EUROPEAN CENTRAL BANK

EUROSYSTEM

WORKING PAPER SERIES

NO 825 / OCTOBER 2007

WHAT CAN PROBABILITY FORECASTS TELL US ABOUT INFLATION RISKS?

by Juan Angel García

and Andrés Manzanares

EUROPEAN CENTRAL BANK

EUROSYSTEM

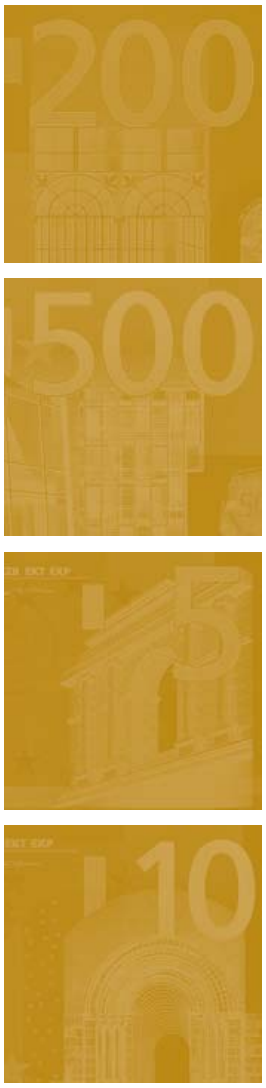

WORKING PAPER SERIES

NO 825 / OCTOBER 2007

WHAT CAN PROBABILITY FORECASTS TELL US ABOUT INFLATION RISKS?[1]

b Ja A ge Ga ca[2]

ad Ad 1 Ma a a e[3]

In 2007 all ECB publications feature a motif taken from the €20 banknote.

T a e ca be d aded c a e f ,
:// .ecb.e a.e f , e S c a Sc e ce Re ea c Ne
e ec c b a a :// . c , /ab ac d=1020964.

1 We are very grateful to Geert Bekaert, Gonzalo Camba-Méndez, Refet Gürkaynak, Gabriel Perez-Quirós, Thomas Werner, an anonymous referee and seminar participants at the ECB and in the conference "Price Stability and Monetary Policy" organized by the Federal Reserve Banks of Dallas and Cleveland in Dallas, May 2007, for useful comments and suggestions. We are particularly indebted to Paul Söderlind for many fruitful discussions in the early stages of this project, and to Ken Wallis for very detailed suggestions. Excellent research assistance by Renate Dreiskena is also gratefully acknowledged. Any remaining errors are our responsibility. The views expressed in this paper are those of the authors and do not necessarily reflect the views of the European Central Bank.

2 Capital Markets and Financial Structure Division, DG-E, European Central Bank, Kaiserstrasse 29, 60311 Frankfurt am Main, Germany; e-mail: juan_angel.garcia@ecb.int

3 Risk Management Division, DG-H, European Central Bank, Kaiserstrasse 29, 60311 Frankfurt am Main, Germany; e-mail: andres.manzanares@ecb.int

Address
Kaiserstrasse 29
60311 Frankfurt am Main, Germany

Postal address
Postfach 16 03 19
60066 Frankfurt am Main, Germany

Telephone
+49 69 1344 0

Website
http://www.ecb.europa.eu

Fax
+49 69 1344 6000

Telex
411 144 ecb d

The statement of purpose for the ECB Working Paper Series is available from the ECB website, http://www.ecb.europa.eu/pub/scientific/wps/date/html/index.en.html

ISSN 1561-0810 (print)
ISSN 1725-2806 (online)

CONTENTS

Abstract

A crucial but often ignored element of inflation expectations is the amount of perceived inflation risk. This paper estimates the degree of uncertainty and asymmetry in the probability forecasts of the Survey of Professional Forecasters (SPF) using a new methodology. The main conclusion from our analysis is that, when monitoring inflation expectations, limiting attention to a point prediction is not sufficient. The analysis of inflation expectations should take into account inflation risks. As an example, we show that our measures of inflation risks can better explain why inflation scares happened in the bond market during the Volcker disinflation.

Keywords: *Inflation risk, inflation expectations, Survey of Professional Forecasters (SPF), skew-normal distribution, power divergence estimators*

JEL Classification: *C16, C42, E31, E47*

Non-technical summary

A crucial but often ignored element of inflation expectations is the amount of perceived inflation risk. Inflation expectations play a central role in economic analysis. All central banks monitor closely private sector's inflation expectations, using both indicators derived from financial instruments and survey of inflation expectations. Yet, researchers and policymakers restrict attention to point predictions and neglect inflation risks.

The analysis of inflation expectations appears to be restricted to point predictions for two reasons. First, most surveys, which are a key source of information about inflation expectations, only gather information about a point prediction. Quantitative evidence on the perceived inflation risks is therefore missing. Second, the analysis of inflation expectations often uses linear models that neglect inflation risks.

In this context, the purpose of this paper is twofold. First, we provide quantitative evidence of the perceived inflation risks over the last four decades. Our quantitative evidence comprises two key metrics to assess the perceived risks in inflation expectations, namely the degree of forecast uncertainty and the degree of forecast skewness (i.e. the asymmetry in the forecasts). We then illustrate that the analysis of inflation risks could play a key role in the analysis of central bank credibility. As an example we show that our measures of inflation risks in the 1980s shed new light on the bond market inflation scares faced by the Federal Reserve under chairman Volcker.

We use the probability forecasts of the Survey of Professional Forecasters (SPF) to measure perceived inflation risks. The SPF is a quarterly survey of macroeconomic expectations that, although launched in 1968 by the American Statistical Association and the NBER, is nowadays conducted by the Federal Reserve Bank of Philadelphia. The SPF allows to estimate the risks surrounding inflation expectations because forecasters report probabilities of inflation falling into pre-specified intervals, i.e. a density forecast in the form of a histogram. The literature has proposed several alternative approaches to analyze the SPF histograms, but extracting risk measures from those histograms, either by parametric or non-parametric methods has proved quite challenging. To obtain reliable estimates of inflation risks for the SPF data, this paper introduces a new methodology.

We measure inflation risks by fitting a density to the SPF histograms. Our methodology departs from existing literature in two fundamental aspects, namely the choice of a fitting criterion and the choice of underlying density. As fitting criterion we propose a small departure from maximum likelihood estimation. As underlying density we employ a potentially asymmetric distribution, the skew-normal. Skewness is a crucial feature of any

forecast, and, often provides valuable information. For example, most central bank policy statements include references to the "assessment of inflation risks". By using the skew normal-density, we "*let the data speak*" about perceived asymmetries in inflation risks.

The main conclusion from our analysis is that, when monitoring inflation expectations, limiting attention to a point prediction is not sufficient. A thorough analysis of inflation expectations should take into account inflation risks. We find that (mean) inflation expectations tend to co-move with inflation uncertainty, but move little with inflation skewness. We document strong swings in inflation risks in the last few decades, swings that were not shared by standard measures of mean inflation expectations and that can improve our understanding of many economic events.

Inflation risks offer additional dimensions to assess the credibility of monetary policy. Central bank credibility should not only anchor mean inflation expectations but also reduce perceived inflation risks. Goodfriend and King [2005] provide an excellent historical analysis of the Volcker disinflation and persuasively argued that one of the crucial elements to understand that disinflation is the imperfect credibility of the Volcker Fed, which led to four "inflation scares" in the 1980s. Our analysis shows that, for most of the 1980s, perceived inflation risks remained high, even when mean inflation expectations (and actual inflation) were declining. Moreover, to understand the movements in long-term interest rates in the 1980s, changes in perceived inflation risks, particularly upside risks to the inflation outlook, are fundamental. A close monitoring of the risks to inflation expectations is therefore useful to gauge central bank credibility.

Policymakers and researchers can obtain valuable information about the general macroeconomic situation by monitoring inflation risks. We show that the influence of the business cycle on inflation risks is strong and timely: inflation risks tend to decline ahead of recessions, are strongly influenced by the perceived uncertainty about real GDP growth and also respond to the perceived probability of recession over the forecast horizon. The strength of the relationship between business cycle conditions and inflation risks has however varied over time, and, since the early 1990s, seems to have weakened.

The evidence suggests that a thorough analysis of the probability distributions of the SPF provides additional insights not only about how inflation expectations change over time but also about developments in other macroeconomic and financial variables. We hope that the methodology developed in this paper and our analysis of the dynamics of inflation beliefs during the Volcker disinflation period could be a first step in that regard.

"...we must consider not only what appears to be the most likely outcome, but also the risks to that outlook..."

– Testimony of Fed Chairman Ben Bernanke to the U.S. Senate, July 19, 2006 –

1 Introduction

A crucial but often ignored element of inflation expectations is the amount of perceived inflation risk. Inflation expectations play a central role in economic analysis. Central banks closely monitor inflation expectations, both from financial market indicators and surveys. Yet, that researchers and policymakers restrict attention to point predictions and neglect inflation risks is very surprising. To explain developments in inflation expectations, a point forecast is only sufficient under very restrictive assumptions. In general, understanding inflation expectations requires the analysis of perceived inflation risks.[1]

The analysis of inflation expectations appears to be restricted to point predictions for two main reasons. First, inflation surveys emphasize point predictions, and evidence on perceived inflation risks is either not requested at all or, when requested, not easy to assess.[2] Second, although in recent years research has emphasized time-varying inflation volatility and time-varying inflation risk premia as crucial elements to understand US macroeconomic dynamics, the analysis of inflation expectations often uses linear (or linearized) models that ignore higher-order moments and abstract from inflation risks.

The purpose of this paper is twofold. We first provide quantitative evidence of the perceived inflation risks over the last four decades. Our quantitative evidence comprises two key metrics to assess the perceived risks in inflation expectations, namely the degree of forecast uncertainty and the degree of forecast skewness. Second, we show that the analysis of inflation risks explains the bond market inflation scares in the 1980s.

We use the probability forecasts of the Survey of Professional Forecasters (SPF) to measure perceived inflation risks. The American Statistical Association and the NBER launched the SPF in 1968 as a quarterly survey of macroeconomic expectations, and the Federal Reserve Bank of Philadelphia took over in 1990. We can therefore analyze perceived inflation risks since 1968Q4. The SPF allows to estimate the risks surrounding inflation expectations because forecasters report probabilities of inflation falling into prespecified intervals, i.e. a density forecast in the form of a histogram. The literature has proposed several alternative approaches to analyze the SPF histograms, but extracting

[1] Rather than (knightian) *"uncertainty"* we prefer inflation *"risks"* because we summarize the inflation randomness by the statistical moments of probability forecasts (see Machina and Rothchild [2007]).

[2] Survey data are often critisized because panelists have little incentive to provide their true expectations. The superior accuracy of survey forecasts found by Ang *et al.*[2007] however questions those concerns.

risk measures from those histograms, either by parametric or non-parametric methods has proved quite challenging.[3]

To obtain reliable estimates of inflation risks from the SPF data, this paper introduces a new methodology. We follow recent literature (see for example Giordani and Söderlind [2003], Engelberg, Manski and Williams [2007]) and measure inflation risks by fitting a theoretical density to the SPF histograms. Our methodology however departs from existing literature in two fundamental aspects, namely the choice of a fitting criterion and the choice of underlying density.

As fitting criterion we propose a small departure from maximum likelihood estimation. Extracting reliable risk measures from the SPF histograms requires an efficient and robust estimator. Each SPF histogram is a discretized version of the (true) density forecast. We assume that the discretization reflects how many "draws" from the true density lie within each of the pre-specified intervals in the questionnaire, and therefore interpret the reported probabilities as the realization of a multinomial distribution. In this context, least squares, the fitting criterion usually employed in existing literature on SPF data, is not efficient. In addition, a small departure from maximum likelihood provides additional robustness to zero-probability intervals in the SPF data.

As theoretical density we employ a potentially asymmetric distribution, the Azzalini's [1985] skew-normal family. Recent SPF work neglects inflation skewness (see Giordani and Söderlind [2003], Rich and Tracy [2006] and D'Amico and Orphanides [2006]). Skewness is however a crucial feature of any forecast, and, often provides valuable additional information. For example, most central bank policy statements include references to the "assessment of inflation risks". Moreover, if the SPF probability forecasts are skewed, neglecting asymmetries in the theoretical density leads to biased estimates. By using the skew normal density, we "*let the data speak*" about perceived asymmetries in inflation risks. Monte Carlo evidence confirms that these two methodological contributions lead to significant accuracy gains.

The main conclusion from our analysis is that, when monitoring inflation expectations, limiting attention to a point prediction is not sufficient. A thorough analysis of inflation expectations should take into account inflation risks. We find that (mean) inflation expectations tend to co-move with inflation uncertainty, but move little with inflation skewness. We document strong swings in inflation risks in the last few decades, swings that were not shared by standard measures of mean inflation expectations and that can improve our understanding of many economic events.

[3] The pioneering work of Zarnowitz and Llambros [1987] and Lahiri *et al.* [1988] used non-parametric methods. In search for more reliable estimates, more recent research proposed fitting a distribution to the SPF histograms.

Measures of inflation risks offer additional information to assess the credibility of monetary policy. Central bank credibility implies that (mean) inflation expectations are well-anchored and perceived inflation risks are low. A striking feature of our estimates of inflation risks is their behavior in the 1980s, during the Volcker disinflation period. Goodfriend and King [2005] provide a detailed historical analysis of the Volcker disinflation and argue that a crucial element to understand this disinflation is the imperfect credibility of the Volcker Fed. Indeed, the Fed faced four "inflation scares" in the 1980s. To gauge its credibility, the FOMC focused on long-term bond yields. Our analysis shows that, to understand the movements in long-term interest rates in the 1980s, changes in perceived inflation risks, particularly upside risks to the inflation outlook, are fundamental. For most of the 1980s, perceived inflation risks remained high, even when mean inflation expectations (and actual inflation) were declining. The high levels of perceived inflation risks explain why inflation scares happened in the bond market, in particular in 1983-84, when inflation and inflation expectations appeared contained. A close monitoring of the risks to inflation expectations proves therefore useful to gauge central bank credibility.

Policymakers and researchers can also obtain valuable information about the general macroeconomic situation by monitoring inflation risks. A reason for monitoring inflation expectations is that, besides the inflation outlook, they comprise information about the general macroeconomic outlook. In particular, monetary policy discussions often interpret changes in inflation expectations in the context of expected changes in the general macroeconomic situation, for good reasons. We show that the influence of the business cycle on inflation risks is strong and timely: inflation risks tend to decline ahead of recessions, are strongly influenced by the perceived uncertainty about real GDP growth and also respond to the perceived probability of recession over the forecast horizon. The strength of the relationship between business cycle conditions and inflation risks has however varied over time, and, since the early 1990s, seems to have weakened. These findings provide strong support for the work allowing for time-varying relationships with inflation expectations and opens new lines of research for the role of inflation risks in business cycle analysis.

The remainder of the paper is as follows. Section 2 introduces our methodology to measure the inflation risks, in particular our choice of theoretical density function and fitting criterion. Section 3 presents Monte Carlo results that support the accuracy gains from our proposed methodology. Section 4 provides stylized facts on inflation risks. We also analyze the relationship between different moments of inflation forecasts. Section 5 focuses on the dynamics of private sector beliefs about inflation in the 1980s and interprets the inflation scares during the Volcker disinflation period in the light of our inflation risk measures. Section 6 investigates the relationship between inflation risks and key macroeconomic variables. Finally Section 7 concludes.

2 A new methodology to analyze the SPF histograms

The Survey of Professional Forecasters (SPF) asks panelists to assign probability to future inflation falling within some predetermined intervals.[4] The subjective probability forecasts are therefore reported in the form of histograms. As part of the published survey results, the individual histograms are aggregated across panelists to construct a *combined* probability forecast, which reflects the average probability assigned to each interval in every survey round.

We interpret the SPF histograms as a discretized version of an unknown density forecast, $f_{\pi k}$, of each forecaster $k = 1, \ldots, K$. A thorough analysis of the information content of the SPF probability forecasts requires to elicit the underlying density forecast from the reported frequencies. In theory, the probabilities assigned to each survey interval should correspond to the integrals of the underlying density function $p_{ik} := \int_{\alpha_{i-1}}^{\alpha_i} f_{\pi k}(x)dx$, calculated over each of the intervals (α_{i-1}, α_i), $i = 1, \ldots, I$.[5] In practice, however, it is unlikely that survey participants discretize their subjective density forecast by computing those integrals.

As working assumption we interpret the reported probabilities as the proportion of "random draws" taken from the subjective density forecast that lie within each of the intervals (α_{i-1}, α_i). Without loss of generality, we assume that the unknown density function $f_{\pi k}$ belongs to a suitable parametric family of distributions f_{ϱ_k} where $\varrho_k \in \Theta \subseteq R^r$ and r is the number of parameters characterizing the family. [6] Formally, we interpret the probabilities assigned to each interval by the kth forecaster ($\{\widehat{p}_{ik}\}$, $i = 1, \ldots, I$) as the realization of a multinomial random variable with I classes.

In this multinomial framework, the observed frequencies $\{\widehat{p}_{ik}\}$ are a sufficient statistic for estimating the theoretical probabilities $\{p_{ik}\}$. Formally, denoting by $\{Z_n\}_{n=1,\ldots,N}$ the random draws used to fill in the questionnaire by a panelist, we have:

$$P\left(Z_1 = z_1, \ldots, Z_N = z_N \mid \widehat{p}_1, \ldots, \widehat{p}_I\right) = \frac{\prod_{i=1}^{I}(n\widehat{p}_i)!}{n!}, \forall \{z_n\}_{n=1,\ldots,N}$$

which shows that, as the right-hand side of the expression only depends on the observed frequencies, sufficiency holds.[7]

[4] The SPF also requests probability distributions for output growth but here we restrict the analysis to the inflation density forecasts. See Croushore [1993] for a description of the SPF data.

[5] Except for the first and the last intervals, which are one-side open (i.e. α_0=-∞ and α_I=+∞), all central intervals in each survey round have equal length.

[6] The larger the number of observations drawn, the closer to the integrals the reported frequencies would be. In any case, no assumption about the number of draws is needed for our methodology.

[7] See, for instance, Gouriéroux and Montfort [1995] for a general development on sufficiency for discrete distributions and information theoretical implications.

Interpreting the estimation of the underlying density forecast in the context of the (parametric) multinomial framework offers some advantages. Our framework makes explicit that the link between the unknown density function $f_{\pi k}$ and the reported probabilities is stochastic, which, for example, helps explain the presence of rounded probabilities. More importantly, the multinomial framework lends formal structure to the estimation problem, and that structure provides some additional insights that have been overlooked so far.

2.1 The choice of fitting criterion

We base our choice of fitting criterion on its properties to handle the peculiarities of the SPF data. The inference problem is to find the parameters of the unknown density function to match the reported frequencies of the SPF histograms. Although long-sample properties (consistency, asymptotic normality and asymptotic efficiency) are desirable[8], the robustness of the estimator in small samples is crucial given the particular features of the SPF data.

Recent work estimating parametric densities from the SPF histograms (Giordani and Söderlind [2003], Rich and Tracy [2006], Engelberg *et al.*[2007], D'Amico and Orphanides [2006]) uses least squares as fitting criterion, i.e. minimizing the sum of the squared deviations between the theoretical and the observed probabilities over the set of intervals. Formally, $LS = \sum_{i=1}^{I} \left(\widehat{p}_i - p_i(\varrho)\right)^2$.

For the SPF data, the LS criterion (although consistent) is however not efficient in the standard sense that not all available information is used in the estimation. The LS criterion assigns equal weight to the fitting errors for each interval. An efficient criterion would instead assign different weights to the fitting errors depending on the probability assigned to each interval, thereby exploiting the bell shape structure of the SPF histograms to improve the estimation. Appendix A elaborates on the shortcomings of (unweighted) least squares in the context of the SPF data.

Maximum likelihood estimation is a natural alternative approach in the context of the multinomial framework. The sufficiency of reported probabilities $\{\widehat{p}_i\}$ implies that maximum likelihood would deliver a consistent, asymptotically normal and efficient estimate of the density parameters. Specifically the optimal ϱ maximizes:[9]

[8] Consistency, in the sense that the estimator $\widehat{\varrho}$ converges to the true parameter value in probability as the number of observations (draws from the true density forecast in our multinomial framework) $n \to \infty$. Asymptotic efficiency, in the standard sense that no other estimator has smaller variance, as $n \to \infty$.

[9] The likelihood is $\frac{n!}{\prod_{i=0}^{I}(n\widehat{p}_i)!} \prod_{i=0}^{I} p_i^{n\widehat{p}_i}$.

$$\log \text{lik}(\widehat{p}_0, \ldots, \widehat{p}_I, \varrho) \ \propto \ \sum_{i=1}^{I} \widehat{p}_i \log p_i(\varrho) \tag{1}$$

or, equivalently, minimizes the log-likelihood ratio:[10]

$$G^2 = \sum_{i=1}^{I} \widehat{p}_i \log \widehat{p}_i - \log \text{lik}(\widehat{p}_0, \ldots, \widehat{p}_I, \varrho) \ = \ \sum_{i=1}^{I} \widehat{p}_i \log(\widehat{p}_i / p_i(\varrho)) \tag{2}$$

Applying maximum likelihood estimation to the SPF data is however problematic, because the low probabilities assigned to some intervals are likely to trigger numerical problems. In addition, the optimal estimator for SPF data must be robust to potential model specification (i.e. possible misspecification of the theoretical density function) and to the likely "loose" reporting of the discretization (sloppy handwriting, rounding practices, *etc*).

To improve the robustness of the estimator in the context of multinomial distributions, Cressie and Read [1988] consider departures from maximum likelihood estimation within the family of "power divergence estimators" (PDE henceforth). Indexed by the parameter $\tau \in R$, the family is defined as the estimators obtained by minimizing the following expression with respect to ϱ:[11]

$$I^{\tau}(\widehat{p}, p) \ = \ \frac{1}{\tau(\tau+1)} \sum_{i=0}^{I} \widehat{p}_i \left[\left(\frac{\widehat{p}_i}{p_i(\varrho)} \right)^{\tau} - 1 \right] \tag{3}$$

To see that maximum likelihood can be interpreted as a particular case of this family of criteria consider the limiting case of $\tau = 0$ in equation (3)

$$I^0 \equiv \lim_{\tau \to 0} I^{\tau} \ = \ \sum_{i=0}^{I} \left[\widehat{p}_i \log \frac{\widehat{p}_i}{p_i(\varrho)} + (\widehat{p}_i - p_i(\varrho)) \right]$$

which shows that minimizing I^0 with respect to ϱ is equivalent to minimizing the log-likelihood G^2(see equation (2)) above.[12]

Under Birch's [1964] regularity conditions, all power divergence estimators have the optimal large sample properties of maximum likelihood (Cressie and Read [1984]), and are all first-order efficient (see Lindsay [1994]). More robust power distance estimators however underperform with respect to maximum likelihood estimation in terms of efficiency in small samples (in our context small number of draws). We therefore choose our fitting criterion (i.e. the optimal τ) within that family of estimators taking into account the

[10] The likelihood ratio G^2 is obtained by changing sign and adding the quantity (non dependent on ϱ) $\sum_{i=0}^{I} \widehat{p}_i \log \widehat{p}_i$. Here we follow the notation of Cressie and Read [1988].

[11] The family of power distance estimators encompasses many widely-used estimators such as the chi-square criteria of Pearson and Neyman, the Hellinger distance, and the Kullback-Leibler divergence. See Appendix A for further details.

[12] Note that $\sum_{i=0}^{I} \widehat{p}_i - p_i(\varrho) = 0, \ \forall \, \varrho \in \Theta$.

small sample properties of the power divergence estimators and the characteristics of the SPF data. An inspection of the SPF data suggests that (numerical) robustness to *inliers* (i.e. intervals with much lower observed probability than the theoretical density suggests, for example related to rounding) is fundamental: about 20% of the individual histograms have some probabilities in two bins or less. In the presence of inliers, Lindsay [1994]) suggests a positive, but relatively low, value of the parameter τ, and Cressie and Read [1988] recommend $\tau = 2/3$. In Section 3.1 below we provide some Monte Carlo simulations that confirm that a positive but rather low τ is optimal for the SPF data.

2.2 The choice of theoretical density function

Recent work on the SPF data has focused on the first two moments of inflation expectations and neglected forecasts skewness. The asymmetry in the risks surrounding the central expectation is however a crucial feature of any forecast. In the Bank of Sweden's or the Bank of England's "fan charts",[13] skewness plays a prominent role, and central bank official statements often include references to the "assessment of risks". Moreover, from a purely technical point of view, if the SPF probability forecasts are skewed, neglecting asymmetries in the theoretical density used to fit the histograms is likely to bias estimates of the first two moments of the distribution.

Our goal is to extract as much information as possible from the SPF probability forecasts. In principle, kurtosis could also be of interest, but the SPF histograms provide a limited amount of information about the density forecasts. The need to keep the exercise feasible leads us to focus on the first three moments of the distribution and consider a three-parameter density family.

We depart from existing literature by using a potentially-skewed density function: the skew-normal distribution (see Azzalini [1985]). The skew-normal provides a one-to-one mapping between its three parameters and the mean, variance and skewness.[14] We can therefore "*let the data speak*" about potential asymmetries present in the SPF data without restricting our estimates of the other moments of the distribution. In addition, the skew-normal density always remains unimodal, which is a reasonable premise when dealing with macroeconomic expectations.

The skew-normal class is built by shifting and re-scaling a standard distribution with a density function defined as

$$f_\lambda(z) := 2\varphi(z)\Phi(\lambda z) \qquad z \in \mathbb{R}$$

[13] See Blix and Sellin [1998] and Britton *et al.* [1998].

[14] Other three-parameter distributions (i.e. the two-piece normal) do not provide such a direct mapping. In more parsimonious but potentially skewed distributions, like the beta family, two parameters impose undesired constraints on the estimation of the three moments of interest.

where φ and Φ are respectively the standard normal density and distribution functions, and $\lambda \in \mathbb{R}$ is the *shape* parameter, which determines the skew of the distribution.[15] In particular, the skew-normal distribution nests the standard normal as a particular case (λ=0). In what follows, we use $SN(\lambda)$ for the skew-normal distribution just defined. A general random variable Y is said to be skew-normal distributed when it can be written as

$$Y = \mu + \sigma \left(\frac{Z - E[Z]}{\sqrt{V(Z)}} \right) \qquad Z \sim \mathcal{S}N(\lambda)$$

The first three central moments of Y are then expressed as[16]

$$\begin{aligned} E[Y] &= \mu \\ V(Y) &= \sigma^2 \\ SK(Y) &= \gamma_1 = \left(2b^2 - 1\right) b\delta^3 / \left(1 - b^2\delta^2\right)^{3/2} \\ &\quad \text{where} \quad b = \sqrt{2/\pi} \quad \text{and} \quad \delta = \lambda / \sqrt{\left(1 + \lambda^2\right)}. \end{aligned}$$

The skewness coefficient γ_1 is bounded in the interval $(-0.995, 0.995)$, and is a monotonous function of the shape parameter λ, sharing its sign. Figure 1 illustrates the flexibility of this distribution by depicting four skew-normal densities with zero mean, unit variance and four different values of the shape parameter λ.

3 Assessing our methodology

We choose our fitting criterion (i.e. the optimal τ^*) based on the performance of power divergence estimators for SPF-like data. This section also provides some quantitative evidence on the accuracy gains in the estimation of the first two moments of the distribution and illustrates the advantages of employing a potentially skewed distribution to fit the SPF histograms with an actual example.

3.1 Accuracy gains

The Monte Carlo simulations are designed as follows. Pseudo-random draws are taken from different underlying densities that range from a skew-normal (SN henceforth) with a high degree of asymmetry ($\gamma_1 = 0.9$) to the normal distribution ($\gamma_1 = 0$). SPF-like histograms are formed on the basis of the relative frequencies of the draws falling over a grid of

[15] The sign of λ determines the direction of asymmetry. Moreover, if Z is a $\mathcal{S}N(\lambda)$-distributed random variable, then $-Z$ is a $\mathcal{S}N(-\lambda)$ random variable (see Azzalini [1985] for the properties of the distribution).

[16] This parametrization avoids inference problems stemming from singularities in the Fisher information matrix (Azzalini [1985] and Pewsey [2000]).

intervals similar to the actual survey questionnaire.[17] The parameters of the underlying densities are estimated using different theoretical densities and fitting criteria (Normal *(N)* and Skew-Normal *(SN)* densities, and least squares *(LS)* and our proposed fitting criterion $PDE(\tau^*)$) in order to assess the accuracy gains stemming from the individual or combined use of our two methodological contributions.[18]

The optimal τ should provide the lowest estimation errors for all three density parameters (μ, σ, γ_1), but in practice such optimality is unlikely. Estimators within the $\tau \in (-0.5, 1)$ range show a relatively similar estimation performance.[19] The results also confirm that a low but positive value of τ helps cope with the inliers in SPF data, so we fix $\tau^* = 0.2$. Our approach is therefore fully specified as *PDE(*τ=0.2), *SN.*

Monte Carlo simulations confirm accuracy gains from our approach. Figures 2 to 5 illustrate the improvements in the estimation of the mean(μ) and standard deviation (σ) in terms of the percentage decline in the mean square error (MSE) and mean absolute error (MAE) using 50-draw samples and four different degrees of asymmetry in the distributions. Appendix C provides a more comprehensive sensitivity analysis.

Figure 2 shows that for the mean and for the standard deviation the decline in MSE is of at least 20% in all cases. Using the skew-normal rather than the normal distribution also improves the estimation of the key moments of the distribution (see Figure 3). Gains increase with the asymmetry in the data but the *SN* also performs well in the case of symmetry in the true data generating process (DGP), despite being overparametrized. The flexibility of the SN therefore appears to overcome the shortcoming of overparametrization.

Accuracy gains are maximized by combining the $PDE(\tau^*)$ fitting criterion and *SN* as underlying distribution: Figures 4 and 5 show the accuracy gains from using our criterion rather than least squares fitting assuming a normal distribution in terms of MSE and MAE. For both the mean and the standard deviation the improvement in terms of the MSE is about 40%. In terms of the MAE, the improvement is also substantial (above 20% in all cases).

The Monte Carlo results also confirm the reliability of our methodology for the estimation of the skewness parameter. Mean square errors are small (below 0.5 even for very small samples), and decrease fast with sample size (see Table 1). In most cases the estimation bias is also small (below 0.1), and only rises for the extreme hypothetical case

[17] The intervals are $(-\infty, 0)$,$[0, 1)$,$[1, 2) \ldots$,$[7, 8)$,$[8, +\infty)$.To obtain draws from the SN distribution we follow Henze [1986].

[18] To compute the theoretical probabilities for each interval, $p_i(\mu, \sigma, \gamma) = F_Y(\alpha_i) - F_Y(\alpha_{i-1})$, we use that $F_Y(\alpha) = F_{Z,\lambda}\left(\left(1 - (b\delta)^2\right)^{1/2}\left(\frac{\alpha-\mu}{\sigma}\right) + b\delta\right)$, and $F_{Z,\lambda}(z) = \Phi(z) - 2T(z, \lambda)$, where $T(z, \lambda)$ is the so-called Owen function, for which precise approximations exist (see Azzalini [1985] and Patefield and Tandy [2000]).

[19] Appendix C provides results for different sample sizes. The qualitative conclusions are robust to alternative values of the mean and the standard deviation of the true distributions.

of a rather small sample size (n=20) and highly skewed distribution (γ=0.9). The biases decrease with higher skewness, since errors on each side are bounded (see section 2.2 on the admissible range of γ), leaving less scope for error on one side of the true parameter. Pinning down skewness however appears to be more challenging for highly skewed distributions (γ=0.9): the bias is negligible for symmetric (γ=0) distributions even in rather small samples, but tends to rise with the asymmetry in the true data. For highly skewed densities, the estimates tend to be pushed towards the upper and lower bounds for the skewness index, which leads to a higher bias for extreme skewness. Such a extreme skew in macroeconomic forecasts, in particular in inflation forecasts, is however unlikely. Therefore, the impact of that estimation bias in our qualitative assessment of the skewness embodied in the SPF inflation forecasts appears minimal.

In general, the first two moments are easier to pin down than the skewness parameter. Despite having a higher (finite sample) MSE than the lower moments, the skewness estimates asymptotically converge to the true values at a linear rate, in line with existing results (Pewsey [2000]).

3.2 Fitting densities to the SPF histograms: practical considerations

We here illustrate the differences of fitting a normal and a skew-normal density to an actual SPF histogram. The histogram has probability assigned only to three intervals, and the distribution of those probabilities is quite asymmetric. Our fitting criterion uses the information in those three intervals and also considers the surrounding empty bins (inliers) to better identify the shape of the density. In addition our approach exploits the flexibility of the skew-normal density to handle skewed forecasts.

Figure 6 depicts the normal and the skew-normal estimation for an actual SPF histogram with asymmetric features using our fitting criterion. This example also illustrates that the increased flexibility stemming from the endogenous skew of our theoretical distribution is very useful to extract reliable estimates of the first two moments of the probability forecast.

The identification of the three parameters of the skew-normal distribution in the case of two or less bins containing probability can be understood by reference to the case for a normal distribution. The parameters of a normal distribution are identified for histograms with three or more active intervals. Heuristically, an additional interval with positive probability should then suffice to identify the additional parameter implied by the skew-normal distribution. Even if probability is only assigned to two intervals, the survey response in fact conveys information for the probability contained in four intervals, which guarantees identification for the skew-normal. The remaining case, i.e., when all the probability mass is assigned just to one bin – that is information about three intervals

– is paradoxically trivial: the assumption of uniform distributions within intervals for the computation of initial parameter values yields in this case a zero value estimate for the skewness parameter γ_1(as well as a mid-point estimate for the mean of the distribution), with only a (low) variance to be identified.[20] In practise, numerical results confirm the stability of the mean and skewness parameters at those initial values because with all the probability mass in one single interval, the SPF data provide no information to explore potential asymmetries in the distribution.

Our methodology, being robust to data inliers, can be therefore applied across the panel of replies. Previous research has often made additional assumptions for the estimation of the densities concentrated in a low number of intervals, thereby employing several different approaches across the panel of replies.[21]

4 Four decades of inflation risks: stylized facts

One of the purposes of the SPF requesting probability forecasts is to gauge information about the risks to the point forecast. We use almost four decades, between 1969Q1 and 2006Q1, of probability forecasts to calculate their uncertainty (i.e. the variance) and their risk assessment (i.e. skewness). For each survey round, we estimate the key moments of both the individual density forecasts and the *combined* (or aggregate) distribution calculated as a simple average of responses across forecasters.

The SPF requests probability forecasts for inflation over the current (since 1981Q3 also the next) calendar year for four consecutive quarters.[22] With each survey round the forecast horizon therefore decreases. To assess how forecasts have evolved over time, we use the current-year probability forecasts from the first-quarter surveys of each year, and focus perceived inflation risks over a constant horizon of four quarters ahead. As a robustness check for our correlations and regressions, we use together current and next calendar year forecasts to construct a quarterly series of forecasts about four quarters ahead. Specifically we take the current-year forecast in the Q1 (i.e. about four quarters ahead) and in the Q2 survey rounds (i.e. about three quarters ahead) with the next year

[20] This case only arises in less than 5% of replies for the current year inflation forecasts.

[21] For instance, Engelberg *et al.* [2007] employ a triangular distribution for those histograms with less than three active bins and a beta for the rest. Rich and Tracy [2006] propose using information from the point estimates reported by the forecasters in the framework advocated by Wallis [2005], while D'Amico and Orphanides [2006] propose direct fitting of the cumulative distribution of inflation uncertainty measures.

[22] There were occasional mistakes in the horizons in certain surveys. In 1968Q4, 1969Q4, 1970Q4, 1971Q4, 1972Q3 & Q4, 1973Q4, 1975Q4, 1976Q4, 1977Q4, 1978Q4, and 1979Q2, Q3 and Q4 the probability variables referred to the following year, rather than the current year. The opposite situation appeared to occur in the 1985Q1 and 1986Q1 surveys, but for those two survey rounds four quarters ahead forecasts are still available in the dataset. See http://www.philadelphiafed.org/econ/spf/index.html.

forecasts in the Q3 (i.e. about six quarters ahead) and in the Q4 survey rounds (i.e. about five quarters ahead).

Our analysis of the SPF inflation risks provides some stylized facts. We first assess the dynamics of inflation uncertainty and inflation skew. Next we check the link between (mean) expected inflation and perceived inflation risks. Finally, we investigate the explanatory power of key macroeconomic variables on our measures inflation risks.

4.1 Measuring inflation uncertainty

Forecast uncertainty is directly linked to the probabilities that a forecaster assigns to the possible values of the predicted variable: the tighter the distribution, the lower the uncertainty. The variance (or standard deviation) of a forecasts is a direct measure of the uncertainty surrounding the forecast. The estimated individual variances show significant variations across forecasters that use the same number of intervals, reflecting large discrepancies in probability mass assigned to those intervals (see Figure 7).

Measuring uncertainty for the panel of forecasters as a whole is however more problematic. Aggregate measures of uncertainty can be estimated from the combined SPF probability distribution, or can be obtained by averaging the variances of the individual forecasts. Moreover, given that many surveys only request a single point prediction, a long stream of literature focuses on the dispersion of point forecasts –the so-called disagreement– as alternative indicator.[23] The variance of the combined distribution, by construction, reflects both the average of the uncertainty surrounding the individual forecasts and the dispersion of the individual mean forecasts. Formally:

$$Var(\pi_c) = \frac{1}{K}\sum_{k=1}^{K} Var(\pi_k) + \frac{1}{K}\sum_{k=1}^{K}[E(\pi_k) - E(\pi_c)]^2 \qquad (4)$$

Equation (4) holds in population terms independently of distributional assumptions. The left-hand side can be estimated from the combined probability forecast and the right-hand side can be calculated from the estimated individual density functions.

Over the last 40 years inflation uncertainty has exhibited significant fluctuations (see Figure 8). It rose sharply following the two oil price crises in the 1970s, and during the 1980s there were also some other spikes before a significant decline since early 1990s.The three measures of uncertainty from equation (4) experienced those fluctuations. The decomposition of the variance of the combined distribution depicted in Figure 8, however, shows that average uncertainty across panelists is the main component; disagreement plays a minor role. The correlation with the variance of the combined distribution is also higher for average uncertainty than for disagreement (see Table 2), a result that appears

[23] For a critical assessment see Rich and Butler [1998].

robust to data frequency and period of study. The decomposition depicted in Figure 8 however makes evident that the independent estimation of the three terms in equation (4), although accurate in general, is occasionally subject to errors, which call for caution when interpreting changes in one single measure of inflation uncertainty.

4.2 Asymmetries in inflation risks

The assessment of risks surrounding the baseline scenario is an important characteristic of any forecast. One of the novelties of our methodology is the estimation of the skewness of the SPF probability forecasts. Evidence on the changes in the balance of inflation risks can help better understand a number of macroeconomic phenomena. Moreover, our findings provide two key features of the assessment of inflation risks that theories of expectations formation should match.

First, we find that the SPF probability forecasts show strong asymmetries. More than 60% of the SPF individual probability forecasts have skewness index γ_1 higher than 0.3 (in absolute value). Moreover, more than 60% of the combined probability forecasts also do so (see Figures 9 and 10). These results make accounting for asymmetry a must to obtain reliable estimates of the first two moments of the SPF probability forecasts, and therefore render support to our methodological approach.

A second key finding about those asymmetries is their sign: U.S. inflation forecasts are positively skewed; only in a few occasions over the last four decades downside risks predominated (see Figure 11). The difference between the mean and the mode of the combined probability forecasts, a simple measure of asymmetry, also corroborates those stylized facts.

In terms of their evolution over time, the balance of risks surrounding inflation forecasts has fluctuated significantly in the last four decades. Such fluctuations are however more difficult to interpret and identify with specific historical episodes than the ones of inflation uncertainty. Notwithstanding this, the strong upside risks to inflation throughout the 1980s, particularly in the mid-1980s, stand out. We provide an interpretation for those fluctuations in the next section, with the help of additional historical evidence.

The skewness of the combined probability forecast is also linked to the moments of the individual densities as follows (see Appendix B for details):[24]

$$SK(\pi_c) := E\left[(\pi_c - E(\pi_c))^3\right] = E\left[SK\left(\pi_k\right)\right] + SK[E(\pi_k)] + 3E[Var\left(\pi_k\right)\left(E(\pi_k) - E(\pi_c)\right)]$$

[24] For the sake of exposition, in this subsection we will use the term *skewness* $SK(\pi)$ to denote the non-normalized central third moment, instead of the standard skewness coefficient used in the previous sections.

The first and second components are the average individual skewness and the skewness of the distribution of individual means, in a similar fashion to the variance decomposition. The skew decomposition however has an additional component: the covariance between the distance of the individual forecaster's mean to the combined mean in that survey round and the variance of that forecaster. Intuitively, this third component reflects the impact that the different degrees of uncertainty among forecasters can have via its correlation with the disagreement with respect to the combined forecast. For instance, the term would vanish if all forecasters had the same degree of uncertainty.

In contrast to the variance decomposition, the average skew of the individual forecasts does not appear to be a reliable proxy for the asymmetries in the combined probability forecast. The correlation between the skewness of the combined distribution and its three components is relatively limited and suggests that inferring the sources of the risk assessment embodied in the combined distribution is quite challenging (see Table 3).

4.3 Inflation risks and the central tendency of inflation expectations

We here investigate how changes in (mean) inflation expectations affect inflation risks (uncertainty and the risks assessment), thereby assessing the evolution of inflation density forecasts over the last four decades. Figure 12 depicts mean inflation expectations together with inflation uncertainty and the degree of skewness. The relation between inflation uncertainty and the level of expected inflation is positive, while perceived skewness and inflation exhibit a weak negative relationship.

As regards the volatility of the first three moments of inflation expectations between 1969 and 2006, mean expectations have been the most volatile (see variances in the principal diagonal in Table 4),[25] and inflation uncertainty have been much less volatile than the perceived asymmetry in inflation risks. In terms of co-movement, movements in mean inflation expectations appear to be more related to changes in inflation uncertainty than to changes in the perceived asymmetry in inflation risks. Uncertainty and skewness in inflation forecasts have however shown some co-movement, possibly reflecting the upward movement both higher-order movements experienced in the 1980s and the decline thereafter.

To assess the statistical significance of these correlations, Table 5 reports some bivariate linear regressions. As robustness checks, we use the quarterly time series introduced above for the sample 1981Q1-2006Q1, as well as an additional central tendency proxy, the average of individual means. The level of expected inflation significantly contributes to explaining

[25] Results using the average and the median of the means of the individual probability forecasts, or the quarterly time series that combines current and next year forecasts were qualitatively similar. They are therefore omitted here but available upon request.

movements in inflation uncertainty, which confirms the positive relationship between the level of a variable and the uncertainty surrounding its forecast. Such a relationship, is however limited (the $\overline{R}^2$ in the bivariate regressions are about 0.4), which indicates that a significant part of the variation in the uncertainty surrounding inflation expectations is explained by other factors. Asymmetries in inflation risks appear to be more related to measures of inflation uncertainty, particularly average uncertainty, than to the level of inflation expectations.

The key message from these results is that, although the first three moments of inflation expectations do co-move, such a co-movement is limited. To characterize changes in inflation expectations the mean, or any central tendency measure, is therefore not sufficient.

5 Inflation scares and inflation risks in the 1980s

A thorough analysis of inflation expectations must comprise not only point predictions but also perceived inflation risks. Previous sections show that our measures of inflation risks appeared to be much higher in the 1980s that the level of expected (or actual) inflation suggests. In the 1980s the U.S. bond market also experienced some turbulent episodes that have been interpreted as inflation scares.[26] Point forecasts offer little information about why bond premia rose. We exploit the link between perceived inflation risks and the premia embodied in bond yields, and show that a thorough assessment of inflation expectations sheds new light on the bond market inflation scares in the 1980s.

Under Chairman Volcker, the U.S. Federal Reserve pursued a protracted disinflationary policy during the 1980s. The Volcker disinflation is a crucial episode in the history of U.S. monetary policy, and remains the subject of much research and discussion. Goodfriend and King [2005] document that the FOMC saw credibility as fundamental for the success of the disinflation and monitored long-term interest rates as indicators of that credibility.[27]

It is surprising that a thorough analysis of perceived inflation risks during the Volcker disinflation remains missing. This paper aims at filling that void. Although long-term interest rates are often used as indicators of inflation expectations, changes in nominal yields may reflect other factors beyond inflation expectations. The SPF forecasts are a more direct source of information about private sector's inflation expectations, and offer information about both point predictions and inflation risks.

Our analysis provides an assessment of the evolving credibility of the Fed´s disinflation-

[26] See Goodfriend [1993], Gürkaynak *et al.* [2005] and Orphanides and Williams [2005].

[27] Ball [1994] established the need for imperfect credibility of a disinflation to cause output losses even in the presence of nominal rigidities. See also Ball [1995].

ary policy. Modeling the interaction between the Fed´s policy and private sector beliefs about inflation is however beyond the scope of this paper. Nor do we provide a thorough historical account of all the events surrounding monetary policy decisions (see Goodfriend [1993] and Goodfriend and King [2005] are excellent contributions). We combine information from both surveys and long-term interest rates, and interpret the inflation scares in the bond market as particular episodes in the process of gaining credibility.[28] The changes in inflation risks around the time of the bond market inflation scares suggests a close link between the two, and a taxonomy of those inflation scares on the basis of direct evidence on private sector inflation beliefs emerges as a by-product of our interpretative historical analysis.

Figure 13 shows the policy and long-term interest rates, actual and expected inflation and our measures of inflation risks over the 1980s. The inflation scares are represented by shadowed areas (see top chart). Figure 13 (middle chart) depicts survey measures of actual and expected inflation, both over short-term (four quarters ahead) and long-term horizons. Mean expected inflation rose sharply at the end of the 1970s and stood above realized inflation during most of the 1980s, which lends support to the thesis of a very slow improvement in Fed's credibility. Short and long-term inflation expectations however declined in the first half of the 1980s and remained broadly unchanged for most of the second half. This dynamics suggests that the inflation scares in the 1980s need to be explained by other features of inflation expectations. Inflation uncertainty and the inflation skewness exhibited significant fluctuations in the 1980s, and rose sharply around the time of the inflation scares in the bond market. We analyze their movements in each of the inflation scare episodes.

Early stages of the disinflation: the 1979-80 and the 1981 inflation scares

In the summer of 1979 oil price developments had greatly worsened the inflation outlook. The Fed switched to non-borrowed reserve targeting and raised rates at its October meeting to fight the rise in inflation, but the tightening stopped at the end of the year with the federal funds rate around 13.5% and the ten-year rate at 10.5% (see Figure 13). By February 1980 the long-term rate however edged up above 13%, triggering further interest rate hikes to control the rise in inflation expectations and related premia. Higher policy rates helped contain inflation but also triggered a recession that halted the Fed tightening. By late 1980, the ten-year bond rate had reversed the 2 percentage point decline associated with the recession and was again above 12%. Moreover, despite having the federal funds rate at 19% in 1981, the Fed faced a second inflation scare: long-term interest rates rose

[28] It has to be borne in mind that matching movements of financial and survey data is quite difficult given their different frequency.

by 3 percentage points between January and October 1981 to a peak of 15%.

Inflation expectations did rise significantly over that period, thereby explaining the inflation scares of 1979-80 and 1991 (see Figure 13). (Mean) expected inflation steadily rose between 1979 and 1981, reaching two digits in 1981 despite the decline in realized inflation since 1980. Inflation uncertainty rose sharply to a historical peak in 1980, to moderate somewhat later on. The variance decomposition shown in Figure 8 suggests that rising inflation uncertainty was accompanied by a rise in both disagreement about the inflation outlook and the individual uncertainty surrounding that outlook. Inflation risks were also tilted towards higher inflation outcomes, as indicated by the upward shift in the skewness of the combined probability forecast.

The 1983-84 inflation scare

Two years after the start of the disinflationary policy in 1980, inflation was coming gradually down and the Fed somewhat relaxed its policy stance. By February 1983, the funds rate was around 8.5% and the ten-year rate remained steady at around 10.5%. A few months later, however, with economic activity gaining momentum, the Volcker Fed was confronted with another inflation scare. Long-term rates started to rise again, although annualized (quarterly) inflation was below 5% throughout most of 1983 and 1984, that is about 4 or 5 percentage points lower than in 1981. Long-term rates reached 13.5% in mid-1984, only about a percentage point lower than their October 1981 peak. The Fed raised rates to resist the on going inflation scare. Long-term rates, however, did not start falling until the second half of 1984, and did so very gradually, returning below the 10% level only in the last months of 1985. The containment of the 1983-84 inflation scare was arguably one of the most remarkable achievements of the Volcker Fed (Goodfriend [2005]).

Our analysis reveals that the 1983-84 inflation scare was significantly different from the two previous scares in several crucial aspects. Movements in inflation expectations around the time of the 1983-84 bond market inflation scare were somewhat surprising. Inflation expectations (as measured by the mean of the combined distribution) actually *declined* during most of the 1980s and particularly between the first quarters of 1983 and 1986 (by more that 1.5 percentage points, see Figure 13). However, despite the decline in the (mean) inflation expectations, inflation risks edged strongly upwards since early 1983 and remained high over most of the period of the bond market scare. The rise in inflation uncertainty and the risk assessment suggests that private agents believed that the reduction in inflation was temporary, and a return to high inflation likely. Therefore, in sharp contrast to the two previous inflation scares, the then imperfect credibility of the Fed was reflected not in the usually considered point estimates of inflation expectations but in the perceived inflation risks (uncertainty and in particular asymmetry).

The 1987 inflation scare

In 1987, the Volcker Fed faced a fourth inflation scare. With inflation rising for the first time in the decade, this inflation scare was marked by a sharp but relatively brief, rise in long-term bond yields by 2 percentage points between March and October 1987. (Mean) inflation expectations reacted little to the rise in actual inflation rates, allowing for the gap between the two to close by early 1988. Longer-term inflation expectations appeared to move only slightly upwards around 1987. At the same time, perceived inflation risks rose again from early 1986, and only declined significantly once the scare was contained.

Assessment

The analysis of central tendency and inflation risk measures during the inflation scares of the 1980s illustrates that changes in inflation expectations may be quite complex. Moreover, to correctly interpret movements in inflation expectations, the analysis of higher-order moments of inflation forecasts is necessary. Our discussion of the four inflation scares confronted by the Volcker Fed in the 1980s shows that such episodes were triggered by distinct changes in inflation expectations. The first two inflation scares reflected a rise in both central tendencies of and the risks surrounding inflation expectations. The information provided exclusively by the point estimates is however insufficient, if not misleading, to explain the 1984 scare; the rise in long-term rates around 1984 was mainly triggered by higher perceived inflation risks, most likely affecting risk premia. Although the relationship between the higher-order moments of inflation forecasts and the risk premia embodied in bond yields is likely to be rather complex, this section shows that a thorough analysis of probability forecasts can shed new light in this regard.

6 Inflation risks and the macroeconomy

A reason for monitoring inflation expectations is that, besides the inflation outlook, they comprise information about the general macroeconomic outlook. We here show that policymakers and researchers can obtain valuable information about the general macroeconomic situation by monitoring inflation risks.

We now ask how macroeconomic variables influence inflation risks. Inflation uncertainty and actual inflation display the expected positive relationship (see Figure 14). Given the observed inflation rate, a noticeable feature of Figure 14 is that, both for a brief period (around 1980 following the second oil crisis) and over a more protracted period in the second half of the 1980s, inflation uncertainty was abnormally high. Such outliers correspond to the spikes in inflation uncertainty shown in Figure 8 and analyzed in the

previous section.

The relation between perceived inflation skewness and actual inflation, as for expected inflation, instead appears to be weak (see Figure 14). Actual inflation does not shed any light on two puzzling features of our inflation skewness measures (see also Section 4.2 and Figure 11). The SPF density forecasts have only exhibited downside risks on a few occasions: with high inflation levels (for example in 1974-75 and in 1980) but also at more moderate rates of inflation (for example in 1990-91), which suggests that negative skew in inflation forecasts may be capturing other factors beyond the level of actual or expected inflation.

The assessment of inflation risks appears to be more related to economic activity than to the level of inflation. Figure 15 shows our measures of inflation risks together with two standard measures of excess capacity: (i) an HP-filtered output gap measure for the (log of) actual chain-weighted real output; (ii) episodes of economic contraction according to the NBER Business Cycle Dating Committee, represented as shaded areas.

Ahead of a recession, inflation forecasts tend to exhibit negative skewness. This tendency indicates a stronger business cycle response of inflation expectations than considered so far: point forecasts of inflation tend to decline with lower expected growth but business cycle conditions also influence the assessment of inflation risks. Moreover, in the three available episodes of contracting real GDP in the sample, inflation uncertainty also declined ahead of the recession. During recession periods uncertainty about future inflation however appears to rise. Outside recession periods, over the 1980s a widening of the output gap seemed to have a positive impact on both inflation uncertainty and upside risks, but the protracted expansion of the late 1990s appeared to have little impact on perceived inflation risks. The business cycle influence on inflation risks therefore seems to have significantly weakened.

To investigate more formally the link between inflation risks and macroeconomic variables, we regress the variance of the combined distribution and its two components, namely average uncertainty and disagreement, and the two measures of risk assessment, skewness and the distance between the mean and the mode of the combined distribution, on several macroeconomic variables . Specifically, together with the level of inflation, as nominal variables we consider the (squared) change of inflation as a proxy for the volatility of inflation, and as real economy variables we consider the output gap measure used above, and the probability of recession and the average uncertainty about real GDP growth (four quarters ahead) as reported by SPF panelists (see Table 6). For each of those variables Panel A reports the coefficients from bivariate regressions. To assess the robustness of those estimates, Panels B and C respectively control for actual inflation and (mean) expected inflation.

The squared change of inflation appears to be a robust determinant of average uncertainty, even after controlling for actual and expected inflation. This result lends some support to the link between uncertainty and inflation volatility. Such a link is however less clear for the variance of the combined distribution, and even less so for disagreement, particularly after controlling for inflation or inflation expectations the relationship weakens considerably or becomes insignificant. The volatility of inflation seems to have no effect in the asymmetry of the perceived inflation risks either.

As regards economic activity, inflation risks show no relationship with the output gap measure. Inflation risks instead respond to uncertainty about growth and the probability of recession over the horizon of the inflation forecast, which underscores the forward-looking nature of the SPF inflation risks.

7 Concluding remarks

This paper provides stylized facts about perceived inflation risks. We analyze inflation uncertainty and inflation skewness in forty years of probability forecasts from the Survey of Professional Forecasters (SPF) using a new methodology. Our estimation comprises two methodological contributions. First, our fitting criterion is a small departure from maximum likelihood that provides the needed robustness to cope with the particular features of the SPF data. Second, the use of a parsimonious but flexible theoretical density function, Azzalini's [1985] skew-normal. Compared to standard approaches for the estimation of SPF density forecast, both methodological contributions lead to substantial accuracy gains.

The main result from our analysis is that moving beyond point estimates of inflation expectations is necessary. To provide a thorough account of private sector beliefs about the inflation outlook, inflation risks provide additional insights. As an example, we illustrate how the movements in our inflation risk measures in the period of the Volcker disinflation shed some new light on the inflation scares in the U.S. bond market in the 1980s.

Our results suggest that a thorough analysis of the SPF probability forecasts may provide additional insights not only about how inflation expectations may change over time but also about developments in other macroeconomic and financial variables. We hope that the methodology developed in this paper and our analysis of the dynamics of inflation beliefs in the Volcker disinflation period could be a first step in that process. Our findings point to business cycle analysis and the estimation of inflation risk premia in bond yields as interesting avenues to explore.

A. On the optimal fitting criterion for the SPF histograms

A.1. Shortcomings of (unweighted) least squares

Most literature fitting parametric densities to the SPF histograms uses least squares as fitting criterion, i.e. minimizing the sum of the squared deviations of the observed probabilities with respect to the theoretical ones over the set of intervals. More formally, $LS = \sum_{i=0}^{I} \left(\widehat{p}_i - p_i(\varrho)\right)^2$. Minimizing the LS criterion (although consistent) is not efficient, for it does not use all available information. Generalized least squares (GLS), a standard (efficient) approach, would suggest minimizing $GLS = \left(\widehat{\mathbf{p}} - \mathbf{p}(\varrho)\right)' S^{-1} \left(\widehat{\mathbf{p}} - \mathbf{p}(\varrho)\right)$ with respect to ϱ where S^{-1} denotes the inverse of the covariance matrix of the frequencies vector, $\widehat{\mathbf{p}}$. The multinomial framework outlined in the main text provides a direct identification of the structure of the matrix S without *ad-hoc* assumptions. Specifically, the structure of the covariance in a multinomial model is of the form

$$S = \begin{pmatrix} p_0(1-p_0) & -p_0p_1 & \cdots & -p_0p_I \\ -p_0p_1 & p_1(1-p_1) & \cdots & -p_1p_I \\ \vdots & & \ddots & \vdots \\ -p_0p_{I-1} & \cdots & p_{I-1}(1-p_{I-1}) & -p_1p_{I-1} \\ -p_0p_I & \cdots & \cdots & p_I(1-p_I) \end{pmatrix}$$

The matrix S is singular (as the sum of frequencies p_i must add up to one). Its Moore-Penrose inverse, S^-, has the following structure [29]

$$S^- = \begin{pmatrix} 1/p_0 & 0 & \cdots & 0 \\ 0 & \ddots & \ddots & \vdots \\ \vdots & \ddots & \ddots & 0 \\ 0 & \cdots & 0 & 1/p_I \end{pmatrix}$$

This evidence suggests using the weighted least squares (WLS) estimator. Minimizing WLS is algebraically equivalent to minimizing the Pearson chi-square criterion

$$X^2 = \sum_{i=0}^{I} \frac{\left(\widehat{p}_i - p_i(\varrho)\right)^2}{p_i(\varrho)} \tag{5}$$

The unknown vector $\mathbf{p}(\varrho)$ in the denominator could be replaced by its estimate, $\widehat{\mathbf{p}}$, to obtain the modified chi-square criterion of Neyman [1949].

The fundamental difference between LS and WLS is the weighting of the fitting errors. LS assigns equal weighting to the fitting errors for each interval. As shown in equation (5),

[29] Note that the structures of the S and the S^- matrices come directly from the interpretation of the reported histograms as the realizations of a multinomial variable.

the efficient criterion WLS instead weighs the fitting errors by the inverse of the observed (or fitted) frequency . If the frequencies were constant across the intervals i, both LS and WLS would provide the same results. Formally, in terms of the above expressions, for WLS (or NM^2) and LS to be equivalent, the matrix S^- should be proportional to the identity matrix, thereby providing equal weighting for all the intervals. In the SPF probability distributions the observed relative frequencies in each interval are far from similar, they are relatively bell-shaped. Assigning the same weight to the fitting errors at the center of the distribution and at the tails, as the LS criterion does, does not make efficient use of all the information contained in the SPF data.

Equal weighting is therefore a fundamental shortcoming of the commonly-used LS criterion. In contrast, the chi-square criteria derived above provides consistent, asymptotically normal and efficient estimates for the vector of distribution parameters. Those properties, which are crucial to recover the underlying density forecasts, are provided by all power divergence estimators.

A.2. Power divergence estimators

We here present in greater detail the main properties of the family of power divergence estimators, which encompasses the X^2 and NM^2 criteria discussed above. We focus on the properties that are more relevant for our purpose in this paper, but Cressie and Read [1988] provide a general assessment.

Recall that the family of minimum power divergence estimators is defined in general terms as the estimates obtained by minimizing the expression

$$I^{\tau}(\widehat{p},p) \;=\; \frac{1}{\tau(\tau+1)}\sum_{i=0}^{I}\widehat{p}_i\left[\left(\frac{\widehat{p}_i}{p_i(\varrho)}\right)^{\tau}-1\right] \tag{6}$$

with respect to ϱ. The family is thus indexed by the parameter $\tau \in \mathbb{R}$.

To grasp the broadness of this family, note that, if $\tau = -2$ for instance, equation (6) yields the Neyman chi-square criterion, while for $\tau = 1$ one obtains Pearson's X^2.[30]

Serfling [1980] shows that the estimate obtained from minimization of the likelihood ratio G^2 (see equation (2)) is asymptotically equivalent in probability to the estimate obtained by minimizing the standard Pearson chi-square criterion X^2 discussed above. We have however argued that small sample properties are also fundamental in the choice of the optimal estimator. The figures below show the relative performance of some of the estimators discussed above based on the Monte Carlo experiment detailed in Section 3.1 .

[30] Moreover, the family encompasses other widely-used estimators such as the Hellinger distance $\left(I^{-1/2} = 2\sum_{i=0}^{I}\left[\sqrt{\widehat{p}_i}-\sqrt{p_i(\varrho)}\right]^2\right)$ and the Kullback-Leibler divergence $\left(I^{-1} \equiv \lim_{\tau\to-1} I^{\tau} \;=\; \sum_{i=0}^{I}\left[p_i(\varrho)\log\frac{p_i(\varrho)}{\widehat{p}_i}+(\widehat{p}_i-p_i(\varrho))\right]\right)$.

[Figure AI about here]

B. Decomposing the skewness of the combined forecast

The third moment decomposition used in the main text is derived in a similar way to the standard ANOVA decomposition (with Y being a conditioning variable denoting the forecaster), as follows:

$$
\begin{aligned}
S(X) &= E\left[(X - E[X])^3\right] \\
&= E\left[E\left[(X - E[X])^3 |Y\right]\right] \qquad (7) \\
&= E\left[E\left[(X - E[X|Y] + E[X|Y] - E[X])^3 |Y\right]\right] \\
&= E\left[E\left[(X - E[X|Y])^3 |Y\right]\right] + 3\,E\left[(E[X|Y] - E[X])\,E\left[(X - E[X|Y])^2 |Y\right]\right] \\
&\quad + 3\,E\left[(E[X|Y] - E[X])^2\,E\left[(X - E[X|Y])\,|Y\right]\right] \\
&\quad + E\left[E\left[(E[X|Y] - E[X])^3 |Y\right]\right] \qquad (8) \\
&= E\left[S\left(X|Y\right)\right] + S\left(E[X|Y]\right) + 3\,E\left[V\left(X|Y\right)\left(E[X|Y] - E[X]\right)\right] \qquad (9)
\end{aligned}
$$

Step (7) is obtained applying iterated expectations. Step (8) follows from the expansion of the power and the measurability of $(E[X|Y] - E[X])$ with respect to the σ-algebra induced by Y. Finally, step (9) is obtained by observing that the third term vanishes, since $(E[X|Y] - E[X])^2$ can be taken out of the expectation (due to its measurability) and the remaining factor is equal to zero.

C. Monte Carlo evidence

This Appendix provides additional evidence on the accuracy gains in the estimation of the first two moments of inflation expectations from the methodology presented in this paper. Specifically, the results reported in the columns in each of the tables below correspond to the power distance estimator (PDE) using a skew-normal density (PDE(τ^*), SN)), the power distance estimator using a normal density (PDE(τ^*), N), the standard least squares criterion using the skew-normal density (LS (SN)) and the least squares criterion using the normal density (LS (N)). Hence, the first column represents our preferred methodology while the fourth represents the method most frequently used in the literature. The relative accuracy of the different methodologies is measured by computing mean square errors

(MSE), mean absolute errors (MAE), and the empirical bias obtained from 1000 simulated samples of the different sizes (i.e. the different number of random draws) in order to check the robustness of the proposed methodology. We highlight some key results below. In general they are consistent across samples sizes so we omit specific references to the sample size and focus on additional comparisons.

Our estimator PDE(τ^*=0.2) outperforms (unweighed) LS, thereby providing some quantitative support for our theoretical argumentation that both kinds of estimators are consistent but only power divergence estimators are efficient. Moreover, these gains appear to be also robust to the choice of model distribution to fit the data: either the *SN* (compare columns 1 and 3 for the mean and columns 5 and 7 for the variance; see also Figure 2) or the *N* distribution (compare columns 3 and 4 for the mean and columns 7 and 8 for the variance).

The gains from using the *SN* shown in Figure 3 were more limited than for the PDE(τ^*=0.2), but when a less efficient method, such as LS, is used, the relative gains from employing the *SN* are much larger (compare columns 3 and 4 for the mean and columns 7 and 8 for the variance). Another important results is that even in small samples (i.e. discretizations based on relatively small number of draws) the simulations suggest that the *SN* is capable of accommodating deviations from symmetry without compromising the estimation of the first two moments.

Table C.1. DGP: *SN* (mean μ=3 , Standard deviation σ=1 , Skew γ_1=0.6)

		Mean μ				Standard deviation σ			
Sample size		PDE, SN	PDE, N	LS, SN	LS, N	PDE , SN	PDE, N	LS, SN	LS, N
n=20	MSE*	5.32	5.53	7.19	9.84	3.92	4.09	4.64	6.50
	MAE	0.19	0.19	0.22	0.26	0.16	0.16	0.18	0.21
	bias*	1.47	-0.26	-4.41	-13.2	-3.20	-1.88	-5.87	-12.3
n=50	MSE*	2.16	2.26	2.59	4.91	1.38	1.65	2.05	2.84
	MAE	0.12	0.12	0.13	0.18	0.09	0.10	0.12	0.14
	bias*	1.26	0.83	-1.99	-13.5	-1.35	0.44	-2.71	-8.12
n=100	MSE*	1.06	1.06	1.40	3.85	0.70	0.80	1.08	1.56
	MAE	0.08	0.08	0.10	0.16	0.07	0.07	0.08	0.10
	bias*	0.99	0.76	-0.82	-15.0	-0.53	1.82	-1.98	-6.15
n=200	MSE*	0.54	0.62	0.72	3.00	0.45	0.81	0.52	1.02
	MAE	0.06	0.06	0.07	0.15	0.06	0.07	0.06	0.08
	bias*	2.23	1.66	0.61	-15.0	2.55	5.07	-0.23	-6.31

*Reported figures are scaled up by 100.

Table C.2. DGP: *SN* (mean μ=3 , Standard deviation σ=1 , Skew γ_1=0.3)

		Mean μ				Standard deviation σ			
Sample size		PDE, SN	PDE, N	LS, SN	LS, N	PDE, SN	PDE, N	LS, SN	LS, N
n=20	MSE*	5.74	5.46	6.21	8.11	3.00	3.61	4.57	5.62
	MAE	0.19	0.19	0.20	0.23	0.14	0.15	0.17	0.19
	bias*	1.29	-0.11	-1.46	-6.79	-3.08	-1.89	-3.45	-8.70
n=50	MSE*	2.07	2.33	2.90	3.79	1.40	1.39	1.86	2.26
	MAE	0.11	0.12	0.14	0.16	0.09	0.09	0.11	0.12
	bias*	0.49	-0.14	0.50	-6.81	-0.85	-0.13	-1.03	-4.98
n=100	MSE*	1.01	1.05	1.66	2.19	0.65	0.69	0.98	1.19
	MAE	0.08	0.08	0.10	0.12	0.07	0.07	0.08	0.09
	bias*	0.47	0.17	-0.36	-7.28	-0.70	0.26	0.16	-3.04
n=200	MSE*	0.61	0.56	0.78	1.41	0.47	0.57	0.48	0.60
	MAE	0.06	0.06	0.07	0.10	0.06	0.06	0.06	0.06
	bias*	0.95	1.26	0.61	-7.50	2.30	3.51	0.03	-2.43

*Reported figures are scaled up by 100.

Table C.3. DGP: *SN* (mean μ=3 , Standard deviation σ=1 , Skew γ_1=0)

		Mean μ				Standard deviation σ			
Sample size		PDE, SN	PDE, N	LS, SN	LS, N	PDE, SN	PDE, N	LS, SN	LS, N
n=20	MSE*	5.39	5.15	6.11	8.33	3.51	3.47	4.49	5.49
	MAE	0.18	0.18	0.20	0.23	0.15	0.15	0.17	0.20
	bias*	1.45	-0.53	0.70	0.18	-2.78	-1.58	-2.23	-5.89
n=50	MSE*	2.19	2.25	3.08	3.42	1.17	1.19	1.92	2.10
	MAE	0.12	0.12	0.14	0.15	0.09	0.09	0.11	0.11
	bias*	0.88	0.23	0.14	0.96	-1.26	-0.52	0.13	-3.11
n=100	MSE*	1.07	1.20	1.49	1.64	0.62	0.57	1.00	1.07
	MAE	0.08	0.09	0.10	0.10	0.06	0.06	0.08	0.08
	bias*	-0.02	0.38	-0.03	-0.32	-0.94	-0.24	1.10	-1.40
n=200	MSE*	0.60	0.56	0.82	0.90	0.41	0.46	0.52	0.52
	MAE	0.06	0.06	0.07	0.08	0.05	0.05	0.06	0.06
	bias*	0.50	-0.30	-0.06	-0.38	1.17	2.52	0.78	-0.98

*Reported figures are scaled up by 100.

References

Azzalini, A., "A Class of Distributions which Includes the Normal Ones ", *Scandinavian Journal of Statistics*, 12, 171-178, 1985.

Ang, A., G. Bekaert, and M. Wei, "The Term Structure of Real Rates and Expected Inflation?", *Journal of Finance, forthcoming.*

Ang, A., G. Bekaert, and M. Wei, "Do Macro Variables, Asset Markets or Surveys Forecast Inflation Better?", *Journal of Monetary Economics*, 54 (4), 1163-1212, 2007.

Ball, L., "Credible Disinflation with Staggered Price-Setting", *American Economic Review*, 84, 282-89, March, 1994.

Ball, L., "Disinflation with Imperfect Credibility", *Journal of Monetary Economics* 35, 5–23.

Boero, G., J. Smith and K. Wallis, "Uncertainty and Disagreement in Economic Prediction: The Bank of England Survey of External Forecasters", *Economic Journal, forthcoming.*

Birch, M. V. ,"A New Proof of the Fisher-Pearson Theorem", *Annals of Mathematical Statistics*, 35, 817-824, 1964.

Blix, M. and P. Sellin, "Uncertainty Bands for Inflation Forecasts", Sveriges Riksbank Working Paper Series 65, September, 1998.

Britton, E., P. Fisher and J. Whitley, "The Inflation Report Projections: Understanding the fan chart", Bank of England Quarterly Bulletin, February, 1998.

Cressie, N. and T.R.C. Read. "Multinomial goodness-of-fit tests", *Journal of the Royal Statistical Society B*, 46, 440-464, 1984.

Cressie, N. and T.R.C. Read, "*Goodness-of-fit statistics for discrete multivariate data*", New York, Springer Verlag, 1988.

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

D'Amico, S. and A. Orphanides, "Uncertainty and Disagreement in economic forecasting", Board of Governors of the Federal Reserve, *manuscript,* 2006.

Engelberg, J., C. F. Manski, and J. Williams, "Comparing the Point Predictions and Subjective Probability Distributions of Professional Forecasters", *Journal of Business and Economic Statistics*, *forthcoming.*

García, J.A. "An introduction to the ECB's Survey of Professional Forecasters", ECB Occasional Paper series, No. 8, September, 2003.

García, J.A. and A. Manzanares, "Reporting biases and survey results: evidence from European professional forecasters", ECB Working Paper Series, *forthcoming.*

Giordani, P., and P. Söderlind, "Inflation Forecast Uncertainty", *European Economic Review*, 47, 1037-1059, 2003.

Goodfriend, M., "Interest Rate Policy and the Inflation Scare Problem:1979–1992", *Federal Reserve Bank of Richmond Economic Quarterly,* 79(1), Winter, 1993.

Goodfriend, M. and R. G. King, "The Incredible Volcker Disinflation", *Journal of Monetary Economics*, 52, 981–1015, 2005.

Gouriéroux, C. and A. Montfort, "*Statistics and Econometric Models*", Cambridge University Press, 1995.

Henze, N. "A Probabilistic Representation of the Skew-Normal Distribution", *Scandinavian Journal of Statistics*, 13, 271-275, 1986.

Gürkaynak, R., B. Sack and E. Swanson, "The Sensitivity of Long-Term Interest Rates to Economic News: Evidence and Implications for Macroeconomic Models", *American Economic Review* 95(1), 425-36, 2005.

Lahiri, K., C. Teigland, and M. Zaporowski. "Interest Rates and the Subjective Probability Distribution of Inflation Forecasts", *Journal of Money, Credit and Banking*, 20, 233-248, 1988.

Lindsay, B.G., "Efficiency versus Robustness: the Case for Minimum Hellinger Distance and Related Methods", *The Annals of Statistics*, 22, 1081-1114, 1994.

Machina M..and M. Rothschild, "Risk", The New Palgrave Dictionary of Economics, 2nd Edition, edited by S. N. Durlauf and L. E. Blume, Macmillan (Basingstoke and New York), *forthcoming.*

Neyman, J., "Contributions to the Theory of the Chi-Square Test", *First Berkeley Symposium on Mathematical Statistics and Probability*, 239-273, 1949.

Orphanides, A. "Comment on: The Incredible Volcker Disinflation", *Journal of Monetary Economics*, 52, 1017–1023, 2005.

Orphanides, A. and J. Williams, "Inflation Scares and Forecast-based Monetary Policy", *Review of Economic Dynamics*, 8, 498-527, 2005.

Patefield, M., and Tandy, D., "Fast and Accurate Computation of Owen's T-Function", *Journal of Statistical Software* 5 (5), 2000.

Pearson, K.,"On the Criterion that a Given System of Deviations from the Probable in the Case of a Correlated System of Variables is such that It can be Reasonably Supposed to Have Arisen from Random Sampling", *Philosophy Magazine*, 50, 157-172, 1900.

Pewsey, A., "Problems of Inference for Azzalini's Skew-Normal Distribution", Journal of Applied Statistics, 27, 859-870, 2000.

Wallis, K. F., "Combining Density and Interval Forecasts: A Modest Proposal", *Oxford Bulletin of Economics and Statistics*, 67, 983-994, 2005.

Zarnowitz, V. and L. A. Llambros, "combined and Uncertainty in Economic Prediction", *Journal of Political Economy*, 95(3), 591-621, 1987.

Rich, R. W. and J. S. Butler, "Disagreement as a Measure of Uncertainty: A Comment on Bomberger", *Journal of Money, Credit and Banking*, 30 (3), Part 1, 411-419, 1998.

Rich, R., and J. Tracy, "The Relationship between Expected Inflation, Disagreement, and Uncertainty: Evidence from Matched Point and Density Forecasts", Federal Reserve Bank of New York Staff Reports, No. 253, July 2006.

Serfling, R.J., "Approximation Theorems of Mathematical Statistics", *J. Wiley. New York*, 163-165, 1980.

Table 1: Estimation errors for the skewness parameter

Degree of skewness		Sample size n=20	n=50	n=100	n=500	n=1000
γ=0	Mean Sq. Error*	50.6	23.9	11.0	5.9	2.6
	Bias	-0.01	-0.01	0.03	0.02	0.03
	Monte Carlo St. Dev.	0.71	0.49	0.33	0.24	0.16
γ=0.3	Mean Sq. Error*	48.5	22.0	10.9	5.4	2.1
	Bias	-0.05	0.04	0.02	0.03	-0.01
	Monte Carlo St. Dev.	0.69	0.47	0.33	0.23	0.14
γ=0.6	Mean Sq. Error*	41.8	16.0	8.8	6.7	5.0
	Bias	-0.13	0.02	0.04	0.03	0.02
	Monte Carlo St. Dev.	0.63	0.40	0.29	0.26	0.22
γ=0.9	Mean Sq. Error*	30.3	6.7	2.4	1.5	0.9
	Bias	-0.20	-0.05	0.03	0.06	0.08
	Monte Carlo St. Dev.	0.51	0.26	0.15	0.11	0.05

*Reported figures are scaled up by 100.

Results are based on our proposed fitting criterion $PDE(\tau^*=0.2)$.

Table 2 : Correlation between measures of inflation uncertainty

Q1surveys (1969Q1-2006Q1)	Variance (Agg. dist.)	Average uncertainty	Disagreement
Variance (Agg. dist.)	1.00		
Average uncertainty	0.90	1.00	
Disagreement	0.74	0.54	1.00

Agg. Dist denotes the average of the individual probability forecast across panelists.

Average uncertainty denotes the average of the variances of the individual probability forecasts.

Disagreement is measured by the variance of the means of the individual probability forecasts.

Table 3 : Co-movement between measures of skewness

Q1surveys (37obs)	Skewness (Agg. dist.)	Average skewness	Skew of means	Cov. mean and variance
Skewness (Agg. dist.)	1.00			
Average skewness	0.28	1.00		
Skew of means	0.55	0.03	1.00	
Cov. mean and variance	0.78	0.39	0.24	1.00

Table 4: Correlation between the moments of inflation forecasts

	Agg. Mean	Agg. Variance	Average uncertainty	Agg. Skew	Mean-Mode
Agg. Mean	(5.45)				
Agg. Variance .	0.59	(0.37)			
Average uncertainty	0.48	0.88	(0.15)		
Agg. Skew	0.07	0.54	0.60	(1.12)	
Mean-Mode	-0.02	0.42	0.53	0.98	(0.40)

Q1 surveys, 38 observations 1969Q1-2006Q1. Figures in main diagonal denote variances of the series.

Table 5: Level, uncertainty and risk assessment in inflation expectations

	Uncertainty measures		Risk assessment measures	
	Variance (Agg. Dist.)	Average uncertainty	Skew (Agg. Dist.)	Mean-Mode
Mean (Agg. Dist.)	0.24^{**}	0.18^{**}	0.10	0.03
Average of means	0.24^{**}	0.18^{**}	0.10	0.03
Variance (Agg. Dist.)			0.85^{**}	0.40^{*}
Average uncertainty			1.12^{**}	0.54^{**}

Each cell reports a regression of the risk measures on the first column variables, qurarterly data.1981Q1-2006Q1.

** and *denote significance at 1% and 5% level (Newey-West corrected standard errors upto a year).

Table 6: Inflation risks and macroeconomic variables

	Variance	Average uncertainty	Disagreement	Skewness	Mean-Mode
Panel A: Bivariate regressions (each cell represents a separate regression)					
Inflation	0.19*	0.16**	0.05**	0.04	-0.04
Δinflation-squared	0.89*	0.63*	0.26*	0.41	0.13
Output gap	-0.82	-1.06	0.33	0.82	-3.20
Prob. recession	0.04**	0.04**	0.01**	0.04*	0.02*
Growth uncertainty	0.98**	0.83**	0.20**	1.06**	0.53**
Panel B: Controlling for actual inflation					
Δinflation-squared	-0.43	-0.41*	-0.04	0.34	0.31
Output gap	5.47	3.78	1.66	-0.57	-3.27
Prob. recession	0.04**	0.03**	0.01**	0.04*	0.02
Growth uncertainty	0.78**	0.70**	0.15**	1.26**	0.69**
Panel C: Controlling for expected inflation					
Δinflation-squared	-0.55*	-0.51**	-0.09	-0.34	-0.15
Output gap	4.94	3.47	1.57	1.62	-1.59
Prob. recession	0.02	0.02*	0.01	0.03	0.02
Growth uncertainty	0.65**	0.62**	0.10*	1.17**	0.63**

** and *denote significance at 1% and 5% level (Newey-West corrected standard errors upto a year)

Prob. recession refers to the probability of negative growth four quarters ahead reported in the SPF.

Growth uncertainty refers to the estimated average uncertainty from the SPF probabilistic forecast.

Figure 1: Examples of skew-normal densities

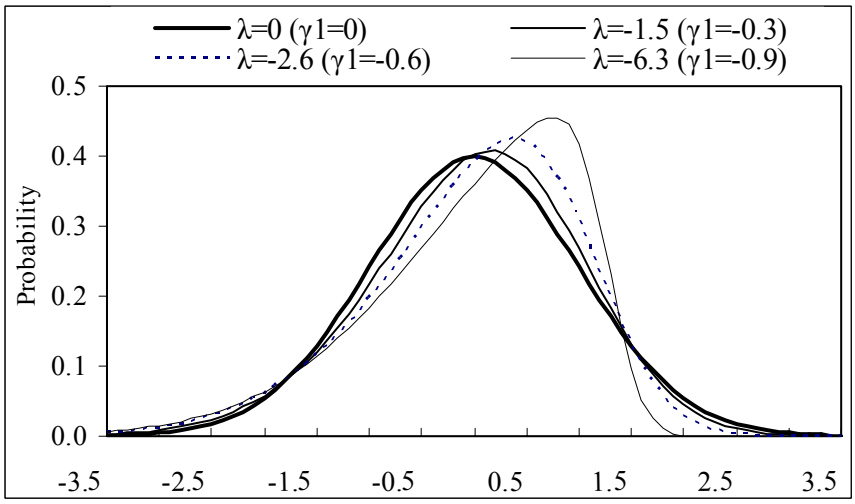

The densities depicted have zero mean and unit variance

Figure 2: Fitting criterion gains (MSE)

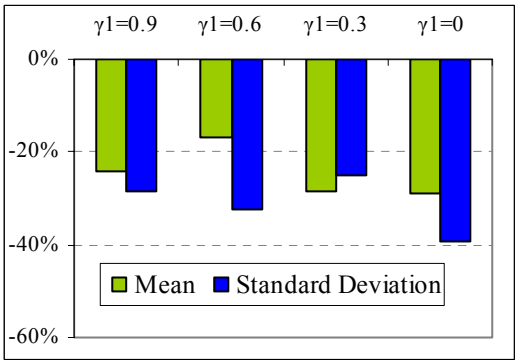

Figure 3: Density gains (MSE)

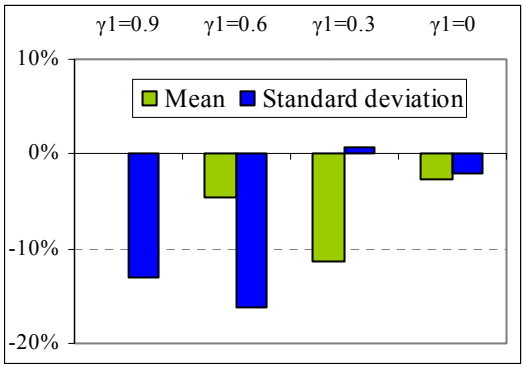

Figure 4: Overall gains (MSE)

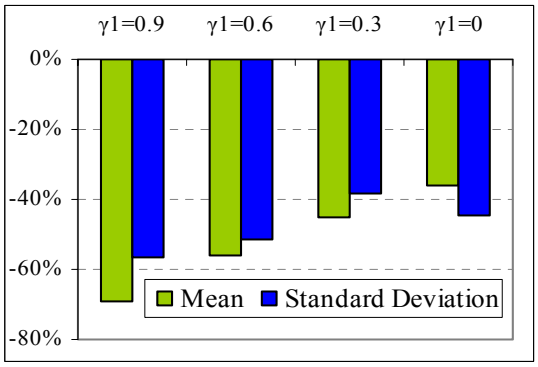

Figure 5: Overall gains from LS,N (MAE)

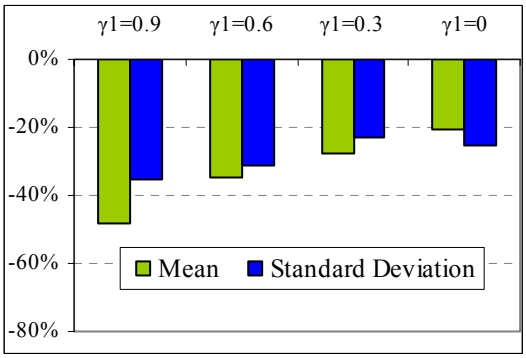

Note: For both the estimated mean and variance of the distribution, the charts depict the percentage reduction in the estimation error (mean absolute error (MAE) and mean square error (MSE)) of our methodology with respect to the least squares (LS) fitting criterion and/or the standard normal (N) density. Specifically, Figure 2 reports the reduction in the MSE from using our fitting criterion instead of LS for the Skew-Normal (SN) distribution, Figure 3 for using the SN instead of the N with our fitting criterion, and Figures 4 and 5 from using our fitting criterion and the SN instead of LS and the N. Results are based on 1,000 simulations from SN distributions with four different degrees of skewness γ_1.

Figure 6: Skew-normal versus normal fitting

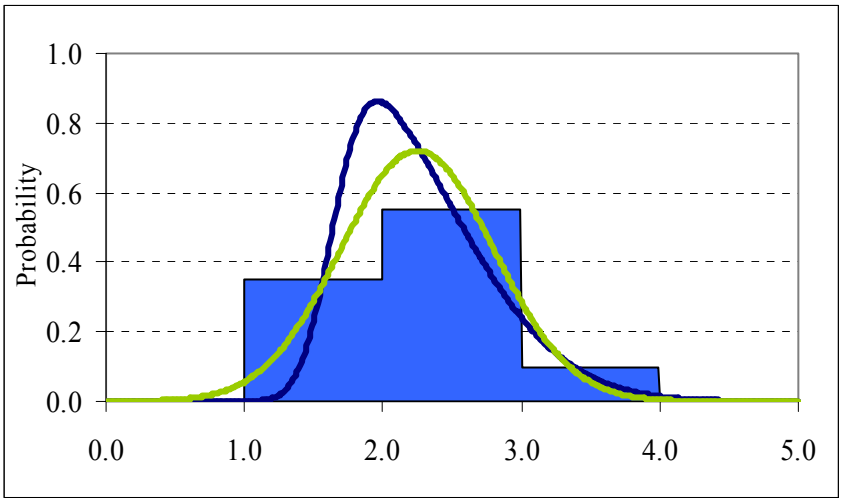

The histogram is from forecaster 541 in the 2005 Q3 SPF

Figure 7: Estimated uncertainty and number of active intervals reported

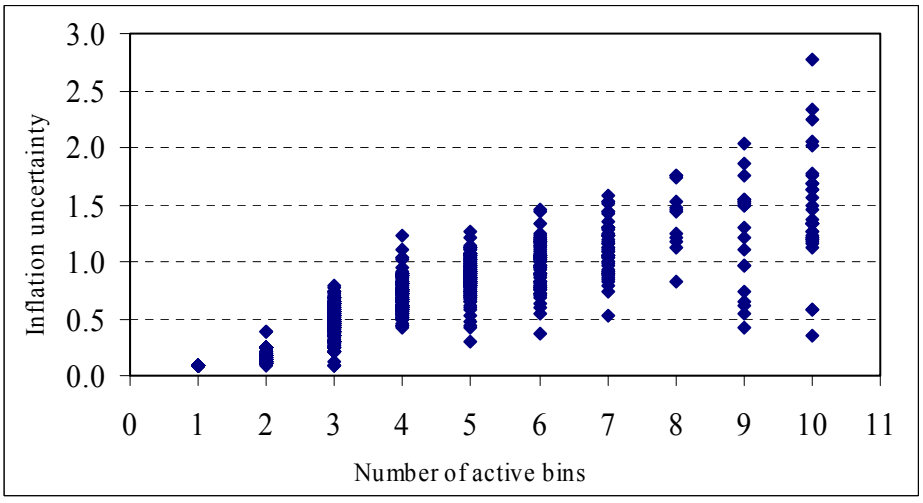

Estimation results for individual probability forecasts between 1992Q1 and 2006Q1

Figure 8: Variance of the combined distribution, average uncertainty and disagreement

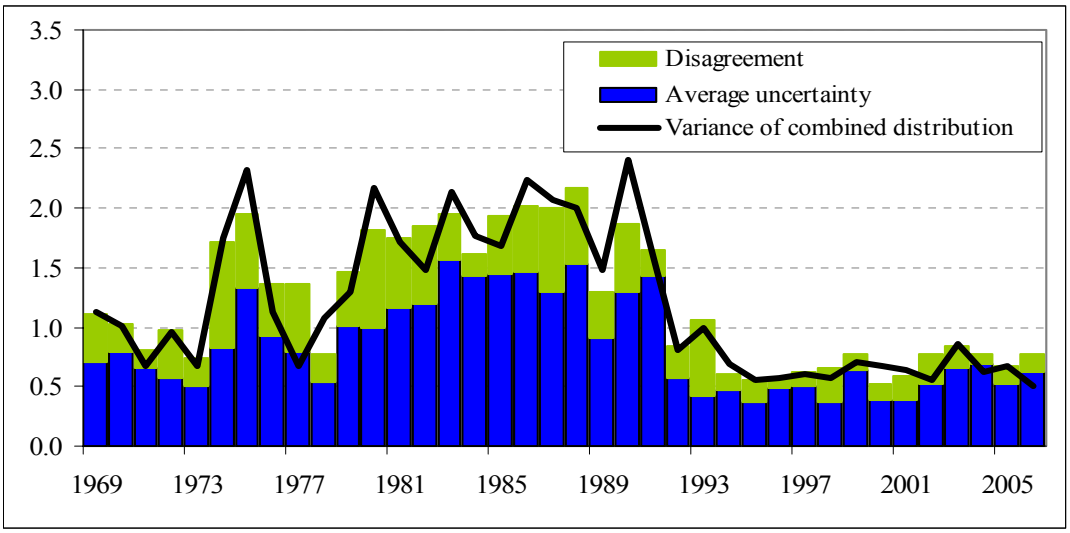

Decomposition of the variance of the combined distribution (current-year forecasts in Q1)

Figure 9: Skewness (combined forecasts)

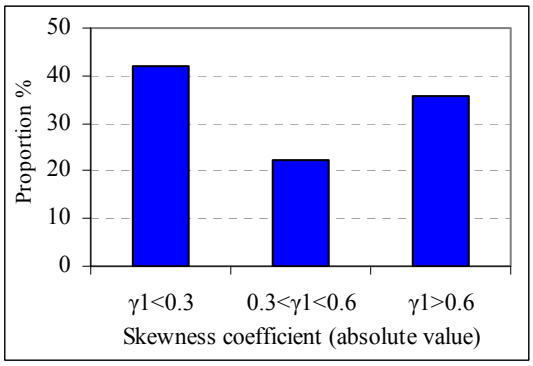

Based on 247 estimated combined densities, 1968Q4-2006Q1

Figure 10: Skewness (individual forecasts)

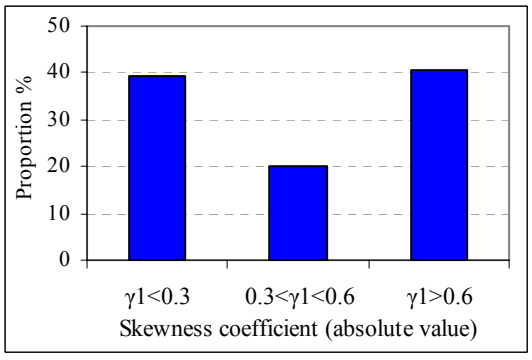

Based on 8,204 estimated individual densities, 1968Q4-2006Q1

Figure 11: Alternative measures of skew (combined distribution)

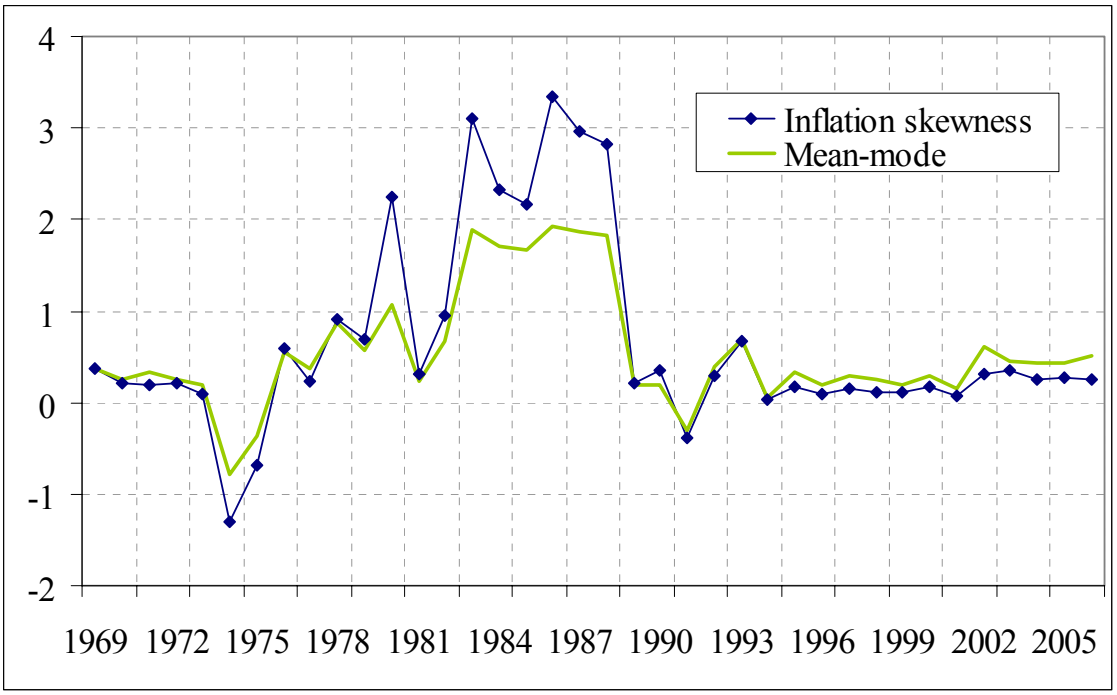

Figure 12: Expected inflation and inflation risks

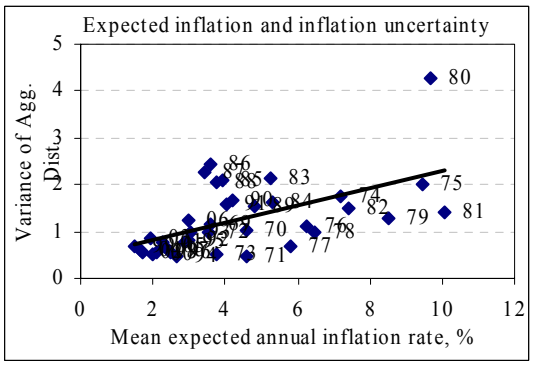

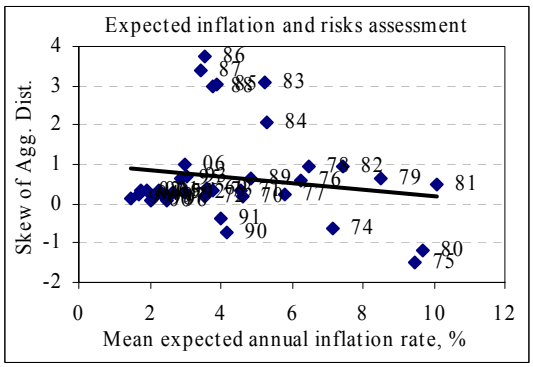

Figure 13: Interest rates, inflation expectations and inflation scares

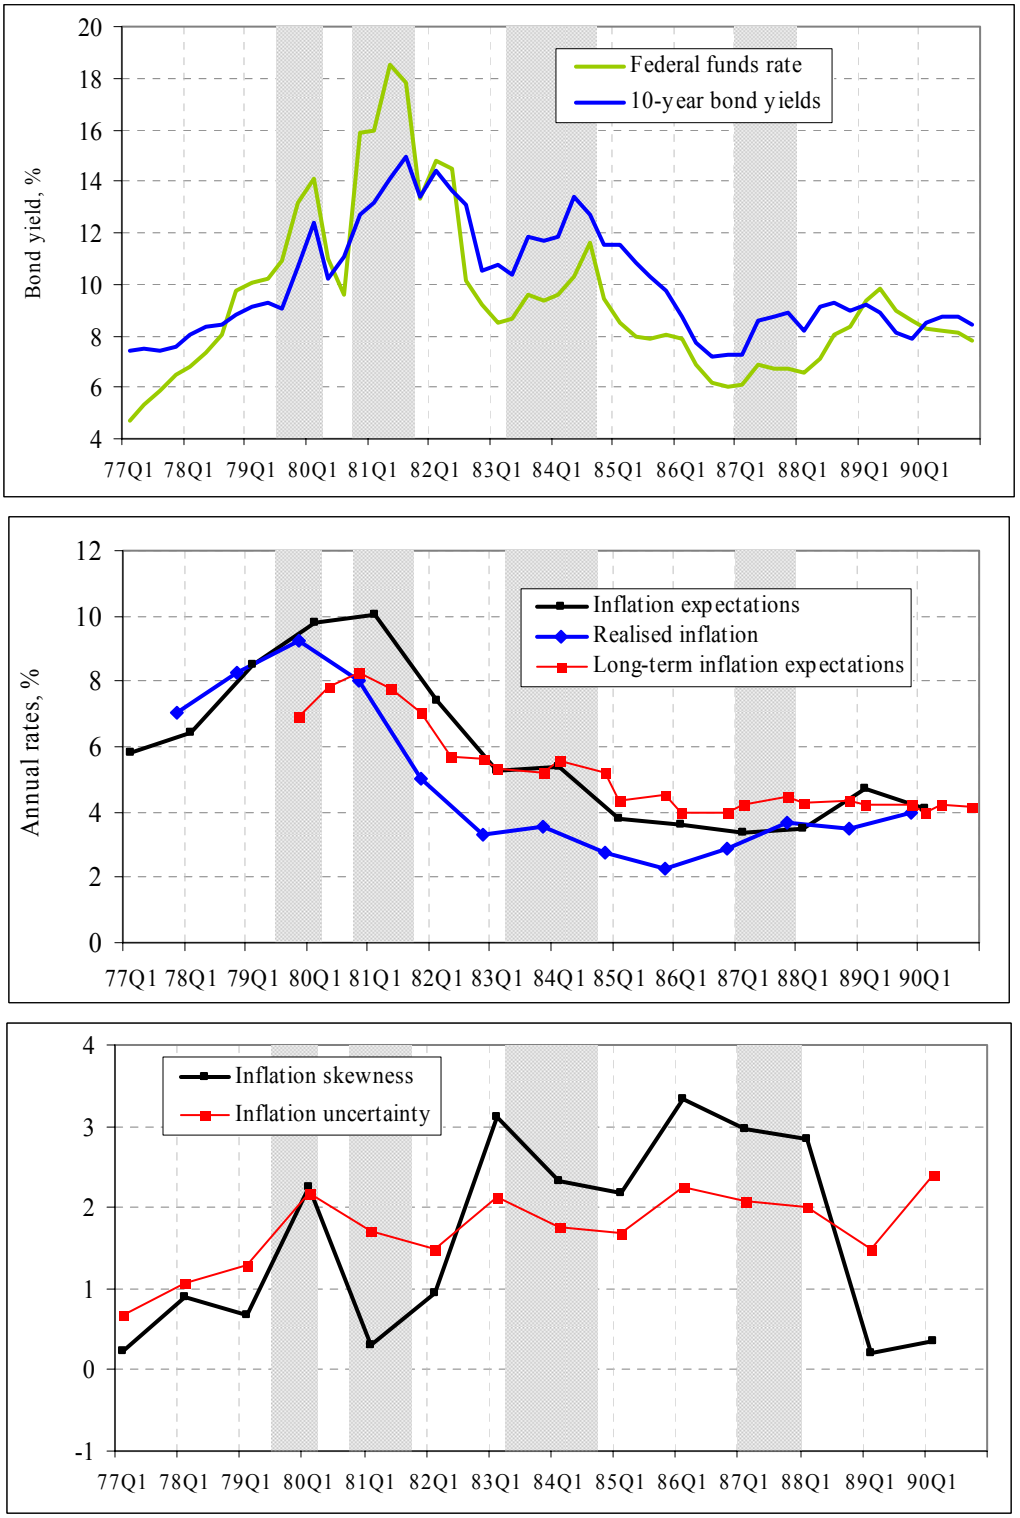

Note: Shadowed areas reflect periods of inflation scares in the bond market. Inflation expectations denote the mean of the combined distribution for current-year inflation; realized inflation is the year-end rate of growth in the GDP deflator; long-term inflation expectations are from the Philadelphia Fed. The bottom chart depicts the variance and skewness from the combined distribution for current-year inflation.

Figure 14: Inflation and inflation risks

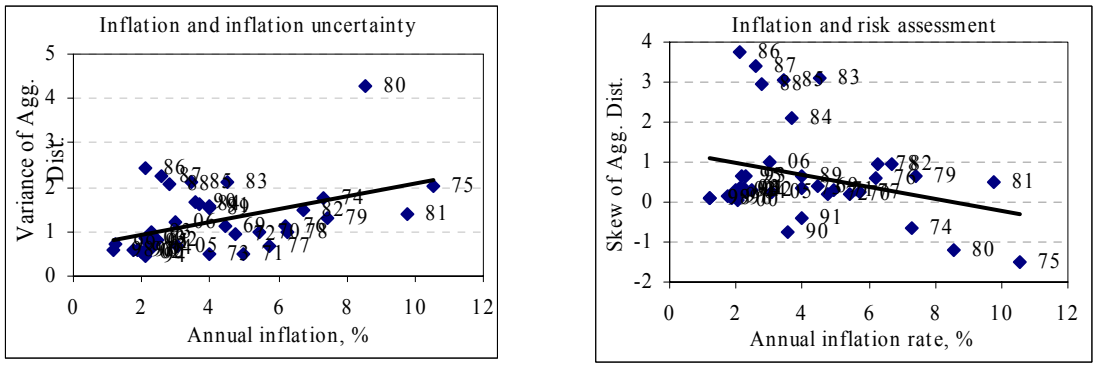

Figure 15: Inflation risks and economic activity

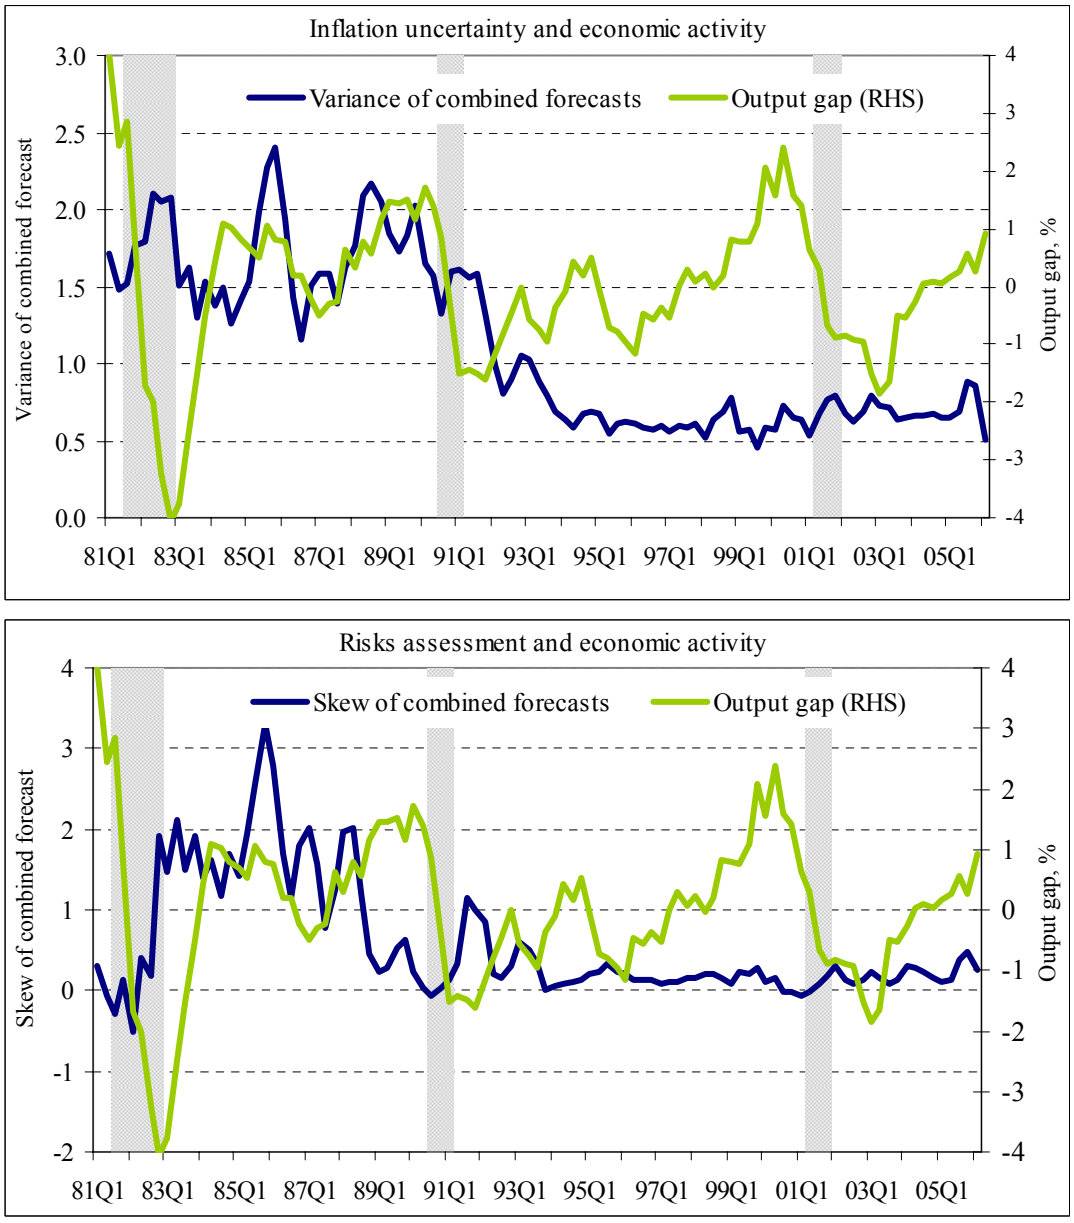

Note: output gap is HP-filtered. Inflation uncertainty and skew are smoothed (five quarter centred moving averages)

Figure AI: Relative performance of alternative estimators, MSE

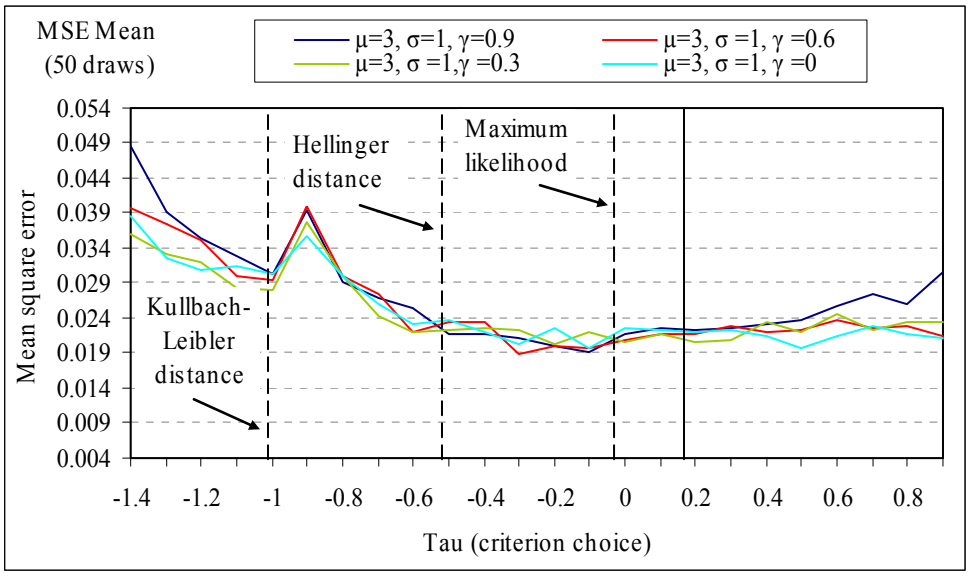

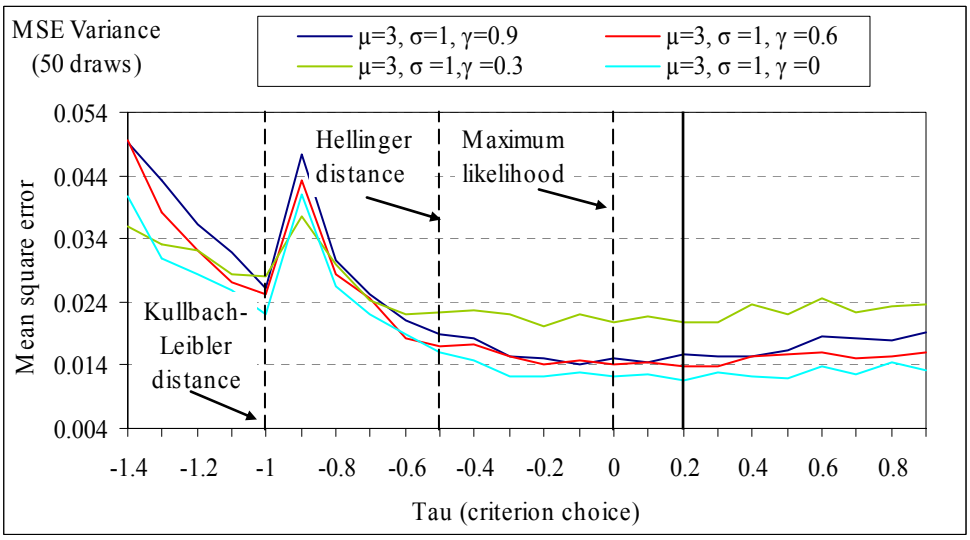

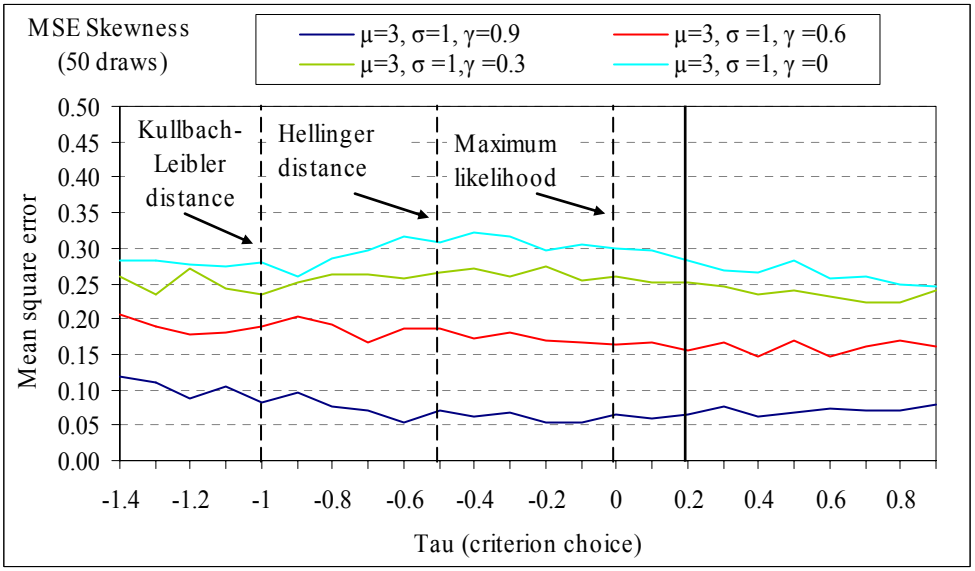

European Central Bank Working Paper Series

For a complete list of Working Papers published by the ECB, please visit the ECB's website (http://www.ecb.europa.eu)

773 "Exchange rate volatility and growth in small open economies at the EMU periphery" by G. Schnabl, July 2007.

774 "Shocks, structures or monetary policies? The euro area and US after 2001" by L. Christiano, R. Motto and M. Rostagno, July 2007.

775 "The dynamic behaviour of budget components and output" by A. Afonso and P. Claeys, July 2007.

776 "Insights gained from conversations with labor market decision makers" by T. F. Bewley, July 2007.

777 "Downward nominal wage rigidity in the OECD" by S. Holden and F. Wulfsberg, July 2007.

778 "Employment protection legislation and wages" by M. Leonardi and G. Pica, July 2007.

779 "On-the-job search and the cyclical dynamics of the labor market" by M. U. Krause and T. A. Lubik, July 2007.

780 "Dynamics and monetary policy in a fair wage model of the business cycle" by D. de la Croix, G. de Walque and R. Wouters, July 2007.

781 "Wage inequality in Spain: recent developments" by M. Izquierdo and A. Lacuesta, July 2007.

782 "Panel data estimates of the production function and product and labor market imperfections" by S. Dobbelaere and J. Mairesse, July 2007.

783 "The cyclicality of effective wages within employer-employee matches: evidence from German panel data" by S. Anger, July 2007.

784 "Understanding the dynamics of labor shares and inflation" by M. Lawless and K. Whelan, July 2007.

785 "Aggregating Phillips curves" by J. Imbs, E. Jondeau and F. Pelgrin, July 2007.

786 "The economic impact of merger control: what is special about banking?" by E. Carletti, P. Hartmann and S. Ongena, July 2007.

787 "Finance and growth: a macroeconomic assessment of the evidence from a European angle" by E. Papaioannou, July 2007.

788 "Evaluating the real effect of bank branching deregulation: comparing contiguous counties across U.S. state borders" by R. R. Huang, July 2007.

789 "Modeling the impact of external factors on the euro area's HICP and real economy: a focus on pass-through and the trade balance" by L. Landolfo, July 2007.

790 "Asset prices, exchange rates and the current account" by M. Fratzscher, L. Juvenal and L. Sarno, August 2007.

791 "Inquiries on dynamics of transition economy convergence in a two-country model" by J. Brůha and J. Podpiera, August 2007.

792 "Euro area market reactions to the monetary developments press release" by J. Coffinet and S. Gouteron, August 2007.

793 "Structural econometric approach to bidding in the main refinancing operations of the Eurosystem" by N. Cassola, C. Ewerhart and C. Morana, August 2007.

794 "(Un)naturally low? Sequential Monte Carlo tracking of the US natural interest rate" by M. J. Lombardi and S. Sgherri, August 2007.

795 "Assessing the impact of a change in the composition of public spending: a DSGE approach" by R. Straub and I. Tchakarov, August 2007.

796 "The impact of exchange rate shocks on sectoral activity and prices in the euro area" by E. Hahn, August 2007.

797 "Joint estimation of the natural rate of interest, the natural rate of unemployment, expected inflation, and potential output" by L. Benati and G. Vitale, August 2007.

798 "The transmission of US cyclical developments to the rest of the world" by S. Dées and I. Vansteenkiste, August 2007.

799 "Monetary policy shocks in a two-sector open economy: an empirical study" by R. Llaudes, August 2007.

800 "Is the corporate bond market forward looking?" by J. Hilscher, August 2007.

801 "Uncovered interest parity at distant horizons: evidence on emerging economies & nonlinearities" by A. Mehl and L. Cappiello, August 2007.

802 "Investigating time-variation in the marginal predictive power of the yield spread" by L. Benati and C. Goodhart, August 2007.

803 "Optimal monetary policy in an estimated DSGE for the euro area" by S. Adjemian, M. Darracq Pariès and S. Moyen, August 2007.

804 "Growth accounting for the euro area: a structural approach" by T. Proietti and A. Musso, August 2007.

805 "The pricing of risk in European credit and corporate bond markets" by A. Berndt and I. Obreja, August 2007.

806 "State-dependency and firm-level optimization: a contribution to Calvo price staggering" by P. McAdam and A. Willman, August 2007.

807 "Cross-border lending contagion in multinational banks" by A. Derviz and J. Podpiera, September 2007.

808 "Model misspecification, the equilibrium natural interest rate and the equity premium" by O. Tristani, September 2007.

809 "Is the New Keynesian Phillips curve flat?" by K. Kuester, G. J. Müller and S. Stölting, September 2007.

810 "Inflation persistence: euro area and new EU Member States" by M. Franta, B. Saxa and K. Šmídková, September 2007.

811 "Instability and nonlinearity in the euro area Phillips curve" by A. Musso, L. Stracca and D. van Dijk, September 2007.

812 "The uncovered return parity condition" by L. Cappiello and R. A. De Santis, September 2007.

813 "The role of the exchange rate for adjustment in boom and bust episodes" by R. Martin, L. Schuknecht and I. Vansteenkiste, September 2007.

814 "Choice of currency in bond issuance and the international role of currencies" by N. Siegfried, E. Simeonova and C. Vespro, September 2007.

815 "Do international portfolio investors follow firms' foreign investment decisions?" by R. A. De Santis and P. Ehling, September 2007.

816 "The role of credit aggregates and asset prices in the transmission mechanism: a comparison between the euro area and the US" by S. Kaufmann and M. T. Valderrama, September 2007.

817 "Convergence and anchoring of yield curves in the euro area" by M. Ehrmann, M. Fratzscher, R. S. Gürkaynak and E. T. Swanson, October 2007.

818 "Is time ripe for price level path stability?" by V. Gaspar, F. Smets and D. Vestin, October 2007.

819 "Proximity and linkages among coalition participants: a new voting power measure applied to the International Monetary Fund" by J. Reynaud, C. Thimann and L. Gatarek, October 2007.

820 "What do we really know about fiscal sustainability in the EU? A panel data diagnostic" by A. Afonso and C. Rault, October 2007.

821 "Social value of public information: testing the limits to transparency" by M. Ehrmann and M. Fratzscher, October 2007.

822 "Exchange rate pass-through to trade prices: the role of non-linearities and asymmetries" by M. Bussière, October 2007.

823 "Modelling Ireland's exchange rates: from EMS to EMU" by D. Bond, M. J. Harrison and E. J. O'Brien, October 2007.

824 "Evolving U.S. monetary policy and the decline of inflation predictability" by L. Benati and P. Surico, October 2007.

825 "What can probability forecasts tell us about inflation risks?" by J. A. García and A. Manzanares, October 2007.

ISSN 1561-0810