NBER WORKING PAPER SERIES

COMPARING THE POINT PREDICTIONS AND
SUBJECTIVE PROBABILITY DISTRIBUTIONS
OF PROFESSIONAL FORECASTERS

Joseph Engelberg
Charles F. Manski
Jared Williams

Comparing the Point Predictions and Subjective Probability Distributions of Professional
Forecasters
Joseph Engelberg, Charles F. Manski and Jared Williams
NBER Working Paper No. 11978
January 2006
JEL No. C42, E27, E47

ABSTRACT

We use data from the Survey of Professional Forecasters to compare point forecasts of GDP growth
and inflation with the subjective probability distributions held by forecasters. We find that SPF
forecasters summarize their underlying distributions in different ways and that their summaries tend
to be favorable relative to the central tendency of the underlying distributions. We also find that
those forecasters who report favorable point estimates in the current survey tend to do so in
subsequent surveys. These findings, plus the inescapable fact that point forecasts reveal nothing
about the uncertainty that forecasters feel, suggest that the SPF and similar surveys should not ask
for point forecasts. It seems more reasonable to elicit probabilistic expectations and derive measures
of central tendency and uncertainty, as we do here.

Joseph Engelberg
Kellogg School of Management
Northwestern University
2001 Sheridan Road
Evanston, IL 60208
j-engelberg@northwestern.edu

Jared Williams
Kellogg School of Management
Northwestern University
2001 Sheridan Road
Evanston, IL 60208
williamsjm@northwestern.edu

Charles F. Manski
Department of Economics
Northwestern University
2001 Sheridan Road
Evanston, IL 60208
and NBER
cfmanski@northwestern.edu

1 Introduction

Professional forecasters have long given point predictions of future events. Financial analysts offer point forecasts of the profit that firms will earn in the quarter ahead, macroeconomic forecasters give point predictions of annual GDP growth and inflation, and film critics conjecture which actors will win Academy Awards. A notable exception is that meteorologists commonly report the chance that it will rain during the next day. However, meteorologists give point rather than probabilistic predictions of the next day's high and low temperatures.

Thoughtful forecasters rarely think that they have perfect foresight. Hence, their point predictions can at most convey some notion of the central tendency of their beliefs, and nothing at all about the uncertainty they feel. Suppose, as often seems reasonable, that forecasters actually have subjective probability distributions for the events they predict. Then their point predictions should somehow be related to their subjective distributions. But how? This is the question that we address. Our particular concern is with prediction of real-valued outcomes such as firm profit, GDP growth, or temperature.[1]

In general, forecasters making point predictions of real outcomes are not asked to provide the mean, median, mode or any other specific feature of a subjective probability distribution. Instead, they are simply asked to "predict" the outcome or to provide their "best prediction," without definition of the word "best." In the absence of explicit guidance, forecasters may report different distributional features as their point predictions. Hence, it is possible that forecasters who hold identical probabilistic beliefs provide different point predictions and that forecasters with dissimilar beliefs provide identical point predictions. If so, comparison of point predictions across forecasters is problematic. Variation in predictions need not imply disagreement among forecasters and homogeneity in predictions need not imply agreement.

To perform our analysis requires a data source in which forecasters provide both point predictions and subjective distributions for specific events. Perhaps uniquely among existing forecasting instruments, the Survey of Professional Forecasters (SPF) provides such data (www.phil.frb.org/econ/spf/). The SPF respondents are macroeconomic forecasters who are queried about future GDP growth and inflation in the USA. The forecasters provide both point and probabilistic predictions for these quantities; the latter are the probabilities that the outcome will lie in each of ten intervals. Another important feature of the SPF is that the survey has a longitudinal component, with many forecasters providing multiple predictions over the course of several years. Section 2 describes the SPF data in detail.

Although the SPF probabilistic prediction data do not completely identify the subjective distributions that respondents hold, they do imply fairly tight bounds on the means, medians, and modes of these distrib-

---

[1] When the forecasting problem is to predict the occurrence of a binary event, the idea that point predictions should somehow be related to subjective probability distributions was suggested forty years ago by Juster (1966). Considering the case in which consumers are asked to give a point prediction of their buying intentions (buy or not buy), Juster wrote (page 664): "Consumers reporting that they 'intend to buy A within X months' can be thought of as saying that the probability of their purchasing A within X months is high enough so that some form of 'yes' answer is more accurate than a 'no' answer." Thus, he hypothesized that a consumer facing a yes/no intentions question responds as would a statistician asked to make a best point prediction of a future random event. Subsequent analysis formalizing Juster's idea appears in Manski (1990).

utions. Hence, we are able to compare respondents' point predictions with these features of their subjective probability distributions. This is done in Section 3.

We find that reporting practices differ substantially across forecasters. Whereas most SPF forecasters give point predictions that are consistent with the means/medians/modes of their subjective distributions, many do not. Moreover, the heterogeneity in reporting is not symmetric but rather tends towards presentation of point forecasts that are favorable relative to the central tendencies of respondents' subjective distributions. In the case of GDP growth, respondents who give point predictions outside the bounds of their subjective means/medians/modes are more likely to report point predictions that are above the upper bound than below the lower bound. In the case of inflation, the point predictions are more often below the lower bounds than above the upper bounds. We also find persistence over time in the reporting practices of individual respondents: forecasters who give point predictions above/below their subjective means/medians/modes in one sample period are more likely to do the same the next period.

The findings described above are nonparametric — we use the SPF interval probability data in their raw form to bound subjective means/medians/modes rather than to fit the precise subjective probability distributions that respondents hold. To enable further analysis of the data, we go on in Section 4 to suppose that each subjective distribution placing positive probability on at least three intervals has the generalized Beta form, and we use the interval probability data to fit the parameters. In those cases where a forecast places positive probability on only one or two intervals, we suppose that the distribution has the shape of an isosceles triangle and we fit its parameters. This done, we are able to compare SPF point forecasts with the fitted probability distributions. The parametric analysis corroborates and sharpens our nonparametric finding of heterogeneity across forecasters, with a tendency for favorable point predictions. The parametric analysis also corroborates our nonparametric finding of persistence over time in individual reporting practices.

Our findings suggest that efforts to aggregate point predictions across forecasters face considerable difficulty. In the finance literature, researchers have used the mean point prediction of a group of forecasters to summarize the mean beliefs held by these forecasters and have interpreted cross-forecaster dispersion in point predictions to indicate disagreement in their beliefs (e.g., Diether, Malloy, and Scherbina, 2002; Mankiw, Reis, and Wolfers, 2003). These research practices confound variation in forecaster beliefs with variation in the manner that forecasters make point predictions.

Our work also suggests that it may be problematic to interpret the observed time-series variation in SPF "consensus" forecasts for the economy. So-called consensus forecasts are simply the median point predictions of respondents. The SPF panel of forecasters evolves over time, with some new members entering each quarter and some old ones exiting. Time-series variation in consensus forecasts need not reflect real changes in beliefs about the economy. It may instead reflect changes in the reporting practices of the evolving panel.

Even if point forecasts were fully successful in describing the central tendencies of SPF forecaster beliefs, they inherently would reveal nothing about the uncertainty that forecasters feel when predicting GDP growth

or inflation. Probabilistic forecasts are well-suited to this task.[2] In Section 5, we use the subjective probability distributions fitted in Section 4 to describe the uncertainty of SPF forecastsers. We also show how the fitted distributions can be used to describe the extent to which different forecasters disagree in their predictions.

A separate matter, which we do not feel the need to address in the body of our paper, is the longstanding use of cross-sectional dispersion in point predictions to measure forecaster uncertainty about future outcomes. See, for example, Cukierman and Wachtel (1979), Levi and Makin (1979, 1980), Makin (1982), Brenner and Landskroner (1983), Hahm and Steigerwald (1999), and Hayford (2000). This research practice is suspect on logical grounds, even if all forecasters make their point predictions in the same way. Even in the best of circumstances, point predictions provide no information about the uncertainty that forecasters feel. This point was made forcefully almost twenty years ago by Zarnowitz and Lambros (1987). Nevertheless, some researchers have continued since then to use the dispersion in point predictions to measure forecaster uncertainty.

Even though Zarnowitz and Lambros (1987) explicitly recognized the logical fallacy in using the dispersion in point predictions to measure forecaster uncertainty, they nevertheless thought it useful to determine the empirical relationship between such dispersion and uncertainty amongst the SPF respondents. More recently, Giordani and Soderlind (2003a) have performed a similar empirical analysis. As far as we are aware, the Zarnowitz and Lambros (1987) and Giordani and Soderlind (2003a) articles are the only research before our own that has compared the SPF point and probabilistic predictions. However, their analyses differ fundamentally from ours. They sought to describe the empirical relationship between the cross-sectional distribution of point predictions and the uncertainty evident in a typical forecaster's subjective probability distribution. In contrast, our concern is to understand the empirical relationship between *individual forecasters*' point predictions and their subjective probability distributions.

2 Data

Our data are from the Survey of Professional Forecasters, administered since 1990 by the Federal Reserve Bank of Philadelphia. The SPF was begun in 1968 by the American Statistical Association and the National Bureau of Economic Research; hence, it was originally called the ASA-NBER survey. The panel of forecasters, who include university professors and private-sector macroeconomic researchers, are asked to predict American GDP, inflation, unemployment, interest rates, and other macroeconomic variables.[3] The survey, which is performed quarterly, is mailed to panel members the day after government release of quarterly data on the national income and product accounts. The composition of the panel changes gradually over time, with individual members providing forecasts for about six years on average.

---

[2] Manski (2004) reviews recent research eliciting probabilistic expectations in surveys and assesses the state of the art.

[3] A partial list of respondents is posted in the Philadelphia Fed's quarterly release at http://www.phil.frb.org/econ/spf/

2.1 Question Format

Each quarter, the SPF asks panel members to make point and probabilistic forecasts of annual real GDP and inflation. To analyze the responses, it is important to understand the specific format of the questions. We explain the point and probabilistic questions in turn.

Point forecasts: The SPF instrument provides the value of GDP in the previous calendar year and asks for a point forecast of GDP in the current and the next year. Thus, in the four quarterly surveys administered during calendar year t, respondents are told the value of real GDP in year t - 1 and are asked to give point forecasts of annual GDP in years t and t + 1. Similarly, the instrument provides the average value of the GDP price index in the previous calendar year and asks for a point forecast of the average GDP index in the current and the next year.

Probabilistic forecasts: In the four quarterly surveys administered during calendar year t, respondents are asked to forecast the percentage change in annual real GDP between years t - 1 and t and, likewise, the percentage change in GDP between years t and t + 1. They are similarly asked to forecast year-to-year percentage changes in the average GDP price index. In each case, the SPF instrument partitions the real number line into intervals and asks respondents to report their subjective probabilities that the variable of interest will take a value in each interval. For GDP growth, the intervals are $(-\infty, -2\%), [x\%, x + 1\%)$ for $x = -2, -1, ..., 5$, and $[6\%, \infty)$. For inflation, they are $(-\infty, 0\%), [x\%, x+1\%)$ for $x = 0, 1, ..., 7$ and $[8\%, \infty)$.

Observe that panel members are asked to give point forecasts of the levels of GDP and the GDP price index, but probabilistic forecasts of the year-to-year percentage changes in these quantities. To make the point forecasts comparable to the probabilistic ones, we must convert the point forecasts of levels into point forecasts of percentage change. This is straightforward to do between years t - 1 and t, because the SPF instrument tells respondents the actual value of the year t - 1 level.[4] Thus, to obtain a forecaster's point forecast of the percentage change in GDP between years t - 1 and t, we calculate the percentage change between the SPF-specified level of year t - 1 GDP and the respondent's point forecast of year t GDP. This calculation is correct provided that panel members accept as accurate the SPF figure for year t - 1 GDP[5].

It is not straightforward to produce point forecasts of percentage change between years t and t + 1. In this case, we only know forecasters' point forecasts of the levels for the two years, and these point forecasts are related in unknown ways to their subjective distributions. Hence, the percentage change in point forecasts between years t and t + 1 need not equal the point forecast for percentage change that a panel member would have stated had he been asked this question. Given this, our analysis restricts attention to comparisons of

---

[4] The year t-1 figures provided to respondents are posted in the Philadelphia Fed's "Real-Time Data Set" at http://www.phil.frb.org/econ/forecast/reaindex.html

[5] During year t, information about year t-1 GDP and the GDP price index continues to be collected by the government. Therefore, year t-1 values for GDP and the price index often are revised from quarter to quarter. The SPF updates the figures it gives forecasters accordingly. These revisions were generally small during our 1992-2004 sample period, the one notable exception being a large revision in both accounts due to an accounting overhaul of the National Income and Product Accounts (NIPA) in quarter 4 of 1999. Otherwise, the average absolute value of the revision from quarter 1 to quarter 2 was 0.05% for GDP and 0.03% for the price index, from quarter 2 to quarter 3 was 0.37% for GDP and 0.22% for the price index, and from quarter 3 to quarter 4 was 0.12% for GDP and 0.02% for the price index. There were only four cases for GDP and no cases for the price index where the change between quarters was larger than 1 percent.

year t - 1 and year t.

2.2   Sample for Analysis

Although the SPF began in 1968, we restrict attention to data collected from 1992 on.   There are several reasons for this:

1. The survey only began asking forecasters for their annual (rather than quarterly) point forecasts in the third quarter of 1983. This is important since the probabilistic forecasts are for annual changes.

2. After the third quarter of 1983, the survey intended to ask forecasters to report point forecasts for current and next year's GDP. However, the Philadelphia Fed has identified a few surveys between 1985 and 1990 in which it appears that forecasters were mistakenly asked about the previous and current year GDP instead. The Fed took over the survey in Quarter 2 of 1990 and is sure that none of the surveys since then have these errors.

3. The intervals in which respondents place probabilities have changed over the years. From Quarter 1 of 1983 through the end of 1991 there were 6 intervals, and after 1991 there were 10 intervals.

4. Prior to Quarter 3 of 1990, the SPF sometimes did not provide forecasters with figures for the previous year's GDP and GDP price index.[6]   We need these figures to convert point forecasts of levels into point forecasts of percentage changes, as described above.

 Because of these issues and inconsistencies, we restrict our analysis to the survey responses between Quarter 1 of 1992 and Quarter 4 of 2004. We also exclude Quarter 1 of 1996 since the previous-year values were unknwon at the time because of a delay in the release of the data caused by the federal government shutdown. As Table 1 demonstrates, even after this restriction our sample is large, with 3173 observations provided by 116 unique forecasters over the 13 year period.

---

[6]For example, the Quarter 1 survey of 1990 (which can be viewed at http://www.phil.frb.org/files/spf/form90.pdf) does not have previous year data

Table 1: Descriptive Statistics

Variable	Quarter	Observations	Missing Observations	Unique Forecasters	Mean Forecasters per Survey
GDP growth	1	365	46	91	30.4
GDP growth	2	430	49	105	33.1
GDP growth	3	414	47	103	31.8
GDP growth	4	414	54	97	31.8
Inflation	1	347	64	88	28.9
Inflation	2	412	67	100	31.7
Inflation	3	395	66	99	30.4
Inflation	4	396	72	94	30.5
ALL	-	3173	465	116	31.1

We count an observation as missing if (1) the forecaster does not provide a current year level point estimate or (2) the forecaster does not provide values for his subjective distribution.

3   Nonparametric Analysis

It is reasonable to conjecture that, when asked for point forecasts of quantities that are uncertain, forecasters respond with some measure of the central tendency of their subjective distributions. In probability theory, the three most prominent measures of central tendency are the mean, median, and mode of a distribution. The SPF data enable us to study the relationship between forecasters' point forecasts and these measures. Our analysis is in two parts. In this section, we perform a nonparametric analysis that does not assume subjective distributions to have any specific shape. In Section 4 we suppose that each subjective distribution has a Beta or isosceles-triangle shape, and we re-analyze the data from that perspective.

3.1   Bounding Means, Medians, and Modes

Recall that the SPF probabilistic questions ask forecasters to report their subjective probabilities that GDP growth and inflation will lie in given intervals on the real line. The responses to these questions do not fully reveal the subjective distributions that respondents hold, but they do bound these distributions. The data directly imply bounds on subjective means and medians. These bounds are most easily explained by way of examples. We assume throughout that the forecasters' subjective distributions are continuous.

Consider the SPF questions concerning inflation. Respondents are asked for ten subjective probabilities: the probability that prices will decline, the probabilities that the percentage inflation rate, denoted i, will lie in the interval [x, x+1) for x = 0, 1, . . . , 7 percent; and the probability that the inflation rate will be 8 percent or higher. Suppose that a forecaster gives these positive responses: $P(0 \leq i < 1) = 0.2$, $P(1 \leq i < 2) = 0.2$, $P(2 \leq i < 3) = 0.3$, $P(3 \leq i < 4) = 0.2$, $P(4 \leq i < 5) = 0.1$, with all other responses being zero. Then we can immediately conclude that this forecaster's subjective median for inflation lies in the interval [0.02, 0.03). Lower and upper bounds on his subjective mean are obtained by placing all of each

interval's probability mass at the interval's lower and upper endpoint, respectively. The resulting bound is [0.018, 0.028).

These bounds on the median and mean are finite intervals of width 0.01, but there are other examples in which the bounds are infinite. Consider a forecaster with strong bimodal expectations who states that $P(i < 0) = 0.4$ and $P(8 \leq i) = 0.6$. In this case, we can only conclude that the subjective median is larger than 0.08 and we cannot conclude anything at all about the subjective mean. Fortunately for our analysis, cases with bounds of infinite width occur only occasionally. Of the 3173 probabilistic forecasts that we observe, none yield a bound of infinite width on the forecaster's subjective median, and 363 yield a bound of infinite width on the subjective mean. In all other cases, the bound on the median and mean has width 0.01.

Using the SPF data to bound the mode of a subjective distribution is more subtle because the mode is a local property that cannot be inferred from interval data without imposing some assumption. Our analysis assumes that the mode is contained in the interval with the greatest probability mass. Thus, we conclude in the first example above that the mode lies in the interval [0.02, 0.03) and in the second that it lies in the interval [0.08, ∞). In some cases, multiple intervals have the same greatest probability mass; for example this would have occurred in the first example if the forecaster had placed probability 0.3 on both of the intervals [0.02, 0.03) and [0.03, 0.04). Of the 3173 forecasts, none yield a bound of infinite width on the forecaster's subjective mode and 243 yield a finite bound of width greater than 0.01. In all other cases, the bound on the mode has width 0.01.

Although the bounds cannot pinpoint the mean/median/mode of each subjective distribution, they typically have width 0.01 and, hence, are very informative. Examination of the bounds shows that in most forecasts, the three measures of central tendency are fairly close to one another. Calculations not reported here show that in three quarters of the cases, the distance separating the three measures must be no larger than 0.015.

3.2   Consistency of Point Forecasts with the Bounds

Having computed the above bounds, now consider a forecaster's point forecast of some quantity. If the point forecast lies within the bound for the median, then we cannot reject the hypothesis that the point forecast is the median. If the point forecast does not lie within the bound for the median, we can reject this hypothesis. The same reasoning applies to the mean and mode. Thus, we can determine the frequency with which point forecasts are and are not consistent with the three measures of central tendency.

Table 2 gives the findings, aggregated across the years 1992-2004 but disaggregated by quarter. Each entry in the table gives the percentage of cases in which a panel member's point forecast of a given quantity in a given quarter lies within his bound for a given measure of central tendency. For example, the upper left entry for GDP growth, which is 83.01%, means that 83.01% of all first-quarter GDP point forecasts lay within the bound obtained for the forecaster's subjective mean. We call this an "Upper Bound on Mean

Providers" because consistency of a point forecast with a bound does not imply that the forecaster must have given his subjective mean as his point forecast. It only implies that he may have given the subjective mean.

Table 2: Upper Bounds on mean/median/mode Providers

	GDP Growth			
	Quarter 1	Quarter 2	Quarter 3	Quarter 4
Mean	83.01%	82.09%	88.89%	93.72%
Median	75.34%	77.44%	87.68%	89.86%
Mode	85.21%	86.74%	89.86%	92.03%

	Inflation			
	Quarter 1	Quarter 2	Quarter 3	Quarter 4
Mean	82.42%	84.22%	84.30%	88.89%
Median	79.25%	82.29%	84.81%	87.63%
Mode	84.44%	85.68%	87.09%	89.90%

The table shows that most point forecasts are consistent with the hypotheses that SPF panel members report their subjective means, medians, or modes. However, there are many panel members whose point forecasts are inconsistent with these hypotheses. This is especially evident in Quarter 1, where (17%, 25%, 15%) percent of the respondents give point forecasts that cannot be their subjective mean/median/mode.

Observe that, in each row of the table, the entries in the table increase markedly from Quarter 1 to Quarter 4. This is reasonable to expect, because the forecast horizon shrinks as the year goes on – while most of the current year lies ahead in the Quarter 1 surveys, most of it lies behind when the Quarter 4 surveys are conducted. Thus, the SPF forecasters should tend to have sharper subjective distributions in Quarter 4 than in Quarter 1, implying that alternative measures of central tendency should draw nearer to one another.

Table 3 shows that subjective distributions do tend to sharpen as the year goes on. For each quarter and quantity of interest, the table shows the percentage of forecasters whose positive probability assessments are concentrated in N or less intervals, where N ranges from 1 to 10. Observe how the entries increase with Quarter. Consider, for example, the column for N = 2 when the quantity is GDP growth. In Quarter 1, only 6.9% of the forecasters have subjective distributions that are concentrated in two or less intervals. In Quarter 4, beliefs have sharpened so much that 58.9% of the subjective distributions are this concentrated.

GDP Growth										
	1	2	3	4	5	6	7	8	9	10
Quarter 1	0.0%	6.9%	32.3%	54.0%	71.8%	81.9%	90.4%	92.9%	95.1%	100.0%
Quarter 2	0.5%	15.1%	43.5%	62.1%	78.1%	85.1%	90.0%	93.5%	95.4%	100.0%
Quarter 3	1.9%	22.9%	62.3%	77.0%	86.2%	90.3%	94.2%	96.1%	96.8%	100.0%
Quarter 4	11.1%	58.9%	82.6%	91.1%	94.7%	96.1%	97.3%	97.8%	98.5%	100.0%

Inflation										
	1	2	3	4	5	6	7	8	9	10
Quarter 1	0.6%	15.0%	51.0%	74.1%	86.5%	91.9%	97.1%	97.4%	98.0%	100.0%
Quarter 2	1.7%	21.9%	55.4%	74.8%	85.7%	92.2%	94.9%	96.6%	98.3%	100.0%
Quarter 3	4.1%	33.4%	70.1%	84.1%	92.7%	94.9%	97.7%	98.0%	98.7%	100.0%
Quarter 4	14.7%	59.6%	84.6%	94.7%	97.0%	97.7%	98.7%	99.0%	99.0%	100.0%

3.3 Inconsistencies Tend To Present Favorable Scenarios

Now consider the SPF panel members whose point forecasts are not consistent with their subjective means, medians, or modes. Table 4 report the percentage of such cases in which the point forecast lies below or above the bound. A clear finding emerges. Most such point forecasts give a view of the economy that is favorable relative to the central tendencies of respondents' subjective distributions. Thus, forecasters who skew their point forecasts tend to present rosy scenarios.

The table shows that, when forecasting GDP, the point forecasts that are inconsistent with measures of central tendency are much more often above the upper bounds than below the lower bounds on means, medians, and modes. Symmetrically, when forecasting inflation, the point forecasts are much more often below the lower bounds than above the upper ones. We do not know why forecasters skew their point forecasts in this way. One might conjecture that the answer lies in strategic consideration or herd phenomena. However, panel members ostensibly are unaware of each others forecasts when they respond to the survey and individual forecasts are not later identified in the public release of the SPF.

Table 4: Evidence of Favorable Point forecasts

GDP Growth			
	N	Above Bound	Below Bound
Mean	211	73.46%	26.54%
Median	280	73.93%	26.07%
Mode	186	58.60%	41.40%

Inflation			
	N	Above Bound	Below Bound
Mean	232	6.47%	93.53%
Median	254	22.44%	77.56%
Mode	204	12.25%	87.75%

3.4 Persistence of Favorable (Unfavorable) Inconsistencies

The SPF attaches an ID number to each forecaster, so we are able to analyze the behavior of individual forecasters across time. We find that forecasters show some persistence in providing favorable (unfavorable) point forecasts relative to their beliefs: a forecaster whose point forecast for inflation (GDP growth) lies above the upper bound for his mean/median/mode one period is more likely to provide a point forecast above the upper bound for his mean/median/mode in later periods than is a forecaster whose point forecast is within or below the bounds for his mean/median/mode. This persistence is also found for forecasters whose point forecasts lie below the lower bound for their means/medians/modes.

Table 5 presents evidence of short term persistence, comparing forecasts in adjacent quarters. The table considers point forecasts in relation to the bounds for the mean; the results for the median and mode are similar and are omitted for brevity. Table 6 presents evidence of longer term persistence. Here we compare forecasts in Quarter t with those made one, two, and three years later; that is, in Quarters $t + 4k$, $k = 1, 2, 3$.

Table 5: Short Term Persistence (Mean)

GDP Growth

| | | Quarter t+1 | | |
Quarter t	N	Above Bound	Inside Bound	Below Bound
Above Bound	113	24.8%	73.5%	1.8%
Inside Bound	1036	8.5%	89.2%	2.3%
Below Bound	48	6.3%	75.0%	18.8%

Inflation

| | | Quarter t+1 | | |
Quarter t	N	Above Bound	Inside Bound	Below Bound
Above Bound	10	10.0%	80.0%	10.0%
Inside Bound	966	1.1%	89.2%	9.6%
Below Bound	150	0.0%	67.3%	32.7%

Table 6: Long Term Persistence (Mean)

GDP Growth

| | | Quarter t + 4 | | | | Quarter t + 8 | | | | Quarter t + 12 | | |
Quarter t	N	Above	Inside	Below	N	Above	Inside	Below	N	Above	Inside	Below
Above Bound	92	20.7%	77.2%	2.2%	73	21.9%	78.1%	0.0%	62	25.8%	74.2%	0.0%
Inside Bound	907	8.2%	89.6%	2.2%	707	8.5%	89.3%	2.3%	560	7.3%	90.2%	2.5%
Below Bound	35	14.3%	80.0%	5.7%	25	16.0%	80.0%	4.0%	17	11.8%	88.2%	0.0%

Inflation

| | | Quarter t + 4 | | | | Quarter t + 8 | | | | Quarter t + 12 | | |
Quarter t	N	Above	Inside	Below	N	Above	Inside	Below	N	Above	Inside	Below
Above Bound	5	0.0%	100.0%	0.0%	7	0.0%	100.0%	0.0%	4	0.0%	100.0%	0.0%
Inside Bound	832	0.8%	88.0%	11.2%	664	0.5%	90.1%	9.5%	539	0.7%	90.4%	8.9%
Below Bound	125	0.8%	80.8%	18.4%	84	1.2%	84.5%	14.3%	59	0.0%	83.1%	16.9%

4 Parametric Analysis

The analysis of Section 3 has the great advantage of using the SPF data alone, without any assumptions about the shapes of forecasters' subjective distributions for GDP growth and inflation. The accompanying disadvantage is that we can draw only limited conclusions about the relationship of point forecasts to probabilistic beliefs. Imposing assumptions enables sharper empirical analysis, albeit subject to the credibility of the assumptions imposed.

In this section, we report a parametric analysis whose most basic assumption is that probabilistic beliefs are unimodal. When an SPF probabilistic forecast assigns positive probability to three or more intervals, we furthermore assume that the subjective distribution is a member of the generalized Beta family and use the data to fit the Beta parameters. The generalized Beta distribution, which uses two parameters to describe the shape of beliefs and two more to give their support, is a flexible form that permits a distribution to have different values for its mean, median, and mode.

It is possible to fit a unique Beta distribution to the SPF data only when a forecast assigns positive probability to at least three intervals. This occurs in 2234 of the 3173 forecasts that we observe. The remaining 939 forecasts only assign positive probability to one bounded interval or to two adjacent intervals. In these cases, we assume instead that the subjective distribution has the shape of an isosceles triangle whose base includes all of the interval with greater probability mass and part of the other interval, if there is one.[7] This assumption gives one parameter to be fit, which fixes the center and height of the triangle. An isosceles triangle is symmetric, so the mean, median, and mode of the fitted subjective distribution necessarily coincide in these cases. This feature of the fit is not much of a practical concern because the actual means, medians, and modes of forecasts that lie entirely within one or two intervals must lie relatively close to one another in any event.

Section 4.1 describes the fitting methods in more detail. Section 4.2 compares the point forecasts with the fitted probability distributions.

4.1 Fitting Methods

Case 1 - The forecaster uses 1 bounded interval: We assume that the subjective distribution takes the shape of an isosceles triangle whose support is the interval. For example, if a forecaster places all probability in the interval [3%, 4%], we assume that the support of his distribution is [.03, .04]. Then the base of the triangle has its center at .035 and has length .01. The height of the triangle is 200, yielding area equal to 1.

Case 2 - The forecaster uses 2 intervals: Every forecaster who uses exactly 2 intervals uses adjacent intervals. If the forecaster places equal probability in these intervals, we assume the support is the union of the intervals and we fit an isosceles triangle whose base has length .02 and whose height is 100. If a forecaster places more probability in one interval than the other, we assume that the support of the subjective

---

[7]The SPF data contain no cases in which a probabilistic forecast assigns all positive probability to two non-adjacent intervals. Nor does it contain cases in which the majority of the probability mass lies in an unbounded tail interval. Hence, we do not need a fitting method for these situations.

distribution contains the entirety of the more probable interval. This restricts one endpoint of the support. With this restriction and our assumption that the subjective distribution is an isosceles triangle, we are able to completely specify the distribution: Suppose that a forecaster places probability α and $1 - \alpha$ in the intervals $[y\%, (y+1)\%]$ and $[(y+1)\%, (y+2)\%]$ respectively, where $\alpha < \frac{1}{2}$. Letting $t = \frac{\sqrt{\frac{\alpha}{2}}}{1-\sqrt{\frac{\alpha}{2}}}$, it is straightforward to show that the isosceles triangle with height $\frac{200}{t+1}$ and endpoints $(y+1-t)\%$ and $(y+2)\%$ defines a subjective probability density function that is consistent with the respondent's beliefs.[8] Figure 1 illustrates our construction.

Figure 1: Illustration of 2 Interval Case
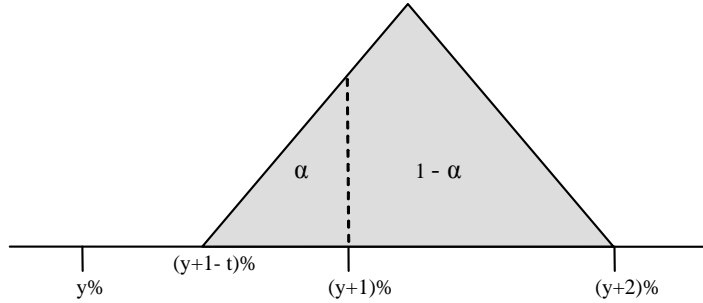
Case 3 - The forecaster uses 3 or more intervals: In general, the probabilities that a forecaster reports for the ten intervals in an SPF forecast reveal points on the cumulative distribution function (CDF) of his beliefs. For example, if a forecaster reports a 0.3 chance that GDP growth will be in the interval $[2\%, 3\%)$, a 0.6 chance for the interval $[3\%, 4\%)$ and a 0.1 chance for the interval $[4\%, 5\%)$, we can infer these points on the forecaster's CDF: $F(.02) = 0$, $F(.03) = 0.3$, $F(.04) = 0.9$, and $F(.05) = 1$. When the forecast places positive probability on three or more intervals, we fit a unimodal generalized Beta distribution to the data. The generalized beta distribution has this CDF:

$$Beta(t; a, b, l, r) = \left\{ \begin{array}{ll} 0 & \text{if } t \leq l \\ \frac{1}{B(a,b)} \int_l^t \frac{(x-l)^{a-1}(r-x)^{b-1}}{(r-l)^{a+b-1}} dx & \text{if } l < t \leq r \\ 1 & \text{if } t > r \end{array} \right\}$$

where $B(a,b) = \frac{\Gamma(a)\Gamma(b)}{\Gamma(a+b)}$ and $\Gamma(a) = \int_0^\infty x^{a-1}e^{-x}dx$. We choose a generalized Beta distribution because the two shape parameters a and b give considerable flexibility and the two location parameters l and r allow us to specify the support of the distribution. (The standard Beta distribution assumes that the support is the interval $[0,1]$). To enforce unimodality, we maintain the restriction that $a > 1$ and $b > 1$.

Consider a forecaster whose observed points on his CDF are $F(t_1), ...F(t_{10})$, where $t_1, ...t_{10}$ are the right endpoints of the ten intervals. If the forecaster does not place positive probabilities on the two tail intervals,

[8]To see this, note that $h = \frac{2}{t+1}$ and $t = \frac{\sqrt{\alpha/2}}{1-\sqrt{\alpha/2}}$ solve the equations $\frac{1}{2}(t+1)h = 1$ and $\frac{1}{2}t[h\frac{t}{(t+1)/2}] = \alpha$. The left hand side of the first equation gives the area of the entire isosceles triangle. The left hand side of the second equation gives the area of the left sub-triangle in the interval with smaller probability mass.

we take the support of the distribution to be the left and right endpoints of the intervals with positive probability. This fixes l and r. For example, if a forecaster only places mass in the [2%, 3%), [3%, 4%0, and [4%, 5%) intervals, then $l = .02$ and $r = .05$ We then minimize over the shape parameters a and b:

$$\min_{a>1,b>1} \sum_{i=1}^{10} [Beta(t_i; a, b, .02, .05) - F(t_i)]^2$$

When a forecaster places mass in the lower tail interval, which is unbounded from below, we let l be a free parameter in the minimization. Likewise, if a forecaster places mass in the upper tail, we let r be a free parameter in the minimization. However, we restrict the support parameters l and r to lie within the most extreme values that have actually occurred in the United States since 1930. Thus, for change in GDP, we restrict l and r to the range $-.13 < l < r < .19$. For inflation, we restrict l and r to the range $-.12 < l < r < .12$. For example, if a forecaster reports a 0.3 chance that GDP growth will be in the interval [4%, 5%], a 0.6 chance for the interval [5%, 6%] and a 0.1 chance for the interval [6%, ∞%), the minimization problem becomes:

$$\min_{a>1,b>1,r<.19} \sum_{i=1}^{n} [Beta(t_i; a, b, .04, r) - F(t_i)]^2$$

4.2 Comparing the Point Forecasts with the Fitted Probability Distributions

Having fitted a CDF to each SPF probabilistic forecast, we can compare the point forecasts with the fitted CDFs. This is accomplished in Figures 2 and 3, the former for GDP growth and the latter for inflation. In each figure, the x-axis gives a percentile of the fitted probability distribution, ranging from 0 to 100. The left-hand y-axis gives the number of cases in which a point forecast equals that percentile and the right-hand y-axis gives the height of a kernel density forecast fitted to the empirical distribution of cases.[9] For example, Figure 2 shows that, of the 1623 point forecasts of GDP growth, 9 were at the 25th percentile of the fitted probability distribution, 51 were at the 50th percentile, and 14 were at the 75th percentile.

Figure 2 shows that forecasters tend to report point forecasts for GDP growth that are high percentiles of the fitted probability distributions. In contrast, Figure 3 shows that most forecasters report point forecasts for inflation that are low percentiles. Overall, 41.16% of point estimates were equal to or below their fitted medians for the GDP growth questions while 71.16% of point estimates were equal to or below their fitted medians for the inflation questions. These results corroborate our earlier nonparametric finding, discussed in Section 3.3, that forecasters tend to provide favorable point forecasts relative to their probabilistic beliefs.

---

[9] When a forecaster reports a point estimate off the lower (upper) bound of his support, we place this observation in the 0 (100) percentile bin. In all but a few cases, the point forecast was within the support of the fitted distribution. For the output question, the point forecast lay outside the support in only 18 of 1623 forecasts. For the inflation question, this occurred in 27 of the 1550 forecasts.

The kernel density estimate uses Silverman's rule of thumb to compute the bandwidth.

Moreover, we find that the distance between the point forecast and measures of central tendency are positively correlated with forecaster uncertainty. Letting $DIST$ be the distance between a forecaster's point forecast and his fitted median and IQR be the interquartile range of the fitted distribution, the correlation between $DIST$ and IQR is .25 in the GDP growth case and .29 in the inflation case. This correlation seems reasonable if forecasters are giving favorable percentiles of their subjective distributions as point estimates; for example, if a forecaster always gave the 90th percentile of his distribution as a point forecast the distance between this point forecast and a measure of central tendency will shrink as the the distribution tightens.

Figure 2: Empirical Distribution of Point-Forecast Percentiles (GDP Growth)
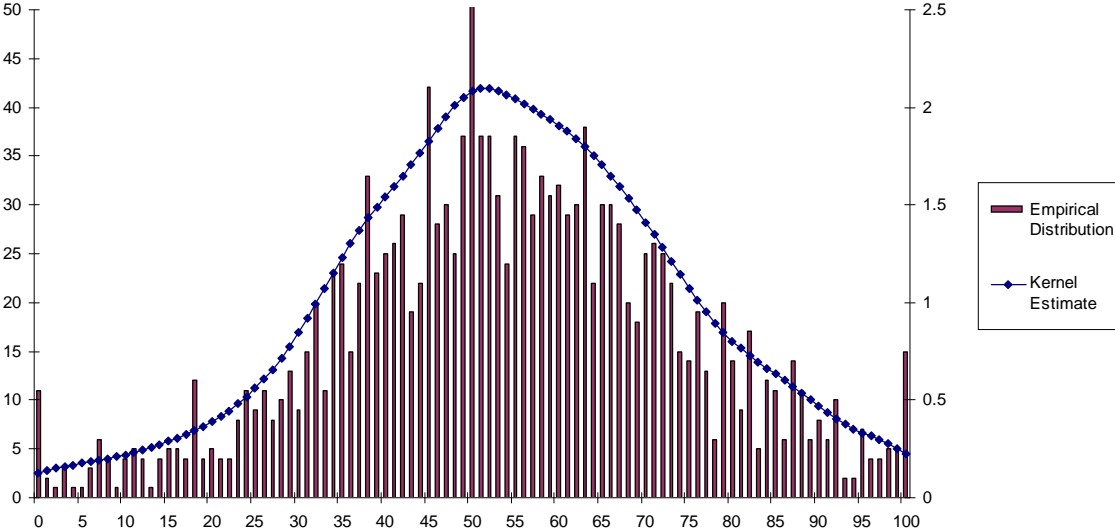
Figure 3: Empirical Distribution of Point-Forecast Percentiles (Inflation)
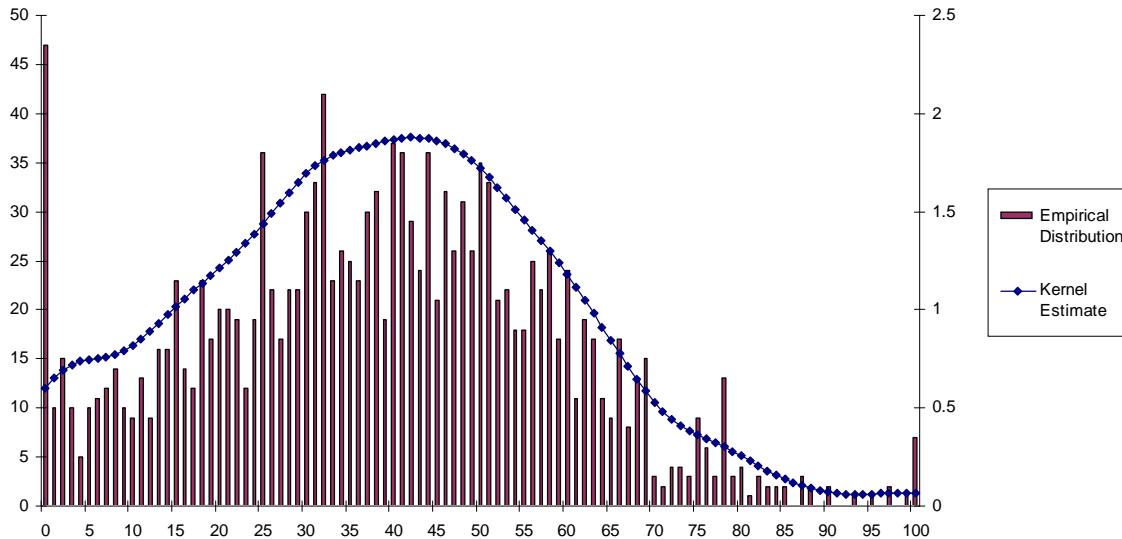
In Section 3.4, we showed nonparametrically that SPF forecasters exhibit some time-series persistence, with those who give favorable point forecasts in one quarter tending to do the same in subsequent quarters. Table 7 examines persistence parametrically, by calculating the serial correlation of the point-forecast percentiles. We find considerable short-term persistence but less long-term persistence, especially in the inflation forecasts.

Table 7: Persistence of Point-Forecast Percentiles

	GDP Growth							
	Quarter $t + 1$		Quarter $t + 4$		Quarter $t + 8$		Quarter $t + 12$	
Quarter t	N	Correlation	N	Correlation	N	Correlation	N	Correlation
Point-Forecast Percentile	1197	0.2553	1034	0.0847	805	0.1196	639	0.1490

	Inflation							
	Quarter $t + 1$		Quarter $t + 4$		Quarter $t + 8$		Quarter $t + 12$	
Quarter t	N	Correlation	N	Correlation	N	Correlation	N	Correlation
Point-Forecast Percentile	1126	0.2157	962	0.0496	755	0.0243	602	0.0144

5 Measuring Forecaster Uncertainty and Disagreement

The analysis of Sections 3 and 4 has shown that the point forecasts of SPF forecasters tend to be favorable relative to their probabilistic beliefs. We have also shown evidence of time-series persistence in the way that particular forecasters form their point forecasts. These findings suggest that elicitation of subjective

distributions provides clearer insights into forecasters' beliefs than does collection of point forecasts. The argument for probabilistic measurement of expectations is further enhanced when one recognizes that point forecasts provide no sense of the uncertainty that forecasters feel.

Indeed, present practices in reporting SPF data conflate forecaster uncertainty and disagreement. In the quarterly reports released by the Philadelphia Fed, the SPF data are aggregated across forecasters in two ways: (1) for the questions that elicit point forecasts, the median point forecast is reported and (2) for the questions that elicit distributions by having forecasters place probability mass in 10 intervals, the mean probability mass in each interval is reported. The median point forecast reveals no information about forecaster uncertainty. Reporting the mean probability mass in each interval mixes together forecaster uncertainty and disagreement. This fact has previously been recognized by Giordani and Soderlind (2003).

To illustrate, suppose that there are two forecasters, labeled A and B, and consider two scenarios. In one scenario, Forecaster A places all probability mass in the GDP growth interval [2%,3%) while Forecaster B places all probability mass in the GDP growth interval [3%, 4%). In the other scenario, both forecasters place half their probability mass in the interval [2%,3%) and half in the interval [3%,4%). In the first scenario, the forecasters are individually quite certain about the outcome but they completely disagree with one another. In the second scenario, the forecasters are individually uncertain about the outcome but they completely agree. Reporting the mean probability mass in each interval makes these two scenarios indistinguishable.

We think it important to separately report forecaster uncertainty and disagreement. In this section we show how to do so. Section 5.1 uses the probability distributions fitted in Section 4 to describe the uncertainty that SPF forecasters feel. Section 5.2 uses the fitted distributions to describe disagreements in the central tendency of forecasts. Finally, Section 5.3 jointly portrays the uncertainty and central tendency of SPF forecasts.

5.1   Measuring Uncertainty

The interquartile range (IQR) of a forecaster's fitted subjective probability distribution provides an informative scalar measure of the uncertainty that this forecaster perceives. Forecasters vary in the uncertainty they feel so, at any point in time, the IQR varies across the members of the SPF panel. A simple way to summarize the cross-sectional distribution of forecaster uncertainty is to calculate the cross-sectional lower quartile, median, and upper quartile of the IQR. This is done in Figures 4 and 5 for GDP growth and inflation respectively.

To illustrate, consider the entries in Figure 4 for the first quarter of 2004. The figures shows the cross-sectional lower quartile of IQR to be 0.011, the median to be 0.013 and the upper quartile to be 0.017. This means that, in the first quarter of 2004, the fitted subjective probability distributions of (25, 50, 75) percent of the then-active SPF forecasters had an IQR no larger than (0.011, 0.013, 0.017) respectively. The first quarter of other years shows similar variation.

Observe that, within each year, the IQRs of SPF forecasters tend to fall sharply as the year progresses

from quarter 1 to quarter 4. In the nonparametric analysis of Section 3.2, we reported (see Table 3) that subjective distributions tend to sharpen as the year goes on. Using the fitted probability distributions, Figures 4 and 5 show that this phenomenon occurs each and every year. The reason presumably is that the forecast horizon shrinks as the year progresses.

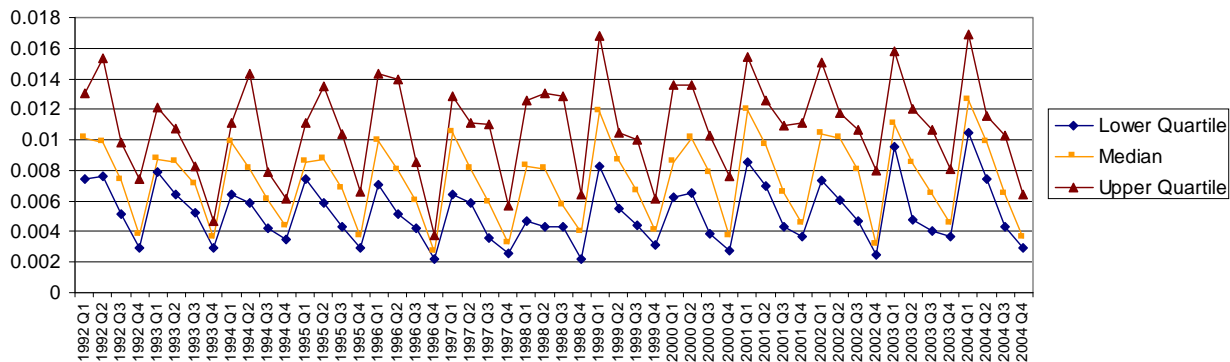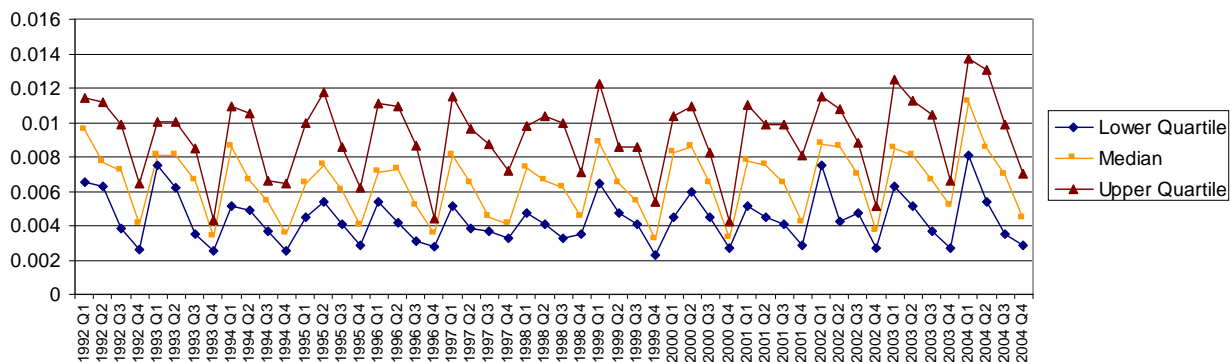**Figure 4: Lower Quartile, Median and Upper Quartile of IQR (GDP growth)**
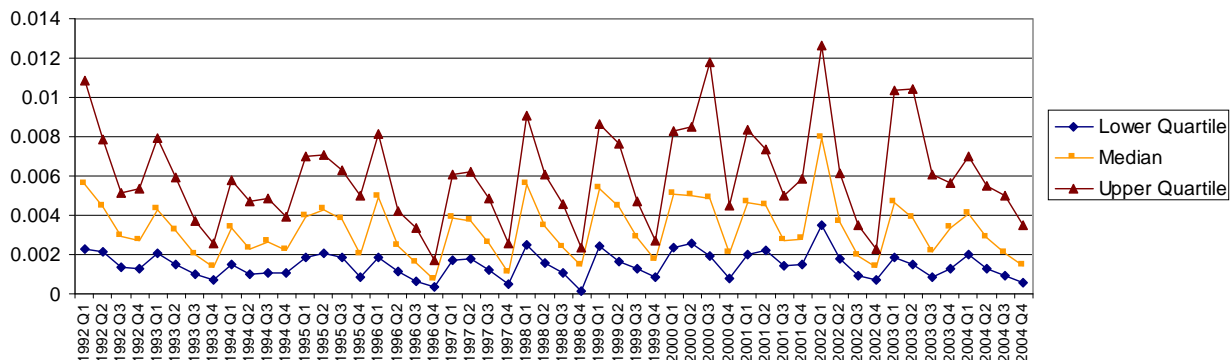

Figure 5: Lower Quartile, Median and Upper Quartile of IQR (Inflation)

## 5.2 Measuring Disagreement in Central Tendency

For the sake of specificity, let us use the median of a fitted subjective probability distribution to measure the central tendency of an SPF forecast. Then the absolute difference in medians (ADM) of two forecasts measures the degree to which these forecasts disagree in central tendency. If N SPF forecasters provide a forecast in a given quarter, there are $\binom{N}{2}$ pairs of medians to be compared. A simple way to summarize the extent to which forecasts disagree in central tendency is to calculate the cross-sectional lower quartile, median, and upper quartile of the ADM. This is done in Figures 6 and 7 for GDP growth and inflation respectively.

Observe that the majority of disagreements in central tendency are smaller than 0.01 and few exceed 0.012. Within each year, the ADMs of SPF forecasters tend to fall as the year progresses from quarter 1 to quarter 4. Again, the reason presumably is that the forecast horizon shrinks as the year progresses.

Figure 6: Lower Quartile, Median and Upper Quartile of ADM (GDP Growth)

Figure 7: Lower Quartile, Median and Upper Quartile of ADM (Inflation)
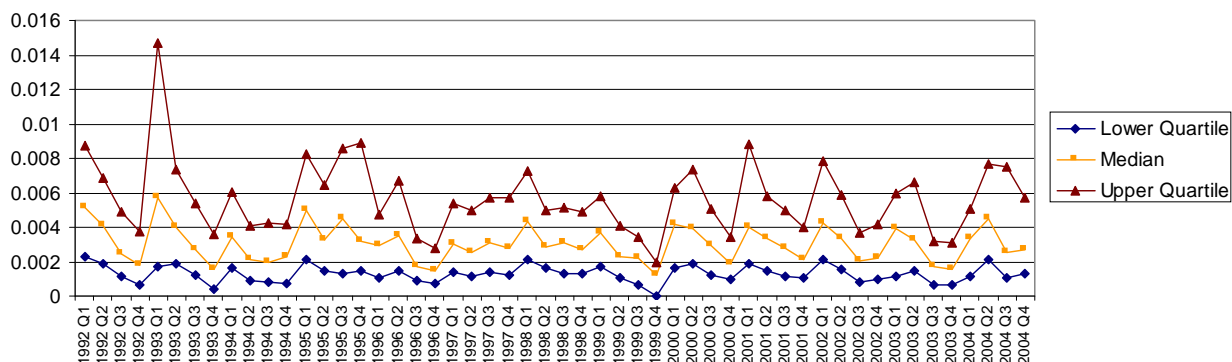
5.3 Joint Portrayal of the Central Tendency and Uncertainty of SPF Forecasts

We have thus far discussed separately the central tendency and uncertainty of SPF probabilistic forecasts. When aggregating the beliefs of forecasters, we recommend joint portrayal of these two basic features of the fitted subjective probability distributions. This can be done in a two-dimensional plot having the median on the x-axis and the IQR on the y-axis. To illustrate, Figures 8 and 9 present such plots for forecasts of inflation in the third quarter of 1999 and the first quarter of 2002.

The plots are simple to interpret. Each point in a plot represents a unique forecaster. When the points cluster towards the top, forecasters tend to feel much uncertainty. When the points are dispersed horizontally, disagreement in the central tendency of forecasts is high. Comparing the 1999 Q3 and 2002

Q1 plots, we see that the cross-sectional distribution of uncertainty is similar in these two periods, but disagreement in central tendency is more pronounced in the latter period.

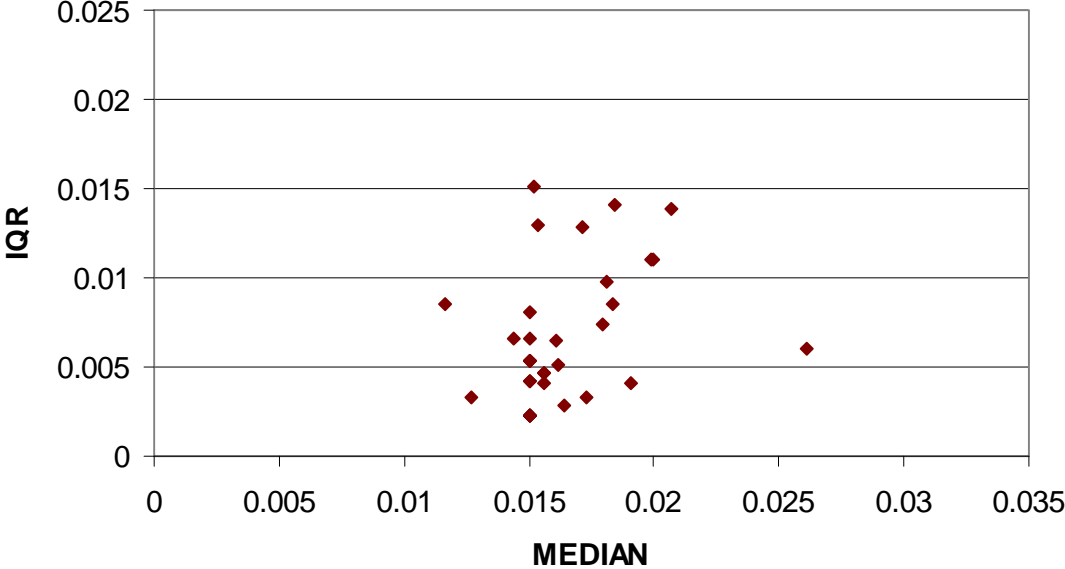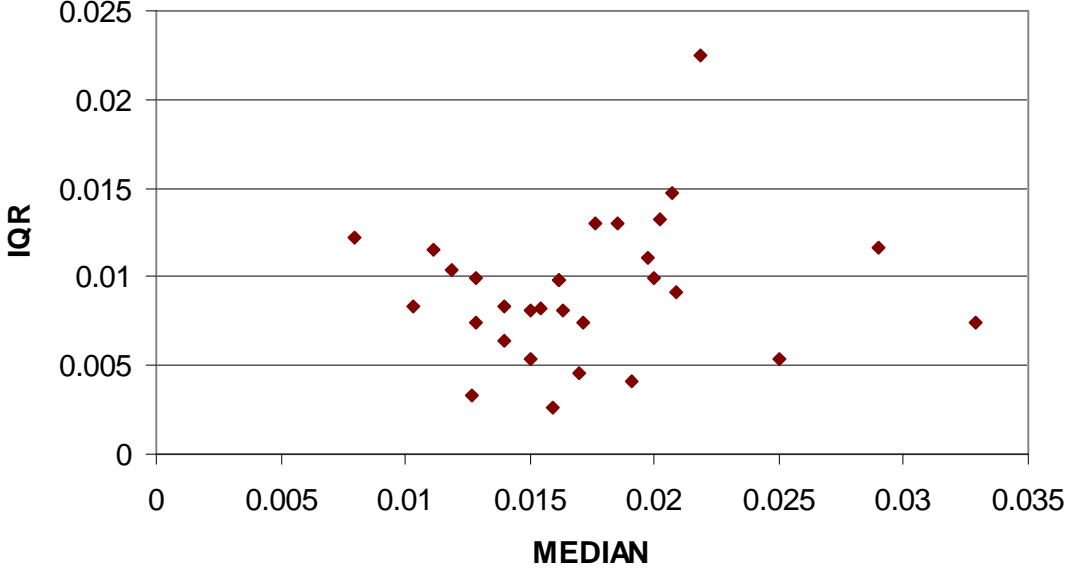**Figure 8: Median/IQR Plot for 1999 Q3 Inflation Forecasts**

Figure 9: Median/IQR Plot for 2002 Q1 Inflation Forecasts

6   Conclusion

If people think probabilistically, as economists generally assume, their point forecasts should somehow be related to their underlying subjective distributions. We have shown that SPF forecasters summarize their underlying distributions in different ways and that their summaries tend to be favorable relative to the central tendency of the underlying distributions.

This finding has implications for research relying on point forecasts. For example, the accounting literature has documented a positive relationship between the optimism of analysts' earnings forecasts and the forecast horizon.[10] In particular, longer horizon earnings forecasts tend to be higher than the realized earnings, but the bias disappears for shorter horizon forecasts; in fact, there is some evidence that short horizon forecasts are slightly pessimistic. Some authors, for example Richardson, Teoh, and Wysocki (2004), have argued that the observed pattern of bias reflects firms' efforts to "guide" analysts' forecasts down to beatable levels as the horizon shortens. Our findings suggest another explanation. It could be that analysts have rational expectations with respect to future earnings but that their point forecasts are simply high percentiles of their underlying distributions. Recall from Section 4.2 that the distance between SPF point forecasts and subjective medians are positively correlated with forecaster uncertainty. Presumably, analysts become more certain of their forecasts as the earnings announcement date approaches. Hence, we should expect to find the positive relationship between point forecast optimism and forecast horizon that is documented in the accounting literature.

Our findings also have important implications outside academia. It is standard practice for experts to express their beliefs and predictions using point forecasts.[11] These point forecasts may have significant consequences: the Congressional Budget Office's (CBO) (www.cbo.gov) point forecasts for the costs of legislative bills may affect legislators' voting decisions, and financial analysts' earnings forecasts and price targets may affect people's investment decisions. The SPF evidence suggests that point forecasts may have a systematic, favorable bias. This, plus the inescapable fact that point forecasts reveal nothing about the uncertainty that forecasters feel, suggest that the agencies who commission forecasts should not ask for point forecasts. Instead, they should elicit probabilistic expectations and derive measures of central tendency and uncertainty, as we have done here.

---

[10]Ramnath, Rock, and Shane (2005) provide a review of the literature.

[11]One notable exception is the Bank of England (www.bankofengland.co.uk), which offers "fan charts" for its projections of future inflation and GDP. These fan charts display the 10th, 20th, ..., 90th percentile from the Bank's probabilistic forecasts for future inflation and GDP.

.

References

[1] Ajinkya, Bipin B., Rowland K. Atiase, and Michael J. Gift (1991): "Volume of Trading and the Dispersion in Financial Analysts' Earnings Forecasts," *The Accounting Review*, 66, 389-401.

[2] Bamber, Linda Smith, Orie E. Barron, and Thomas L. Stober (1997): "Trading Volume and Different Aspects of Disagreement Coincident with Earnings Announcements," *The Accounting Review*, 72, 575-597.

[3] Barron, Orie E (1995): "Trading Volume and Belief Revisions that Differ Among Individual Analysts," *The Accounting Review*, 70, 581-597.

[4] Bassett, William F. and Robin L. Lumsdaine (2001): "Probability Limits: Are Subjective Assessments Adequately Accurate?" *The Journal of Human Resources*, 36, 327-363.

[5] Bomberger, William A (1996): "Disagreement as a Measure of Uncertainty," *Journal of Money, Credit, and Banking*, 28, 3, 381-392.

[6] Brenner, Menachem and Yoram Landskroner (1983): "Inflation Uncertainties and Returns on Bonds," *Economica*, 50, 463-468.

[7] Butler, Kirt C. and Larry H. P. Lang (1991): "The Forecast Accuracy of Individual Analysts: Evidence of Systematic Optimism and Pessimism," *Journal of Accounting Research*, 29, 150-156.

[8] Cukierman, Alex and Paul Wachtel (1979): "Differential Inflationary Expectations and the Variability of the Rate of Inflation: Theory and Evidence," *American Economic Review*, 69, 595-609.

[9] Diether, Karl B., Christopher J. Malloy, and Anna Scherbina (2002): "Differences of Opinion and the Cross Section of Stock Returns," *The Journal of Finance*, 57, 2113-2141.

[10] Giordani, Paolo and Paul Soderlind (2003a): "Inflation Forecast Uncertainty," *European Economic Review*, 47, 1037-1059.

[11] Giordani, Paolo and Paul Soderlind (2003b): "Is there Evidence of Pessimism and Doubt in Subjective Distributions? A Comment on Abel," *A Research Report from Stockholm Institute for Financial Research.*

[12] Granger, Clive and Ramu Ramanathan (1984): "Improved Methods of Combining Forecasts," *Journal of Forecasting*, 3, 197-204.

[13] Hahm, Joon-Ho and Douglas Steigerwald (1999): "Consumption Adjustment under Time-Varying Income Uncertainty," *The Review of Economics and Statistics*, 81, 32-40.

[14] Hayford, Marc (2000): "Inflation Uncertainty, Unemployment Uncertainty, and Economic Activity," *The Journal of Macroeconomics*, 22, 315-329.

[15] Juster T. (1966), "Consumer Buying Intentions and Purchase Probability: An Experiment in Survey Design," *Journal of the American Statistical Association*, 61, 658-696.

[16] Lahiri, Kajal, Christie Teigland, and Mark Zaporowski (1988): "Interest Rates and the Subjective Probability Distribution of Inflation Forecasts," *Journal of Money, Credit, and Banking*, 20, 2, 233-248.

[17] Levi, Maurice and John H. Makin (1979): "Fisher, Phillips, Friedman and the Measured Impact of Inflation on Interest," *Journal of Finance*, 34, 1, 35-52.

[18] Levi, Maurice and John H. Makin (1980): "Inflation Uncertainty and the Phillips Curve: Some Empirical Evidence," *American Economic Review*, 70, 5, 1022-1027.

[19] Makin, John (1982): "Anticipated Money, Inflation Uncertainty, and Real Economic Activity," *The Review of Economics and Statistics*, 64, 126-134.

[20] Mankiw, N. Gregory, Ricardo Reis, and Justin Wolfers (2003): "Disagreement about Inflation Expectations," *NBER Macroeconomics Annual* 2003.

[21] Manski, Charles F. (2004): "Measuring Expectations," *Econometrica*, 72, 1329-1376.

[22] Manski, Charles F. (1990): "The Use of Intentions Data to Predict Behavior: A Best Case Analysis," *Journal of American Statistical Association*, 85, 412, 934-940.

[23] Mullineaux, Donald (1980): "Unemployment, Industrial Production, and Inflation Uncertainty in the United States," *The Review of Economics and Statistics*, 62, 163-169.

[24] Ramnath, Sundaresh, Steve Rock, and Philip Shane (2005): "A Review of Research Related to Financial Analysts' Forecasts and Stock Recommendations," Working Paper.

[25] Rich, R. W., J. E. Raymond, and J. S. Butler (1992): "The Relationship Between Forecast Dispersion and Forecast Uncertainty: Evidence from a Survey Data–ARCH Model," *Journal of Applied Econometrics*, 7, 131-148.

[26] Rich, Robert, and Joseph Tracy (2004): "Uncertainty and Labor Contract Durations," *The Review of Economics and Statistics*, 86, 270-287.

[27] Richardson, Scott, S.H. Teoh, and Peter Wysocki (2004): "The Walkdown to Beatable Analyst Forecasts: The Role of Equity Issuance and Insider Trading Incentives" *Contemporary Accounting Research*, 21 (4), pp 885-924.

[28] Zarnowitz, Victor and Louis A. Lambros (1987): "Consensus and Uncertainty in Economic Prediction," *Journal of Political Economy*, 95, 3, 591-621.