DO MARKETS RESPOND MORE TO MORE
RELIABLE LABOR MARKET DATA?
A TEST OF MARKET RATIONALITY

Alan B. Krueger

Working Paper 5769

DO MARKETS RESPOND MORE TO MORE
RELIABLE LABOR MARKET DATA?
A TEST OF MARKET RATIONALITY

ABSTRACT

Since 1979, the Bureau of Labor Statistics (BLS) has nearly quadrupled the size of the
sample used to estimate monthly employment changes. Although first-reported employment
estimates are still noisy, the magnitude of sampling variability has declined in proportion to the
increase in the sample size. A model of rational Bayesian updating predicts that investors would
assign more weight to the BLS employment survey as it became more precise. However, a
regression analysis of changes in interest rates on the day the employment data are released finds no
evidence that the bond market's reaction to employment news intensified in the late 1980s or early
1990s. For the time period as a whole, an unexpected increase of 200,000 jobs is associated with
an 8 basis point increase in the interest rate on 30 year Treasury bonds, and a 9 basis point increase
in the interest rate on 3 month bills, all else equal. Additionally, announced hourly wage increases
are associated with higher long-term interest rates, while announced changes in the unemployment
rate and revisions to past months' employment estimates have a statistically insignificant effect on
long-term interest rates.

Alan B. Krueger
Woodrow Wilson School
Princeton University
Princeton, NJ 08544
and NBER
akrueger@pucc.princeton.edu

Much existing work on the efficient markets hypothesis tests whether "new information" affects market valuations, and whether "old information" has no impact on market valuations. For example, Pearce and Roley (1985) find that deviations between the announced and expected money supply affect the S&P 500 on the announcement day, while the expected money supply has no effect. Related work on inflation announcements has been done by Schwert (1981), Urich and Wachtel (1985), McQueen and Roley (1993), and others. This approach to testing the efficient markets hypothesis, however, leaves open the possibility that markets may under or over react to new information.[1]

This paper uses an alternative approach to test whether markets respond efficiently to new information. The test is based on the fact that the survey the Bureau of Labor Statistics (BLS) uses to estimate payroll employment was greatly expanded and improved in the 1980s. The BLS employment survey provides the first government economic statistics each month, and is widely reported on and watched by financial players. Stevenson (1996), for example, recently noted, "In the markets, the monthly employment report has become the single most important indicator of economic strength, potential inflation and Fed strategy." Most importantly for our purposes, the sample size of the BLS survey increased nearly four-fold between 1979 and 1995. Because the sample size was increased, data from the survey became more reliable over time. The Root Mean Square Error (RMSE) of estimated employment fell from 121,000 in 1979 to 71,000 in 1994. The central question raised in this paper is: Did the bond market respond more to unexpected movements in the announced employment data as the data became more reliable?

---

[1]Shiller (1979), of course, provides a test of whether long-term interest rates are too volatile. Other work tests for mean reversion in asset price movements, or for calendar effects, or for weather effects (for examples, see DeBondt and Thaler, 1985, Reinganum, 1981, and Saunders, 1993).

Before examining the market responses to employment data, in the next section I describe the BLS Current Employment Statistics (CES) program in more detail. An unusual feature of this program is that the population mean employment is eventually revealed from a complete count of unemployment insurance payroll tax records, the so called "benchmark employment figures." A comparison of the survey results to the benchmark figures provides evidence that the reliability of the monthly releases has indeed improved as the sample size increased. Another finding is that the median prediction of professional forecasters is highly correlated with the survey estimates and with the benchmark figures, and is close to an unbiased estimate of the actual employment change. Interestingly, the median forecast has some ability to predict the ultimate benchmark employment numbers even after conditioning on the survey estimate. This finding suggests that analysts can derive a better estimate of true employment movements by combining the forecast and survey estimates.

The remainder of the paper examines the relationship between the employment data releases and interest rates. Section II presents a theoretical model showing that rational Bayesian investors would place more weight on the employment survey as the size of the sample increases. Another reason the employment data might cause a greater response in the market over time is that many observers believe the Federal Reserve Board shifted its focus away from money supply targets and toward employment monitoring in the mid 1980s.[2] Indeed, most forecasters did not begin making employment forecasts until 1985, and there is no "consensus" estimate of

---

[2]Cook and Korn (1991) provide some evidence of a switch in the short-term interest rate's reaction to employment news after the change in Fed policy. One anomalous finding in their results, however, is that the strongest reaction to employment news occurred in the 1980-82 period.

employment growth prior to 1985. If the shift in the Fed's emphasis was not motivated by improvements in the employment data, this change in monetary policy may confound the relationship between employment data and market reactions. Thus, a finding of increased market responses to the employment numbers as the sample size increased may reflect a rational response to more precise data, or a reaction to the change in the Fed's policy, or both. But a finding of a constant response to employment news would present a puzzle for two reasons.

As expected, results presented in Section III and IV indicate that movements in the BLS survey data have a significant and sizable effect on the 30 year Treasury bond yield and on short term Treasury bill yields on the day the employment numbers are released. An unanticipated increase of 200,000 jobs is associated with about a 4 to 8 basis point increase in the long bond yield, and a larger increase in the short-term bond yield. *Contrary to expectations, however, the effect of reported employment changes on interest rates was at least as strong in the early 1980s as it was in the mid 1990s.* This finding is quite surprising in view of the widely held belief that financial markets only followed the employment releases after the Fed shifted its focus away from the money supply in the mid 1980s.

Four additional findings are noteworthy. First, interest rates are not significantly affected by announcements of revisions to past months' employment data; only the latest month employment change seems to matter. Second, announced increases in hourly pay have a statistically significant, positive impact on interest rates. Third, unexpected changes in unemployment are insignificantly related to long-term interest rates, but significantly related to short-term rates. Fourth, the long-term interest rate is significantly related to the forecasted employment growth conditional on the deviation between the forecast and the employment

survey. This latter finding results from a positive correlation between the forecast error and the employment forecast, as the forecast has an insignificant effect on the bond yield when the forecast error is not conditioned on. Because investors cannot condition on the employment news in advance of the employment release, this finding is not evidence of an inefficient market.

I. Evaluation of BLS Employment Data

The BLS conducts the CES survey of business establishments each month to make timely estimates of nonfarm employment, hours of work, and pay.[3] The survey pertains to the pay period covering the 12th day of each month. The CES data for the previous month are typically released at 8:30 AM on the first Friday of each month, although the data may be released on other days if the first Friday falls on a holiday (e.g., July 4th) or if there is insufficient time between the survey and first Friday of the month (e.g., March). On the release date, the BLS reports the first estimate of employment for the previous month, as well as the first revision to the employment estimate from two months ago and the second revision to the estimate from three months ago. In addition, the BLS reports information from the household survey (the Current Population Survey), which includes the monthly unemployment rate.

The sample of establishments underlying the CES survey is drawn from unemployment insurance (UI) tax records. The sample is a stratified sample, with strata consisting of geographic location, establishment size, and industry cells. Sampled establishments may stay in the survey for several years and even decades. The BLS uses the survey data to make a "linked relative" estimate of employment, in which only establishments that are in the survey in two

---

[3]For a detailed description of the CES survey, see BLS (1996).

adjacent months are used to estimate the change in employment. This method was selected to take advantage of the high month-to-month covariance in employment within establishments. To adjust for births of establishments that are not captured in the sampling scheme, the BLS adds a "bias" factor to the survey estimate each month.[4] The bias factor is based in part on a model of births of establishments. Additionally, after about a 7 month lag the BLS receives a census of employment based on the universe of UI payroll tax records. This count also is used to adjust the bias factor. Each June the complete count of employment from the UI files is used to make a "benchmark adjustment" to the previously released employment data; in principle, the benchmark data are the population mean for nonagricultural employment.[5]

Beginning in the late 1970s, the BLS greatly increased the sample of establishments that were surveyed to improve geographic and industry coverage. Additionally, steps were taken to increase the response rate; most importantly, the BLS moved to an automated system in which respondents could answer the survey by pushing the buttons of a touch-tone telephone. Both of these efforts led to a higher sample size. Figure 1 shows the sample size used to estimate the first employment report each month since 1979. The sample increased from an average of 68,000 respondents in the first quarter of 1979 to an average of 239,000 in the first quarter of 1996. Under random sampling, this increase in the sample size would be expected to reduce the

---

[4]The bias factor is set at the beginning of each quarter, and remains fixed for the quarter.

[5]The benchmark adjustment is made to the March data, and wedged in uniformly for the previous 12 months. In making the benchmark adjustment, the BLS also uses additional sources to count the small number of nonagricultural workers who are not covered by unemployment insurance. This study uses benchmark data through February 1995. In June 1996, the BLS released more recent benchmark data, but also changed its seasonal adjustment method to take account of 4 or 5 week intervals between survey dates. Because forecasts are based on the older seasonal adjustment method, the most recent benchmark data are not used.

Figure 1

Sample Size at First Closing, 79:1–96:3
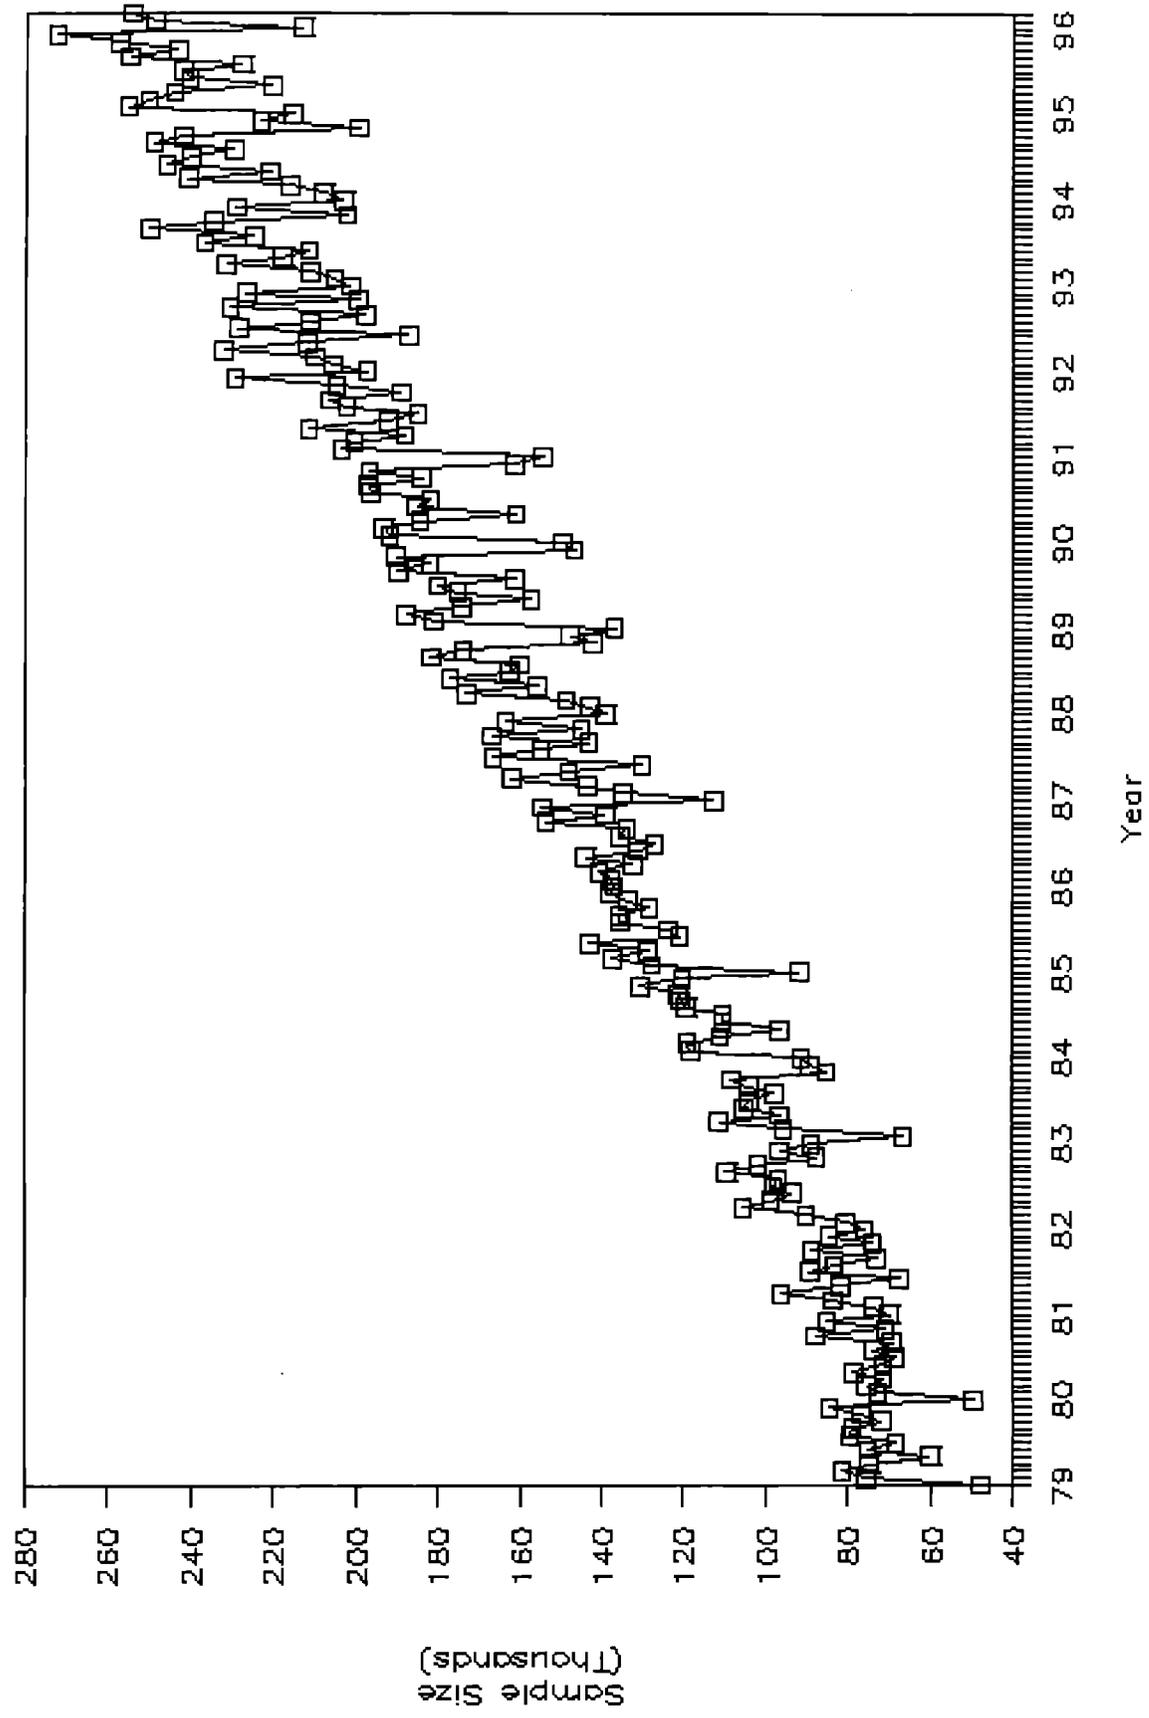
standard error of the estimates by 47 percent.

Table 1 reports summary statistics for seasonally adjusted monthly employment changes.[6] The first-reported employment change, denoted e_1, is the difference between the first report of employment for the latest month and the second report of employment for the preceding month. For example, if the data are released in September, the first-reported change would be the first estimate for August minus the second estimate for July. The second report is the revised August estimate minus the third estimate for July (denoted e_2), and this figure would be released in October. The third report would be available in November, and is the third estimate of August's employment minus the third estimate of July's (denoted e_3). The revised employment reports incorporate data from sampled establishments that responded late. The benchmark employment change is the final estimate from the UI files (denoted μ), which I take to be the population mean.

The forecasted employment change (denoted f) is the median employment change forecasted by market specialists surveyed by Money Market Services International (MMS). The MMS employment data only are available from January 1985 forward. Consequently, I have constructed my own monthly forecast (denoted f_c) for 1979-96 by recursively estimating a regression of the first-reported employment change on lags of the preceding six months of employment changes and the number of new unemployment insurance claims in the week containing the 12th of the month. The sample used to estimate the 1979 forecasting equation consists of monthly observations from December 1969 to December 1978. The coefficients of

---

[6]Because market forecasts are for seasonally adjusted data, and because most discussion focuses on the seasonally adjusted data, all of the data used in this paper are seasonally adjusted.

Table 1

Descriptive Statistics for Monthly Employment Changes

Variable	Mean Value	Mean Absolute Value	Std. Dev.	Period
First Report Employment (e_1)	135.8	210.0	208.1	79:1-96:4
Second Report Employment (e_2)	130.4	213.8	216.4	79:1-96:2
Third Report Employment (e_3)	147.8	222.4	216.8	79:1-96:2
Benchmark Employment (μ)	142.1	212.8	205.5	79:1-95:2
MMS Median Forecast (f)	156.1	168.9	99.7	85:1-96:3
Constructed Forecast (f_c)	122.0	160.7	133.1	79:1-96:4
First Report - Second Report	6.0	52.3	70.3	79:1-96:2
First Report - Third Report	-11.4	67.3	91.4	79:1-96:2
First Report - Benchmark	-5.9	88.9	114.7	79:1-95:2
First Report - MMS Forecast	1.6	95.7	124.5	85:1-96:4
First Report - Constr. Forecast	5.0	109.0	148.1	85:1-96:4

Notes: All numbers are in 1,000s. MMS forecast is the median of professional forecasters. The constructed forecast is the predicted employment change from a recursive regression of the first-reported employment change on six of its lags and seasonally adjusted initial unemployment insurance claims for the week including the 12th of each month. All employment data are seasonally adjusted.

the forecasting equation were re-estimated each year, rolling the sample forward to include the latest 12 months of data. The most current coefficients are used to forecast employment each year. The sample used to construct each month's forecast covers a time period before that month's employment data were released.

Several conclusions are evident from Table 1. First, the mean employment change is fairly close among all of the estimates. Second, the data are quite noisy. The average absolute revision between the first-reported employment change and the final benchmark figure is 89,000; the standard deviation of the revisions is 115,000. Third, over the period when the MMS data are available, the average of the median forecasts is fairly close to the average first report of employment growth, so the median forecast provides essentially an unbiased estimate. Fourth, the standard deviation of the forecast errors (i.e., the first-reported employment change minus the MMS forecast) is substantial, suggesting a good deal of news is revealed on the day the employment numbers are released.

The MMS forecast is reasonably strongly correlated with the BLS estimates of employment growth. Table 2 shows that the MMS forecast has a correlation of .73 with both the first report of employment and the final benchmark number. A regression of the first report of the employment change on the forecasted change yields the following coefficient estimates, with standard errors shown in parentheses:[7]

$$(1) \quad \hat{e}_1 = -38.45 + 1.25 \, f \qquad R^2 = .54 \qquad \rho = -.03 \, .$$
$$\phantom{(1) \quad \hat{e}_1 = } (18.64) \quad (.10) \qquad\qquad\qquad (.09)$$

_____

[7]The equation was estimated using the Cochrane-Orcutt procedure to correct for possible first-order serial correlation. The sample size is 121 monthly observations.

Table 2 -- Correlations Among Employment Change Measures

	1st Rept (e_1)	2nd Rept (e_2)	3rd Rept (e_3)	Benchmark (μ)	MMS Forecast (f)	Constr. Forecast (f_c)
1st Rept	1.00					
2nd Rept	0.94	1.00				
3rd Rept	0.91	0.98	1.00			
Benchmark	0.82	0.88	0.88	1.00		
MMS Forecast	0.73	0.74	0.72	0.73	1.00	
Constr. Forecast	0.63	0.65	0.64	0.63	0.82	1.00

Sample size is 122 months. MMS forecast is the median of professional forecasters. Constr. forecast is the constructed employment forecast from a recursive regression equation based on six lags of e_1 and seasonally adjusted initial unemployment insurance claims for the week including the 12th of each month; see text.

Although the median forecast is virtually <u>mean</u> unbiased, movements in the forecast tend to understate movements in the employment survey. This result generates a positive correlation between the forecast error and the forecast itself.

To further explore the predictive power of the MMS forecast, I regressed the benchmark employment (population mean) on the first report and MMS forecast. The results, which are presented in equation (2) below, indicate that the forecast has surprising explanatory power, even conditional on the survey data. Combining the survey data and the MMS forecast yields a more accurate prediction of the benchmark data. Equation (2) suggests that when the survey estimate of employment growth deviates substantially from the forecasted amount, the survey data are likely to be adjusted toward the forecasted amount.[8] Indeed, in 65% of the months when the survey data exceeded the forecast the survey data were adjusted downward in the benchmark revision, and in 64% of the months when the survey data were less than the forecast, the survey data were adjusted upward.

$$(2) \qquad \hat{\mu} = \underset{(28.92)}{20.33} + \underset{(.06)}{.59e_1} + \underset{(.14)}{.33f} \qquad R^2 = .71 \qquad \underset{(.08)}{\rho = -.47}$$

According to Table 2, the correlation between the MMS forecast and the constructed forecast is .83, suggesting that, to some extent, professional forecasters base their predictions on

---

[8]Further evidence that the survey may exaggerate swings in employment comes from regressing the benchmark employment change on e_1. This regression yields a coefficient of .76 on e_1, with a standard error of .05, suggesting that large movements in first reported employment growth tend to be adjusted toward the mean in the benchmark revision. This finding and the positive effect of the forecast in equation (2) may partially result from smoothing due to revised seasonal adjustment factors in the benchmark data. But a qualitatively similar set of results is found when the third reported employment change is used instead of the benchmark data, so ex-post seasonal adjustments are not the entire story.

a model similar to the one used to derive the constructed forecast. The constructed forecast is also positively correlated with the BLS data, but not as strongly as the median of the professional forecasters. In contrast to the results for the MMS forecast, when equation (1) is estimated using the constructed forecast, the coefficient on the constructed forecast is less than 1. It is also worth noting that when equation (2) is estimated using the constructed forecast instead of the MMS forecast, the coefficient on the survey employment rises, and that on the forecast falls.[9] In sum, the constructed forecast is not as strong a predictor of employment growth as the median MMS forecast, but nonetheless is reasonably correlated with the BLS data.

Another issue concerns revisions to previous months' data. As far as I know, forecasters do not systematically try to predict BLS revisions to earlier months' employment data, even though the revisions are often substantial (see Table 1). To explore the time-series properties of the revisions, I regressed the second report of the employment change for a given month on that month's first-reported change, and the third-reported change on that month's second-reported change. In both cases, the intercept was insignificantly different from zero, and the coefficient on the previous report was very close to one.[10] Because the successive reports appear to follow a random walk, the deviation between the second and first report of employment growth (or the deviation between the third and second report) probably provides a reasonable estimate of the unanticipated revision.

---

[9]If the sample is restricted to the same months used for equation (2), the coefficient on the constructed forecast is .20.

[10]A regression of e_1 on e_2 yields a coefficient of .98 with a standard error of .026, and a regression of e_2 on e_3 yields a coefficient of .98 with a standard error of .014.

Improved Precision?

Table 3 examines whether the precision of the first-reported employment data improved as the sample size increased. Although it is a simplification of the BLS's estimation procedure, the functional form used in these models was derived under the assumption that a univariate mean was estimated from a randomly selected sample. Specifically, write the variance of the mean $(\bar{x})$ of a sample of n observations, $x_1,\ldots x_n$, as $var(\bar{x}) = E(\bar{x}-\mu)^2 = \sigma^2/n$, where μ is the population mean and σ^2 is the variance of x. Taking logarithms of each side yields:

$$(3) \quad \log E(\bar{x}-\mu)^2 = \log(\sigma^2) - \log(n).$$

Intuitively, the sampling variance should decline in proportion with the increase in the sample size. In the CES program, the benchmark employment can be thought of as the population mean, and the first report as the sample mean.

Column 1 of Table 3 presents a regression of the log of the squared deviation between the benchmark and first-reported employment change on the log of the sample size at the time of the first report. As equation (3) predicts, the coefficient on the log of the sample size is insignificantly different from -1, indicating that the sampling variability has declined at about the rate one would expect (with random sampling) as the sample size increased.[11] To check whether the results just reflect a coincidental trend toward less variable employment data, in column 2 a linear time trend is added to the model. The coefficient on the log sample size

---

[11]In results not reported here, I estimated the model in column 1 using the log of the squared deviation between the first report and the second or third report as the dependent variable. The coefficient on the log sample size in these equations was also insignificantly different from -1, indicating a decline in the magnitude of revisions as the sample size increased.

Table 3: Tests of Employment Data Reliability as Sample Size Grows

Explanatory Variable	**Dependent Variable**					
	$\log[e_1-\mu]^2$		$\log[e_1-f]^2$		$\log[f-\mu]^2$	
	(1)	(2)	(3)	(4)	(5)	(6)
Constant	23.69 (4.10)	36.05 (14.24)	35.52 (9.98)	73.97 (24.97)	-17.80 (10.65)	-11.91 (25.09)
Log n	-1.31 (0.35)	-2.42 (1.28)	-2.25 (0.82)	-5.67 (2.20)	2.14 (0.88)	1.61 (2.21)
Trend (months)	---	0.008 (0.009)	---	0.021 (0.013)	---	0.003 (0.013)
R^2	.06	.07	.05	.07	.05	.05
Sample Size	193	193	135	135	121	121

Notes: e_1 is the first employment change report; μ is the final benchmark employment change; f is the median forecast of the employment change from MMS. Log n is the log of the number of establishments used by BLS to estimate the first employment change. Sample period is 1979:1-1995:3 in columns 1 and 2, 1985:1-96:3 in columns 3 and 4, and 1985:1-95:2 in columns 5 and 6. One observation was dropped in columns 5 and 6 because the deviation between μ and f is 0; if this is set to 1, the results are qualitatively similar. Standard errors are in parentheses.

increases in magnitude in this model, although its standard error increases substantially as well. Furthermore, the coefficient on the time trend suggests that the sampling variability increased over time, although this coefficient is not statistically significant.

As a further check on these results, I regressed the squared deviation between e_1 and μ on $1/n$ and an intercept. The coefficient on $1/n$ in this regression is an estimate of σ^2, and the intercept in this regression should be 0. The coefficient on $1/n$ was large (1.4×10^9) and statistically significant, whereas the intercept was small and statistically insignificant.

If the employment survey became less noisy over time, then one would expect the forecast errors to decline. Columns 3 and 4 of Table 3 test this hypothesis with the MMS data. As expected, the results indicate that the forecast error has declined with the increase in the sample size.[12] An alternative interpretation, however, is that forecasts have improved as forecasters have become more sophisticated. This hypothesis is explored in Columns 5 and 6, which use the log of the squared deviation between the forecast and the benchmark employment numbers as the dependent variable. The motivation for estimating this model is that, if the forecasts have improved, they should do a better job predicting the population mean as well as the sample mean. Contrary to this prediction, however, the forecast errors around the population mean have actually increased as the sample size increased.

As a final test of the improved reliability of the data, I have re-estimated equation (2) also including an interaction between the survey employment and its underlying sample size. This interaction term was statistically significant ($t=2.5$) and positive, indicating that the survey data became a stronger predictor of the population mean as the sample size increased, conditional on

---

[12]Similar results were obtained using the constructed forecast instead of the MMS forecast.

the forecast. Taken together, these results suggest that the noise in the survey estimates declined as the sample size increased.

II. A Model of Bayesian Investors with More Precise Data

Intuitively, one would expect rational investors to place more credence in the BLS survey of employment growth as the survey became more reliable. This intuition can easily be formalized for the case of Bayesian updating. For example, assume that the underlying employment change data are normally distributed, with variance σ^2, and unknown mean μ. Also assume that the distribution of priors about μ is normally distributed, with variance v^2 and mean f. Suppose a random sample of n observations is drawn from the population, and the average of this sample is denoted e. The mean of the posterior distribution (μ) after the sample mean is observed is:

$$(4) \qquad \hat{\mu} = \frac{\sigma^2}{\sigma^2 + nv^2} f + \frac{nv^2}{\sigma^2 + nv^2} e .$$

Equation (4) specifies the posterior estimate of the mean as a weighted average of the mean prior expectation and the sample average (see DeGroot, 1975; pp. 269-271). For fixed values of σ^2 and v^2, the relative weight assigned to the sample mean increases as the sample size increases.

Notice that one can re-write the posterior estimate as:

$$(4') \qquad \hat{\mu} = \frac{nv^2}{\sigma^2 + nv^2} (e\text{-}f) + f = \frac{1}{1 + \psi/n} (e\text{-}f) + f$$

where ψ is the ratio of σ^2 to v^2. Deviations between the survey estimate and the prior expectation receive more weight as the sample size increases, and as ψ decreases.

It seems reasonable to take the consensus estimate of professional forecasters as a measure of the mean of the prior distribution, and the first-reported employment change as the sample mean. If n were fixed, the regression results in equation (2) would provide an estimate of the optimal "Bayesian weights" to assign to the prior and to the sample average (e_1). However, n is not fixed. The optimal Bayesian weight to apply to the forecast error in equation (4') can be estimated directly by Nonlinear Least Squares (NLS). Specifically, I used NLS to estimate the parameter ψ in equation (4'), using the benchmark data as the dependent variable (μ), e_1 as the sample average, the MMS forecast as f, and the BLS sample size as n.[13] Plugging in actual values of n, the results imply that the optimal Bayesian weight to apply to deviations between the forecast and employment survey was 0.32 in 1979, 0.47 in 1985 and 0.66 in 1996.

Next consider how the precision of the employment survey might affect the bond market. An increase in employment is typically interpreted as a sign that the labor market is tightening, and that wage-push inflation may follow. Because the bond yield is positively related to expectations of future inflation, any news raising the probability of higher inflation would be expected to increase the bond yield. Moreover, equation (4) indicates that an increase in surveyed employment growth of a given magnitude will lead investors to revise their expectations of true employment growth by a greater margin if the sample size is larger. Consequently, the bond market would be expected to react more to unexpected blips from the survey if the survey is based on a larger sample size.

Formally, assume the bond yield is a function of expected employment growth, $g(\hat{\mu})$, and that $g' > 0$. The market's reaction to a given change in survey employment as the sample size

---

[13]The coefficient estimate of ψ equalled 127,101, with a t-ratio of 3.81.

increases is given by $\delta^2 g(\hat{\mu})/\delta n\delta e$, which is:

(5) $\quad \delta^2 g(\hat{\mu})/\delta n\delta e = g''(\delta\hat{\mu}/\delta e)(\delta\hat{\mu}/\delta n) + g'(\delta^2\hat{\mu}/\delta n\delta e) > 0.$

Notice that if priors are unbiased, the first term drops out in expectation because $\delta\hat{\mu}/\delta n = 0$ if

$e = f$. From (4), the second term is clearly positive, indicating a larger reaction of the bond

yield to a given increase in the survey estimate of employment growth, when the survey estimate

is based on a larger sample size. If $g()$ is linear, the market reaction to the employment surprise

would increase in proportion to $\dfrac{1}{1 + \psi/n}$, all else equal. The next two sections test whether

interest rate reactions to employment news have varied with the precision of the employment

survey.

III. Estimating Market Reactions

Table 4 presents several Ordinary Least Squares (OLS) regression models simply relating

the close-to-close change in the benchmark 30 year Treasury bond yield to the announced change

in survey employment (e_1) on the day the employment news is released. Because a consensus

forecast of employment growth is not available until 1985, results are first presented for a larger

sample without subtracting forecasted employment growth from announced employment growth.

To facilitate comparison to models that adjust for expectations, two time periods are used:

January 1979 to April 1996, and February 1985 to April 1996. Although some of the

employment news was anticipated, the results indicate that an increase in employment is

associated with a statistically significant increase in the bond yield. In column 1 or 4, for

example, an increase in employment of 200,000 jobs is associated with a 4 basis point increase

Table 4: OLS Estimates of Effect of Employment Releases on Bond Yields

Dependent Variable: Change in 30 year T-Bond Yield (x 100)

Explanatory Variable	1979:2 - 1996:4			1985:1 - 1996:4		
	(1)	(2)	(3)	(4)	(5)	(6)
e_1	.019 (.004)	.030 (.017)	.026 (.018)	.018 (.005)	.056 (.051)	.070 (.053)
$e_1 \times \dfrac{1}{1 + \psi/n}$	...	-.023 (.034)	-.017 (.036)	...	-.063 (.086)	-.088 (.089)
Δ hourly rate	...	...	32.392 (25.919)	...	...	35.933 (29.822)
Δ unemployment rate	...	...	-3.871 (3.836)	...	...	-3.947 (5.343)
P-value for F-test of year, year squared, Wednesday and Thursday dummies	.563	.496	.551	.462	.396	.485
P-value for F-test of e_1 and $e_1 \times 1/(1 + \psi/n)$	...	.000	.000	...	.002	.003
R^2	.161	.163	.175	.141	.144	.156

Notes: Standard errors are in parentheses. e_1 is the first-reported employment change, n is the underlying BLS sample size, and ψ is an estimate of σ^2/v^2 (see text for details). All equations also include year, year squared, a dummy indicating whether the release date is a Wednesday, and a dummy indicating whether the release date is a Thursday. Sample consists of 202 monthly observations in columns 1 and 2, 201 in column 3, and 132 in columns 4-6.

in the bond yield. As shown below, the response is about twice as large if one looks at unanticipated employment changes.

To test whether the market reactions have become stronger as the CES sample size increased, the models in columns 2 and 5 also include an interaction between the first-reported employment change and the optimal Bayesian weight, defined as $e_1 \times \dfrac{1}{1 + \psi/n}$, where ψ was estimated by NLS as previously described. Contrary to what one would expect with Bayesian updating, this interaction term is statistically insignificant and slightly negative. The Bayesian weight is a nonlinear function of the sample size. To explore the robustness of this interactive effect, I estimated two alternative functional forms: (1) interact the employment change with the linear sample size; (2) interact the employment change with the log sample size. In both specifications, the interaction was negative and statistically insignificant. Further evidence that the market reaction did not intensified as the sample size increased comes from estimating the model in column (1) on two subsamples, one covering 1979-84, and one covering 1985-96. The coefficient (and standard error) on e_1 for the 1979-84 sample is .021 (.006), compared to .018 (.005) for the post-1984 period.

The other variables in the model are individually and jointly statistically insignificant. The change in the unemployment rate and change in the hourly wage of production/ nonsupervisory workers, which are also announced on the day the employment data are released, have small and statistically insignificant effects. A 0.5 percentage point increase in the unemployment rate is associated with a 1.8 basis point decline in the yield rate, while a 6 cent increase in the hourly wage is associated with a 2.0 basis point increase in the yield rate. Additionally, dummy variables indicating the day of the week the data were released and a

quadratic time trend are jointly insignificant.

Only unanticipated news should affect financial markets. To adjust for expectations, results in Table 5 use the MMS data to calculate the employment change forecast error $(e_i\text{-}f)$. The estimates are based on data for 1985 to 1996. Comparing column (1) of Table 5 to column (4) of Table 4 indicates that an unanticipated increase in employment has a larger impact on the bond yield than does the total increase in employment. An unanticipated increase of 200,000 jobs is associated with a 7.6 basis point increase in the bond yield. If the market applies the optimal Bayesian weights to the employment survey, the interaction between employment growth and the (nonlinear function of the) sample size in column (2) would be significant and positive, and the forecast error itself would be insignificant. Again, however, the interaction term is negative and statistically insignificant, providing no evidence of a stronger reaction to more precise data.[14]

Figure 2 provides a scatter diagram of the change in the bond yield versus the unexpected change in employment. An upward sloping relationship is apparent.[15] There is no obvious change in the relationship over time. A bivariate regression of the change in the interest rate on the forecast error using the 1985-89 sample yields a slope of .053 (s.e. = .011), whereas the same regression using the 1990-96 sample yields a slope of .033 (s.e. = .007). As a rough check on the power of the estimates, suppose the effect of employment surprises had increased by 40

---

[14]Similar results are found if the linear sample size or log of the sample size is interacted with the employment surprise.

[15]Notice that Figure 2 displays little evidence of a nonlinear relationship between the change in the interest rate and employment surprises. More formal statistical tests based on fitting a quadratic in the forecast error and a linear spline that allows for differential effects of positive or negative employment shocks also supported a linear relationship.

Table 5: OLS Estimates of Effect of Deviations from Forecasts on the Long Bond Yield, February 1985-April 1996
Dependent Variable: Change in 30 Year T-Bond Yield (x 100)

Variable	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
$e_1 - f$	.038 (.006)	.109 (.067)	.040 (.007)	.040 (.007)	.124 (.067)	.042 (.007)	....	.040 (.006)	....
$(e_1-f) \times \dfrac{1}{1+\psi/n}$	....	-.120 (.114)	....	....	-.144 (.114)	....	....	....	....
Revision of last month's employment change			....	.016 (.013)	.016 (.013)	.022 (.013)	.015 (.015)	....	....
Revision of 2 months ago employment change			....	.024 (.026)	.020 (.026)	.019 (.026)	.030 (.029)	....	....
Δ hourly wage	....	....	55.807 (26.956)	53.542 (27.184)	55.511 (27.162)	53.908 (26.589)	18.184 (29.986)	....	....
Deviation of unemployment rate change from MMS forecast	....	....	-5.233 (4.574)	-4.396 (4.634)	-5.105 (4.656)	-5.055 (4.540)	-11.493 (5.107)	....	....
Median employment forecast	....	....	....	....	....	-.022 (.008)	-.016 (.010)	-.018 (.008)	-.013 (.010)
P-value for F-test of year, year squared, Wednesday and Thursday dummies	.400	.400	.591	.525	.523	.147	.051	.185	.063
P-value for F-test of e_1 and $(e_1-f) \times 1/(1 + \psi/n)$	....	.000	....	....	.000	....	....	....	....
P-value for F-test of the two employment revisions	....	....	....	.393	.422	.232	.427	....	....
R^2	.271	.278	.300	.311	.320	.346	.123	.297	.068

Notes: e_1 is the first employment change; f is the MMS median employment forecast; n is the sample size of the BLS employment survey; and ψ is an estimate of σ^2/v^2 (see text for details). Standard errors are in parentheses. Sample size is 132.

Figure 2: Change in 30 Year Treas. Bond Yield vs. Forecast Error
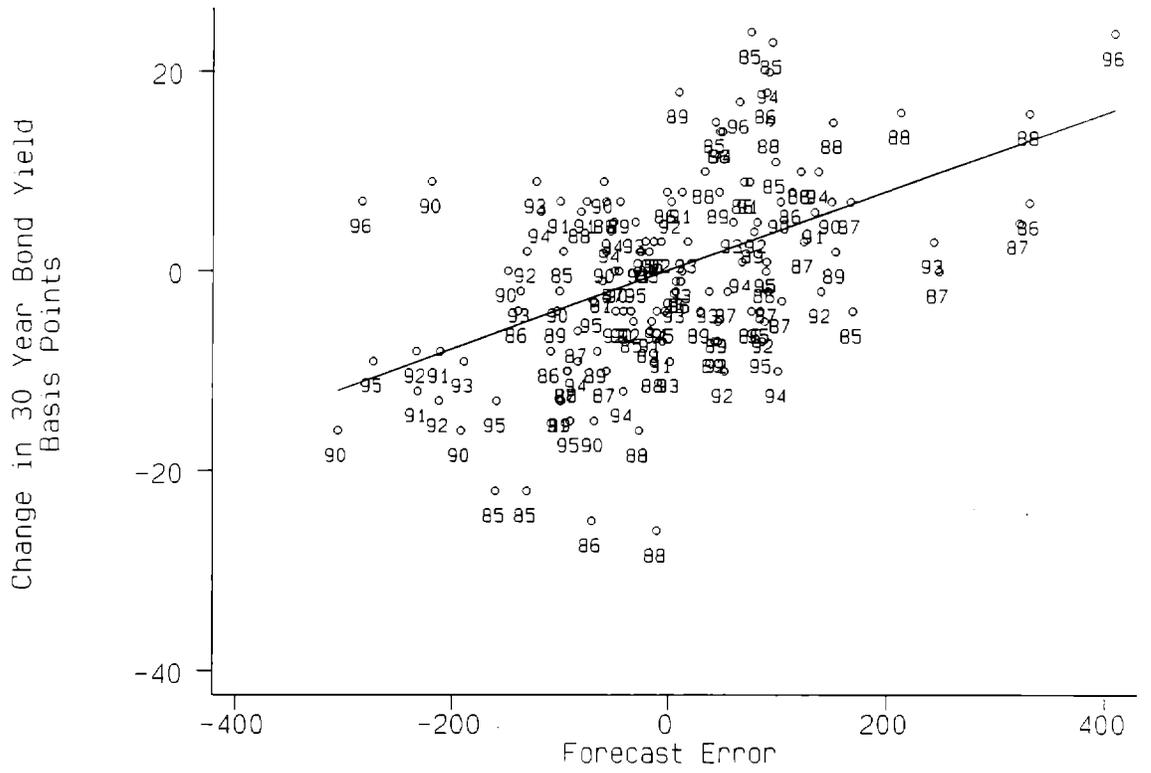
percent since 1985-89, i.e., to a coefficient of .074 (=1.4 x .053). The 1990-96 estimate is statistically different from .074, so the data likely would have the ability to discriminate between effects of this magnitude.

To further explore changes in the responsiveness of interest rates to unexpected employment news over time, bivariate regressions of the change in rates on the employment surprise were estimated for each two year period. Figure 3 illustrates the predicted change in rates associated with an unexpected increase of 200,000 jobs based on these regressions. If anything, the responsiveness of the interest rate to employment news declined in the late 1980s, and may have rebounded somewhat in 1995-96. But since the standard error of each of these estimates is about 3 basis points, too much should not be made of individual points.

McQueen and Roley (1993) find that interest rate reactions to real economic news are invariant to the state of the business cycle, in contrast to stock market returns. A similar pattern holds for employment surprises in these data: an interaction between the unemployment rate and forecast error is insignificant if it is added to the model in column 1. Thus, business cycle effects are unlikely to confound any effect of more precise data.

Column 3 of Table 5 indicates that deviations between the unemployment rate and the MMS consensus forecast of the unemployment rate are statistically insignificant. Previous studies that have found a significant relationship between the bond rate and unexpected changes in unemployment generally have not controlled for the effects of employment changes (e.g., Hardouvelis, 1988 and Prag, 1994). If the model in column 3 is estimated without the employment forecast error, the coefficient on unexpected changes in unemployment becomes -14, with a t-ratio of -3.5. Thus, the negative correlation between unexpected employment growth

Figure 3

Long Bond Mkt. Response to Forecast Error

Basis Point Incr. for 200k Extra Jobs

Basis Point Change

1985-86   1987-88   1989-90   1991-92   1993-94   1995-96
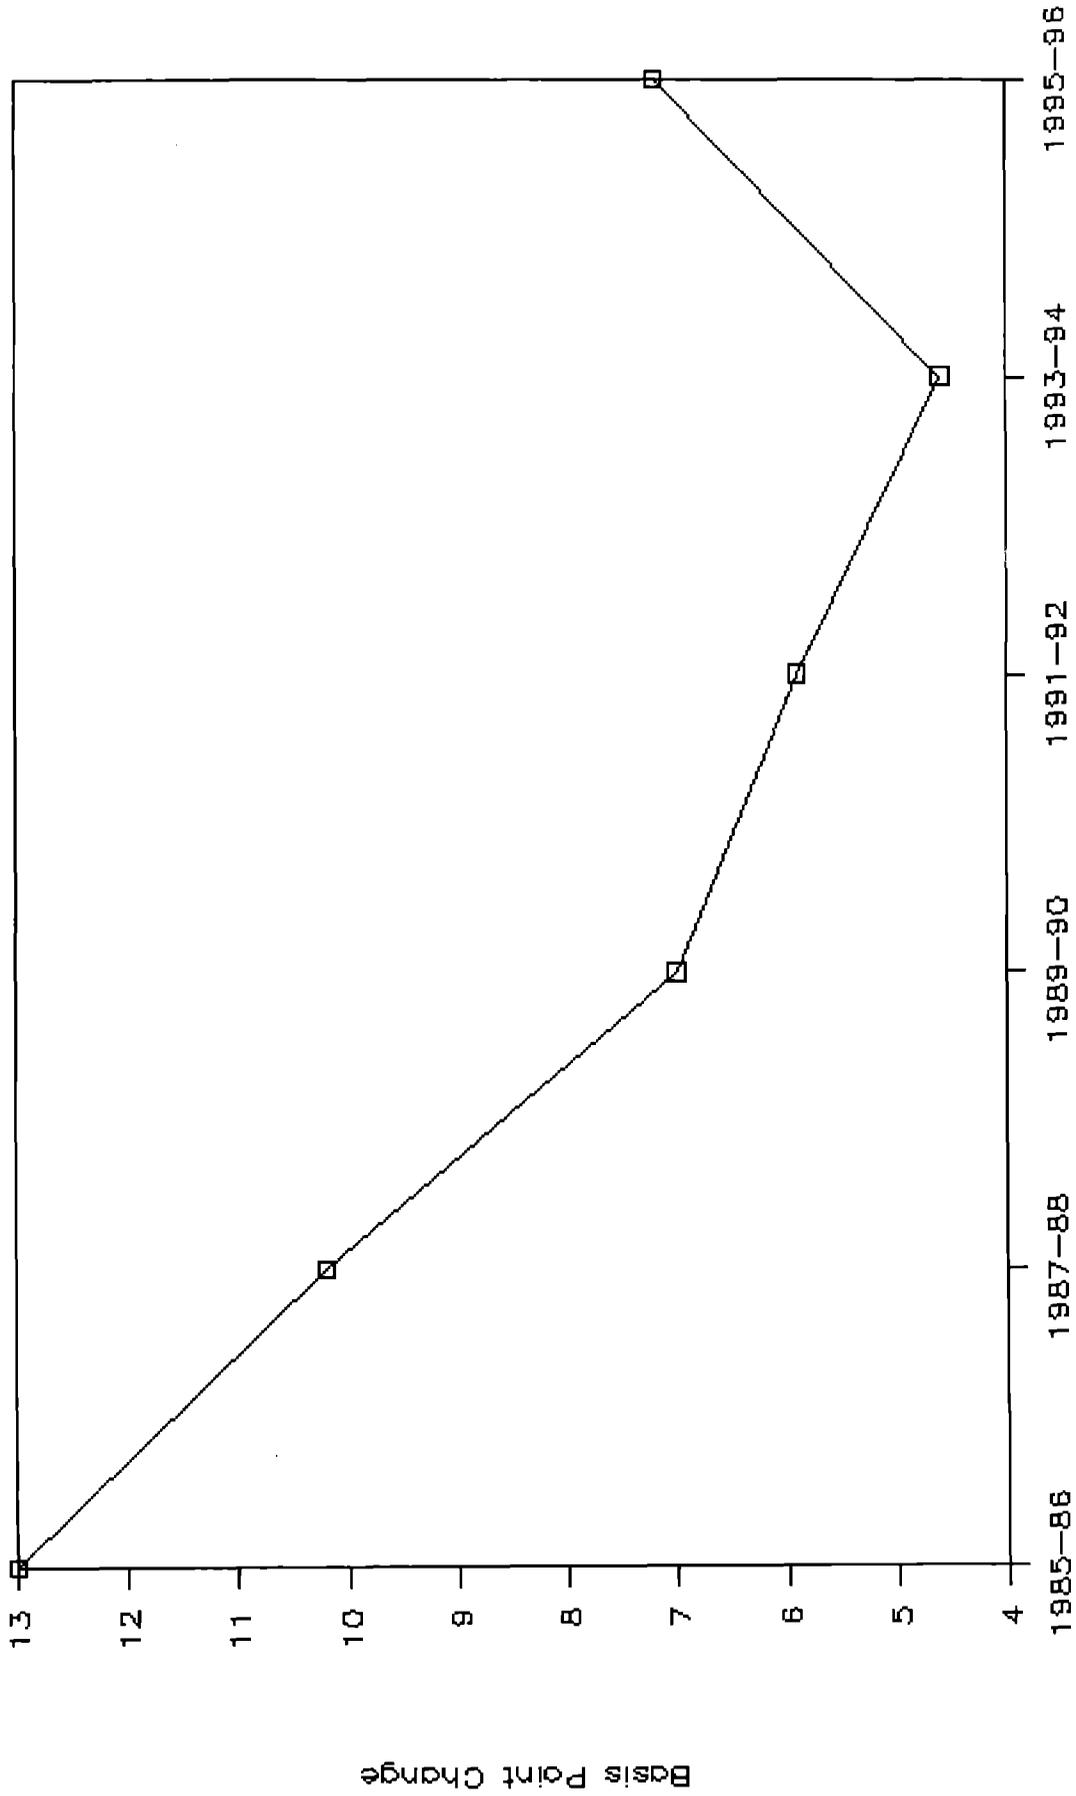
and unexpected unemployment rate changes (r = -.26) may partially drive earlier findings of a significant effect of unemployment shocks on the bond rate. As discussed below, another issue involves bond maturity: short-term rates are more sensitive to the unemployment rate.

Because forecasts of wage changes are not available before the last few years, the wage change is included as a regressor in column 3 without subtracting off expectations.[16] Announced changes in the hourly wage of production and nonsupervisory workers have a statistically significant effect in these models. A 6 cent increase in the hourly wage is associated with about a 3 basis point increase in the 30 year bond rate. Notice, however, that when the employment surprise is excluded from the model (see column 7), wages changes have an insignificant and small effect.

Revisions to the two previous months' employment numbers, which are released along with the unemployment rate and the latest employment data, also have a statistically insignificant effect. Nonetheless, the magnitude of the effects of the revisions are sizable, about half the size of the effect of the latest month's data. As mentioned earlier, the second and third revisions to the employment changes appear to follow a random walk, so the revisions could be viewed as largely unanticipated.

In column (6) the MMS employment forecast is included in the model along with the forecast error. Surprisingly, the forecast has a statistically significant effect. This result is also found in the more parsimonious model in column (8). When the forecast error is omitted from the model, however, the forecast is insignificant. Thus, the positive correlation between the

---

[16]Experimentation with modelling wage changes as an autoregressive process, and using residuals from this process as an explanatory variable in Table 5, yielded similar results.

forecast error and the forecast drives this result.[17] Because the forecast error is not known in advance of the employment release, this finding does not conflict with the efficient markets hypothesis.

To look over a longer time period, Table 6 uses the constructed forecast in place of the MMS median forecast. The models were estimated with data from 1979 to 1996. The forecast error based on the constructed forecast has a smaller impact on the bond yield than does the MMS forecast error, but is nonetheless statistically significant.[18] An unexpected increase in employment of 200,000 jobs is associated with a 4 basis point increase in these models. The interaction between the constructed forecast error and the sample size again yields an insignificant effect. If the model in column 4 is estimated with a sample limited to the post-1984 period, the coefficient on the forecast error is insignificantly different from that estimated with the pre-1985 sample, though larger.

A curious result emerges when the constructed forecast is included as a regressor in columns 6 and 8. Here, the forecast has a statistically significant, positive impact on the bond yield. Recall that the MMS forecast had a negative and significant effect in this specification. The reason for the opposite signs on the two forecasts is that the constructed forecast is negatively correlated with its forecast error, while the MMS forecast is positively correlated with

---

[17]To see this, write $\hat{y} = bf$ and $\hat{y} = c(e_1 - f) + df$, where constants have been suppressed for simplicity. By the omitted-variable-bias formula, $b = d + c\pi$, where π is the coefficient from an auxiliary regression of $(e_1 - f)$ on f. If $b = 0$, then $d = -c\pi$. In this case, d is negative because c and π are positive.

[18]Unlike in Table 5, the deviation between the announced unemployment rate and the MMS unemployment rate forecast is statistically significant in these models. This results in part because the constructed employment forecast is a weaker predictor than the MMS forecast, and in part because the unexpected unemployment rate has a smaller effect in the post-1984 period.

Table 6: OLS Estimates of Effect of Deviations from Constructed Forecasts on the Long Bond Yield, February 1979-April 1996

Dependent Variable: Change in 30 Year T-Bond Yield (x 100)

Variable	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
e_1-f_c	.020 (.004)	.025 (.021)	.018 (.004)	.018 (.004)	.017 (.022)	.019 (.004)	....	.022 (.004)	....
$(e_1-f_c) \times \dfrac{1}{1+\psi/n}$	....	-.010 (.041)	....	....	-.002 (.043)	....	....	....	....
Revision of last month's employment change	....	....	....	-.010 (.010)	-.010 (.010)	-.010 (.010)	-.010 (.011)	....	....
Revision of 2 months ago employment change	....	....	....	.014 (.017)	.014 (.017)	.014 (.017)	.013 (.018)	....	....
Δ hourly wage	....	....	36.676 (25.582)	38.936 (25.732)	39.240 (26.679)	38.079 (25.505)	38.063 (26.774)	....	....
Deviation of unemployment rate change from MMS forecast	....	....	-9.896 (3.909)	-9.825 (3.930)	-9.819 (3.944)	-8.053 (3.990)	-13.169 (4.016)	....	....
Median employment forecast	....	....	....	....	....	.012 (.006)	.007 (.006)	.015 (.006)	.011 (.006)
P-value for F-test of year, year squared, Wednesday and Thursday dummies	.245	.247	.180	.117	.119	.325	.120	.506	.303
P-value for F-test of e_1-f_c and $(e_1-f_c) \times 1/(1+\psi/n)$	....	.000	....	....	.000	....	....	....	....
P-value for F-test of the two employment revisions	....	....	....	.406	.408	.400	.456	....	....
R^2	.137	.137	.172	.181	.181	.199	.113	.166	.049

Notes: e_1 is the first employment change; f_c is the constructed employment forecast; n is the sample size of the BLS employment survey; and ψ is an estimate of σ^2/v^2 (see text for details). Standard errors are in parentheses. Sample size is 202 in columns 1-2 and 8-9, 201 in column 3, and 200 in columns 4-7.

its forecast error. Notice that another surprising result is that the constructed forecast borders on statistical significance when the forecast error is excluded from the equation (see column 9).

One possible explanation for the frequent statistical significance of the forecast is that the standard errors of the estimates are understated, perhaps because homoskedasticity is assumed in calculating the OLS standard errors. To explore this possibility, White (1980) standard errors were calculated for all the models in Tables 5 and 6. The heteroskedasticity-consistent standard errors for the forecast and forecast error tended to be slightly _smaller_ than the OLS standard errors, however, so the more conservative OLS standard errors are reported in the tables.

IV. Short-term Treasury Bills

Although the focus of this paper is on responses to announced employment news in the long-term bond market, I have also examined responses in the market for short-term Treasury bills. Table 7 presents results using as the dependent variable the yield change for 3 month or 1 year Treasury bills. The sample covers 1985-1996, and the estimated models are comparable to those in Table 5 (i.e., they use the MMS employment change forecast).

The interest rate changes associated with unanticipated employment news is greater for the short term maturities than for the 30 year bond. An unexpected increase of 200,000 jobs is associated with a 9.4 basis point increase in the interest rate for 3 month bills, and a 14 basis point increase for 1 year bills. There is no evidence that the market reaction to unanticipated employment changes has increased as the BLS sample size increased; indeed, for the 3 month Treasury bill the interaction term between the (nonlinear function of the) sample size and the employment forecast error is negative and statistically significant. If the constructed forecast is

Table 7: Effect of Deviations from Forecasts on the Short-term Bond Market, 1985-96
Dependent Variable: Close-to-Close Change T-Bill Yield (x 100)

Variable	3 Month T-Bill Yield			1 Year T-Bill Year		
	(1)	(2)	(3)	(4)	(5)	(6)
$e_1 - f$	.047 (.006)	.219 (.055)	.047 (.006)	.070 (.008)	.191 (.078)	.071 (.008)
$(e_1 - f) \dfrac{1}{1 + \psi/n}$	....	-.294 (.094)	....	....	-.207 (.132)	....
Revision of last month's employment change	.012 (.011)	.012 (.011)	.011 (.012)	.018 (.015)	.018 (.015)	.021 (.016)
Revision of 2 months ago employment change	.041 (.022)	.034 (.022)	.041 (.022)	.024 (.030)	.019 (.030)	.022 (.030)
Δ hourly wage	72.809 (23.226)	76.830 (22.470)	72.787 (23.319)	89.028 (31.772)	91.858 (31.637)	89.212 (31.751)
Deviation of unemployment rate change from MMS forecast	-11.724 (3.959)	-13.172 (3.852)	-11.684 (3.982)	-13.002 (5.416)	-14.022 (5.424)	-13.335 (5.421)
Median employment forecast	....	....	.001 (.007)	....	....	-.011 (.010)
P-value for F-test of year, year squared, Wednesday and Thursday dummies	.725	.823	.757	.917	.897	.864
P-value for F-test of e_1-f and (e_1-f) x n	....	.000	....	....	.000	....
P-value for F-test of the two employment revisions	.157	.226	.159	.452	.483	.384
R^2	.492	.530	.492	.489	.499	.493

Notes: e_1 is the first employment change; f is the MMS median employment forecast; n is the sample size of the BLS employment survey; and ψ is an estimate of σ^2/ν^2 (see text for details). Standard errors are in parentheses. Sample size is 132.

used to look at a longer time period, this interaction term is negative but insignificant. Thus, like long-term interest rates, short-term rates do not exhibit a more intense reaction to employment news after the news became more precise.

Three results differ somewhat from the long-term market. First, the MMS forecast has a statistically insignificant and small effect on the short term interest rate. Cook and Korn (1991) find a similar result for short-term rates. Second, changes in the hourly wage have a larger effect on short-term rates: a 6 cent increase in the hourly wage is associated with a 4 basis point increase in the interest rate on the 3 month bills and a 5 basis point increase in the interest rate on 1 year bills. Third, as others have found, unexpected changes in the unemployment rate have a statistically significant effect on short-term rates: a 0.5 percentage point increase in unemployment is associated with roughly a 6 basis point increase in short-term rates.

V. Extensions

A number of extensions of the basic results were explored. First, to test whether bond markets react to the employment news with a lag, the effect of the employment announcement on the change in the next trading day's yield was examined. Specifically, I estimated the models in Tables 4-6 using as the dependent variable the change in the bond yield for the next trading day. These results provided no indication that the employment news had a lagged effect on long-term bond yields: neither the employment numbers nor their deviation from forecasts had a statistically significant impact on bond yields in these models. Curiously, however, the change in the hourly wage had a statistically significant, positive effect in these models.

Second, it is often alleged that markets over react to the employment data on the day the

data are released, and then correct for this over reaction on the following trading day. Although the results described in the previous paragraph are inconsistent with this view, to further test this hypothesis I correlated the residuals from the models in Tables 4-6 with the next day's yield changes. If larger than predicted movements in yields on announcement days are corrected the next day, this correlation would be negative. The results of this exercise yielded numerically small, statistically insignificant and typically positive correlations, however.

Third, if some market participants have advance knowledge of the BLS data, the news may affect the markets prior to the announcement date. To test for this possibility, I regressed the change in the yield <u>on the day before</u> the employment announcement on the forecast errors.[19] When the models in Table 5 were estimated using the post-1984 sample, the unexpected soon-to-be-announced changes in employment and unemployment had statistically insignificant and small effects. Surprisingly, however, when the models in Tables 4 and 6 were estimated for the 1979-96 sample using the previous day's yield change as the dependent variable, the changes in the unemployment rate, employment and hourly wage often were statistically significant. These findings do not necessarily imply that the BLS data were leaked to some investors prior to their announcement date, however. It is possible, for example, that the Fed -- which is notified of the employment data prior to their release -- acts on this information, and that the markets respond to the Fed's actions without direct knowledge of the BLS employment data.

Lastly, it is worth comparing these results to a recent study by Jones, Lamont, and

---

[19]The employment data typically are known to the BLS only one or two days before the announcement date.

Lumsdaine (1996), which examines the effect of CES announcement dates on volatility in the 5, 10 and 30 year bond markets. Jones, Lamont, and Lumsdaine find that the variance of the daily excess return is about 50 percent greater in all three markets on announcement days than on non-announcement days. My results indicate that the R-square of the employment surprises, revisions to past months data, wage change and unemployment rate, is about 30 percent (using the MMS forecast data) for long term bonds, and about 50 percent for short-term bills, suggesting that the news conveyed by the employment report can account for a good deal of the excess volatility on announcement dates.

VI. Summary and Conclusion

The three main findings of this paper are: (1) the precision of the BLS's first estimate of nonfarm employment growth improved as the sample size increased; (2) announcements of unanticipated employment changes strongly affected daily interest rate movements from 1979 to 1996; (3) the effect of unanticipated employment changes on interest rates does not appear to have increased over time, as the BLS employment data became more reliable and as the Federal Reserve Board shifted its focus to employment indicators. Although the latter finding may seem surprising, there is additional evidence suggesting that financial markets paid at least some attention to the BLS employment announcements before the mid 1980s. Table 8 reports the percent of times an article in the New York Times cited the BLS employment release data the day after the announcement date in a story regarding the bond or stock market. After 1983, the employment data were noted as influencing the financial markets virtually every day after they were released. In the 1979-83 period, the employment news was cited as influencing the

Table 8

Percent of times the New York Times cited the
BLS employment release in connection to the
bond or stock market on the day after the release

Period	Bond Mkt.	Stock Mkt.	Either Bond or Stock Market
1979-83	27%	33%	45%
1984-88	97	78	98
1989-93	100	92	100
1994-96*	97	97	100

*1994-96 data are through June of 1996.

financial markets less frequently, but was still mentioned on 45 percent of post-announcement days.

One possible explanation for the apparent constant market reaction to more precise news is that the markets were not aware of the increased precision of the BLS employment survey. This is certainly possible, but there were news articles written about the improvements in the BLS data in the 1980s. Moreover, the BLS reports the sample size and sampling variability of the employment survey in every issue of Employment and Earnings, so it is not difficult to learn that the series improved. Market forecasters have a tremendous amount of information regarding the construction of the BLS data; for example, many are knowledgeable of the magnitude of the quarterly bias adjustment. It would be somewhat surprising if they were not aware of the expansion of the survey.

Another possibility is that the amount of additional information available to market participants increased during the same time period the precision of the employment survey increased, so the employment survey provided less new information than otherwise would be the case.[20] There is little evidence that the quantity or quality of other labor market data improved over this period, however. For example, the BLS's closely-watched survey of manufacturing turnover was eliminated in the early 1980s, and the sample size of the CPS was cut from 65,500 to 53,600 households in the 1980s. More generally, Abraham (1996) reports that the BLS's budget was constant in real terms since 1978, so it is unlikely that there was a major increase in other labor market data in this period. Finally, the finding that MMS employment forecasts have not improved relative to the benchmark data suggests that private forecasters were unable

---

[20]In terms of equation (4), this hypothesis is equivalent to a decrease in v^2.

to make more precise forecasts over time.

At this stage, it seems anomalous that the bond market did not respond more to more reliable employment data. Kahneman and Tversky (1982) document that, compared to optimal Bayesian updating, it is common for individuals to place too much weight on recent information and too little weight on their prior data. The findings of this study may reflect this broader phenomenon. From the present results, however, it is not possible to determine if the market over or under reacts to employment news. It is possible that the market over reacted to the news initially, and that the current reaction is efficient. Alternatively, it is possible that the market over reacted initially, and continues to over react today. But in either scenario, the results suggest that investors did not rationally respond to changes in the precision of relevant information.

Before the conclusion that financial markets do not rationally respond to more accurate information is generally accepted, it would be useful to see this hypothesis tested further by examining the effects of changes in other data series. For example, there have been (or are planned) changes in the quality of the Employment Cost Index, Consumer Price Index, Produce Price Index, and National Income and Product Accounts. Other countries also have changed the quality of their economic statistics over time. The present type of analysis can be performed using the experiments provided by changes in these other data series.

Schwert, G. William. 1981. "The Adjustment of Stock Prices to Information About Inflation." The Journal of Finance 36(1): 15-29.

Shiller, Robert J. 1979. "The Volatility of Long-Term Interest Rates and Expectations Models of the Term Structure." Journal of Political Economy 87 (6), pp. 1190-1219.

Stevenson, Richard W. 1996. "In Eye of the Storm, the Fed Governors Stay Calm." The New York Times. (June 8): 31.

Urich, Thomas and Paul Wachtel. 1985. "The Effects of Inflation and Money Supply Announcements on Interest Rates." The Journal of Finance 39 (4), September, pp. 1177-88.

White, Halbert. 1980. "A Heteroskedasticity-Consistent Covariance Matrix Estimator and a Direct Test for Heteroskedasticity." Econometrica 48, 817-38.