T. Valev, 2003, Sources of Dispersion in Consumer Inflation Forecasts, *Applied Economic Letters*, 10(2) 77-81.

[7] Carroll, C., 2003, Macroeconomic Expectations of Households and Professional Forecasters, *Quarterly Journal of Economics*, 118(1), 269-298.

[8] Christoffersen. P.F., and F.X. Diebold, 1996, Further Results on Forecasting and Model Selection Under Asymmetric Loss, *Journal of Applied Econometrics*, 11, 561-572.

[9] Christoffersen. P.F., and F.X. Diebold, 1997, Optimal Prediction Under Asymmetric Loss, *Econometric Theory*, 13, 808-817.

[10] Clarida, R., J. Galí, and M. Gertler, 1999, The Science of Monetary Policy: A New Keynesian Perspective, *Journal of Economic Literature*, 36, 1661-1707.

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

[12] Croushore, D., 1998, Evaluating Inflation Forecasts, *Federal Reserve Bank of Philadelphia Working Paper #98-14*.

[13] Croushore, D., and T. Stark, 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.

[14] Cukierman, A. and P. Wachtel, 1979, Differential Inflationary Expectations and the Variability of the Rate of Inflation: Theory and Evidence, *American Economic Review*, 69(4), 595-609.

[15] Elliott, G., I. Komunjer, and A. Timmermann, 2005, Estimation and Testing of Forecast Rationality under Flexible Loss, *The Review of Economic Studies*, 72, 1107-1125.

[16] Ehrbeck, T. and R. Waldmann, 1996, Why Are Professional Forecasters Biased? Agency versus Behavioral Explanations, *Quarterly Journal of Economics*, 111, 21-40.

[17] Engle, R.F., 1982, Autoregressive Conditional Heteroskedasticity with Estimates of the Variance of U.K. Inflation, *Econometrica*, 50, 987-1008.

[18] Engle, R.F., and D. Kraft, 1983, Multiperiod Forecast Error Variances of Inflation Estimated from ARCH Models, in A. Zellner (ed.), *Applied Time Series Analysis of Economic Data*, 293-302. Washington, D.C.: Bureau of the Census.

[19] Evans, G.W., and S. Honkapohja, 2001, *Learning and Expectations in Macroeconomics*, Princeton, NJ: Princeton University Press.

[20] Fama, E., 1991, "Efficient Markets: II" Fiftieth Anniversary Invited Paper,*Journal of Finance*, 46, 1575-1617.

[21] Granger, C.W.J., 1969, Prediction with a Generalized Cost Function, *Operational Research*, 20, 199-207.

[22] Granger, C.W.J., and P. Newbold, 1986, *Forecasting Economic Time Series*, 2nd edition. Orlando: Academic Press.

[23] Grier, K., and M. J. Perry, 1998, On inflation and Inflation Uncertainty in the G-7 Countries, *Journal of International Money and Finance*, 17.

[24] Hommes, C.H., 2005, Heterogeneous Agent Models in Economics and Finance, forthcoming in: L. Tesfatsion and K. Judd. (eds.), *Handbook of Computational Economics II: Agent-Based Computational Economics*.

[25] Ito, T., 1990, Foreign Exchange Rate Expectations: Micro Survey Data, *American Economic Review*, 80(3), 434-449.

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

[27] Kiefer, N.M., and T.J. Vogelsang, 2002, Heteroskedasticity-Autocorrelation Robust Standard Errors Using the Bartlett Kernel without Truncation, *Econometrica*, 70, 2093-2095.

[28] Laster, D., P. Bennett, and I.S. Geoum, 1999, Rational Bias In Macroeconomic Forecasts, *Quarterly Journal of Economics*, 114(1), 293-318.

[29] Leitch, G. and J.E. Tanner, 1991, Economic Forecast Evaluation: Profits Versus the Conventional Error Measures, *American Economic Review*, 81, 580-90.

[30] Logue, D.E., and T.D. Willett, 1976, A Note on the Relation Between the Rate and the Variability of Inflation, *Economica*, 43, 151-158.

[31] Lucas, R.E., Jr., 1973, Some International Evidence on Output-Inflation Tradeoffs, *American Economic Review*, 63, 326-334.

[32] Mankiw, G.N., and R. Reis, 2002, Sticky Information Versus Sticky Prices: A Proposal to Replace the New Keynesian Phillips Curve, *Quarterly Journal of Economics*, 117(4), 1295-1328.

[33] Mankiw, G.N., R. Reis, and J. Wolfers, 2003, Disagreement about Inflation Expectations, *NBER Macroeconomics Annual*.

[34] Olds, E.G., 1938, Distribution of Sums of Squares of Rank Differences for Small Samples, *Annals of Mathematical Statistics*, 9.

[35] Ottaviani, M., and P.N. Sørensen, 2005, The Strategy of Professional Forecasting, *LBS mimeo.*

[36] Patton, A. and A. Timmermann, 2004, Properties of Optimal Forecasts under Asymmetric Loss and Nonlinearity, *UCSD mimeo.*

[37] Pesaran, M.H., and M. Weale, 2005, Survey Expectations, forthcoming in: G. Elliott, C.W.J. Granger, and A. Timmermann (eds.), *Handbook of Economic Forecasting.*

[38] Pigou, A., 1927, *Industrial Fluctuations*, London: Macmillan.

[39] 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(1), 39-72.

[40] Timmermann, A, 1993, How Learning in Financial Markets Generates Excess Volatility and Predictability in Stock Prices, *Quarterly Journal of Economics*, 108, 1135-1145.

[41] Townsend, R.M., 1983, Forecasting the Forecasts of Others, *The Journal of Political Economy*, 91(4), 546-588.

[42] Varian, H., 1974, A Bayesian Approach to Real Estate Assessment, in S.E. Fienberg and A. Zellner (eds.), *Studies in Bayesian Econometrics and Statistics in Honor of L.F. Savage*, 195-208. Amsterdam: North-Holland.

[43] 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.

[44] Weber, E.U., and B. Kirsner, 1997, Reasons for Rank-Dependent Utility Evaluation, *Journal of Risk and Uncertainty*, 14, 41-61.

[45] Zarnowitz, V., and P. Braun, 1992, Twenty-two Years of the NBER-ASA Quarterly Economic Outlook Surveys: Aspects and Comparisons of Forecasting Performance, *NBER Working Paper No. 3965.*

[46] 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	158	157
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.

Statistic	Revised Data	Real-time Data
Quantile 0.25	-0.85	-0.74
Quantile 0.50	-0.20	-0.17
Quantile 0.75	0.98	0.76
Inter-quartile Range	1.83	1.50
Maximum	3.25	3.18
Minimum	-2.15	-1.80
Mean	0.06	0.08

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

Figure 2: Histogram of linex asymmetry parameter estimates, Survey of Professional Forecasters (One-step-ahead forecasts, real-time data).


Figure 3: Degree of asymmetry.

Figure 4: Histogram of first order autocorrelations, Survey of Professional Forecasters (One-step-ahead forecasts, real-time data).


Figure 5: Histogram of mean forecast errors, Livingston Survey (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.