How the FOMC Weighs Its Dual Mandate

Hyuk An*

September 2026

Abstract

We estimate the weight on unemployment relative to inflation in the Federal Reserve's quadratic loss function over 2012–2020. Large language models count the negative and positive words in each member's remarks in the FOMC transcripts. Taking net negativity as a proxy proportional to the loss, we regress it on the Tealbook forecasts. The results show that the FOMC places greater weight on the inflation gap than on the unemployment gap. A weight recovered from the observed federal funds rate target, using the FOMC's projections in the Summary of Economic Projections together with the Board's own models, EDO and LINVER, is consistent with the transcript estimate.

Keywords: FOMC, loss function, monetary policy objectives, text analysis, large language models.
JEL: E31, E52, E58, C55.

*Korea Development Institute. hyukan@kdi.re.kr

1 Introduction

Price stability and maximum employment constitute the dual mandate of the Federal Reserve (Fed). Because the two trade off in the short run, the Federal Open Market Committee (FOMC) must choose the relative weight it places on each objective. However, the FOMC's commitment to a balanced approach puts no number on it. Earlier studies have estimated the weight indirectly, through a macroeconomic model in which the policy rate minimizes a loss function. However, the weight they estimate is on output relative to inflation, not on the employment side of the mandate. The models are written in the output gap rather than the unemployment gap (Favero and Rovelli, 2003; Ozlale, 2003; Salemi, 2006; Dennis, 2006; Ilbas, 2012; Givens, 2012; Debortoli and Lakdawala, 2016; Lakdawala, 2016).

A growing body of research on central banking uses non-traditional data, but none of it estimates that weight. Chappell et al. (2005) recover member-level reaction functions from transcripts and voting records. Pierdzioch et al. (2016) recover an asymmetric forecast loss from the projections members submit. Tsang and Yang (2025) score transcripts for the disagreement that formal votes conceal. Shapiro and Wilson (2022) estimate the inflation objective through the loss function, assuming that loss is proportional to the net negativity of the FOMC transcripts. They regress net negativity, counted with the Loughran and McDonald (2011) dictionary, on the macroeconomic state. Haavio et al. (2026) extend the approach to the ECB's press conferences and find an asymmetric loss function.

The FOMC transcripts through 2020 were released in January 2026, by which time large language models (LLMs) could classify each speaker's sentiment close to human accuracy (Hansen and Kazinnik, 2024). The 2012–2020 period matters for the inflation gap, because the target was explicitly 2% throughout it. The FOMC announced the target in January 2012 (FOMC, 2012) and loosened the framework in August 2020 to seek inflation that averages 2% over time (FOMC, 2020), known as flexible average inflation targeting (FAIT).

This paper estimates a quadratic loss function from the 2012–2020 transcripts, using LLMs as text classifiers. Two models from different providers, OpenAI's gpt-5.6-luna and Google's gemini-3.6-flash, each score the 1,147 speaker-by-meeting cells ten times. The mean over the ten runs of their net negativity is the proxy for the loss. We regress that proxy on the forecasts of inflation and the unemployment gap in Tealbook A (the Greenbook until June 2010), the staff forecast prepared before each meeting. We find that net negativity responds more to the inflation gap than to the unemployment gap, which puts the estimated weight well below one. We then recover the weight from the observed federal funds rate target through EDO and LINVER, two macroeconomic models the Board of Governors uses in its own policy work. Both models give a weight consistent with the transcript estimate.

This paper makes three contributions. First, to the work that recovers the Fed's objective (Favero and Rovelli, 2003; Ozlale, 2003; Salemi, 2006; Dennis, 2006; Ilbas, 2012; Givens, 2012; Debortoli and Lakdawala, 2016; Lakdawala, 2016; Shapiro and Wilson, 2022), we add evidence on the relative weight from the 2012–2020 transcripts. Since it comes from the record rather than from a model, it offers a reference point for models that would otherwise assume the weight.

Second, work on non-traditional data such as text has mostly turned it into a particular index. That literature has measured how hawkish or dovish FOMC statements read (Lucca and Trebbi, 2009), decomposed statements and transcripts with topic models (Hansen and McMahon, 2016; Hansen et al., 2018), predicted the rate decision from the staff's own documents and read the residual as a shock (Aruoba and Drechsel, 2024), detected emotion in the chair's voice (Gorodnichenko et al., 2023), and

classified central bank text with language models (Hansen and Kazinnik, 2024; Dunn et al., 2024; Gambacorta et al., 2026). We carry that index further and use it to estimate a parameter of macroeconomic models.

Third, a measure built from text is usually checked for robustness within the assumptions it is built on. For example, Shapiro and Wilson (2022) score with the Loughran and McDonald (2011) dictionary and Hansen and Kazinnik (2024) with an LLM, and both check their measures against a human reading of a sample. However, a human audit checks the score, not what it proxies for. Our estimate depends on the assumption that net negativity is proportional to the loss. We therefore estimate the weight again from a structural model, outside that assumption. The two estimates depend on different assumptions and give complementary evidence.

The rest of this paper is organized as follows. Section 2 sets out the loss function and the estimating equation. Section 3 describes how net negativity and the two gaps are measured. Section 4 estimates the weight from the transcripts. Section 5 recovers the weight from the funds-rate target. Section 6 concludes.

2 The loss function and the estimating equation

We use the standard quadratic loss function of the monetary policy literature (Svensson and Tetlow, 2005; Svensson, 2011; Shapiro and Wilson, 2022),

$$L_t \;=\; (\pi_t - \pi^*)^2 \;+\; \varphi\,(u_t - u_t^*)^2, \tag{1}$$

where π_t is inflation, π^* the inflation target, u_t the unemployment rate, and u_t^* the natural rate of unemployment. $\pi_t - \pi^*$ is the inflation gap and $u_t - u_t^*$ the unemployment gap. The FOMC held the target at 2% over the whole of the sample, so we set $\pi^* = 2$ (FOMC, 2012). The parameter φ captures the relative weight placed on a squared unit of the unemployment gap against a squared unit of the inflation gap.

The loss L_t in equation (1) is not observed. We assume that the net negativity N_{it} of member i at meeting t is proportional to the loss (Shapiro and Wilson, 2022),

$$N_{it} \;=\; f_i \;+\; \delta\, L_t \;+\; \varepsilon_{it}, \tag{2}$$

where f_i is a speaker fixed effect to control for individual differences in speaking style, and δ is the unknown constant of proportionality. Substituting equation (1) into equation (2) gives the equation we estimate,

$$N_{it} \;=\; f_i \;+\; c_1(\pi_t - 2)^2 \;+\; c_2(u_t - u_t^*)^2 \;+\; \gamma_1(\Delta i_t)^2 \;+\; \gamma_2 P(i_t)^2 \;+\; \varepsilon_{it}, \qquad c_1 = \delta, \quad c_2 = \delta\varphi, \tag{3}$$

so that δ multiplies both coefficients and cancels in their ratio. The specification also includes the two instrument terms of the objective in the Board's FRB/US optimal-control program, with the program's functional forms and constants.[1] Δi_t penalizes moving the rate from one meeting to the next, and $P(i) = 3\big(\underline{i} - i\big)\mathbf{1}\{i < \underline{i}\} + 0.10\,e^{10(i-\underline{i})}\mathbf{1}\{i < \underline{i}\} + 0.10\,e^{-20(i-\underline{i})}\mathbf{1}\{i \geq \underline{i}\}$ penalizes a rate close to or below its lower bound, with i_t the midpoint of the target range set at meeting t and $\underline{i}$ the bound of 12.5 basis points.

[1] The program is `ocpolicy.prg` in the EViews package of FRB/US.

3 Measuring net negativity and the two gaps

3.1 Transcripts and net negativity

We parse the FOMC transcripts into individual remarks and aggregate the speech of the Chair, governors, and Reserve Bank presidents into speaker-by-meeting cells, excluding speakers who appear at fewer than eight meetings. The estimation sample covers the 69 meetings with a Tealbook forecast from January 2012 to July 2020 and holds 1,147 cells from 31 speakers.

Each cell is scored by gpt-5.6-luna from OpenAI (luna) and by gemini-3.6-flash from Google (gemini), each at its provider's middle reasoning setting. The two models are developed independently by different providers, and scoring with both checks that the result is not specific to one LLM. The prompt asks the model to return the words conveying the speaker's own positive and negative reaction to economic conditions, each with the number of times it occurs.[2] The LLMs occasionally list words absent from the text or overreport occurrences, so every reported word is credited only up to its verified occurrences. This step trims 1.5% and 0.3% of the reported negative counts for the two models, and 0.8% and 0.1% of the positive ones.

Unlike a dictionary such as Loughran and McDonald (2011), an LLM reads a word in the context it stands in. For example, at the January 2012 meeting Charles Plosser, then president of the Federal Reserve Bank of Philadelphia, expected "the unemployment rate to continue to decline another half a percentage point this coming year." Both 'unemployment' and 'decline' are negative entries in the Loughran and McDonald (2011) dictionary, so the dictionary counts two negative signals in a forecast of falling unemployment. Reading the sentence, neither LLM returns either word as negative for that cell.

Net negativity, the proxy for the loss in equation (2), is defined at the cell level as

$$N_{it} \;=\; \frac{1}{10}\sum_{r=1}^{10}\frac{\mathrm{Neg}_{itr}-\mathrm{Pos}_{itr}}{\mathrm{Total}_{it}}\times 100, \tag{4}$$

where Neg_{itr} and Pos_{itr} are the verified counts of negative and positive words returned on run r and Total_{it} is the number of words in the cell. Net negativity is therefore the number of negative words minus the number of positive words per hundred words spoken. Since LLMs are probabilistic, each model scores every cell ten times and equation (4) averages the ten runs to smooth the variation across them. Table 1 summarizes the net negativity the two models return. Mean net negativity is close to zero in both models. Its standard deviation is 0.38 for luna and 0.16 for gemini, since the two models draw the negativity threshold differently. However, the two series move together, correlating at 0.67 across cells and at 0.81 across meeting averages.

3.2 Tealbook forecasts

The inflation and unemployment forecasts come from Tealbook A, which the staff prepares in the days preceding each meeting. Because they are settled before the meeting opens, they are predetermined with respect to the remarks in the transcript. Forecasts over longer horizons reflect the funds-rate path the staff assumes (Romer and Romer, 2004), so we average both forecasts over the current and next

[2] Figure A.1 shows the prompt with the request parameters. It adapts the instruction Shapiro and Wilson (2022) gave the research assistants who scored their human audit, which asks for the speaker's own sentiment and not for factual statements about the economy.

Table 1. LLM Net-Negativity Measurement, by Model

	luna	gemini
Negative words	0.66 (0.33)	0.27 (0.15)
Positive words	0.66 (0.26)	0.29 (0.17)
Net negativity N	−0.00 (0.38)	−0.02 (0.16)

Notes: Shares of cell words in percent, mean with standard deviation in parentheses, over the 1,147 cells of the estimation sample. Means over the ten runs.

quarters (Shapiro and Wilson, 2022).[3] What members say can then be read as a response to the state of the economy.

The inflation forecast π_t is the Tealbook projection of core PCE inflation in annualized percentage terms. We use core rather than headline inflation, since in the short term headline inflation moves with transitory changes in food and energy prices. The unemployment gap forecast $u_t - u_t^*$ is the Tealbook unemployment rate forecast minus the staff's real-time estimate of the natural rate (NAIRU).

Figure 1 plots net negativity against the left axis and the two squared gaps against the right. From 2012Q1 to 2013Q4 the squared unemployment gap averages 4.17 against 0.29 for the squared inflation gap, and net negativity averages 0.09 for luna and 0.03 for gemini. The squared unemployment gap falls to 0.83 over 2014 and net negativity falls with it, to -0.08 and -0.05. From 2015Q1 to 2016Q4 the squared unemployment gap is below 0.07 at every meeting, and yet net negativity rises again, to 0.02 and -0.03. The gap that remains in those quarters is the inflation gap, whose square averages 0.38. From 2018Q1 to 2019Q2 the squared inflation gap averages 0.03 while the squared unemployment gap is between 0.49 and 1.00, and net negativity is at its lowest, -0.09 and -0.08. When the pandemic hits in 2020, both squared gaps and net negativity rise. Read on its own, the figure suggests that net negativity responds more to the inflation gap than to the unemployment gap.

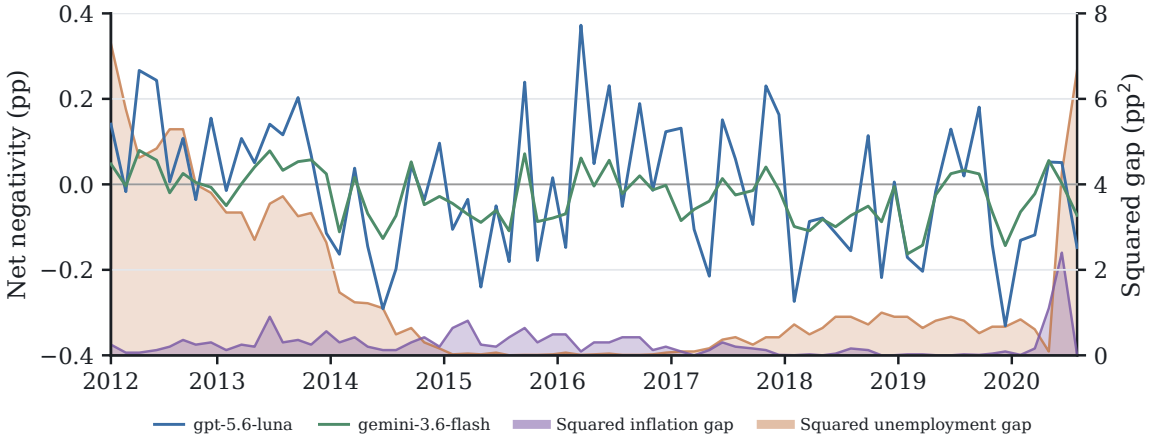


Figure 1. Net Negativity and the Two Squared Gaps

Notes: Net negativity on the left axis and the two squared gaps on the right, across the 69 FOMC meetings from January 2012 to July 2020. The gaps are Tealbook forecasts averaged over the quarter of the meeting and the next.

[3] Table A.1 reports columns (1)–(2) of Table 2 with both forecasts averaged through two, three and four quarters ahead. Across these horizons the weight lies between 0.04 and 0.23, consistently below one.

4 The weight estimated from the transcripts

Table 2 reports the estimates, computed by a Huber M-estimator (Huber, 1964), which limits the influence of outliers,[4] with standard errors two-way clustered by meeting and speaker (Cameron et al., 2011). The table holds three specifications, each estimated for both models, luna in the odd columns and gemini in the even. Columns (1)–(2) estimate equation (3). Columns (3)–(4) drop the instrument terms of equation (3). Columns (5)–(6) add two financial controls, the S&P 500 return over the 90 days ending on the meeting day and the VIX averaged over the three trading days before it, since the stock market is discussed at these meetings (Cieslak and Vissing-Jorgensen, 2021).

Both gaps enter positively in every column. A one-unit increase in the squared inflation gap raises net negativity by between 0.060 and 0.187 for luna and between 0.033 and 0.067 for gemini, and a one-unit increase in the squared unemployment gap by between 0.012 and 0.037 and between 0.008 and 0.014. However, the gap coefficients do not estimate how far the loss itself moves, since the constant of proportionality is unknown.

Our interest is the relative weight φ, the ratio of the two coefficients, in which that constant cancels. A null of $\varphi = 0$ would test whether the unemployment gap enters the loss at all. We test $\varphi = 1$ instead, equal weight on the two mandates and the weight set in the Board's optimal-control program and in the optimal-control calculation of Yellen (2012).

Every column puts the estimate between 0.195 and 0.232 and rejects $\varphi = 1$ at the 1% level. The relative weight φ is 0.222 and 0.232 in columns (1)–(2).[5] Without the two instrument terms φ is 0.205 for luna and unchanged at 0.232 for gemini. The financial controls raise both gap coefficients, while φ remains 0.195 and 0.210.

Table 2. Loss-Function Estimates

	(1) luna	(2) gemini	(3) luna	(4) gemini	(5) luna	(6) gemini
Relative weight						
φ	0.222***	0.232***	0.205***	0.232***	0.195***	0.210***
	(0.210)	(0.201)	(0.289)	(0.215)	(0.078)	(0.084)
Coefficients						
$(\pi_t - 2)^2$	0.087	0.036	0.060	0.033	0.187***	0.067***
	(0.053)	(0.025)	(0.043)	(0.023)	(0.054)	(0.021)
$(u_t - u_t^*)^2$	0.019	0.008*	0.012	0.008*	0.037**	0.014***
	(0.014)	(0.005)	(0.013)	(0.005)	(0.014)	(0.005)
S&P 500 (3m)					-0.886***	-0.300***
					(0.319)	(0.114)
VIX					-0.010	-0.003
					(0.006)	(0.002)
Instrument terms	Yes	Yes	No	No	Yes	Yes
Speaker FE	Yes	Yes	Yes	Yes	Yes	Yes
Observations	1147	1147	1147	1147	1147	1147
R^2	0.154	0.108	0.150	0.108	0.165	0.116

Notes: Huber M-estimator, two-way clustered standard errors in parentheses. The standard error under φ is the delta-method error on the same clustered covariance, and the stars under φ test $\varphi = 1$. */**/***: 10/5/1%.

To check whether the estimate is robust, we vary three features of the specification in columns (1)–(2) of Table 2, and Table 3 reports the result. Columns (1)–(2) add a dummy for the 2020 meetings.

[4] Table A.2 reports least squares estimates of the same specifications, where the weight lies between 0.23 and 0.31.

[5] Table A.3 estimates the two gaps on each run separately, where φ ranges from 0.136 to 0.335 for luna and from 0.133 to 0.394 for gemini.

The unemployment gap returns to its 2012 level within a single quarter in 2020, the jump at the right edge of Figure 1, so the last few meetings could account for the estimate. The dummy puts the relative weight at 0.188 and 0.213. Columns (3)–(4) replace both squared terms with absolute gaps,

$$N_{it} \;=\; f_i \;+\; c_1\,|\pi_t - 2| \;+\; c_2\,|u_t - u_t^*| \;+\; \varepsilon_{it}, \tag{5}$$

which puts it at 0.296 and 0.398. The two specifications are different loss functions, so the levels are not comparable, but the inflation gap still carries the larger weight. Column (5) scores the same cells with the Loughran and McDonald (2011) dictionary in place of the two LLMs. It puts the ratio at 0.180, though it is estimated less precisely, the standard error on c_1 running 0.100 against 0.047 for luna in column (1).

Table 3. Alternative Specifications

	COVID dummy		Absolute gaps		Dictionary
	(1) luna	(2) gemini	(3) luna	(4) gemini	(5) LM
Relative weight					
φ	0.188***	0.213***	0.296***	0.398***	0.180***
	(0.119)	(0.142)	(0.206)	(0.212)	(0.220)
Coefficients					
$(\pi_t - 2)^2$	0.133***	0.044*	0.174**	0.066**	0.155
	(0.047)	(0.024)	(0.073)	(0.031)	(0.100)
$(u_t - u_t^*)^2$	0.025*	0.009**	0.052	0.026**	0.028
	(0.014)	(0.005)	(0.035)	(0.012)	(0.024)
COVID dummy	-0.156**	-0.030			
	(0.068)	(0.031)			
Instrument terms	Yes	Yes	Yes	Yes	Yes
Financial controls	No	No	No	No	No
Speaker FE	Yes	Yes	Yes	Yes	Yes
Observations	1147	1147	1147	1147	1147
R^2	0.160	0.109	0.157	0.112	0.172

Notes: Huber M-estimator, two-way clustered standard errors in parentheses. Standard errors and stars as in Table 2. */**/***: 10/5/1%.

5 The weight recovered from the funds rate

5.1 The prescribed rate at each weight

Earlier studies obtain the weight from a model the researcher specifies, so the weight they report is conditional on that model. Their identification also depends on the slope of the Phillips curve, which is hard to identify over 2012–2020. We estimate the Phillips curve in the unemployment gap, motivated by the New Keynesian model with unemployment of Blanchard and Galí (2010),

$$\pi_t - 2 \;=\; \rho\,(\pi_{t-1} - 2) \;+\; \kappa\,(u_{t-1} - u_{t-1}^*) \;+\; e_t, \tag{6}$$

On the Tealbook forecasts of Section 3.2, the slope is $\kappa = 0.009$ with a standard error of 0.038. We therefore compute the loss from the projections of the FOMC and the impulse responses of the Board's own models. Barnichon and Mesters (2023, 2026) show that these two objects are sufficient statistics for evaluating a policy.

Suppose the FOMC sets the rate that minimizes its loss function, in which the weight on the unem-

ployment gap is unknown. Let i_t range over the candidate rates the FOMC could have set in quarter t, let $L_t(i_t;\varphi)$ be the loss under i_t and φ, and let $i_t^*(\varphi)$ be the prescribed rate under φ,

$$i_t^*(\varphi) \;=\; \arg\min_{i_t}\; L_t(i_t;\varphi), \qquad t = 1,\dots,T. \tag{7}$$

Let $\{i_t^{\text{obs}}\}_{t=1}^{T}$ be the observed funds-rate target, the midpoint of the target range, over the $T = 35$ rounds of the Summary of Economic Projections (SEP) from 2012Q1 to 2020Q3. We derive the prescribed rate at each weight and measure how far its path lies from the observed funds-rate target,

$$\widehat{\varphi} \;=\; \arg\min_{\varphi}\; D(\varphi), \qquad D(\varphi) \;=\; \sqrt{\frac{1}{T}\sum_{t=1}^{T}\left(i_t^*(\varphi) - i_t^{\text{obs}}\right)^2} \quad \text{or} \quad \frac{1}{T}\sum_{t=1}^{T}\left|i_t^*(\varphi) - i_t^{\text{obs}}\right|. \tag{8}$$

The weight whose prescribed path comes closest to the observed funds-rate target, $\widehat{\varphi}$, serves as a cross-check on the transcript estimate.

Section 4 measures the loss meeting by meeting, because a member's remarks respond to the economic conditions prevailing at the meeting. In this section the weight is recovered from the rate, whose effect on inflation and unemployment is spread over subsequent quarters, so the loss L_t in equation (7) is summed over the H quarters that follow,

$$L_t(i_t;\varphi) \;=\; \sum_{h=0}^{H-1}\beta^h\Big[\left(\pi_{t+h|t}(i_t) - 2\right)^2 \;+\; \varphi\left(u_{t+h|t}(i_t) - u_t^*\right)^2 \;+\; \nu\left(\Delta i_{t+h|t}(i_t)\right)^2 \;+\; \omega\, P\left(i_{t+h|t}(i_t)\right)^2\Big], \tag{9}$$

at a candidate rate i_t, where $\pi_{t+h|t}(i_t)$, $u_{t+h|t}(i_t)$ and $i_{t+h|t}(i_t)$ are the paths of inflation, unemployment and the funds rate expected at t under that rate. The four terms are those of equation (3), with $\nu = 1$, $\omega = 10$ and $\underline{i} = 12.5$ basis points, and each is discounted at $\beta = 0.99$ over $H = 20$ quarters.[6] The loss has no closed-form minimizer, and we minimize it over the grid of candidate rates

$$i_t \;\in\; \left\{\, i_t^{\text{obs}} + 0.01k \;:\; k = -600,\dots,600 \,\right\}. \tag{10}$$

We write each gap in equation (9) as the sum of two terms, what the FOMC projected and what a different rate would have done.[7] Each gap and the rate path move as

$$\begin{aligned}
\pi_{t+h|t}(i_t) - 2 \;&=\; \left(\pi_{t+h|t} - 2\right) \;+\; \left(i_t - i_t^{\text{obs}}\right) a_h^{\pi},\\
u_{t+h|t}(i_t) - u_t^* \;&=\; \left(u_{t+h|t} - u_t^*\right) \;+\; \left(i_t - i_t^{\text{obs}}\right) a_h^{u},\\
i_{t+h|t}(i_t) \;&=\; i_{t+h|t} \;+\; \left(i_t - i_t^{\text{obs}}\right) a_h^{r}.
\end{aligned} \tag{11}$$

In each SEP round, participants project headline PCE inflation, the unemployment rate and the funds rate, the midpoint of the target range at the end of each year, for the current year and the following two years. In the September and December rounds, they also project a third subsequent year. They also provide longer-run projections for all three variables. We take the median across participants. In equation (11) the first term of each gap is the projection minus its longer-run value, and the first

[6] Table A.4 reports the distances at $H = 10$ and 40 and at $\beta = 0.975$ and 0.995. The best fits stay between 0.12 and 0.36.

[7] Barnichon and Mesters (2023) set $\varphi = 1$ in their application and read the deviation as the distance of the observed rate from the optimum.

term of the rate path is the projected rate itself. Figure 2 plots the medians. The longer-run inflation projection is 2.0 in every round and the longer-run unemployment projection falls from 5.50 in January 2012 to 4.05 in September 2020. We interpolate linearly (Barnichon and Mesters, 2023) from the latest observation through each annual projection in the fourth quarter of its year to the longer-run value five years ahead.

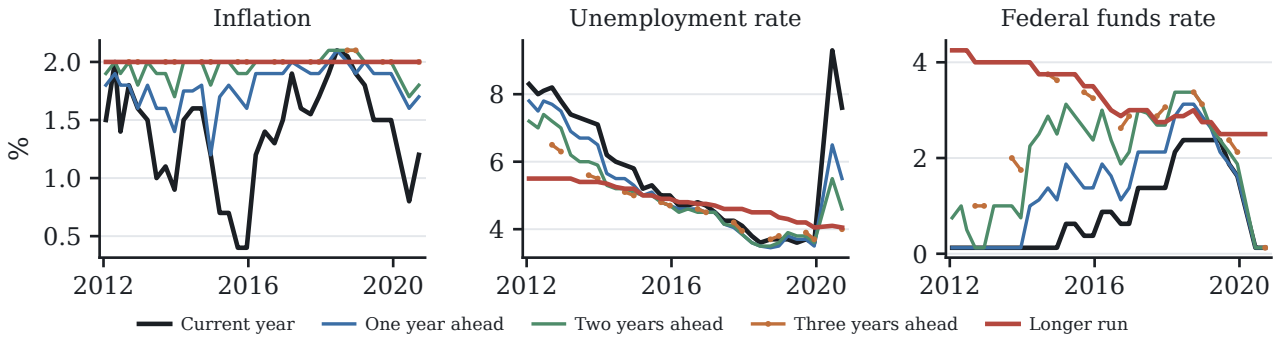


Figure 2. The Summary of Economic Projections, 2012Q1–2020Q3

Notes: The third year ahead appears only in the September and December rounds.

The second term, the impulse response, measures the additional gap that a candidate rate would have produced, relative to the observed funds-rate target i_t^{obs}. The responses of inflation, the unemployment rate and the funds rate to a one percentage point increase in the funds rate are denoted a_h^{π}, a_h^{u} and a_h^{r}. The responses come from two models the Board publishes and uses in its own policy work, the Estimated Dynamic Optimization (EDO) model (Chung et al., 2010) and LINVER, a linear version of FRB/US.[8] The shock to the models is the difference $i_t - i_t^{\text{obs}}$. Both models are linearized, so that difference moves each path by the response multiplied by it. Figure 3 plots the responses. In both models the increase raises the unemployment rate and lowers inflation at first, and the effects fade as the funds rate returns to its path. The unemployment response peaks at 0.10 percentage points after three quarters in LINVER and at 0.22 after one quarter in EDO. Half of the peak has faded by the twelfth quarter in LINVER and by the sixth in EDO. The inflation response is smaller, reaching −0.01 percentage points after seven quarters in LINVER and −0.02 after three quarters in EDO. One set of responses is used for every round, so the transmission of policy is taken as stable over the sample.

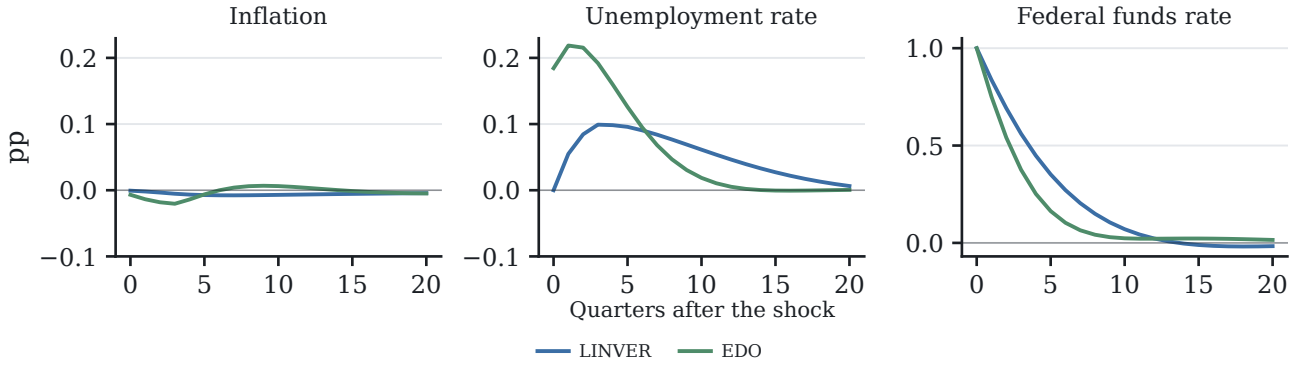


Figure 3. Response to a One Percentage Point Increase in the Federal Funds Rate

[8]LINVER is taken from the Board's July 2024 package, whose model files are those of the June 2022 release. EDO is taken from the Board's October 2016 package.

5.2 The weight closest to the rate decisions

Figure 4 plots the distance in both models. The distance is close to a quadratic function of the weight. The best fit in LINVER is 0.20 by root mean square and 0.15 by mean absolute distance, and in EDO it is 0.16 by both. The transcript estimate, roughly 0.2 across Table 2, lies close to each minimum.

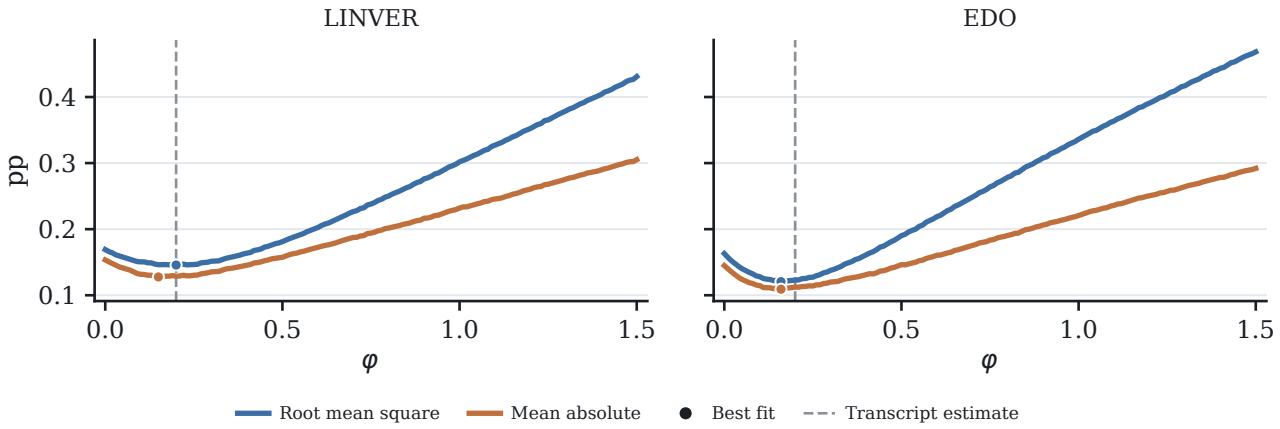


Figure 4. Distance to the Observed Funds-Rate Target

Notes: The dashed line marks the transcript estimate of 0.2.

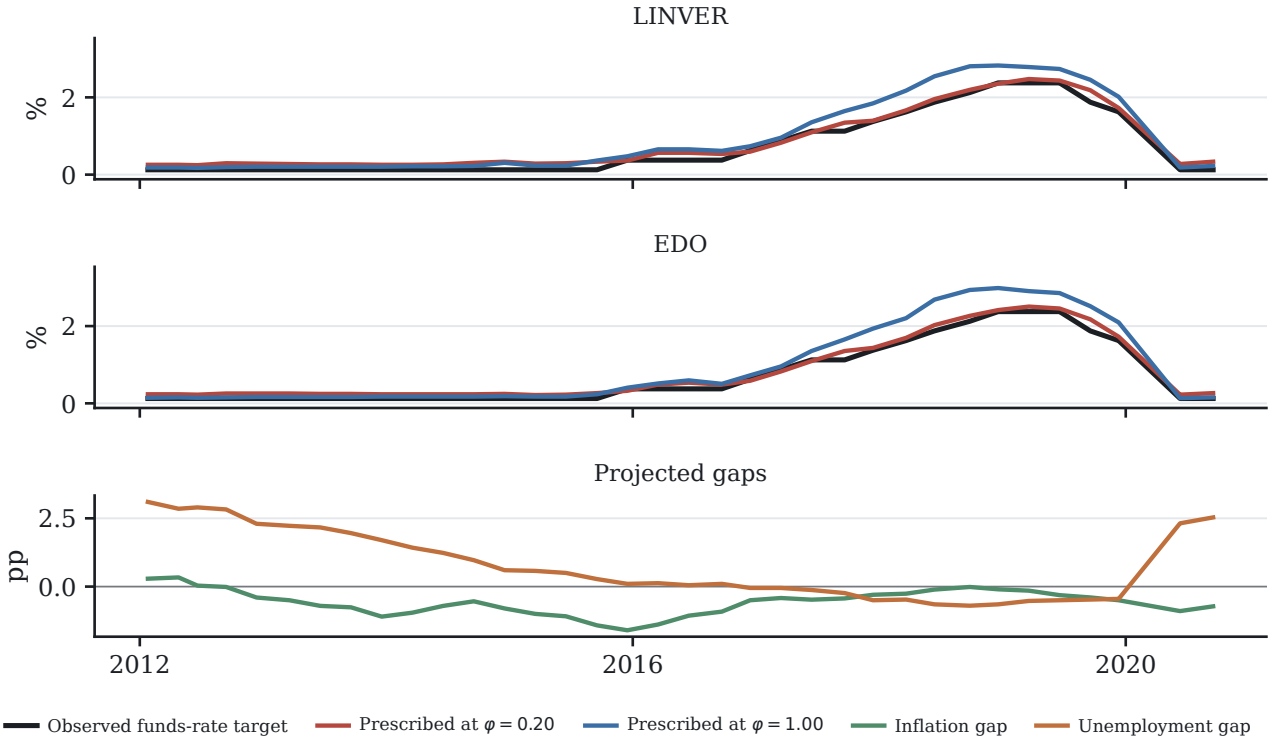


Figure 5. Prescribed Rate and Observed Funds-Rate Target

Notes: The gaps in the lowest panel are interpolated SEP projections for the quarter of each round.

Figure 5 plots the observed funds-rate target with the rates prescribed at $\varphi = 0.2$, the transcript estimate rounded from Table 2, and at equal weights, $\varphi = 1$. Its lowest panel plots the two gaps projected for the quarter of each SEP round. Through 2014 the projected unemployment gap is large and positive, and from the September 2012 round the projected inflation gap is negative. Both gaps then imply a lower rate, but the prescribed rate is held near its lower bound at both weights. From the December

2016 round until the pandemic, the projected unemployment gap is negative and implies a higher rate, while the projected inflation gap, also negative, implies a lower one. At $\varphi = 1$ the unemployment gap carries more weight, so the prescribed rate runs above the observed funds-rate target. At $\varphi = 0.2$ the prescribed rate stays close to that target.

6 Conclusion

We estimate the weight on the unemployment gap relative to the inflation gap from the 2012–2020 transcripts, under the assumption that net negativity is proportional to the loss. The relative weight is well below one. Across the specifications it lies between 0.18 and 0.23, and the weight recovered from the observed funds-rate target in the Board's own models lies between 0.15 and 0.20. The estimate holds when the two instrument terms are dropped, when the financial controls enter, and when a dummy for the pandemic meetings enters.

The weight helps people anticipate how the FOMC weighs its two mandates. Since it comes from the transcripts rather than from a structural model, it offers a reference point for models that would otherwise assume it, in line with a broader empirical approach to constructing macroeconomic models (Coibion and Gorodnichenko, 2015; Nakamura and Steinsson, 2018).

Two extensions remain. The decade studied had inflation below target and a funds rate at its lower bound for four years, so the weight may differ in periods such as the late 1970s or 2022. Word counts are not the only possible proxy for the loss, and other proxies can be compared with this estimate.

References

Aruoba, S. B. and T. Drechsel (2024). Identifying monetary policy shocks: A natural language approach. NBER Working Paper 32417. Forthcoming, *American Economic Journal: Macroeconomics*, doi:10.1257/mac.20240184.

Barnichon, R. and G. Mesters (2023). A sufficient statistics approach for macro policy. *American Economic Review 113*(11), 2809–2845.

Barnichon, R. and G. Mesters (2026). Policy evaluation with sufficient macro statistics: A primer. *The Econometrics Journal 29*(1), 27–68.

Blanchard, O. and J. Galí (2010). Labor markets and monetary policy: A New Keynesian model with unemployment. *American Economic Journal: Macroeconomics 2*(2), 1–30.

Cameron, A. C., J. B. Gelbach, and D. L. Miller (2011). Robust inference with multiway clustering. *Journal of Business & Economic Statistics 29*(2), 238–249.

Chappell, Jr., H. W., R. R. McGregor, and T. A. Vermilyea (2005). *Committee Decisions on Monetary Policy: Evidence from Historical Records of the Federal Open Market Committee*. MIT Press.

Chung, H., M. T. Kiley, and J.-P. Laforte (2010). Documentation of the estimated, dynamic, optimization-based (EDO) model of the U.S. economy: 2010 version. Finance and Economics Discussion Series 2010-29, Board of Governors of the Federal Reserve System.

Cieslak, A. and A. Vissing-Jorgensen (2021). The economics of the Fed put. *Review of Financial Studies 34*(9), 4045–4089.

Coibion, O. and Y. Gorodnichenko (2015). Information rigidity and the expectations formation process: A simple framework and new facts. *American Economic Review 105*(8), 2644–2678.

Debortoli, D. and A. Lakdawala (2016). How credible is the Federal Reserve? a structural estimation of policy re-optimizations. *American Economic Journal: Macroeconomics 8*(3), 42–76.

Dennis, R. (2006). The policy preferences of the US Federal Reserve. *Journal of Applied Econometrics 21*(1), 55–77.

Dunn, W., E. E. Meade, N. R. Sinha, and R. Kabir (2024). Using generative AI models to understand FOMC monetary policy discussions. FEDS Notes, Board of Governors of the Federal Reserve System, December 6.

Favero, C. A. and R. Rovelli (2003). Macroeconomic stability and the preferences of the Fed: A formal analysis, 1961–98. *Journal of Money, Credit and Banking 35*(4), 545–556.

FOMC (2012). Statement on longer-run goals and monetary policy strategy. Adopted effective January 24; issued January 25, 2012. Board of Governors of the Federal Reserve System.

FOMC (2020). Statement on longer-run goals and monetary policy strategy. Adopted effective January 24, 2012; as amended effective August 27, 2020. Board of Governors of the Federal Reserve System.

Gambacorta, L., B. Kwon, T. Park, P. Patelli, and X. S. Zhu (2026). CB-LMs: Language models for central banking. *Journal of Financial Stability 87*, 101585.

Givens, G. E. (2012). Estimating central bank preferences under commitment and discretion. *Journal of Money, Credit and Banking 44*(6), 1033–1061.

Gorodnichenko, Y., T. Pham, and O. Talavera (2023). The voice of monetary policy. *American Economic Review 113*(2), 548–584.

Haavio, M., J. Heikkinen, P. Jalasjoki, J. Kilponen, M. Paloviita, and I. Vänni (2026). Reading between the lines—uncovering asymmetry in the central bank loss function. *Journal of Money, Credit and Banking*. forthcoming, doi:10.1111/jmcb.70053.

Hansen, A. L. and S. Kazinnik (2024). Can ChatGPT decipher Fedspeak? SSRN Working Paper 4399406.

Hansen, S. and M. McMahon (2016). Shocking language: Understanding the macroeconomic effects of central bank communication. *Journal of International Economics 99*, S114–S133.

Hansen, S., M. McMahon, and A. Prat (2018). Transparency and deliberation within the FOMC: A computational linguistics approach. *Quarterly Journal of Economics 133*(2), 801–870.

Huber, P. J. (1964). Robust estimation of a location parameter. *Annals of Mathematical Statistics 35*(1), 73–101.

Ilbas, P. (2012). Revealing the preferences of the US Federal Reserve. *Journal of Applied Econometrics 27*(3), 440–473.

Lakdawala, A. (2016). Changes in Federal Reserve preferences. *Journal of Economic Dynamics and Control 70*,

124–143.

Loughran, T. and B. McDonald (2011). When is a liability not a liability? textual analysis, dictionaries, and 10-Ks. *Journal of Finance 66*(1), 35–65.

Lucca, D. O. and F. Trebbi (2009). Measuring central bank communication: An automated approach with application to FOMC statements. NBER Working Paper 15367.

Nakamura, E. and J. Steinsson (2018). Identification in macroeconomics. *Journal of Economic Perspectives 32*(3), 59–86.

Ozlale, U. (2003). Price stability vs. output stability: Tales of Federal Reserve administrations. *Journal of Economic Dynamics and Control 27*(9), 1595–1610.

Pierdzioch, C., J.-C. Rülke, and P. Tillmann (2016). Using forecasts to uncover the loss function of Federal Open Market Committee members. *Macroeconomic Dynamics 20*(3), 791–818.

Romer, C. D. and D. H. Romer (2004). A new measure of monetary shocks: Derivation and implications. *American Economic Review 94*(4), 1055–1084.

Salemi, M. K. (2006). Econometric policy evaluation and inverse control. *Journal of Money, Credit and Banking 38*(7), 1737–1764.

Shapiro, A. H. and D. J. Wilson (2022). Taking the Fed at its word: A new approach to estimating central bank objectives using text analysis. *Review of Economic Studies 89*(5), 2768–2805.

Svensson, L. E. O. (2011). Practical monetary policy: Examples from Sweden and the United States. *Brookings Papers on Economic Activity 42*(2), 289–332.

Svensson, L. E. O. and R. J. Tetlow (2005). Optimal policy projections. *International Journal of Central Banking 1*(3), 177–207.

Tsang, K. P. and Z. Yang (2025). Agree to disagree: Measuring hidden dissent in FOMC meetings. *Journal of Economic Dynamics and Control 180*, 105197.

Yellen, J. L. (2012). Perspectives on monetary policy. Speech at the Boston Economic Club, June 6. Board of Governors of the Federal Reserve System.

Appendix

Table A.1. Forecasts Averaged over Longer Horizons

	luna			gemini		
Forecast horizon (Q)	c_1	c_2	φ	c_1	c_2	φ
F0–F1	0.087	0.019	0.222*** (0.210)	0.036	0.008*	0.232*** (0.201)
F0–F2	0.241***	0.032**	0.131*** (0.074)	0.088***	0.014***	0.164*** (0.071)
F0–F3	0.306**	0.022*	0.073*** (0.049)	0.091**	0.011***	0.119*** (0.063)
F0–F4	0.330**	0.013	0.040*** (0.035)	0.089*	0.007**	0.078*** (0.052)

Notes: Columns (1)–(2) of Table 2 with both forecasts averaged from the quarter of the meeting. Estimator, standard errors and stars as in Table 2. */**/***: 10/5/1%.

Table A.2. Least-Squares Estimates

	(1) luna	(2) gemini	(3) luna	(4) gemini	(5) luna	(6) gemini
Relative weight						
φ	0.285***	0.274***	0.307*	0.279***	0.230***	0.233***
	(0.254)	(0.206)	(0.413)	(0.235)	(0.105)	(0.087)
Coefficients						
$(\pi_t - 2)^2$	0.084	0.041	0.051	0.035	0.166***	0.075***
	(0.053)	(0.026)	(0.040)	(0.022)	(0.060)	(0.023)
$(u_t - u_t^*)^2$	0.024*	0.011**	0.016	0.010**	0.038**	0.018***
	(0.014)	(0.005)	(0.013)	(0.005)	(0.015)	(0.006)
Instrument terms	Yes	Yes	No	No	Yes	Yes
Financial controls	No	No	No	No	Yes	Yes
Observations	1147	1147	1147	1147	1147	1147
R^2	0.180	0.128	0.176	0.127	0.189	0.138

Notes: Ordinary least squares on the specifications of Table 2. Standard errors and stars as in Table 2. */**/***: 10/5/1%.

Table A.3. Run-by-Run Estimates

	luna			gemini		
Run	c_1	c_2	φ	c_1	c_2	φ
1	0.059	0.020	0.335 (0.420)	0.044*	0.008*	0.177*** (0.118)
2	0.090	0.024*	0.268*** (0.218)	0.041	0.005	0.133*** (0.144)
3	0.085*	0.019	0.228*** (0.233)	0.037	0.009*	0.228*** (0.173)
4	0.062	0.011	0.173*** (0.300)	0.054**	0.010	0.183*** (0.139)
5	0.084*	0.023	0.271*** (0.253)	0.030	0.008	0.247*** (0.247)
6	0.070	0.017	0.246*** (0.288)	0.036	0.008	0.231*** (0.264)
7	0.113*	0.022	0.192*** (0.175)	0.024	0.009*	0.394 (0.440)
8	0.118**	0.025	0.215*** (0.178)	0.049**	0.009*	0.178*** (0.120)
9	0.084	0.020*	0.242*** (0.213)	0.027	0.010**	0.351 (0.435)
10	0.110**	0.015	0.136*** (0.135)	0.026	0.007	0.278*** (0.272)

Notes: Columns (1)–(2) of Table 2 estimated on one scoring run per row. Estimator, standard errors and stars as in Table 2. */**/***: 10/5/1%.

Table A.4. Prescribed Rate versus Observed Funds-Rate Target by Horizon and Discount Factor

		φ								Best fit
		0.00	0.10	0.20	0.25	0.50	0.75	1.00	1.50	$\widehat{\varphi}$
$H = 10, \beta = 0.99$										
LINVER	RMSE	0.188	0.167	0.155	0.151	0.153	0.176	0.209	0.288	0.36
	MAE	0.171	0.150	0.135	0.131	0.136	0.154	0.174	0.219	0.31
EDO	RMSE	0.165	0.130	0.123	0.127	0.177	0.243	0.310	0.431	0.16
	MAE	0.148	0.116	0.111	0.115	0.139	0.172	0.207	0.271	0.16
$H = 40, \beta = 0.99$										
LINVER	RMSE	0.165	0.149	0.146	0.147	0.187	0.247	0.314	0.448	0.17
	MAE	0.150	0.130	0.129	0.131	0.163	0.200	0.239	0.316	0.17
EDO	RMSE	0.163	0.128	0.123	0.128	0.190	0.265	0.338	0.471	0.16
	MAE	0.144	0.115	0.112	0.115	0.146	0.184	0.222	0.293	0.16
$H = 20, \beta = 0.975$										
LINVER	RMSE	0.163	0.145	0.139	0.139	0.166	0.217	0.273	0.391	0.20
	MAE	0.149	0.130	0.123	0.124	0.145	0.177	0.211	0.279	0.20
EDO	RMSE	0.159	0.124	0.120	0.123	0.179	0.250	0.320	0.448	0.17
	MAE	0.143	0.112	0.109	0.111	0.139	0.174	0.210	0.279	0.17
$H = 20, \beta = 0.995$										
LINVER	RMSE	0.171	0.153	0.148	0.151	0.187	0.246	0.311	0.443	0.21
	MAE	0.155	0.134	0.130	0.134	0.163	0.200	0.238	0.313	0.12
EDO	RMSE	0.165	0.129	0.124	0.130	0.193	0.268	0.341	0.475	0.16
	MAE	0.146	0.115	0.113	0.117	0.148	0.186	0.224	0.297	0.16

Notes: The distance D of equation (8), in percentage points. The weight grid runs from 0.00 to 1.50 in steps of 0.01. The columns shown are a selection.

Figure A.1. Scoring Prompt

```
Below are all remarks made by one speaker at one FOMC meeting, in the
order they were made. Remarks are separated by blank lines. Other
participants spoke in between.

List the words that convey negative sentiment, with how many times each
occurs. Do the same for words that convey positive sentiment. Select
words that carry the speaker's own reaction, not words that merely
describe economic data.

The lists should reflect your best guess of how the speaker would answer
if asked, "how negative are you today?" and "how positive are you
today?". Use only the remarks below.

Respond only in this JSON structure:
{"negative": [{"word": "...", "count": 0}],
 "positive": [{"word": "...", "count": 0}]}

REMARKS:
{<full text of the cell>}
```

Notes: Both models are called through their providers' batch endpoints, gpt-5.6-luna through `batches` (chat) at reasoning effort medium with a strict `json_schema`, gemini-3.6-flash through `batchGenerateContent` at `thinkingLevel` medium with a `responseSchema`.