Asymmetry in Unemployment Rate Forecast Errors

John W. Galbraith[a,c,d], Simon van Norden[b,c,d,*]

[a] *McGill University*
[b] *HEC Montréal*
[c] *CIRANO*
[d] *CIREQ*

Abstract

Asymmetries in unemployment dynamics have been observed in the time series of a number of countries, including the United States. This paper studies asymmetries in unemployment rate forecast errors. We consider conditions under which optimal forecasts will display asymmetrically-distributed errors and how the degree of asymmetry might vary with forecast horizon. Using data from the U.S. Survey of Professional Forecasters and the Federal Reserve Greenbook, we find substantial evidence of forecast error asymmetry, which tends to increase with the forecast horizons; we also find noteworthy differences in forecasts from these two sources. The results give insight into the ability of professional forecasters to adapt their forecasts to asymmetry in underlying processes.

Keywords:
asymmetry, forecast efficiency, Survey of Professional Forecasters, FOMC, Greenbook, unemployment rate

*Corresponding author. We thank Todd Clark, Domenico Giannone and participants in the St. Louis Fed Conference on Central Bank Forecasting for valuable comments on earlier drafts, and Mathieu Roussel-Lewis for excellent research assistance. CIRANO *(Centre Interuniversitaire de recherche en analyse des organisations)* and *Fonds québécois de la recherche sur la société et la culture* (FQRSC) provided facilities or financial support of this research. The plots of recessions and unemployment data are taken from the data and plotting tools of the Federal Reserve Bank of St. Louis FRED database and web site.

Email addresses: john.galbraith@mcgill.ca (John W. Galbraith), simon.van-norden@hec.ca (Simon van Norden)

URL: johngalbraith-economics.mcgill.ca (John W. Galbraith), svannorden.org (Simon van Norden)

1. Introduction

This paper examines asymmetries in U.S. unemployment rate forecast errors and how they vary across forecast horizons.

Asymmetry in business cycles has long been observed. Quantifying and forecasting the downside risks for economic activity is of critical importance for central banks and other macroeconomic policymakers.[1] The unemployment rate is in many ways an ideal candidate for quantifying such risks. It is everywhere a key business cycle variable and also an explicit policy target in some economies. It is among the most commonly forecast business cycle variables, with the result that long time series of regular high-quality forecasts are available for study. It is a relatively smooth locally stationary series (in the U.S., at least) that undergoes very little revision after initial publication (unlike the quarterly growth rate of real GDP, for example). This allows business cycle fluctuations to stand out more clearly from possible measurement error 'noise.

Despite its importance and the substantial literature on the asymmetry of business cycles, there have been few studies examining the success of forecasters in modelling and tracking the asymmetry in such processes. This paper attempts to contribute to the literature by documenting the asymmetry in well-known forecast series, and by examining how the asymmetry varies as the forecast horizon increases. We motivate this with two simple models of information arrival. In the first, we assume that linear forecasters learn about upside and downside risks separately. In this case the relationship between forecast error asymmetry and forecast horizon is ambiguous. In the second model, we assume that forecasters learn over time about the symmetric distribution of a forcing variable which is non-linearly related to the unemployment rate. We describe conditions under which this leads to a systematic relationship between forecast error asymmetry and the forecast horizon. We then present evidence consistent with such a relationship in unemployment rate forecast errors from two sets of professional forecasts.

Section 2 briefly reviews business cycle asymmetry and the behavior of changes in U.S. unemployment rates, as background to our investigation of the forecast errors. Section 3 addresses the relationship between asymmetry in forecast errors and forecast horizons in two simple models. Section 4 describes the forecasts we examine, which consist of long time series of unemployment rate forecasts from the Survey of Professional Forecasters as well as from the Federal Reserve Board of Governors Greenbook, both of which have been extensively studied in other respects. Section 5 presents empirical results, and Section 6 concludes.

2. Asymmetry in unemployment data

The suggestion that business cycles may be asymmetric dates back at least to Keynes. If we observe any time series over a long period, the sums of the magnitudes of upward

[1] Related recent literature includes Galbraith and van Norden (2012), Knüppel and Schultefrankenfeld (2012), Adrian et al. (2018), Clark, McCracken and Mertens (2016), and Reifschneider and Tulip (2017), among others.

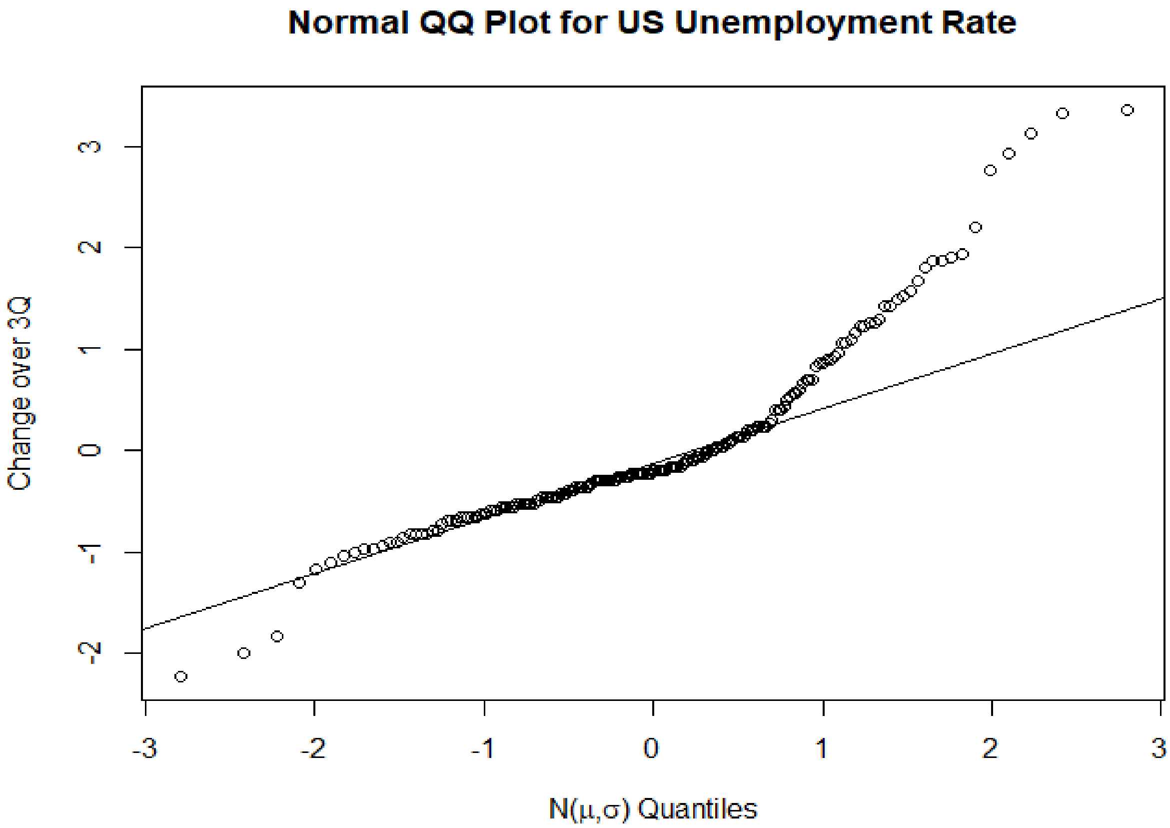

Figure 1: Distribution of Changes in the U.S. Unemployment Rate, 1969/3-2016/4

and downward movements must be approximately equal if the series is to be stationary; otherwise, it will tend to drift. If changes occur more frequently in one direction than in the other, then average sizes of two movements must differ to compensate and keep the unconditional mean unchanged. If a procyclical series increases more often than it decreases (because the economy is more often in an expansion then in a recession), the average magnitude of downward movements must then be greater than that of increases if it is a stationary series. For countercyclical series such as the unemployment rate, the opposite will true and a larger number of decreases than increases in the unemployment rate would imply that the average magnitude of a decline in unemployment is smaller than that of an increase.

This asymmetry can be clearly seen in the distribution of changes in the U.S. unemployment rate. Figure 1 shows a quantile-quantile (QQ) plot of changes in the rate over periods of three quarters.[2] The vertical axis shows the quantiles of the distributions of changes in the unemployment rate while the horizontal axis shows the corresponding quantiles of a

[2]Source data are forecast errors for the "no change" forecast calculated by the Federal Reserve Bank of Philadelphia over the period 1969Q3 to 2016Q4. Quarterly data are used to facilitate comparison with the

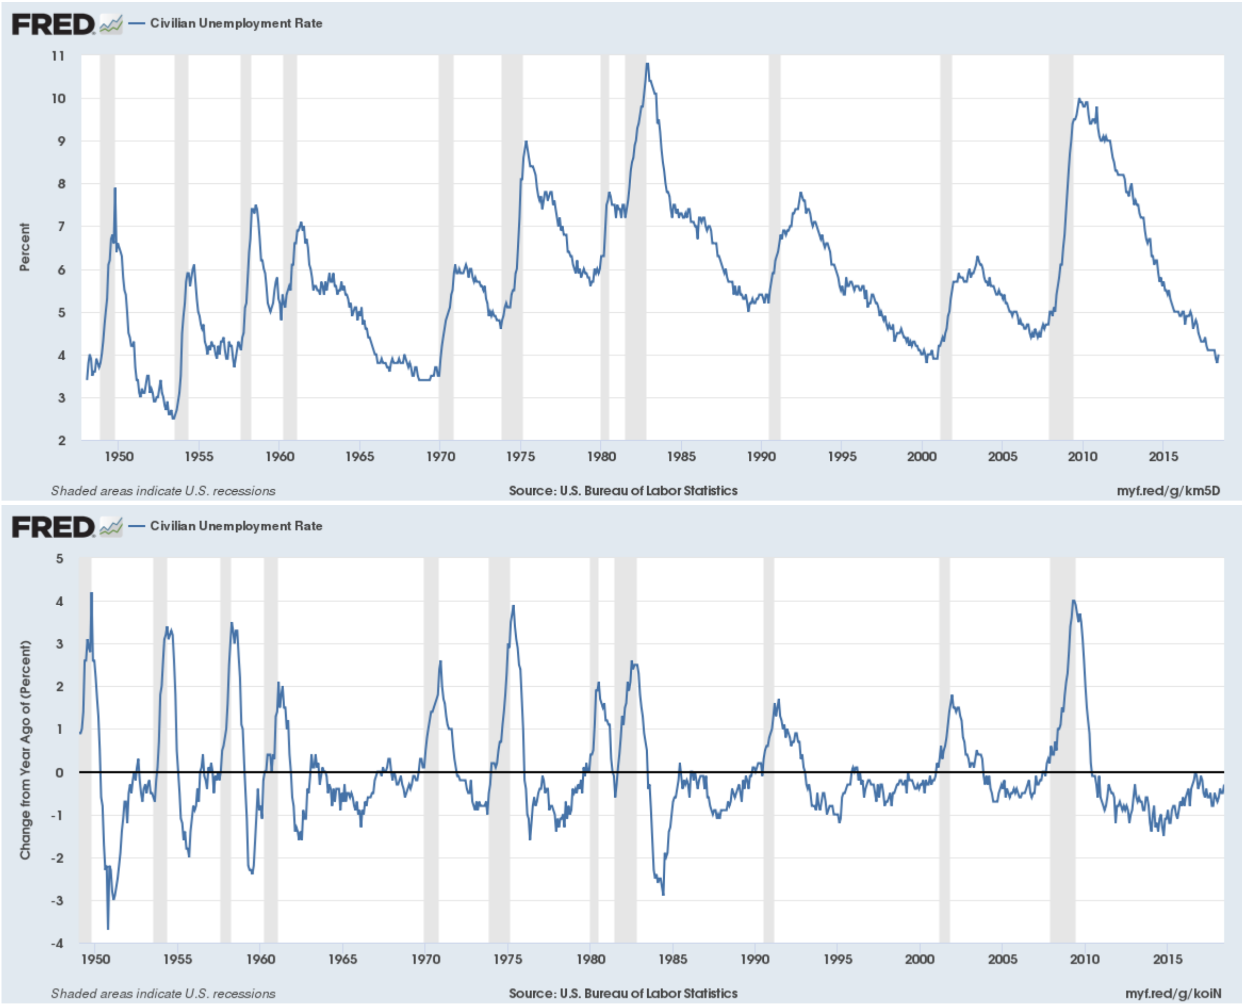

Figure 2: Recessions and Unemployment Rates

standard normal distribution with the same sample mean and variance. The straight line indicates the quantiles of normal distribution with same mean and variance as the observed data. We see that observations falling between $\mu - 2 \cdot \sigma$ and $\mu + \sigma$ seem well described by a normal distribution. However, for values more than one standard deviation above the mean (i.e. large increases in unemployment, corresponding to economic downturns) we see large deviations from the normal distribution, with such increases occurring much more frequently than a symmetric normal distribution would predict, and with larger absolute values in the upper tail. In the raw forecast data of Figures 3 and 4, the same phenomenon is visible as 'steeper' periods of increase than of decrease.

Figure 2 makes clearer the relationship between changes in the unemployment rate and the business cycle. The top panel shows the level of the unemployment rate while the

Survey of Professional Forecasters data considered below. Similar results are found using changes in the unemployment rate over periods ranging from one to five quarters.

shaded vertical bars correspond to NBER-dated recessions. The relatively rapidly increase in the unemployment rate during recessions is clearly visible. The bottom panel shows the corresponding year-over-year change in the unemployment rate. We see that increases in unemployment rates outside of recessions are rare and always small (less than one-half of one percent) while recessions are marked by increases of at least one and one-half percent.

Statistical tests for asymmetry appear as early as Neftçi (1984), who tested the unemployment rate for a difference in the steepness of periods of increase and decrease.[3] Here we provide updated results using data on the U.S. civilian unemployment rate (seasonally adjusted, monthly)[4] from 01/1948 through 12/2016, for which we compute the mean and median of positive and (absolute values of) negative changes in the unemployment rate, conditional on the absolute value exceeding some threshold, either zero or a larger value.[5] That is, we compute $|\Delta(u)|$ conditional on $|\Delta(u)| > \tau,\ \ \tau = \{0, 0.1, 0.2, 0.3\}$. The thresholds are used to exclude either zeroes (which are of course neither positive nor negative) and in the cases with $\tau > 0$, small unemployment rate changes which may not well represent periods of substantial change; note that the inequality is strict so that values at the given threshold are excluded in each case. We then test significance of differences in the mean or median across positive and negative values, using a robust t-type statistic or Wilcoxon rank-sum test respectively. The robust t-statistic and the regression on which it is based (6) are described fully in Section 5 below.

Table 1: Tests of asymmetry in monthly unemployment rate changes, 1948-2016

	no. cases	**mean diff**	**robust t-stat**	**p-value**	**p-value, R-S**
$\tau = 0$	621	-0.028	-2.004	0.017	0.103
$\tau = 0.1$	424	-0.033	-1.962	0.030	0.145
$\tau = 0.2$	267	-0.044	-2.127	0.026	0.031
$\tau = 0.3$	92	-0.055	-1.479	0.187	0.054

Table 1 provides the formal tests on both mean and median effects. For each exclusion threshold, the table provides the available number of observations (column 1), the mean difference between magnitudes of positive and negative changes in the unemployment rate (column 2), a robust $t-$statistic on this mean difference, computed using Newey-West heteroskedasticity and autocorrelation-consistent (HAC) standard errors (column 3), and the

[3]See Sichel (1989) and Falk (1986) for commentary on and development of this work, and Koop and Potter (1999) for somewhat more recent results; Knüppel (2009) provides a clear exposition of the concepts of steepness and deepness and their implications for tests of asymmetry in a class of models with Markov-switching intercept. He also finds statistically significant steepness in U.S. unemployment data. Montgomery et al. (1998) is a rare example of a study attempting to exploit this asymmetry in forecasting models.

[4]Source: U.S. Bureau of Labor Statistics; retrieved as UNEMP from the FRED database, Federal Reserve Bank of St. Louis, February 2017.

[5]Early results of this type, as cited above, were based on data up to approximately 1980 or 1985. One might reasonably have been concerned about whether such results were affected by data mining biases. However, there are now more than forty years of additional data in which the same type of result appears, substantially attenuating such concerns.

p-value (%) for this statistic in column 4. For the rank-sum (R-S) test, the p-value (%) alone is reported, in column 5.

The results tend to suggest that 'steepness' differences between periods of increase and decrease in the unemployment rate are genuine. At the largest threshold, $|\Delta u| > 0.3$, the number of sample points is reduced to 92, and statistical significance at conventional levels is lost. The estimated mean difference nonetheless increases each time we raise the threshold to restrict attention to more extreme changes in the unemployment rate.

Whether such differences can be, or are, exploited by forecasters is the subject of the empirical study below.

3. Theory and hypotheses

The evidence of asymmetries in the distribution of unemployment rate changes suggests that policymakers, and other economic agents, face asymmetric unemployment rate risks. But to understand the nature of these risks, we need to consider how unemployment rate uncertainty (or business cycle uncertainty more generally) is resolved through time. In the case of normally-distributed forecast errors, the variance of these errors captures forecast uncertainty and is easily estimated. Allowing for asymmetric forecast errors implies the possibility that the degree of asymmetry may vary with forecast horizon. This raises the question of what patterns we would expect to see in asymmetry across forecast horizons, if forecasters are exploiting information optimally (in a mean-squared-error sense) and of what we can conclude from empirical observations about the information being exploited by forecasters.

In the remainder of this section we present two simple models of forecasting with asymmetric risks, to help understand how risks may evolve with the forecast horizon when forecasts are optimal. The first model simply serves to illustrate the point that there is no necessary relationship between forecast horizon and the asymmetry of forecast errors. The second treats limiting cases of complete, and complete absence of, information about shocks to forcing variables.

3.1. A simple illustration of asymmetric risk

As a first example of how asymmetries may affect forecast errors across various horizons, we consider a simple 3-period model. A random variable x takes on a value of 1 or -1 at time $t = 0$. At various times $t < 0$ the forecaster makes rational forecasts of the eventual value of x using all information available.[6] The distribution of the forecast errors will in turn depend on the information we have available when a forecast is made.

Suppose that at $t = -2$, with a probability of 1/3, the forecaster receives a signal G. If this signal is received, then with certainty, $x = 1$. If the signal is not received at $t = -2$, then x may take on either value.

[6] By *rational*, we here mean that the forecasts are a linear projection of the target variable on the available information. This is consistent with the approximately mean-zero forecast errors that they display for most other variables.

At $t = -1$, the forecaster may receive a signal B. If so, then $x = -1$ with certainty. Clearly, this implies we can only observe this signal if we have not previously observed the G signal at $t = -2$. Conditional on not having observed G, the probability of observing B at $t = -1$ is 1/2. If neither the G or B signal is received by $t = -1$, then x may take on either value at $t = 0$, with both outcomes equally probable. Figure A.14 (Appendix A) summarizes the information flow and outcomes in this simple model.

Before we derive the distribution of rational forecast errors in this model, we characterize the unconditional and conditional distributions of x, G and B :

- We observe G one-third of the time.
- Conditional on not observing G, we observe B one-half of the time.
  The unconditional probability of observing B is therefore $2/3 \cdot 1/2 = 1/3$,
  the same as G.
- We therefore observe neither G nor B with a probability of 1/3 as well.
- Therefore $Pr(x = 1) = Pr(x = 1|G) \cdot Pr(G) + Pr(x = 1|$ not G, not $B) \cdot Pr($not G, not $B)$ $= 1 \cdot 1/3 + 1/2 \cdot 1/3 = 1/2$

So we may conclude that for $t < -2$, $x = 1$ and $x = -1$ are equally probable.

Next we characterize the rational forecasts and the distribution of their forecast errors in this model:

- If we observe a signal (G or B), we forecast 1 or -1 accordingly and the forecast error will be zero.
- At any point $t < -2$, the best forecast will be the unconditional forecast. This will be $E(x) = 0$ and will have a symmetric forecast error of $\{1, -1\}$.

The forecast errors at $t = -1$ will also be symmetric. Two-thirds of the time we will have observed either G or B, so the forecast error will be zero. The rest of the time, the rational forecast will again be 0, and the forecast error will be $\{1, -1\}$ with equal probability.

In contrast, the forecast errors at $t = -2$ will be asymmetric. With $Pr = 1/3$ we observe G and the forecast error is 0. If we do not observe G, however, we know that the probability that $x = 1$ is only 1/4. (We will observe B next period with $Pr = 1/2$, and if we do not, $x = -1$ with $Pr = 1/2$.) Therefore $E(x|$ not $G) = 1/4 \cdot 1 + 3/4 \cdot -1 = -1/2$. While the forecast error will have a mean of zero by construction, it only takes on the values $\{1/2, 0, -3/2\}$ and so will be asymmetric.

This simple model is sufficient to show that there need not be a monotonic relationship between the forecast horizon and the asymmetry of the forecast errors. In this case, forecast errors at the longest and shortest horizons are symmetrically distributed, but that this is not the case at the intermediate horizon.

3.2. Optimal forecasting in an asymmetric environment

We now consider a process that may resemble more closely the environment that business cycle forecasters face, and ask what we can deduce about asymmetry in limiting cases, if forecasts are generated optimally.

Consider a variable y which responds non-linearly to a symmetrically distributed conditioning variable x and linearly to another conditioning variable z. We may write

$$y_t = \gamma_0 + g(x_t) + \beta z_t + \varepsilon_t \tag{1}$$

or, in the case where we represent the non-linearity with a simple threshold model,

$$y_t = \gamma_0 + \gamma_1 x_t^- + \gamma_2 x_t^+ + \beta z_t + \varepsilon_t. \tag{2}$$

where $x_t^- = x_t(I[x_t < \tau])$ is the set of values of x_t below a threshold τ and $x_t^+ = x_t(I[x_t > \tau])$ is the set of values of x_t above τ. In the present context we might take y to be the change in the employment rate, and x to be the change in GDP. The forecast $F_{t-h}y_t$ is taken to be the conditional expectation of y_t given information on conditioning variables observable at $t-h$, and the forecast error is $y_t - F_{t-h}y_t$.

Two polar cases may be useful to the intuition. In each case we will take model parameters as known to simplify the exposition.

In the first case, $E(x_t|I_{t-h}) = \mu_x$, that is, x_t has no predictable component, so that the optimal forecast $F_{t-h}x_t$ is its unconditional mean. In this case, $F_{t-h}y_t = \gamma_0 + \kappa + \beta F_{t-h}z_t$, where κ is a constant depending on γ_1 and γ_2 and $f(x)$, the unconditional distribution of x. The forecast error is

$$y_t - F_{t-h}y_t = -\kappa + \beta(z_t - F_{t-h}z_t) - \gamma_1 x_t^- - \gamma_2 x_t^+ + \varepsilon_t, \tag{3}$$

and the asymmetry of the underlying process is fully preserved in the forecast errors.

The second polar case is that of a perfectly predictable conditioning variable x_t, so that $F_{t-h}x_t = x_t$. Here, $F_{t-h}y_t = \gamma_0 + \gamma_1 x_t^- + \gamma_2 x_t^+ + \beta F_{t-h}z_t$, and the forecast error is

$$y_t - F_{t-h}y_t = \beta(z_t - F_{t-h}z_t) + \varepsilon_t, \tag{4}$$

which is linear and symmetric given those properties in z_t and in its forecast $F_{t-h}z_t$. Clearly, asymmetry in errors from an optimal forecast is a function of the degree of predictability of the variable to which there is nonlinear response.

In the generic case the forecast of the conditioning variable is neither constant nor identically equal to the outcome, and we have a general nonlinearity $g(x_t)$ in the response. Then $F_{t-h}y_t = \gamma_0 + F_{t-h}g(x_t) + \beta F_{t-h}z_t$, and

$$y_t - F_{t-h}y_t = g(x_t) - F_{t-h}g(x_t) + \beta(z_t - F_{t-h}z_t) + \varepsilon_t, \tag{5}$$

so that the asymmetry is preserved, but its degree depends upon the degree of non-linearity of $g(x_t)$ over the conditional distribution $f(x_t|\Omega_{t-h})$. Intuitively, we might expect this to increase with the impact of a Jensen's inequality term, $[F_{t-h}g(x_t) - g(F_{t-h}x_t)]$. Note that in the special case of the threshold non-linearity, the forecast error depends upon whether or not 'classification' error arises where x_t and $F_{t-h}x_t$ are on opposite sides of the threshold τ, so that the incorrect parameter is applied in computing the forecast of y_t.)

We expect then to see the following patterns with optimal forecasts: at a short horizon a forecaster might be able to exploit information providing a reduction or even elimination of asymmetry in the forecast errors relative to that of the underlying series: short-horizon error asymmetry should not exceed, and may be less than, long-horizon error asymmetry. At some long horizon, any such information ceases to be available, and observed forecast error asymmetry will converge to that of the series being forecast.[7]

4. Data and forecasts

We now turn to an examination of U.S. data on unemployment rates, forecasts and forecast errors. Our forecasts are taken from the Survey of Professional Forecasters (SPF) and the FOMC Greenbook.

4.1. SPF forecasts

The Survey of Professional Forecasters (SPF) conducts quarterly surveys of professional forecasts for many macroeconomic variables, including the U.S. unemployment rate.[8] Surveys are conducted in the middle month of each quarter, asking respondents for their forecasts of the average monthly unemployment rate in the previous and current quarters (backcast and nowcast) and the four subsequent quarters. Our sample from the unemployment survey runs from 1968Q4 through 2016Q4. There is an extensive academic literature studying the properties of the SPF forecasts, which are generally found to be among the best-performing macroeconomic forecasts available.[9]

Figure 3 depicts overlapping sequences of unemployment forecasts from individual forecasters in the SPF. We observe two clear features of these data: first, dispersion across individuals' forecasts increases markedly with forecast horizon; second, these forecasts show evidence of asymmetry of the process previously noted from Figure 1; here, the increase in the rate as the economy enters a period of slowdown appears to be more rapid than in the decrease as the economy enters a period of expansion. For the analysis presented below, we have used the median of survey responses.[10]

---

[7] In between the end points, there is no requirement of strict monotonicity, as the example of the previous section illustrates.

[8] Details concerning the SPF may be found on the Federal Reserve Bank of Philadelphia's web site at *https://www.philadelphiafed.org/research-and-data/real-time-center/survey-of-professional-forecasters.* See also Croushore (1993) and Stark (2016).

[9] See the comprehensive bibliography maintained by the Philadelphia Federal Reserve Bank at *https://www.philadelphiafed.org/research-and-data/real-time-center/survey-of-professional-forecasters/academic-bibliography.*

[10] 'SPF median' refers to the median of individual forecasts recorded in the SPF for any given variable and horizon. As can be seen from Figure 3, the individual SPF forecasts are tightly concentrated for the nowcast and first few forecasts, so that use of the median rather than any individual forecast sequence should have no substantial effects on the results. At longer horizons, forecasts are more dispersed, and asymmetry tests on some individuals could differ noticeably from those on the median.

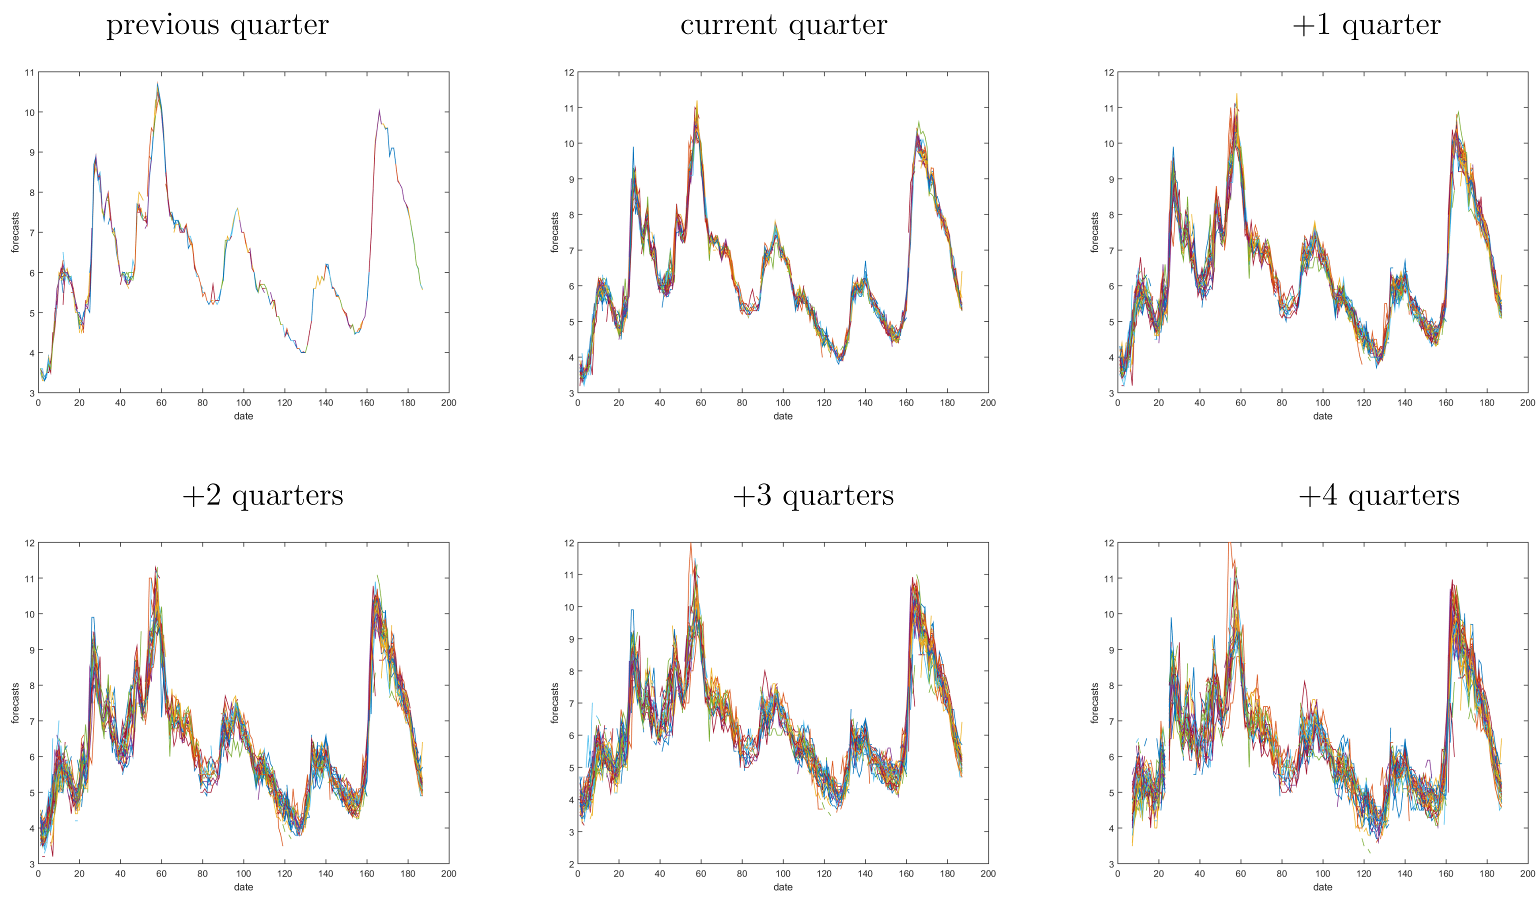

Figure 3: Individuals' unemployment forecast sequences from the SPF

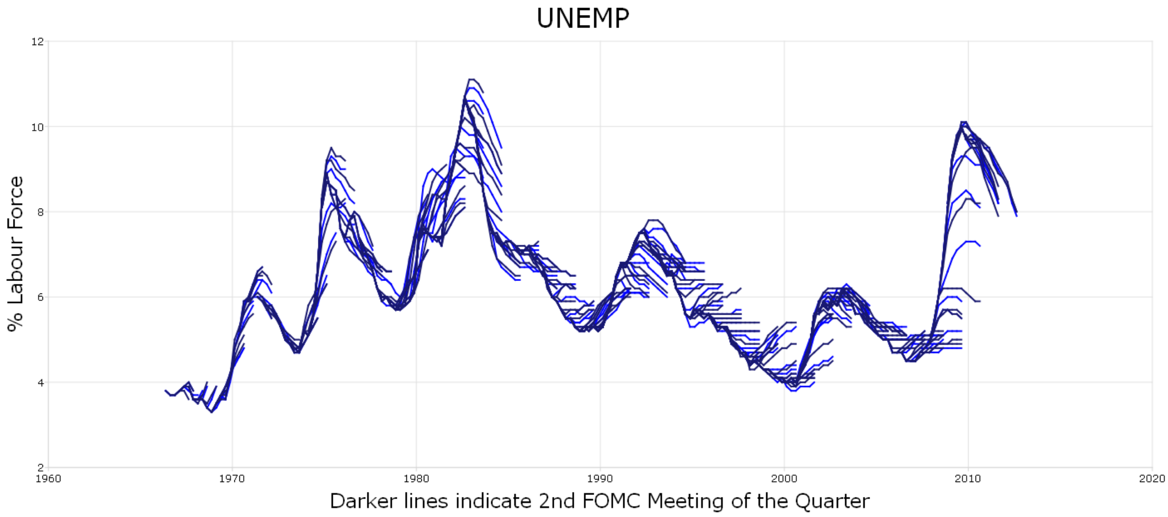

Figure 4: Greenbook Forecasts for Unemployment Rate (UNEMP) 1967-2010

4.2. Greenbook forecasts

The Greenbook is a summary of economic conditions, trends and forecasts prepared by the staff of the U.S. Federal Reserve Board for every meeting of the Federal Open Market Committee (FOMC.) [11] In additional to economic commentary and analysis, it contains the staff's quarterly economic projections. At present, Board policy is to release publicly the Greenbooks, and all other FOMC briefing materials and transcripts, after roughly five years.[12] The ALFRED database of the FRB St. Louis provides the Greenbooks' unemployment rate forecasts made from 1978 through 2008. We instead use the Greenbook unemployment rate forecasts from Croushore and van Norden (2018), which cover forecasts made from August 1967 through December 2010.[13] Figure 4 shows all the Greenbook unemployment rate forecasts and historical estimates. The maximum forecast horizon presented in the Greenbooks varies widely, but typically includes the current quarter and at least the four subsequent quarters. While some forecasts up to nine or more quarters ahead are made, they are few in number, irregularly spaced and tend to be concentrated in the latter half of the sample.

Unlike the SPF, which collects forecasts once per quarter, the FOMC meets (typically) twice per quarter, allowing us to distinguish forecast horizons from the (F)irst and the (L)ast meeting of a given quarter. Like the SPF, there is an extensive literature documenting the

[11] Greenbooks were not produced prior to 1965. In 2010, both the Greenbook and the staff's Bluebook materials were merged into a single report called the Tealbook. For simplicity, we refer to all of these projections as "Greenbook forecasts."

[12] Copies of the Greenbooks are available from the websites of the FRB Philadelphia and the Board of Governors.

[13] Croushore and van Norden (2018) provide additional details on the data collection, properties and construction of forecast errors.

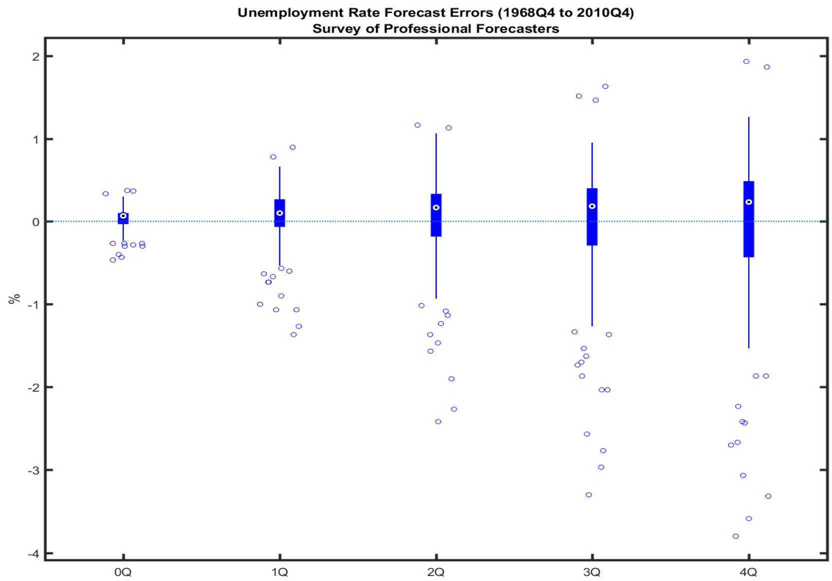

Figure 5: Distribution of SPF Forecast Errors

Note: The box plots above compare forecast error quantiles across forecast horizons for the SPF Unemployment Rate forecasts. The central rectangle spans the interquartile range and the dot within it indicates the median. The whiskers extend an additional 1.5 times the interquartile range; observations beyond this are shown as individual points. (For a normal distribution, the whiskers are equivalent to the 99.3% confidence interval.)

performance of the Greenbook forecasts.[14]

5. Empirical analysis

We now move from description of the forecasts to an examination of the forecast *errors* of the Greenbook and SPF median forecasts. We begin with a visual analysis of the data before turning to formal statistical measures and tests.

5.1. Visual characterization of asymmetries

Figure 5 shows the distribution of the SPF forecast errors at each forecast horizon as a box-whisker plot. The plots show the distribution of forecast errors for the SPF forecast: the

[14]For details, see Croushore and van Norden (2017) and the sources therein. Croushore and van Norden (2014), while primarily concerned with forecasts of fiscal variables, also provides some results on the behaviour of unemployment rate forecast errors for various forecast horizons. In particular, they present (a) box plots of the distribution of forecast errors, and (b) non-parametric tests of the null hypothesis that the median forecast error is zero. These provided visual and statistical evidence of asymmetry, although its relationship to forecast horizon is not clear.

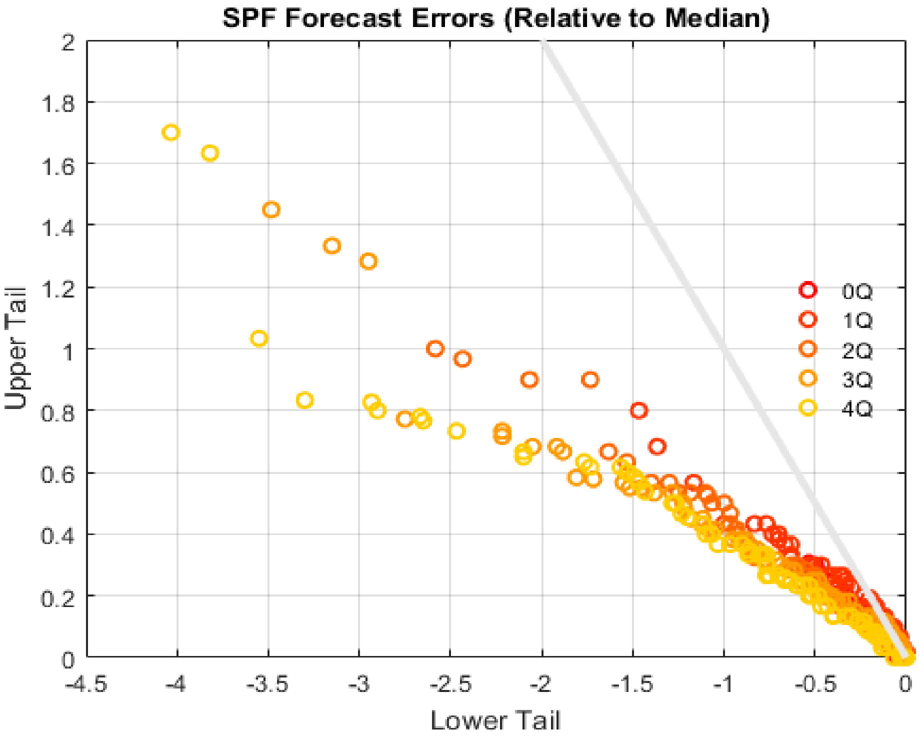

Figure 6: Asymmetry of SPF Forecast Errors

circle in the centre of the vertical bar indicates the median, while the bar itself extends from the first to the third quartile, and the whiskers extend a further 1.5 times the inter-quartile range.[15] All points falling outside the range covered by the whiskers are plotted as circles. It is readily apparent that forecast errors at all horizons have asymmetric distributions; the medians are systematically positive, they are typically much closer to the 3rd quartile than to the 1st, and there are always many more negative errors outside the range of the whiskers than there are positive errors.[16]

An alternative depiction of the asymmetry is given in Figure 6, which presents a modified QQ plot of the SPF forecast errors at various forecast horizons. Each point (x_i, y_i) takes a matched pair of forecast error values, the ith largest and ith smallest, and plots them, subtracting the median. These values are therefore at quantiles α and $1 - \alpha$, and in a perfectly symmetric distribution would lie along the diagonal line of equality, or '45-degree' line, shown in grey. Darker (redder) colors indicate shorter horizons.[17]

Starting in the bottom right corner of Figure 6, we see the median at (0,0). As we

[15]For a normal distribution, the probability of an observation lying outside the whiskers is just under 1.3%.

[16]Forecast errors are calculated as (forecast - actual): negative errors correspond to unexpectedly high rates of unemployment.

[17]It may be useful to think of color in the figure as 'fading' from red to light yellow as the forecaster looks further into the future.

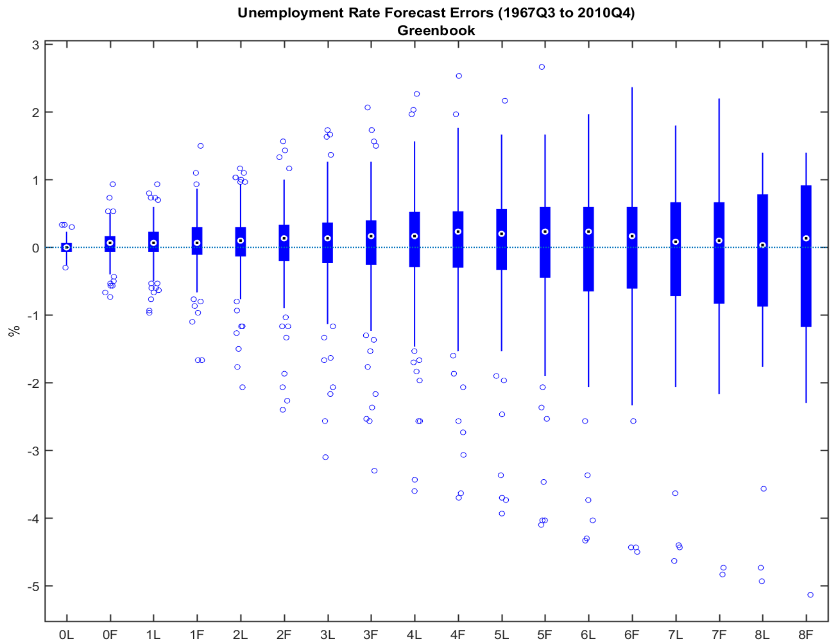

Figure 7: Distribution of Greenbook Forecast Errors

Note: The box plots above compare forecast error quantiles across forecast horizons for the Greenbook Unemployment Rate forecasts. The central rectangle spans the interquartile range and the dot within it indicates the median. The whiskers extend an additional 1.5 times the interquartile range; observations beyond this are shown as individual points. (For a normal distribution, the whiskers are equivalent to the 99.3% confidence interval.)

move upwards and leftwards, we plot percentiles which sum to 100: e.g. the 40th vs the 60th, the 25th vs the 75th, the 10th vs the 90th, etc. We see that almost all points lie well below the diagonal, indicating that the lower tail of the distribution is longer than the upper tail, and therefore that under-predictions of the unemployment rate tend to be larger (in absolute value) than correspondingly frequent over-predictions. The graph also shows that this asymmetry appears to be larger for the 3- and 4-quarter-ahead forecasts than for the 0- to 1-quarter-ahead forecasts, consistent with the results from the box plots.

One might conclude from Figure 6 that capturing the asymmetry in this process is not feasible at any horizon. However, a clear difference is discernible when we examine Greenbook forecasts.

Figures 7 and 8 show the corresponding box-whisker and modified QQ plots for Greenbook forecast errors.[18] In general the forecast errors display a degree of asymmetry similar to that of the SPF forecasts. Figure 7 shows that median forecast errors are again systemat-

[18]Again, we distinguish forecasts from the first (F) and last (L) FOMC meetings in the same quarter; the forecasts in the figures and table below are arranged from the shortest to the longest forecast horizon.

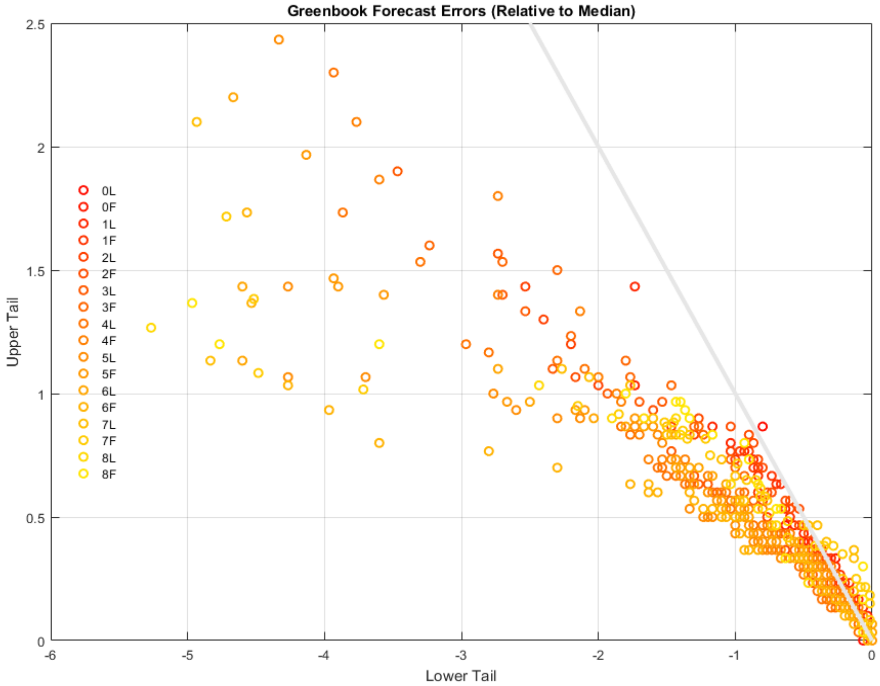

Figure 8: Asymmetry of Greenbook Forecast Errors

ically positive, they are typically much closer to the 3rd quartile than to the 1st, and there are almost always many more negative errors outside the range of the whiskers than there are positive errors. We also see that the preponderance of negative to positive outliers tends to increase with forecast horizon. However, Figure 8 displays some important differences relative to the analogous plot for SPF forecast errors. Although the pattern is broadly similar to that of Figure 6 at intermediate and longer horizons, the shortest horizons (darker, redder) lie quite close to the line of equality: the Greenbook forecasters are, according to this visual evidence, successfully accounting for asymmetry at the shortest horizons and approximately eliminating it from the forecast errors.

Figures 9 and 10, for SPF median and Greenbook forecasts respectively, provide conventional QQ plots, which allow a comparison of the asymmetry in the forecast relative to a random walk ('no-change') forecasts. Each panel simply plots the forecast error quantiles for the random walk forecast (horizontal axis) against those of the SPF or Greenbook forecast (vertical axis), with the red dashed line showing the estimated best linear fit between the two. The nowcast and four-quarter horizons are highlighted in the top and bottom panels of each figure.

Some care must be taken in making comparisons across forecast horizons as (a) the available number of observations drops sharply after about 6Q, (b) the smaller samples are concentrated in the latter part of the sample period, and (c) as forecast horizons increase, forecasts increasingly overlap, reducing the independence between successive forecast errors.

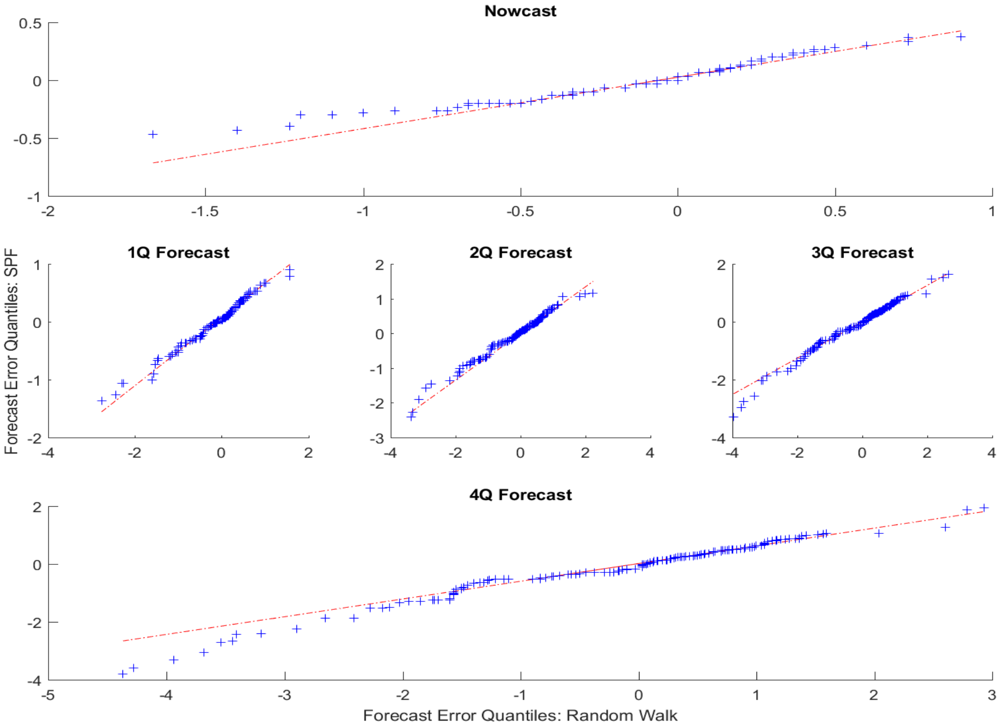

Figure 9: Forecast Error Distributions: SPF versus Random Walk

For the SPF median forecasts, we see from Figure 9 that as we move from the shortest forecast horizon (top panel) to the longest (bottom panel), there is a clear change in the lower tail of the distribution. At the shortest forecast horizon, the random walk forecast has somewhat more frequent large negative errors, while at the longest horizons it has somewhat fewer.

Figure 10 provides similar plots, comparing the distribution of Greenbook forecast errors to those of a random walk forecast. The dashed line in each plot again shows the regression line linking the two series. Comparing the vertical and horizontal scales of the plots confirms that these lines are substantially flatter than the 45-degree, implying that the Greenbook forecasts for the unemployment rate are generally more informative than a random walk. However, comparing the plotted points to the fitted line, we see that as we move from shorter to longer forecast horizons, the degree of improvement in the two tails behaves differently. While forecasts errors in the upper tail (i.e. where forecasts were excessively pessimistic) are generally on or above the regression line, forecasts errors in the lower tail (i.e. where forecasts were excessively optimistic) tend to migrate from above the regression line at shorter horizons to below the line at longer horizons. This implies that while Greenbook forecasts were generally able to improve on random walk forecasts at all horizons, they show relatively less improvement in the face of adverse unemployment surprises at longer horizons (e.g. $> 2Q$) than they do at shorter horizons. This is compatible with the relative absence of useful conditioning information at long horizons.

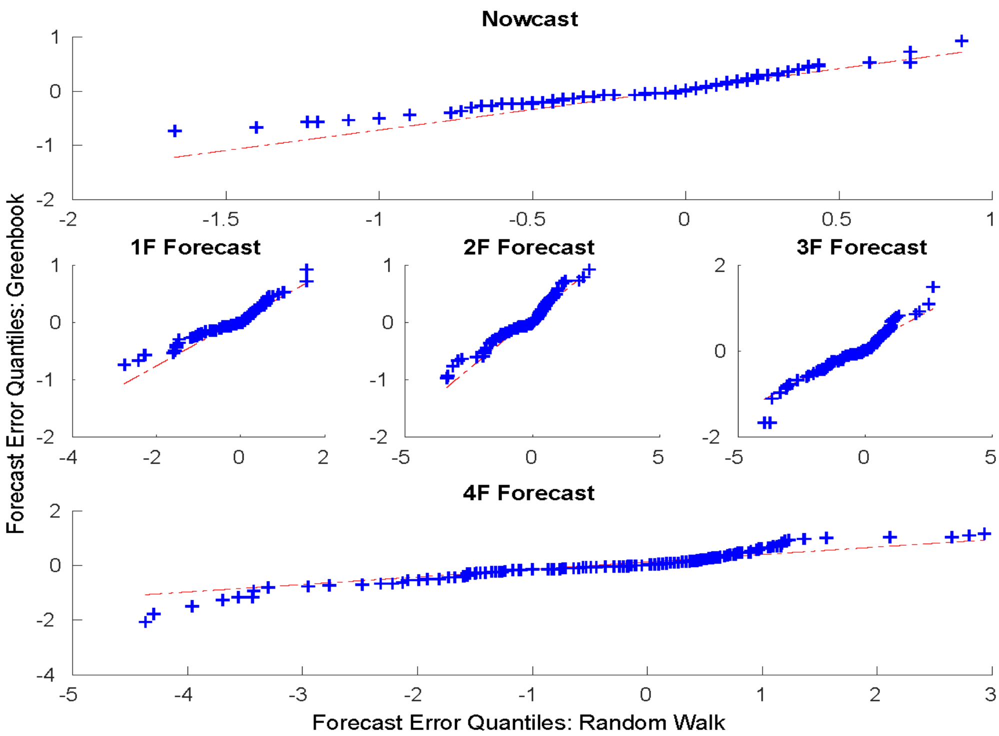

Figure 10: Forecast Error Distributions: Greenbook versus Random Walk

Figure 11 confirms the last point with a slightly different view of the same QQ plots as Figures 9 and 10. With the data points and the axes as the same as before, lines now connect points corresponding to errors with the same forecast horizon, making the comparison across forecast horizons easier. The regression lines have been replaced by a simple 45-degree line through the origin, indicating equality of the two forecast errors. We again see that for small forecast errors (i.e. near the origin), the curves for both the SPF and the Greenbook are flatter than the 45-degree line, implying that both show some forecasting skill.[19] As we might expect, the curve around the origin also appears to be become somewhat steeper as the forecast horizon increases, implying that forecasts become less accurate. However, the slope in the lower tail of the distribution (e.g. for SPF or Greenbook forecast errors $< -1\%$) is roughly the same as that of the 45-degree line. Since this corresponds to periods in which the unemployment rate increased by 2% or more, we know from Figure 2 that this reflects forecast performance during recessions. The fact that the SPF and Greenbook forecast errors in this region plot above the 45-degree line suggests that some of the increase in the unemployment rate was expected. However, when the slope of the curve parallels the 45-degree line, it suggests that the depth of the recession was not.

In general however these conventional QQ plots, unlike the modified plots of Figures 6 and 8, reveal little difference between the two groups of forecasters. We now turn to

[19]Note that for the Greenbook forecasts, the slope for small positive errors $< 1\%$ is much closer to that of the diagonal than is the case for the SPF, before becoming somewhat flatter.

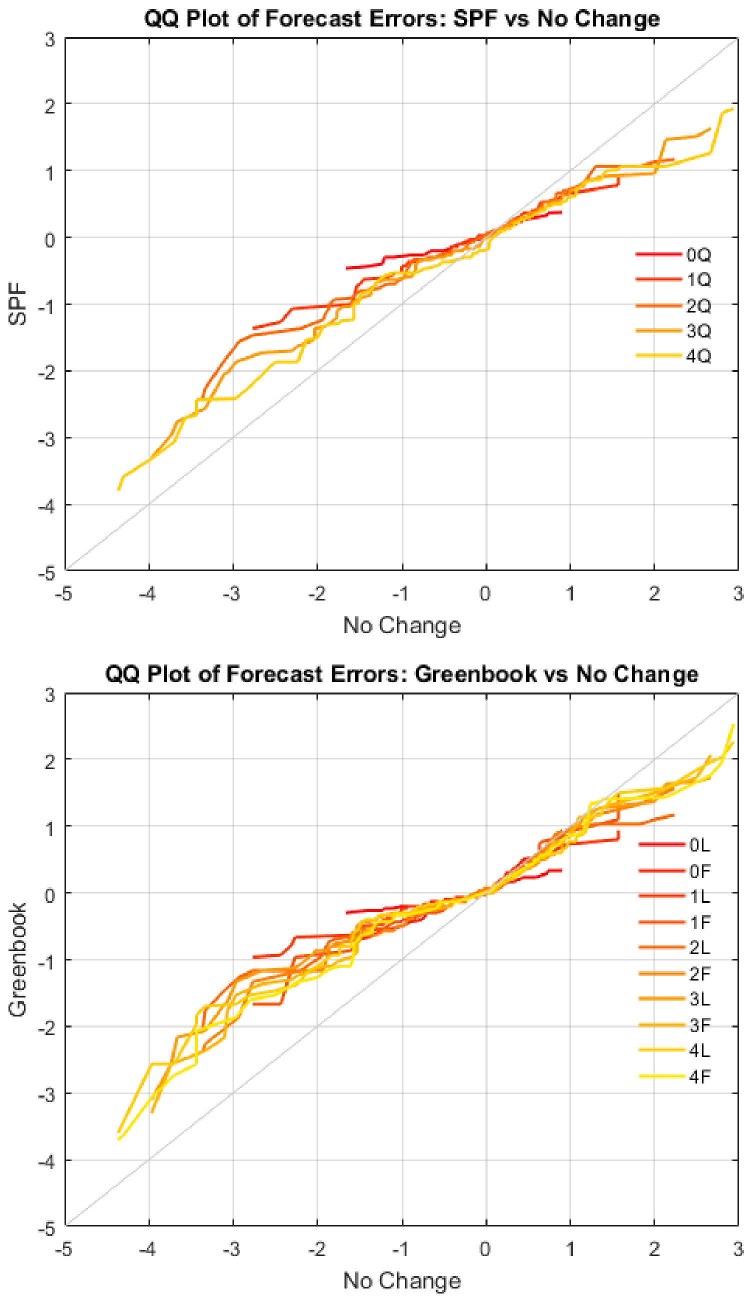

Figure 11: Forecast Error Distributions, versus Random Walk

numerical evidence.

5.2. Numerical measures of asymmetry

Table 2 presents various summary measures of the skewness of the distribution of the SPF forecast errors. The top panel of the table presents results for the median SPF forecast errors, while the middle panel presents those for the No-Change (NC) forecast and the bottom panel presents those for a Direct Autoregressive (DAR) model.[20] The coefficient of skewness is negative at all horizons for all three forecasts, implying that adverse surprises to unemployment tend to be larger than beneficial ones. We are also able to reject the null hypothesis of symmetric forecast errors at all horizons for the SPF and NC forecasts. In the case of the SPF forecasts, the coefficient also becomes increasingly negative as the forecast horizon increases. This is consistent with evidence from the QQ plots, but differs from that of the NC forecast errors (which decrease as the horizon increases) or the DAR forecast errors (which peak at 2Q and then decline slightly.)

We interpret the various measures as follows: the Bowley/Yule-Kendall interquartile measure is not affected by the tails of the distribution; the Pearson measure shows some tail sensitivity via the mean, and the coefficient of skewness based on the third moment shows the greatest tail sensitivity.[21] As we noted earlier, we treat the unemployment serious as approximately stationary with several finite moments, so that the moment-based skewness measures used here are well-defined and estimable.

For the SPF forecast error data, we see that the degree to which the measure changes with forecast horizon corresponds with this classification; the B-Y-K measure shows little discernible pattern across horizon, the Pearson measure appears to show modest increase in asymmetry with horizon, and the coefficient of skewness increases markedly. The results are consistent with asymmetry which arises in the tails of the distributions.

The Kolmogorov-Smirnov test, described below, rejects symmetry at all horizons.

Table 3 compares various quantitative measures of forecast error asymmetry across forecast horizons in the Greenbook.[22] The classification with respect to sensitivity to tails again has explanatory content; the coefficient of skewness is negative at all but the very shortest horizons and shows an almost monotonic increase with forecast horizon from 1L to 6L, after which it stays roughly constant; at the longest horizons, it is roughly ten times larger than at the shortest. Pearson's skewness coefficient shows a more irregular pattern, but is generally increasing with the forecast horizon until about 6Q. The B-Y-K measure again shows positive skewness with no clear relationship to the forecast horizon. We are again able to

[20]Forecast errors for the NC and the DAR models are provided by the Federal Reserve Bank of Philadelphia. The latter uses rolling estimation and lag selection, and forms multi-horizon forecasts by direct projection. See Stark (2016) for details.

[21]Note that Knüppel and Schultefrankenfeld (2011) argue that Pearson's median skewness should be preferred because it is more precisely estimated. However, the fact that this measure is less affected by tail risk than the coefficient of skewness may be the dominant consideration here.

[22]We make no attempt to compare the Greenbook forecasts to NC or DAR forecasts with the same forecast horizon because of the irregular timing of FOMC meetings.

Table 2: Numerical measures of asymmetry, SPF

Forecast Error	Statistic	0Q	1Q	2Q	3Q	4Q
	Skewness	- 0.609	- 1.220	- 1.450	- 1.545	- 1.556
	Pearson	- 0.738	- 0.443	- 0.751	- 0.835	- 0.961
SPF	**B-Y-K**	0.473	-	0.367	0.381	0.448
	KS Test	**0.311**	**0.381**	**0.340**	**0.355**	**0.378**
	p-value	0%	0%	0%	0%	1%
	Skewness	- 1.465	- 1.451	- 1.298	- 1.163	- 1.068
	Pearson	- 0.686	- 0.729	- 0.776	- 0.779	- 0.876
NC	**B-Y-K**	0.333	0.294	0.213	0.308	0.422
	KS Test	**0.237**	**0.247**	**0.284**	**0.397**	**0.340**
	p-value	0%	0%	0%	0%	2%
	Skewness	- 0.634	- 1.015	- 1.165	- 1.148	- 1.132
	Pearson	- 0.210	- 0.422	- 0.551	- 0.675	- 0.732
DAR	**B-Y-K**	0.030	0.158	0.155	0.211	0.172
	KS Test	0.058	0.084	0.168	0.222	0.250
	p-value	90%	49%	24%	23%	33%

SPF refers to SPF forecast errors at the given forecast horizon.
NC refers to No-Change forecast errors.
DAR refers to forecast errors from a Direct Autoregressive model with variable lag lengths. This benchmark is provided by the Federal Reserve Bank of Philadelphia as a benchmark for the SPF.
Skewness is the coefficient of skewness. We also calculated the bias-corrected coefficient of skewness using an adjustment factor of $\frac{\sqrt{T(T-1)}}{T-2}$, which was never substantively different from the coefficient reported above.
Pearson is Pearson's second coefficient of skewness, is given by $3 \cdot (Mean - Median)/StandardDeviation$.
B-Y-K is Bowley's Interquartile Skewness, also called the Yule-Kendall index, and is given by $(Q_3 + Q_1 - 2 \cdot Q_2)/(Q_3 - Q_1)$ where Q_i is the ith quartile.
KS Test is test for symmetry of the forecast errors. The statistic is that of the Kolmogorov-Smirnov 2-sample test of the null hypothesis that x and $-x$ are drawn from the same unknown distribution. For forecast horizons greater than one quarter, the significance levels of the test statistics were calculated using a simple Bonferroni correction. Cases where the null may be rejected at the 5% significance level are shown in **boldface**.
p-value is the marginal significance level of the KS test statistic.

Table 3: Numerical measures of asymmetry, Greenbook

Horizon	**Skewness**	**Pearson**	**B-Y-K**	**KS Test**	p-value	No. Obs.
0L	0.086	0.275	-	0.098	36%	174
0F	- 0.140	- 0.278	0.143	**0.247**	0%	174
1L	- 0.403	- 0.096	- 0.111	**0.337**	0%	174
1F	- 0.732	- 0.080	- 0.120	**0.393**	0%	174
2L	- 0.940	- 0.246	0.077	**0.476**	0%	173
2F	- 1.119	- 0.484	0.250	**0.394**	4%	169
3L	- 1.243	- 0.410	0.222	**0.500**	1%	167
3F	- 1.221	- 0.569	0.300	0.455	10%	165
4L	- 1.287	- 0.466	0.120	**0.526**	5%	165
4F	- 1.345	- 0.685	0.280	0.500	21%	155
5L	- 1.637	- 0.656	0.185	**0.769**	0%	155
5F	- 1.594	- 0.765	0.307	0.546	51%	143
6L	- 1.809	- 0.885	0.421	0.444	100%	130
6F	- 1.693	- 0.700	0.307	0.833	16%	120
7L	- 1.807	- 0.607	0.162	0.600	100%	107
7F	- 1.649	- 0.581	0.249	0.667	100%	85
8L	- 1.625	- 0.659	0.089	1.000	52%	72
8F	- 1.535	- 0.811	0.254	1.000	100%	54

See notes for Table 2. The forecast horizon is shown in quarters, with a suffix F for the first FOMC meeting of the quarter and L for the last meeting of the quarter.

reject the null hypothesis of symmetric forecast error distributions at most forecast horizons up to 4Q.[23]

Having characterized the degree of asymmetry in the distributions of the various forecast errors, we now turn to additional tests of the null hypothesis of symmetry.

5.3. Statistical inference on asymmetry

Different tests for asymmetry may emphasize, and have power against, different forms of departure from symmetry. One widely used test for asymmetry is the general Kolmogorov-Smirnov test, which tests the entire distribution based on the largest deviation between the CDF of a series and the reflection of the CDF about its mean. Specifically, let X be a sequence of forecast errors at the given horizon and $\tilde{X}$ be the de-meaned series; then under the hypothesis of symmetry of the distribution around the mean, the CDF's of $\tilde{X}$ and $-\tilde{X}$ are the same. We can test the hypothesis $H_{0,a} : F(\tilde{X}) = F(-\tilde{X})$ using the Kolmogorov-Smirnov test; while this test has power against the null in most directions, it also requires independent observations, which makes it poorly suited for use with forecast errors of any but the shortest forecast horizons. With the absence of any generally applicable method of correcting for this (see for example Weiss 1978), we simply sampled only non-overlapping forecast errors and applied a Bonferroni correction. While this should deliver a correctly-sized test, it is unlikely to be very powerful. In applying the test to Greenbook forecast errors (see Table 3) we were however able to reject the null hypothesis of symmetry around 0 at the 95% significance level at the 1F, 2L, 2F, 3L, 3F, 4L, 5L, and 6L horizons.

An alternative is to test whether positive and negative deviations (possibly beyond some threshold, τ, to filter out small variations) are on average of the same magnitude. As in Section 1, where such a test was applied to the raw data on changes in the unemployment rate, we can test the estimated difference in mean magnitudes,

$$d = \hat{\mu}_+ - \hat{\mu}_- \equiv (n_+^{-1} \sum_{i=1}^{n_+} [|\tilde{x}_i|_{\tilde{x}_i > \tau}]) - (n_-^{-1} \sum_{j=1}^{n_-} [|\tilde{x}_j|_{\tilde{x}_j < -\tau}]),$$

for some threshold τ and where n_+, n_- are the sample sizes of cases in which $\tilde{X}$ exceeds the threshold magnitude in the positive or negative direction.

Although this is a standard test for the difference in means of two series, as noted in Section 1 we can conduct the test in a regression framework in order to compute robust (HAC) t-type statistics conveniently. Specifically, we

1. standardize each series of forecast errors $(\hat{u}_t^h - u_t)$ to have mean zero and variance
2. select a value of the threshold $\tau = \{0.0, 0.5, 1.0, 1.5, 2.0\}$
3. select only those observations for which $|\hat{u}_t^h - u_t| \geq \tau$

[23]The Bonferroni correction that we used causes the KS statistics to lose power as the forecast horizon increases. Together with the diminishing sample size at long horizons, this could explain some of the failures to reject the null hypothesis of symmetry at long horizons, despite the large KS statistics.

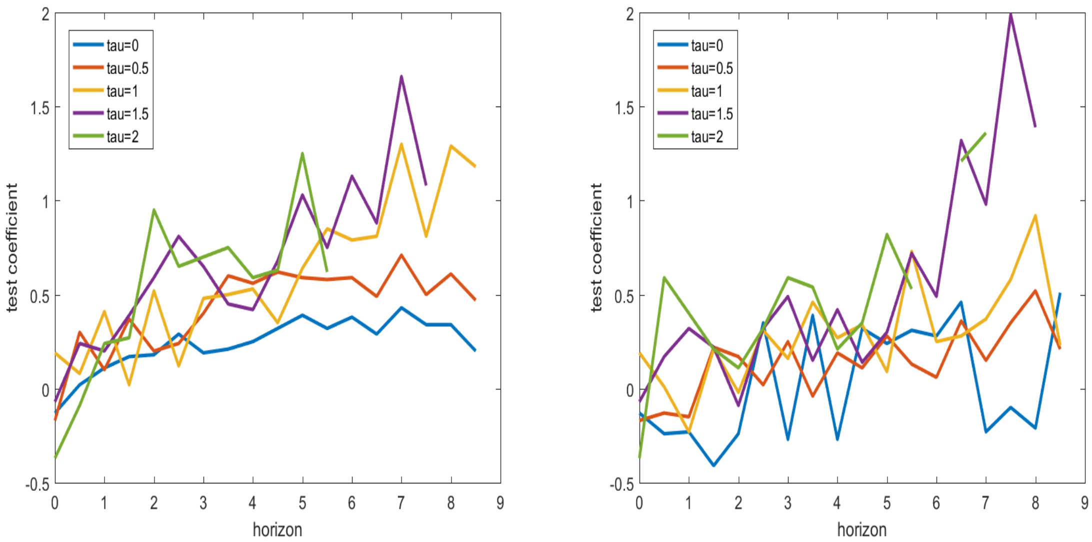

Figure 12: Test coefficients by forecast horizon (0L, 0F, 1L, ... 8L, 8F), Greenbook data; raw errors (left) and revisions (right)

4. estimate (by OLS)

$$|\hat{u}_t^h - u_t)| = \alpha_0 + \alpha_1 \cdot \mathcal{I}(\hat{u}_t^h - u_t) < -\tau) + \epsilon, \tag{6}$$

5. calculate the Newey-West standard errors for α_0 and α_1 using a lag length of $h+1$.
6. The test statistic is then the t-ratio for α_1, where $H_0 : \alpha_1 = 0$ is consistent with forecast errors that are symmetric.

Another way of dealing with autocorrelation would be to examine forecast revisions rather than the raw forecasts. For a set of forecasts of an outcome at t, the initial forecast is $\hat{u}_{t|t-h}$ and the sequence of forecast revisions is $(\hat{u}_{t|t-h+1} - \hat{u}_{t|t-h}), (\hat{u}_{t|t-h+2} - \hat{u}_{t|t-h+1}), \ldots, (\hat{u}_{t|t-1} - \hat{u}_{t|t-2})$. These revisions do not overlap and (assuming forecasts are efficient) standard tests may be used on the sequence without an autocorrelation correction.[24]

Tests using both of these approaches are reported in Tables B.4 – B.6 (Greenbook) and B.7 – B.8 (SPF) in Appendix B. Plots of the test coefficient estimates by horizon and threshold are given in Figures 12 and 13; these are intended to convey a visual impression of the change in the measure of asymmetry (the test coefficient) with forecast horizon. The lines correspond with columns headed 'coef' in Tables B.4 – B.8, separating the results by threshold in each case. Test statistics and p-values are contained in these Appendix tables as well.

To interpret these results, we note the following:

[24]Another alternative is to use the symmetry test for time series data of Bai and Ng (2005), based on the coefficient of skewness. This test requires a nonparametric estimate of spectral density at zero of the time series, with the attendant requirement for choice of a bandwidth parameter.

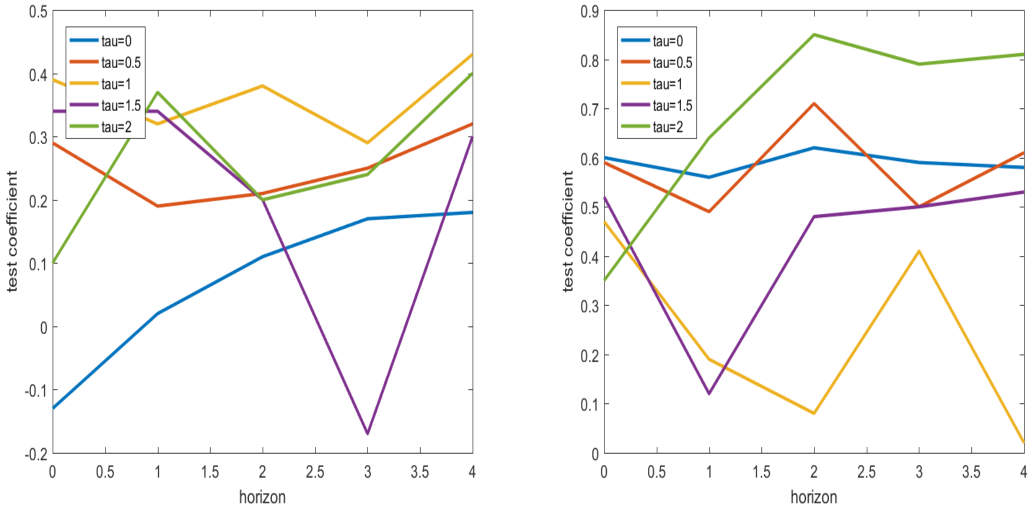

Figure 13: Test coefficients by forecast horizon (0–4), SPF data; raw errors (left) and revisions (right)

- The test coefficient is almost invariably estimated to be positive on the raw forecast errors, indicative of larger average magnitudes of forecast error in the negative direction. In SPF forecasts, this is true in every case; in Greenbook data, there are a few exceptions, always at horizon 0 (nowcast) (Greenbook, Tables B.4 – B.6, column 3; SPF, Table B.7, column 3).

- SPF forecast *revisions* show the same pattern (Table B.8, column 3). The Greenbook revisions have a somewhat larger number of negatives coefficients than in the raw errors (Tables B.4 – B.6, column 7). These mostly occur when $\tau = 0.0, 0.5$ and the horizon is short. Where a higher threshold is used to concentrate on larger changes, the coefficient is more reliably positive (Tables B.4 – B.6, column 7, comparing in particular B.4 and B.6). A number of the negative coefficients are strongly significant, which is consistent with Greenbook forecasters having some ability to reduce (offset) some longer-horizon asymmetries.

- The evidence of asymmetry is typically statistically significant at conventional levels in the SPF results (Tables B.7 and B.8, column 5). For Greenbook forecasts, the evidence is weaker; although the number of significant results substantially exceeds what would be expected by random chance, a majority of tests are nonetheless not significant (Tables B.4 – B.6, columns 5, 10).

- Comparison of the SPF errors with those from no-change (random-walk) or AR forecasts shows similar results (Tables B.7 and B.8, comparing for example column 3 with columns 7 and 11).

- Comparing different horizons is difficult given the sample sizes and corresponding estimation variability. Nonetheless, the overall pattern of change in the coefficients in the Greenbook data (Figure 12) suggests smaller asymmetric effects at shorter horizons, as would be the case if information becomes available closer to the date of interest which allows forecasters to eliminate some of the larger errors. In the SPF data, which show shorter (maximum 4-quarter) horizons, no such pattern is visible (Figure 13).

As we have noted, the primary purpose of this paper is to examine and document any asymmetries in forecast errors, rather than to draw conclusions about forecast efficiency or rationality. Nevertheless, in the commonly considered case in which agents do minimize quadratic loss, asymmetry of forecast errors may be indicative of inefficiency. Globally, the evidence here tends to be compatible with the idea that forecasters generally do not, or are unable to, observe or use information sufficient to allow them to eliminate the asymmetry in the raw data from their forecast errors; however there is evidence that Greenbook forecasters may be successful in doing so at the shortest horizons. This tends to confirm the visual evidence of differences in performance of SPF and Greenbook at the shortest horizons, observed in Figures 8 and 6. That is, the problem of exploiting information about asymmetry for forecasting is difficult, but at least at very short horizons is not impossible.

6. Concluding remarks

It has long been observed that the unemployment rate shows asymmetry over the business cycle, in the sense that periods of increase tend to be shorter and show steeper gradients than periods of decrease. The SPF and Greenbook forecasts therefore provide a test case with which to investigate forecasters' ability to incorporate this asymmetry into their forecasts, and thereby reduce or eliminate asymmetry in their forecast errors.

We find that there is clear evidence of asymmetry in both sets of forecasts, mirroring the asymmetry in the underlying series and implying that forecasters have been unable to eliminate the large negative forecast errors which appear to be the main source of the asymmetry. [25] This evidence appears across forecast horizons; there is some indication of reduced asymmetry at shorter horizons (consistent with availability of relevant information close to the realization date) but this evidence appears primarily in the Greenbook forecasts. Whether this apparent out-performance at short horizons is the result of more timely information available to Greenbook forecasters, or to better modelling of asymmetry and therefore better exploitation of available data, cannot be identified on these observations.

Elimination of forecast-error asymmetry of this type, which entails estimation of the non-linear model of the underlying process as well as timely observation of inputs to such a model, is bound to be challenging.[26] It is interesting that Greenbook forecasts seem to show some success in doing so, but it remains clear nonetheless that there is room for substantial

[25] In analogous tests on SPF forecasts for other macroeconomic time series, including industrial production, CPI, real and nominal GDP and treasury bill rates, we found very little evidence of forecast asymmetry, with the possible exception of treasury bill rates.

[26] See Adrian et al. (2018) for some preliminary work in this direction.

progress in finding and exploiting timely information that could predict periods of rapid unemployment increase.

7. References

Adrian, Tobias, Nina Boyarchenko and Domenico Giannone (2018) "Vulnerable Growth." Forthcoming, *American Economic Review.*

Bai, J. and S. Ng (2005) "Tests for Skewness, Kurtosis and Normality for Time Series Data." *Journal of Business and Economic Statistics* 23, 49-60.

Burns, A.F. and W.C. Mitchell (1946) *Measuring Business Cycles*, NBER, New York.

Clark, Todd E., Michael W. McCracken, and Elmar Mertens (2016) "Modeling Time-Varying Uncertainty of Multiple-Horizon Forecast Errors." mimeo.

Croushore, D. (1993) "Introducing: the Survey of Professional Forecasters" *Business Review,* Federal Reserve Bank of Philadelphia, November/December 1993, 3-15.

Croushore, D. and S. van Norden (2014) "How Well Do Central Bankers Understand Fiscal Policy?" *mimeo.*

Croushore, D. and S. van Norden (2018) "Fiscal Forecasts at the FOMC: Evidence from the Greenbooks." *Review of Economics and Statistics*, forthcoming.

Croushore, D. and S. van Norden (2017) "Fiscal Surprises at the FOMC." *Federal Reserve Bank of Philadelphia Working Paper*, 17-13.

Falk, B. (1986) "Further Evidence on the Asymmetric Behavior of Economic Time Series over the Business Cycle." *Journal of Political Economy* 94, 1096-1109.

Federal Reserve Bank of Philadelphia (2015) "Survey of Professional Forecasters Documentation." Update of September 24, 2015.

Hodrick, Robert J. and Edward C. Prescott (1997) "Postwar U.S. Business Cycles: An Empirical Investigation." *Journal of Money, Credit and Banking*, 29(1), 1-16.

Galbraith, J.W. and S. van Norden, (2012) "Assessing gross domestic product and inflation probability forecasts derived from Bank of England fan charts," *Journal of the Royal Statistical Society: Series A (Statistics in Society)*, 175(3), 713-727.

Knüppel, M. (2009) "Testing Business Cycle Asymmetries Based on Autogregressions with a Markov-Switching Intercept." *Journal of Business and Economic Statistics* 27, 544-552.

Knüppel, M. and G. Schultefrankenfeld (2012) "How Informative Are Central Bank Assessments of Macroeconomic Risks?" *International Journal of Central Banking* 17, 87-139.

Knüppel, M. and G. Schultefrankenfeld (2012) "Evaluating Macroeconomic Risk Forecasts." *Deutsche Bundesbank Discussion Paper* 14/2011.

Koop, G. and S.M. Potter (1999) "Dynamic Asymmetries in U.S. Unemployment." *Journal of Business and Economic Statistics* 17, 298-312.

Koopmans, L.H. (1974), *The Spectral Analysis of Time Series*, San Diego, Academic Press.

Mankiw, N. Gregory and Ricardo Reis (2002) "Sticky Information versus Sticky Prices: A Proposal to Replace the New Keynesian Phillips Curve." *Quarterly Journal of Economics* 117(4) 1295-1328.

Montgomery, A.L., V. Zarnowitz, R.S. Tsay and G.C. Tiao (1998) "Forecasting the U.S. Unemployment Rate." *Journal of the American Statistical Association* 93, 478-493.

Morley, J. and I.B. Panovska (2017) "Is Business Cycle Asymmetry Intrinsic in Industrialized Economies?" *University of New South Wales working paper.*

Neftçi, S.N. (1984) "Are Economic Time Series Asymmetric over the Business Cycle?" *Journal of Political Economy* 92, 307-28.

Reifschneider, D. and P. Tulip (2017) "Gauging the Uncertainty of the Economic Outlook Using Historical Forecasting Errors: The Federal Reserve's Approach" *FEDS Working Paper* 2017-020. DOI:10.17016/FEDS.2017.020

Sichel, D.E. (1989) "Are Business Cycles Asymmetric? A Correction." *Journal of Political Economy* 97, 1255-60.

Sims, C.A. (2003) "Implications of Rational Inattention." *Journal of Monetary Economics* 50(3), 665-90.

Stark, T. (2016) "Error Statistics for the Survey of Professional Forecasters for Unemployment Rate [QA,PPs]. Documentation, Real-Time Data Research Center, Federal Reserve Bank of Philadelphia.

Woodford, M. (2003) "Imperfect Common Knowledge and the Effects of Monetary Policy." in *Knowledge, Information, and Expectations in Modern Macroeconomics: In Honor of Edmund S. Phelps*, edited by Philippe Aghion, Roman Frydman, Joseph Stiglitz and Michael Woodford, 25-58. Princeton: Princeton University Press.

Weiss, M.S. (1978) "Modification of the Kolmogorov-Smirnov Statistic for use with Correlated Data." *Journal of the American Statistical Association* 73, 872-875.

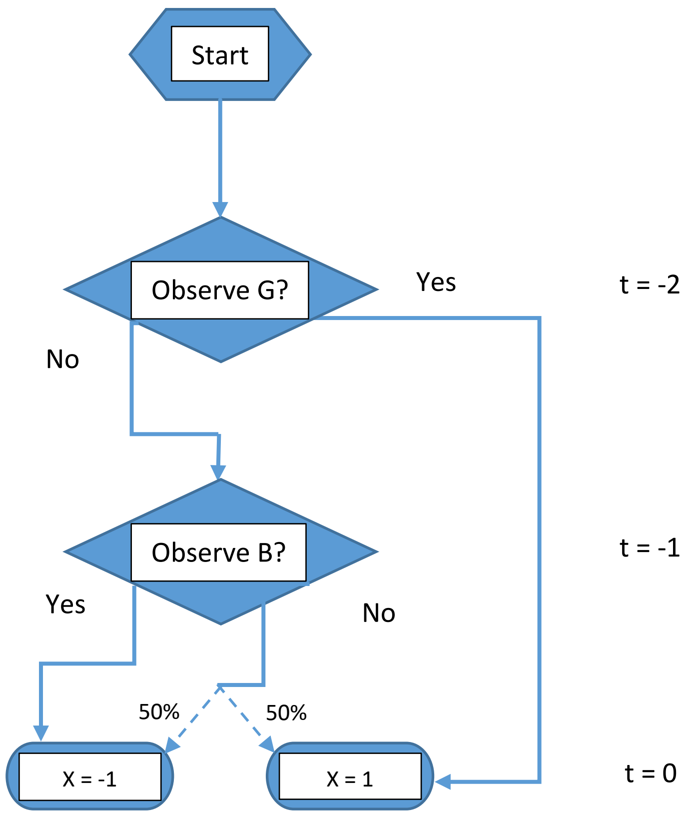

Figure A.14: A Simple Example of Asymmetric Risks

Appendix A. Figure 1

Appendix B. Statistical inference on forecast error symmetry

The following tables present tests of $H_0 : \alpha_1 = 0$ in (6), as described in Section 5.3. Tests on the Greenbook and SPF (median) forecasts are presented on the raw forecast errors at each horizon, and also on the forecast revisions, again as described in 5.3. For Greenbook forecasts, tests on forecast revisions are presented in the right-hand sides of the tables; for SPF, in a separate table following. For the SPF, a comparison is included with analogous tests on the no-change forecast error series (i.e. $\hat{u}_{t+h|t} = u_t$) and on an AR forecast series.

Table B.4: Tests of symmetry in Greenbook forecast errors, part 1

	raw error				revision			
τ	h and F/L	coef	t	p-val (%)	h and F/L	coef	t	p-val (%)
0	0L	-0.13	-1.40	16	0L	-0.13	-1.40	16
0	0F	0.02	0.16	88	0F	-0.24	-2.24	3
0	1L	0.11	0.94	35	1L	-0.23	-2.33	2
0	1F	0.17	1.28	20	1F	-0.41	-3.47	0
0	2L	0.18	1.20	23	2L	-0.24	-2.45	1
0	2F	0.29	1.75	8	2F	0.35	2.67	1
0	3L	0.19	1.10	27	3L	-0.27	-2.56	1
0	3F	0.21	1.13	26	3F	0.39	2.93	0
0	4L	0.25	1.27	20	4L	-0.27	-2.47	1
0	4F	0.32	1.49	14	4F	0.32	2.42	2
0	5L	0.39	1.58	11	5L	0.24	1.62	10
0	5F	0.32	1.25	21	5F	0.31	1.94	5
0	6L	0.38	1.26	21	6L	0.28	1.90	6
0	6F	0.29	1.08	28	6F	0.46	1.77	8
0	7L	0.43	1.22	22	7L	-0.23	-1.54	12
0	7F	0.34	1.14	25	7F	-0.10	-0.45	66
0	8L	0.34	0.98	33	8L	-0.21	-0.84	40
0	8F	0.20	0.55	58	8F	0.51	1.88	6
0.50	0L	-0.17	-1.51	13	0L	-0.17	-1.51	13
0.50	0F	0.30	1.93	5	0F	-0.13	-0.81	42
0.50	1L	0.10	0.62	53	1L	-0.15	-1.14	25
0.50	1F	0.37	1.89	6	1F	0.22	1.17	24
0.50	2L	0.20	0.99	32	2L	0.17	1.43	15
0.50	2F	0.24	1.12	26	2F	0.02	0.10	92
0.50	3L	0.40	1.53	13	3L	0.25	2.07	4
0.50	3F	0.60	2.27	2	3F	-0.04	-0.25	80
0.50	4L	0.56	2.11	4	4L	0.19	1.25	21
0.50	4F	0.62	1.98	5	4F	0.11	0.61	54
0.50	5L	0.59	1.72	8	5L	0.28	1.85	6
0.50	5F	0.58	1.69	9	5F	0.13	0.57	57
0.50	6L	0.59	1.57	12	6L	0.06	0.30	77
0.50	6F	0.49	1.34	18	6F	0.36	1.04	30
0.50	7L	0.71	1.72	9	7L	0.15	0.48	63
0.50	7F	0.50	1.46	14	7F	0.35	1	32
0.50	8L	0.61	1.52	13	8L	0.52	1.15	25
0.50	8F	0.47	1.27	20	8F	0.21	0.74	46

Table B.5: Tests of symmetry in Greenbook forecast errors, part 2

	raw error				revision			
τ	h and F/L	coef	t	p-val (%)	h and F/L	coef	t	p-val (%)
1	0L	0.19	1.28	20	0L	0.19	1.28	20
1	0F	0.08	0.34	73	0F	0.01	0.03	97
1	1L	0.41	1.92	6	1L	-0.23	-1.46	14
1	1F	0.02	0.08	94	1F	0.21	0.94	35
1	2L	0.52	2.14	3	2L	-0.02	-0.17	86
1	2F	0.12	0.44	66	2F	0.31	1.18	24
1	3L	0.48	1.43	15	3L	0.16	1.34	18
1	3F	0.50	1.55	12	3F	0.46	1.73	8
1	4L	0.53	1.44	15	4L	0.27	1.79	7
1	4F	0.35	0.93	35	4F	0.34	1.47	14
1	5L	0.64	1.64	10	5L	0.09	0.53	60
1	5F	0.85	1.99	5	5F	0.73	2.16	3
1	6L	0.79	1.82	7	6L	0.25	1.06	29
1	6F	0.81	1.67	9	6F	0.28	0.70	48
1	7L	1.30	3.56	0	7L	0.37	0.71	48
1	7F	0.81	1.70	9	7F	0.58	0.98	33
1	8L	1.29	4.44	0	8L	0.92	1.10	27
1	8F	1.18	2.65	1	8F	0.23	0.69	49
1.50	0L	-0.07	-0.37	71	0L	-0.07	-0.37	71
1.50	0F	0.24	0.84	40	0F	0.17	0.67	50
1.50	1L	0.20	0.98	33	1L	0.32	1.79	7
1.50	1F	0.39	1.15	25	1F	0.22	0.74	46
1.50	2L	0.59	2.05	4	2L	-0.09	-0.48	63
1.50	2F	0.81	2.43	2	2F	0.32	1.31	19
1.50	3L	0.65	1.98	5	3L	0.49	2.54	1
1.50	3F	0.45	1.45	15	3F	0.15	0.71	48
1.50	4L	0.42	1.06	29	4L	0.42	1.56	12
1.50	4F	0.68	1.53	13	4F	0.14	0.67	50
1.50	5L	1.03	3.49	0	5L	0.30	1.10	27
1.50	5F	0.75	2.11	3	5F	0.72	2.25	2
1.50	6L	1.13	4.53	0	6L	0.49	1.21	23
1.50	6F	0.88	2.53	1	6F	1.32	5.37	0
1.50	7L	1.66	16.41	0	7L	0.98	1.17	24
1.50	7F	1.08	2.67	1	7F	1.99	10.40	0
1.50	8L	-	-	-	8L	1.39	1.38	17
1.50	8F	-	-	-	8F	-	-	-

Table B.6: Tests of symmetry in Greenbook forecast errors, part 3

	raw error				revision			
τ	h and F/L	coef	t	p-val (%)	h and F/L	coef	t	p-val (%)
2	0L	-0.37	-2.89	0	0L	-0.37	-2.89	0
2	0F	-0.09	-0.23	81	0F	0.59	2.43	2
2	1L	0.24	1.21	23	1L	0.40	1.88	6
2	1F	0.27	0.65	51	1F	0.21	0.71	48
2	2L	0.95	4.47	0	2L	0.11	0.53	60
2	2F	0.65	2.51	1	2F	0.32	1.28	20
2	3L	0.70	2.18	3	3L	0.59	3.13	0
2	3F	0.75	3.54	0	3F	0.54	2.61	1
2	4L	0.59	2.31	2	4L	0.21	0.52	60
2	4F	0.63	2.22	3	4F	0.35	1.70	9
2	5L	1.25	4.61	0	5L	0.82	2.21	3
2	5F	0.62	2.38	2	5F	0.53	2.05	4
2	6L	-	-	-	6L	-	-	-
2	6F	1.15	2.87	0	6F	1.21	5.30	0
2	7L	-	-	-	7L	1.36	1.49	14
2	7F	-	-	-	7F	-	-	-
2	8L	-	-	-	8L	-	-	-
2	8F	-	-	-	8F	-	-	-

Table B.7: Tests of symmetry in SPF median forecast errors

	SPF				NC				AR			
τ	h	coef	t	p-val (%)	h	coef	t	p-val (%)	h	coef	t	p-val (%)
0	0	0.34	3.40	0	0	0.41	3.30	0	0	0.13	1.22	22
0	1	0.29	2.20	3	1	0.40	2.63	1	1	0.23	1.87	6
0	2	0.43	2.68	1	2	0.45	2.59	1	2	0.35	2.46	1
0	3	0.46	2.65	1	3	0.54	3.05	0	3	0.39	2.55	1
0	4	0.47	2.37	2	4	0.56	3.23	0	4	0.39	2.40	2
0.50	0	0.29	2.46	1	0	0.78	4.67	0	0	0.23	1.57	12
0.50	1	0.55	3.28	0	1	0.61	3.15	0	1	0.40	2.40	2
0.50	2	0.55	2.75	1	2	0.67	3.40	0	2	0.54	3.20	0
0.50	3	0.69	3.31	0	3	0.64	3.36	0	3	0.60	3.62	0
0.50	4	0.62	2.68	1	4	0.57	3.13	0	4	0.57	3.45	0
1	0	0.22	1.63	10	0	0.55	2.53	1	0	0.42	2.24	3
1	1	0.49	2.39	2	1	0.66	2.49	1	1	0.66	3.02	0
1	2	0.59	2.70	1	2	0.52	1.94	5	2	0.63	3.67	0
1	3	0.65	2.39	2	3	0.35	1.26	21	3	0.67	3.66	0
1	4	0.70	2.42	2	4	0.34	1.32	19	4	0.67	3.76	0
1.50	0	0.17	1.06	29	0	0.50	1.87	6	0	0.13	0.53	60
1.50	1	0.37	1.58	11	1	0.27	0.91	36	1	0.30	1.18	24
1.50	2	0.54	2.02	4	2	0.06	0.28	78	2	0.49	2.23	3
1.50	3	0.47	2	5	3	0.18	0.98	33	3	0.37	2.21	3
1.50	4	0.48	1.96	5	4	0.28	1.31	19	4	0.57	3.65	0
2	0	0.25	1.36	17	0	0.94	3.25	0	0	0.32	1.34	18
2	1	0.64	2.63	1	1	0.44	1.93	5	1	0.14	0.60	55
2	2	0.99	3.33	0	2	0.51	2.50	1	2	0.44	2.37	2
2	3	0.82	3.14	0	3	0.50	2.70	1	3	0.68	5.06	0
2	4	0.96	5.26	0	4	0.55	4.32	0	4	-	-	-

Table B.8: Tests of symmetry in SPF median forecast *revisions*

	SPF				NC				AR			
τ	h	coef	t	p-val (%)	h	coef	t	p-val (%)	h	coef	t	p-val (%)
0	0	0.34	3.40	0	0	0.41	3.30	0	0	0.13	1.22	22
0	1	0.10	0.84	40	1	0.39	2.77	1	1	0.14	1.30	19
0	2	0.29	2.25	2	2	0.39	2.60	1	2	0.20	1.73	8
0	3	0.33	2.40	2	3	0.41	2.66	1	3	0.18	1.50	13
0	4	0.31	2.22	3	4	0.41	2.65	1	4	0.16	1.32	19
0.50	0	0.29	2.46	1	0	0.78	4.67	0	0	0.23	1.57	12
0.50	1	0.60	3.73	0	1	0.53	2.83	0	1	0.33	2.14	3
0.50	2	0.37	2.20	3	2	0.51	2.67	1	2	0.35	2.29	2
0.50	3	0.52	2.95	0	3	0.53	2.75	1	3	0.43	2.49	1
0.50	4	0.43	2.58	1	4	0.52	2.77	1	4	0.48	2.50	1
1	0	0.22	1.63	10	0	0.55	2.53	1	0	0.42	2.24	3
1	1	0.65	3.45	0	1	0.61	2.51	1	1	0.36	1.55	12
1	2	0.78	3.59	0	2	0.61	2.45	1	2	0.39	1.71	9
1	3	0.52	2.42	2	3	0.62	2.57	1	3	0.39	1.51	13
1	4	0.49	2.51	1	4	0.62	2.73	1	4	0.81	2.73	1
1.50	0	0.17	1.06	29	0	0.50	1.87	6	0	0.13	0.53	60
1.50	1	0.46	2.51	1	1	0.46	1.69	9	1	0.15	0.47	64
1.50	2	0.64	3.01	0	2	0.45	1.66	10	2	0.47	1.80	7
1.50	3	0.92	4.39	0	3	0.46	1.70	9	3	0.67	2.40	2
1.50	4	0.80	3.38	0	4	0.46	1.71	9	4	0.72	2.35	2
2	0	0.25	1.36	17	0	0.94	3.25	0	0	0.32	1.34	18
2	1	0.59	3.40	0	1	0.62	2.19	3	1	0.58	1.55	12
2	2	0.31	1.70	9	2	0.76	2.63	1	2	0.42	1.21	23
2	3	-	-	-	3	0.62	2.25	2	3	0.33	1.17	24
2	4	-	-	-	4	0.62	2.11	4	4	0.98	2.66	1