Review of Economic Studies (2022) **0**, 1–38 https://doi.org/10.1093/restud/rdab094
© The Author(s) 2021. Published by Oxford University Press on behalf of The Review of Economic Studies Limited.
Advance access publication 27 December 2021

Downloaded from https://academic.oup.com/restud/advance-article/doi/10.1093/restud/rdab094/6484651 by Federal Reserve Bank of San Francisco Research Library user on 25 Ja

Taking the Fed at its Word: A New Approach to Estimating Central Bank Objectives using Text Analysis

ADAM HALE SHAPIRO and DANIEL J. WILSON
Federal Reserve Bank of San Francisco

First version received November 2019*; Editorial decision July* 2021*; Accepted December* 2021 (*Eds.*)

We propose a new approach to estimating central bank preferences, including the implicit inflation target, that requires no priors on the underlying macroeconomic structure nor observation of monetary policy actions. Our approach entails directly estimating the central bank's objective function from the sentiment expressed by policymakers in their internal meetings. We apply the approach to the objective function of the U.S. Federal Open Market Committee (FOMC). The results challenge two key aspects of conventional wisdom regarding FOMC preferences. First, the FOMC had an implicit inflation target of approximately $1\frac{1}{2}\%$ on average over our baseline 2000–11 sample period, significantly below the commonly assumed value of 2. Second, the FOMC's loss depends strongly on output growth and stock market performance and less so on their perception of current economic slack.

Key words: Central bank preferences, Sentiment analysis, Inflation target

JEL Codes: E52, E58, C80

I'm bemused by the de facto inflation targeters that we have become here [laughter] with the 1.5% goal.

– Ben Bernanke, Chairman of the Federal Reserve
September 2006 FOMC Meeting

1. INTRODUCTION

What is the central bank's objective function? This is a core question of monetary macroeconomics. The canonical model assumes that the central bank has quadratic preferences over inflation, relative to a known target, and economic slack (Walsh, 2017). Although there is broad consensus on what the central bank objective function *should* look like based on the large literature on optimal monetary policy, there has been very little study of what the central bank objective function actually *is* in practice. The dearth of positive analysis of the central bank objective function is surprising considering it is implicitly the foundation underlying monetary policy choices. Moreover, the sparsity is certainly not due to a belief that the objective function is

The editor in charge of this paper was Dirk Krueger.

well understood. For example, even the functional form, much less the parameters, is not widely agreed upon. As Blinder (1997) argued, "academic macroeconomists tend to use quadratic loss functions for reasons of mathematical convenience, without thinking much about their substantive implications. The assumption is not innocuous...practical central bankers and academics would benefit from more serious thinking about the functional form of the loss function."

In this article, we propose a new approach to estimating central bank objectives. Our approach relies on the assumption that the central bank's preferences regarding desired levels of inflation and economic activity are embedded in the words spoken by central bank policymakers at internal meetings. Specifically, we demonstrate that the sentiment expressed by policymakers can be used to directly estimate the central bank's loss function, including the implicit inflation target. This approach could be used internally and externally to study the preferences of any central bank that has transcripts or detailed summaries of their policymaking deliberations. We apply the approach here to estimating the U.S. Federal Open Market Committee's (FOMC) objective function. We construct a proxy for loss using a measure of the negative sentiment expressed in the FOMC's internal discussions. We then estimate how the FOMC's negativity is affected by macroeconomic and financial variables potentially entering the FOMC's loss function, using real-time data on the Federal Reserve staff (Greenbook) forecasts of core PCE inflation and real economic variables along with contemporaneous stock market variables.

The results from this exercise challenge two key aspects of the conventional wisdom on FOMC preferences, the first relating to the inflation target and the second relating to preferences regarding economic slack. Beginning with the first aspect, the analysis indicates that the FOMC had an implicit inflation target of approximately $1\frac{1}{2}\%$ on average over the 2000–11 sample period. The 2000–11 period is when the FOMC both had an agreed-upon preferred inflation measure (core PCE inflation) and had not yet announced an official inflation target.[1] This finding is robust along several dimensions: using alternative measures of negativity, filtering out FOMC discussion unrelated to inflation, conditioning or not on macro-economic forecasts and/or financial variables, and allowing for asymmetric preferences.

Our $1\frac{1}{2}\%$ estimated target over 2000–11 stands in contrast to the commonly held view, both at the time and currently, that the FOMC's inflation target was 2% over this period. Svensson (2015), for example, points to the "widely held perception of an unofficial Federal Reserve inflation target of about 2%" over the 2000–11 period of his study. In addition, many macroeconomic models assume the Federal Reserve had a 2% target over this period.[2] Since both average realized inflation and survey measures of longer-run inflation expectations hovered close to 2% over that period, this perception would imply that the Federal Reserve was keeping inflation at desired levels. By contrast, our results suggest that not only did the FOMC miss hitting its inflation target, but inflation expectations became anchored at a level well above what the FOMC intended.[3] In fact, we see in the internal meeting transcripts that this discrepancy between the FOMC's desired

1. The Greenbook began forecasting PCE inflation in January 2000. The FOMC officially announced an inflation target of 2.0% based on PCE inflation in January 2012. As discussed below, we also use our sentiment approach to separately estimate the inflation target from 2012 to 2019 based on data from FOMC minutes.

2. The Federal Reserve Board's FRB/US model and the New York Fed's DSGE model (Del Negro, Eusepi, Giannoni, Sbordone, Tambalotti, Cocci, Hasegawa and Linder, 2013) set a prior on the inflation target centred at 2% over this time period. The perceived inflation target rate (PTR) variable used in the Federal Reserve Board's FRB/US model and academic papers (e.g. Fuhrer, Olivei and Tootell, 2012; Bauer and Rudebusch, 2017) is based on long-run inflation expectations from the Survey of Professional Forecasters, which was slightly above 2% over our sample period. The 2% target over this period is also assumed in a number of academic papers (e.g. Bianchi, Melosi and Rottner, 2019).

3. An official target was not publicly announced by the FOMC until January 2012.

Downloaded from https://academic.oup.com/restud/advance-article/doi/10.1093/restud/rdab094/6484651 by Federal Reserve Bank of San Francisco Research Library user on 25 Ja

target and the public's perception of the target drew the concern of several committee members, including by St. Louis Federal Reserve President William Poole in the September 2006 meeting:

> I had several conversations at Jackson Hole with Wall Street economists and journalists, and they said, quite frankly, that they really do not believe that our effective inflation target is 1–2%. They believe we have morphed into $1\frac{1}{2}$ to $2\frac{1}{2}$%, and no one thought that we were really going to do anything over time to bring it down to 1 to 2.

This concern of appearing incapable of keeping inflation at its desired levels—or in other words, not seeming credible—was echoed by a multitude of FOMC participants.[4] In fact, Chairman Ben Bernanke expressed concern at the March 2006 meeting that even a 20 or 30 basis point miss of the inflation target could be potentially damaging to the Fed's credibility:

> Like most of you, I am not at all alarmist about inflation. I think the worst that is likely to happen would be 20 or 30 basis points over the next year. But even that amount is a little disconcerting for me. I think it is very important for us to maintain our credibility on inflation and it would be somewhat expensive to bring that additional inflation back down.

The FOMC's strong emphasis on credibility suggests that the central bank would have behaved quite differently had the consensus inflation target been 2% instead of $1\frac{1}{2}$%. The FOMC raised the federal funds rate 17 times in the mid-2000s, in part, from concerns about inflation moving out of the $1\frac{1}{2}$% comfort zone. For instance, at the January 2006 meeting, when the Greenbook Core PCE inflation nowcast sat at 2.0%, Chicago Federal Reserve President Michael Moskow stated:

> With inflation remaining at such rates, we could begin to lose credibility if markets mistakenly inferred that our comfort zone had drifted higher. When we stop raising rates, we ought to be reasonably confident that policy is restrictive enough to bring inflation back toward the centre of our comfort zone, which I believe is $1\frac{1}{2}$%...So for today, we should move forward with an increase of 25 basis points...

Similar comments in favour of raising rates because inflation was moving away from this comfort zone were made by other FOMC participants.[5] In fact, the data corroborate this interest-rate setting behaviour. As we show in Section 5.3, the estimates from a simple Taylor rule with an inflation target of $1\frac{1}{2}$% resemble the "balanced-approach" rule of Taylor (1999). By contrast, a target of 2% produces a worse fit and the wrong sign on the inflation-gap term. Thus, a reaction function where the FOMC tended to raise (lower) interest rates when inflation was above (below) $1\frac{1}{2}$% fits the data better than a 2% target.[6]

4. In Supplementary Appendix D.2, we list similar quotes regarding missing the inflation target and the bank's credibility concerns. Survey evidence shows that this high regard for credibility is ubiquitous across almost all central banks (Blinder, 2000). Numerous papers in the literature have shown that not maintaining central bank credibility can have important welfare implications (e.g. Bomfim and Rudebusch, 2000; Faust and Svensson, 2001; Erceg and Levin, 2003).

5. Additional statements are provided in Supplementary Appendix D.4.

6. In addition to impacting central bank behaviour, an inflation target of $1\frac{1}{2}$ may have inherently different welfare consequences than the commonly held view of 2%. Indeed, some FOMC participants stated they were in favour of an inflation target of $1\frac{1}{2}$% as opposed to something higher, such as 2%, precisely because the lower rate is assumed to substantially improve welfare (see Supplementary Appendix D.3 for participants' statements). While there is disagreement

Downloaded from https://academic.oup.com/restud/advance-article/doi/10.1093/restud/rdab094/6484651 by Federal Reserve Bank of San Francisco Research Library user on 25 Ja

Given the discrepancy between our estimate of an implicit inflation target of $1\frac{1}{2}$ and the conventional prior of 2, we complement our regression analysis with a narrative analysis that identifies and tabulates instances in which FOMC participants explicitly stated their inflation target preferences. Though stated target preferences are conceptually distinct from the implicit target consistent with the overall tone of the committee's discussion, we do find that the consensus target preference was $1\frac{1}{2}\%$ for most of our baseline sample period. However, the narrative analysis also reveals that some participants shifted their preference up to 2% during the Great Recession. As a validation exercise using our sentiment analysis approach, we split our speaker-meeting observations into three sub-samples: those observations where the narrative evidence suggests the speaker had a preference for $1\frac{1}{2}\%$, those with a preference for 2%, and those with no explicit evidence of an inflation target preference. The loss function estimated using the $1\frac{1}{2}\%$ narrative-preference subsample yields an implied inflation target close to $1\frac{1}{2}\%$, while the loss function estimated using the 2% narrative-preference sub-sample yields an implied inflation target close to 2% (and the other group is estimated to have a target in between), thus validating the performance of the sentiment analysis approach.

These results suggest that the majority of participants preferred an inflation target of $1\frac{1}{2}\%$ over the 2000–11 period, despite some FOMC participants shifting to a preference for 2% towards the end of the sample period. It is plausible that these individuals represented the vanguard of a broader shift in inflation preferences among FOMC participants, leading to the announcement of an explicit 2% inflation target in 2012. Indeed, we find corroborating evidence for this broader shift in preferences. Specifically, in an extension in which we apply the sentiment analysis approach using data from FOMC minutes separately for the 2000–11 period and post-2011 (2012–19), we find evidence that the inflation target did indeed shift up to 2.0% after 2011.

These changes in stated preferences highlight an important issue of interpretation. Do changes in the inflation target reflect changing preferences about the desired long-term level of inflation or changing circumstances? In the conventional framework of optimal monetary policy, the loss function and its parameters reflect preferences by definition. If circumstances change that cause policymakers to change their views on the desired long-term level of inflation, for example due to increased salience of the zero lower bound, this would be viewed as a change in their inflation target preferences. We interpret the inflation target estimated in this paper as the target embedded in the central bank loss function, which can be different in different periods of time due to varying circumstances.[7]

The second key aspect of our findings from the sentiment analysis approach is that, in contrast to typical formulations of the central bank loss function in the optimal monetary policy literature, the FOMC's loss was little affected by contemporaneous economic slack. Instead, we find that loss was strongly affected by output growth and financial variables, most prominently recent stock market returns. There are two potential interpretations for this result. The most direct interpretation is that output growth and financial conditions are part of the central bank's objective

in the economics literature regarding the precise welfare costs of inflation, a number of influential papers have shown that relatively small changes in steady-state inflation can have substantial welfare implications (e.g. Dotsey and Ireland, 1996; Di Tella, MacCulloch and Oswald, 2001; Lucas, 2001; Erosa and Ventura, 2002; Coibion, Gorodnichenko and Wieland, 2012).

7. In August 2020, after the sample period studied in this article, the Federal Reserve announced a change in its monetary policy framework, introducing for the first time an explicit distinction between short-term and long-term inflation targets. Specifically, they adopted Average Inflation Targeting, which calls for the Fed to aim to achieve 2% inflation on average over several years. Thus, 2% can be seen as a long-term target whereas the short-term target is the short-run level of inflation needed to achieve that long-run average.

Downloaded from https://academic.oup.com/restud/advance-article/doi/10.1093/restud/rdab094/6484651 by Federal Reserve Bank of San Francisco Research Library user on 25 Ja

function.[8] There is, in fact, considerable external support for this interpretation. Thornton (2011) documents that from 1991 until 2009 the FOMC's policy directive, announced to the public after each FOMC meeting, stated "The Federal Open Market Committee seeks monetary and financial conditions that will foster price stability and promote sustainable *growth in output*" [italics added]. Thornton further notes that neither "maximum sustainable employment nor the unemployment rate" is mentioned in these directives. The finding that stock market performance affects the FOMC's loss is consistent with work by Peek, Rosengren and Tootell (2015) and Wischnewsky, Jansen and Neuenkirch (2019) who argue that the FOMC responds to financial variables.[9]

An alternative, and not mutually exclusive, interpretation is that current financial indicators and near-term forecasts of output growth are signals of future (longer-run) levels of slack. That is, perhaps policymakers do not care about output growth and financial markets *per se*, but do care about future slack, which may be predicted by these variables. We do not take a strong stand on which of these two interpretations is more likely, though we do find some empirical evidence supporting direct inclusion. Specifically, growth has a stronger effect than slack on current loss even when these two variables are measured by longer-horizon Greenbook forecasts.

Direct estimation using text analysis offers a new approach to studying central bank preferences, complementing prior indirect approaches. Prior analyses have relied on indirect inference, deriving central bankers' preferences from either observed interest rate votes[10] or statements about desired interest rates viewed through the lens of an estimated interest-rate rule.[11]

Inferring the objective function from estimated interest-rate rules has two main drawbacks. First, inferring the parameters of the objective function, or even what variables are in it, requires precise knowledge of the structure of the macro-economic model implicitly underlying the central bank's actions.[12] For instance, if the central bank believes the economy behaves according to the New Keynesian framework, backing out the structural parameters of the objective function requires knowledge of the slopes and persistence parameters of the Phillips and IS curves—parameters for which there is very little consensus in the literature. More broadly, the actual structure of the economy is undoubtedly more complex than any available theoretical model (Blanchard, 2018). Second, identification of the objective function in these analyses must be limited to periods in which there are changes to the federal funds rate target. There are plausibly many instances when the FOMC is in some way constrained in its ability to change its policy target to perfectly align with its preferences. The zero lower bound, which was in effect between 2008 and 2015 and from March 2020 to the time of this writing, is the most obvious such constraint. Another is the convention of only moving the fed funds target by quarter-point increments.

In the remainder of the article, we first derive in Section 2 the regression specifications we use to estimate the loss function and discuss identification issues. In Section 3, we describe the textual source data we compiled on FOMC meeting transcripts, minutes, and speeches and how we measured the negativity therein. We also show how the negativity in each of these sources has varied over time. Section 4 presents the main results of the article. Based on the FOMC's

8. Such an objective function has been formulated previously in the literature (e.g. Barro and Gordon, 1983; Walsh, 2017).

9. It is also consistent with Cieslak and Vissing-Jorgensen (2018) who find that negative mentions of the stock market in FOMC minutes and transcripts have high explanatory power for interest rate policy changes.

10. See Favero and Rovelli (2003), Dennis (2006), Surico (2007), Ilbas (2012), and Givens (2012).

11. See Chappell Jr, Havrilesky and McGregor (1997) and Meade and Stasavage (2008).

12. We demonstrate this point with a simple three-equation New Keynesian model in Supplementary Appendix Section A.1. The point is demonstrated in more general models in Rudebusch and Svensson (1999), Dennis (2006), and Surico (2007).

Downloaded from https://academic.oup.com/restud/advance-article/doi/10.1093/restud/rdab094/6484651 by Federal Reserve Bank of San Francisco Research Library user on 25 Ja

private meeting transcripts, we first estimate a standard symmetric loss function for the 2000–11 period and identify the implied inflation target. We then consider asymmetry and robustness to alternative sentiment measures. In Section 5, we provide four types of external validation related to our estimate of the inflation target. Specifically, we provide corroborating evidence from narrative analysis, heterogeneity across different types of FOMC participants, Taylor rule estimation, and estimates based on FOMC minutes. In Section 6, we use our sentiment analysis approach to investigate the inflation target before and after our baseline sample period. We offer concluding remarks in Section 7.

2. DERIVING AN ESTIMABLE LOSS FUNCTION

2.1. *Standard quadratic loss function*

The canonical model of central bank preferences and optimal policy is based on an assumed central bank loss function within a standard New Keynesian (NK) model of the macroeconomy (see e.g. Clarida, Gali and Gertler (1999); Giannoni and Woodford, 2003; Walsh, 2004; Surico, 2007). In Supplementary Appendix A.1, we present a simple version of the model that conveys the role of the loss function in determining interest rate policy and illustrates the difficulty of identifying the structural parameters of the loss function from estimates of the interest rate rule alone. And, of course, more complicated loss functions or economic environments will only introduce more structural parameters into the interest rate rule coefficients.[13] In addition, any instability over time in the structural parameters of the IS or Phillips curves would further complicate efforts to infer the loss function parameters from estimated Taylor rule coefficients. For instance, there has been evidence in recent years that the slope of the Phillips curve has flattened over the past couple of decades (e.g. Leduc and Wilson, 2017).

Therefore, we propose an approach to directly estimate the central bank's loss function. To fix ideas, consider the standard short-run loss function:

$$L_t = \tilde{\pi}_t^2 + \phi \tilde{u}_t^2, \tag{1}$$

where $\tilde{\pi} = \pi - \pi^*$ is the inflation gap, defined as inflation minus the central bank's inflation target, and $\tilde{u} = u - u^*$ represents the unemployment gap and is meant to capture current economic activity or "slack."

Conceptually, there are other variables that could enter the loss function. As discussed in Section 1, previous studies have suggested that monetary policymakers may be concerned with output growth as well as financial conditions. Furthermore, economists have also put forth different functional forms for how economic activity enters the loss function. For instance, Barro and Gordon (1983) and Walsh (2017) formulate a loss function where economic activity is proxied with a linear term in addition to the squared term. For this reason, we consider a more general loss function that nests equation (1):

$$L_t = \tilde{\pi}_t^2 + \mathbf{X}_t \Lambda, \tag{2}$$

where $\mathbf{X}$ is a vector of non-inflation variables potentially entering the loss function such as the unemployment gap, output growth, and financial variables.

13. For instance, a recent paper by Nakamura and Steinsson (2018) argues that Fed interest rate shocks may have an additional channel beyond the conventional IS curve. Interest rate shocks, because they reflect the Fed's economic forecasts, convey information to private agents about future (potential) output growth which can create actual output growth through investment decisions. This additional Fed information channel introduces even more structural parameters into the economic environment, further complicating the ability to infer the Fed's loss function from estimates of their interest rate rule.

Downloaded from https://academic.oup.com/restud/advance-article/doi/10.1093/restud/rdab094/6484651 by Federal Reserve Bank of San Francisco Research Library user on 25 Ja

We proxy for loss using a measure of the negativity, denoted N, expressed by FOMC members in their meetings:

$$N_{it} = \delta L_{it}. \tag{3}$$

Specifically, we assume that the loss of an FOMC member (i), which is unobserved, is proportional to their measured net negativity with a factor of proportionality δ. We measure net negativity for individual members of the FOMC, N_{it}, and hence we allow for speaker-specific fixed effects. These fixed effects allow for heterogeneity across speakers in their general (time-invariant) style/tone of speaking. Allowing for this heterogeneity could be potentially important because it ensures that changes in the composition of the FOMC over time do not lead one to misinterpret changes in negativity over time. Applying these changes to equation (2), and adding an i.i.d. error term to capture measurement error in loss, yields the following regression specification:

$$N_{it} = \alpha_i + \delta \tilde{\pi}_t^2 + \sum_j \omega_j x_{j,t} + \varepsilon_{it}, \tag{4}$$

where α_i is a speaker fixed effect, $x_{j,t}$ is the jth element of the vector $\mathbf{X}_t$, $\omega_j \equiv \delta\lambda_j$, and λ_j is the jth element of the parameter vector Λ.

The implicit target, π^*, is a latent variable that is of wide general interest to the public, economists, and other policymakers. The value of the true π^* guiding monetary policy is unknown, especially prior to the FOMC's formal adoption of a 2.0% inflation target (as measured by PCE price index) in January 2012. Hence, rather than imposing an assumed π^* to measure $\tilde{\pi}_t$, we directly estimate it.

The inflation gap term in equation (4) can be expanded as follows:

$$N_{it} = f_i + \delta \pi_t^2 + \theta \pi_t + \sum_j \omega_j x_{j,t} + \varepsilon_{it}, \tag{5}$$

where $\theta \equiv -2\delta\pi^*$ and $f_i \equiv \alpha_i + \delta\pi^{*2}$. Thus, the implied π^* can be backed out from the estimated coefficients on inflation and inflation squared: $\hat{\pi^*} = -\hat{\theta}/2\hat{\delta}$. The coefficients of equation (5) can be estimated via OLS. We calculate standard errors that are robust to heteroskedasticity and two-way clustering by time (meeting) and speaker. Clustering by time accounts for potential correlation of errors across speakers at a given meeting, while clustering by speaker allows for potential correlation over time within speaker. The standard error for $\hat{\pi^*}$ is calculated using the delta method from the estimated variance-covariance matrix (allowing for heteroskedasticity and two-way clustering).

2.2. *Identification*

Note that the right-hand side of equation (5) should be the FOMC's real-time perceptions of the current and future values of those variables. To measure these real-time perceptions, we use the real-time "Greenbook" forecasts prepared by the Federal Reserve's economic staff in advance of each FOMC meeting. Greenbook forecasts have the appeal that they are the actual numbers discussed by FOMC members at the meetings. Moreover, they incorporate all of the higher-frequency information available between any contemporaneous data releases and the actual meeting. Importantly, because the Greenbook forecasts are prepared prior to when the FOMC meets, they can be considered exogenous with respect to the committee's dialogue.

Following Romer and Romer (2004), Coibion and Gorodnichenko (2011), and others, our baseline regression specifications use *short-run* Greenbook forecasts of these variables and omit

Downloaded from https://academic.oup.com/restud/advance-article/doi/10.1093/restud/rdab094/6484651 by Federal Reserve Bank of San Francisco Research Library user on 25 Ja

longer-run forecasts. Specifically, we use an average of the Greenbook's current-quarter and one-quarter-ahead forecasts. Conceptually, "short-run" forecasts here should be thought of as the time horizon over which the central bank takes current and forecasted inflation and real economic activity as exogenous or outside of their control. As Romer and Romer emphasize, over the longer-run the central bank endogenously determines these variables via interest rate policy in order to minimize its lifetime loss. Hence, longer-run forecasts "are contaminated by assumptions or inside information [on the part of Federal Reserve staff] about the course of monetary policy." Put differently, the FOMC's loss is a function of their perception of economic states over the short-run as well as their perception of states over the longer-run *in the hypothetical scenario where the central bank did not change monetary policy*. However, this latter perception—sometimes referred to as a constant-interest-rate (CIR) forecast (e.g. Rudebusch and Svensson, 1999)—is not observed and thus could be considered an omitted variable(s) in our regressions.[14]

We assess the potential for a resulting omitted-variable bias in two ways. First, in Supplementary Appendix A.2, we discuss how the interpretation of the regression coefficients changes if these short-run forecasts also reflect information about longer-run CIR forecasts due to persistence. As Romer and Romer (2004) note, short-run Greenbook forecasts will capture the central bank's perceived (exogenous) persistence of current variables: "both output growth and inflation are serially correlated enough that forecasts one and two quarters ahead provide a good indication of the likely forecasted path of the economy over longer horizons." We show that such persistence is unlikely to affect our estimate of the estimated inflation target. Second, we show empirically (in Section 4.3) that the results are robust to including farther-ahead Greenbook forecasts in measuring short-run expectations. Specifically, while our baseline uses averages of current and one-quarter-ahead forecasts, we produce alternative results using averages encompassing forecasts as far as four quarters ahead. Forecasts for horizons beyond four quarters are not consistently provided in the Greenbook and, in any case, would increasingly violate the CIR assumption.

Finally, it should be noted that we do not directly consider an interest-rate smoothing term in the objective function (e.g. as specified by Woodford (2003)) because the interest rate target is chosen by the FOMC at the close of the meeting and is therefore endogenous.[15] However, we do include a measure of short-term yield volatility as way to control for the FOMC potentially disliking large interest rate movements.[16]

2.3. *Asymmetric inflation preferences*

Allowing for asymmetry in the relationship between the inflation gap and loss is potentially important both generally for understanding the central bank's preferences but also specifically for estimating the inflation target. Allowing for asymmetry, however, requires a non-linear estimation approach because the ability to solve for π^* as a function of linear least squares coefficients, as we did in Table 1, relied on the symmetric, quadratic specification for the inflation gap. To relax the symmetry constraint in a parsimonious way, we alter the loss function specification to be

14. Greenbooks occasionally include CIR forecasts as "alternative simulations," but these are not consistently available for our full sample period.

15. That is, the federal funds rate is potentially correlated with unobserved factors that also affect the members' language.

16. Supplementary Appendix Tables G.4 and G.5 show the results are similar across all specifications and are also robust to an alternative proxy for interest-rate volatility, the squared first-difference of the 3-month fed funds futures rate.

Downloaded from https://academic.oup.com/restud/advance-article/doi/10.1093/restud/rdab094/6484651 by Federal Reserve Bank of San Francisco Research Library user on 25 Ja

TABLE 1
Estimated U-shaped loss function FOMC meetings: 2000–11

	(1)	(2)	(3)	(4)
π_t, GB	−1.145	−0.697	−1.066*	−1.724***
	(0.716)	(0.705)	(0.562)	(0.422)
π_t^2, GB	0.458**	0.384*	0.367**	0.510***
	(0.203)	(0.197)	(0.164)	(0.127)
$(u_t - u_t^*)$, GB		0.154**		−0.0478
		(0.0713)		(0.0865)
$(u_t - u_t^*)^2$, GB		−0.0170		0.0129
		(0.0143)		(0.0134)
Δy_t , GB			−0.101***	−0.0354
			(0.0184)	(0.0263)
Δy_t^2 , GB			−0.0278***	−0.0345***
			(0.00607)	(0.00674)
S&P return, 3m				−2.208***
				(0.598)
Shiller PE ratio				0.0186*
				(0.0107)
VIX uncertainty				0.00771
				(0.00579)
Short-term yield volatility				0.283
				(0.248)
Observations	1,676	1,676	1,676	1,676
Adjusted R^2	0.043	0.053	0.169	0.227
Type	Meeting	Meeting	Meeting	Meeting
Sample	2000m1–2011m12	2000m1–2011m12	2000m1–2011m12	2000m1–2011m12
$\hat{\pi^*}$	1.249	0.908	1.451	1.691
SE	(0.25)	(0.471)	(0.166)	(0.102)
95% C.I.	[0.759–1.739]	[−0.015–1.831]	[1.125–1.776]	[1.491–1.89]

Notes: The variables y and y* denote log GDP and log potential GDP, respectively. Standard errors (in parentheses) are robust to heteroskedasticity and two-way clustering by sample-month and speaker. Regressions include speaker fixed effects. * $p<0.10$, ** $p<0.05$, *** $p<0.01$.

linear (rather than quadratic) in the inflation gap but with potentially different slopes above and below zero:

$$N_{it} = \delta_1(\pi_t - \pi^*) * \mathbb{1}[\pi_t > \pi^*] + \delta_2(\pi_t - \pi^*) * \mathbb{1}[\pi_t \leq \pi^*] + \mathbf{X}_t\Lambda + \varepsilon_{it}, \quad (6)$$

where $\mathbf{X}_t$ is a vector of non-inflation variables entering the loss function and $\mathbb{1}[]$ denotes an indicator variable equal to one if the condition in the brackets is true and zero otherwise.

We estimate π^* along with the coefficients δ_1, δ_2, and Λ using a grid search over all possible values of π^* from 0 to 3.0, with step size of 0.01. At each iteration, we plug in the corresponding π^* into equation (6) and then estimate the equation via linear OLS. The final estimates of $\hat{\pi}^*$, $\hat{\delta}_1$, $\hat{\delta}_2$, and $\hat{\Lambda}$ are those resulting from the iteration yielding the minimum sum of squared residuals (SSR). The standard errors of $\hat{\delta}_1$, $\hat{\delta}_2$, and $\hat{\Lambda}$ are also obtained from that minimum-SSR regression and allow for two-way clustering by time (sample-month) and speaker (as in the symmetric case above). To obtain the standard error for $\hat{\pi}^*$, we use a block bootstrap procedure that allows for the same two-way clustering.[17]

17. The block bootstrap procedure loops over 100 draws, where each draw pulls all observations from a random sample of $T-1$ months (with replacement) and estimates $\hat{\pi}^*$ using the grid search method described above. The standard error of π^* is calculated as $[\text{var}(\hat{\pi}^*)^{(t)} + \text{var}(\hat{\pi}^*)^{(i)} - \text{var}(\hat{\pi}^*)^{(t,i)}]^{\frac{1}{2}}$ where $\text{var}(\hat{\pi}^*)^{(j)}$ is the variance of the $\hat{\pi}^*$ distribution when sampling with replacement by (j) and the 95% confidence interval is given by plus or minus 1.96 times the standard error (see Cameron and Miller, 2015).

Downloaded from https://academic.oup.com/restud/advance-article/doi/10.1093/restud/rdab094/6484651 by Federal Reserve Bank of San Francisco Research Library user on 25 Ja

3. MEASURING CENTRAL BANK SENTIMENT

We extract a direct measure of the FOMC's loss from publicly available transcripts. Below we outline the data sources used and describe how we quantified the negativity expressed in FOMC meetings, minutes, and speeches. We then show how our negativity measures have varied over time.

3.1. *Data sources*

Our main textual data source is the public archive of FOMC transcripts that can be downloaded directly from the Federal Reserve Board of Governors website: https://www.federalreserve.gov/monetarypolicy/fomc_historical.htm. These include all regular FOMC meetings from 1976 to 2015 (as of the time of this writing), though we make no use of pre-1986 transcripts in this paper given other data constraints. We converted these pdf files to text files and then parsed the text to generate separate text files for every utterance (sentence) in the transcript which is tagged with a speaker. There were 704,499 total utterances. We then apply three filters to remove utterances likely to be especially noisy. First, we drop utterances with less than five words. Second, we drop stand-alone utterances, defined as an utterance that is preceded and followed by utterances from different speakers. Third, we used the Oxford Dictionary of Economics (ODE) to filter out remarks that did not contain at least one economics-related term. We define a remark as a set of consecutive utterances by a single speaker (i.e. comments made by a speaker before discussion turned to another speaker). Because the ODE is quite expansive, containing 3,229 terms, and very little of the discussions at FOMC meetings is unrelated to economics, this filter removed just 3.0% of the 566,696 utterances remaining after the first two filters, leaving 549,649. As we show in Section 4, our results are very similar whether or not we apply the economic-terms filter. We also provide results based on two alternative filters, one that keeps only remarks containing terms related to inflation and one that keeps only remarks related to slack.[18]

Most speakers will make many remarks during a given meeting/transcript. Using historical lists of Federal Reserve governors and bank presidents, and which presidents were voters vs. non-voters at any given meeting, we are also able to perform analyses on particular subsets of remarks, such as those by Fed chairs, governors, presidents, and voters. We exclude remarks from non-FOMC members (mostly Fed research staff members) given that our goal is to estimate the collective loss function of the FOMC. This exclusion also has the benefit of, in effect, filtering out the portion(s) of each FOMC meeting in which the staff presents its assessment of economic and financial conditions to the committee, leaving primarily only the portion of the meeting in which committee members discuss their economic outlooks and policy views.

Using the methodology described below, we calculate a unique negativity score for each speaker-meeting, pooling all remarks by a given speaker at that meeting. This yields 4,223 speaker*meeting observations over all meetings between 1986 and 2015, and 1,676 for our primary sample period of 2000 to 2011.

For additional analyses, we also collected the minutes of the FOMC meetings, which are released to the public 3 weeks after the corresponding meeting. The minutes are a summary of the discussion that took place at the meeting, with no attribution of any specific comments or opinions to individual FOMC members.[19]

18. The inflation terms are: price, prices, pricing, inflation, inflationary, cpi, pcepi, core, deflation, deflationary, disinflation, and disinflationary. The slack terms are: output gap, unemployment, unemployed, jobless, joblessness, nonemployment, U3, U4, U5, U6, nairu, natural rate, potential output, potential GDP, potential GNP, slack, and utilization.

19. PDF files for minutes from 1993 onward are available at: https://www.federalreserve.gov/monetarypolicy/fomc_historical.htm.

Downloaded from https://academic.oup.com/restud/advance-article/doi/10.1093/restud/rdab094/6484651 by Federal Reserve Bank of San Francisco Research Library user on 25 Ja

3.2. *Quantifying negativity in FOMC text*

There is a large and growing literature aimed at quantifying sentiment from text. We use a method known as the "Bag of Words" or "lexical" approach, which relies on pre-defined dictionaries of words that are associated with particular sentiments.[20] We employ an updated version of the dictionaries from Loughran and McDonald (2011) (hereafter, LM), who construct lists of negative and positive words curated to be appropriate for financial text. They show that their dictionaries are superior for classifying economic and financial texts to other dictionaries, for example, the Harvard Psychosociological Dictionary, which tend to mis-categorize words that are neutral in a financial/economic context (e.g. tax, costs, capital, expense, liability, risk, excess, and depreciation).[21] There are 2,355 negative words and 354 positive words in the LM dictionaries. We ignore positive and negative words preceded by "not" or "n't," as recommended by Loughran and McDonald (2016).[22] We show an example of scored text from an FOMC transcript in Supplementary Appendix B.

The net negativity scores, N_{it}, are calculated as the fraction of negative words by speaker i at a given FOMC meeting in month t minus the fraction of positive words by speaker i at the same meeting (subsequent to any text filtering as described in the previous subsection). Using the difference between the negative and positive fractions—as opposed to one or the other—ensures that our measure is independent of overall emotivity (i.e. use of both negative and positive words). Nonetheless, we show later that the results are similar using either the negative or positive fractions alone.[23]

To help validate our sentiment measure, we performed a "human audit" as follows. We randomly selected 10% of the roughly 1,900 speaker-by-meeting observations of FOMC meeting text from 2000 to 2013. We then asked a group of 14 research assistants to read and evaluate the tone of a given observation's text on a 5-point scale (from least to most negative).[24]

20. The other approach uses machine learning (ML) techniques. See Liu (2010) for a detailed description of the ML approach to sentiment analysis and Shapiro, Sudhof and Wilson (2018) for an application of both lexical and ML approaches to measuring news sentiment.

21. Heston and Sinha (2015) measure negativity in news articles about companies and estimate their impact on those companies' stock returns. They use the Harvard Psychosociological Dictionary along with Loughran and McDonald's dictionary. More closely related to our article, Jegadeesh and Wu (2017) use a combination of these same two dictionaries to measure the tone (net positivity) of FOMC meeting minutes, with sentences in the minutes separated by topic. Apel, Grimaldi and Hull (2019) and Picault and Renault (2017) use the lexical approach to classify the hawkishness vs. dovishness of central bank communications for the FOMC and the ECB, respectively. Lastly, the study of economic policy uncertainty by Baker, Bloom and Davis (2016) also uses a lexical-based measure as part of its uncertainty index. That measure is a simple count of news articles containing terms such as "uncertain" and "not certain" along with terms related to economic policy.

22. We first removed stop words and intensifiers and then searched for positive and negative words preceded by "not" and "n't." For example, "not a very bad reading" would be reduced to "not bad reading." We also ran specifications where we removed the words "unemployed" and "unemployment" from the LM negativity dictionary since higher unemployment can be a positive or a negative event depending on the current unemployment gap. We found no change in our results. We reviewed the LM dictionaries for other such terms and did not find any.

23. See Supplementary Appendix Table G.1.

24. Specifically, they were instructed to: "Assign a score from 1 [very positive] to 5 [very negative] of the sentiment/negativity expressed by the speaker at that meeting. This score should be thought of as your best guess of the score that speaker would give if they were asked, how negative are you today? at that meeting. That is, this score should be your assessment of how you think the speaker felt (about economic conditions) as of that meeting date, without imposing how you think he or she should have felt. That is, you should ignore factual statements (e.g. revenue growth is low, unemployment rose last month, inflation has been trending higher) and focus on the speakers sentimental language (e.g. I'm not happy about that, This concerns me, ... which is a troubling development)." The resulting data set consists of 182 human audit scores.

Downloaded from https://academic.oup.com/restud/advance-article/doi/10.1093/restud/rdab094/6484651 by Federal Reserve Bank of San Francisco Research Library user on 25 Ja

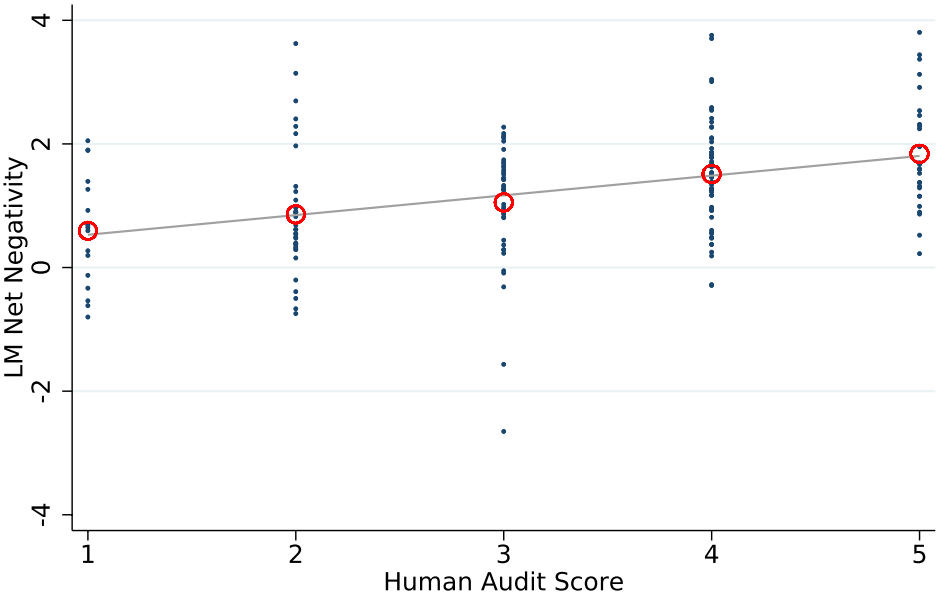

FIGURE 1

LM net negativity vs. human scores

The audit results are shown in the scatterplot in Figure 1. Each dot represents a speaker-meeting observation. The negativity score assigned by a research assistant is shown on the x-axis while the LM net negativity fraction is shown on the y-axis. The mean LM fraction, by human score category, is indicated by the red circles. A linear regression fit line also is provided. While there is a good deal of variation in the LM net negativity measure within a category, we find a strong correlation between the human scores and LM net negativity. Specifically, the Spearman's rank correlation is 0.40, which is statistically significant below the 0.01% level.[25]

Figure 2 plots the estimated time (meeting-month) fixed effects from a simple regression of the negativity scores on time and speaker fixed effects, using data from January 1986 through December 2015. The dotted line shows the actual time series of coefficients on the time dummies, while the solid line shows a smoothed version calculated as an 11-month centred moving average. FOMC meeting negativity has a clear countercyclical pattern, rising in recessions and falling in

25. We also performed the same analysis on two other common computational sentiment analysis measures. The first is a lexical measure analogous to the LM net negativity but using the Harvard Psychosociological Dictionary to define negative and positive words. The rank correlation between this measure and the human scores is just 0.14. The comparatively worse performance of the Harvard Psychosociological Dictionary based lexical measure relative to LM illustrates the importance of domain (economics/finance) specificity in constructing a sentiment measure. The second alternative measure we considered is the widely used, open-source python package, Vader. Vader is a predictive model that provides estimated probabilities that a given set of text is negative, neutral, and positive. The Vader model is based on a machine learning algorithm trained on a combination of general english language words and social media data. Relative to LM, Vader has the advantage of taking into account parts-of-speech, context (to some extent), and other advanced features of current natural language processing tools. However, it is not domain-specific to economics and finance. We find a rank correlation between the Vader net negativity measure (probability negative minus probability positive) and the human scores of 0.41, very similar to that of LM net negativity. As noted later in the article, using Vader to measure net negativity also yields similar results to our baseline LM measure in the loss function regressions.

Downloaded from https://academic.oup.com/restud/advance-article/doi/10.1093/restud/rdab094/6484651 by Federal Reserve Bank of San Francisco Research Library user on 25 Ja

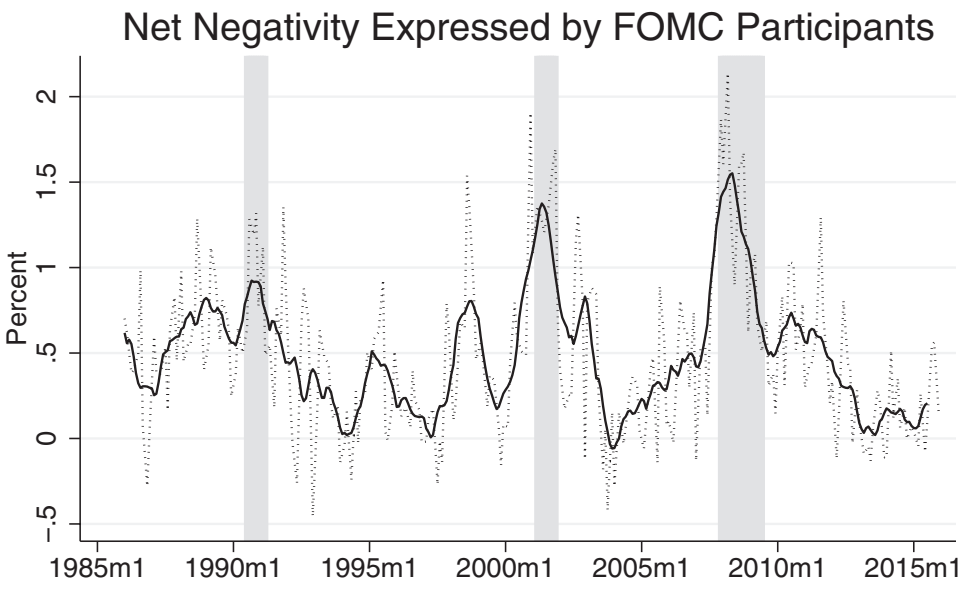

FIGURE 2
Transcripts of FOMC meetings

expansions. Negativity rises especially sharply during the 2007–9 Great Recession. It gradually fell after the Great Recession but was still relatively high by the end of 2015, consistent with the weakness of the recovery during the 2010–15 period.[26]

It also is interesting that negativity was not relatively higher during the late 1980s, when inflation and unemployment were relatively higher, than in later years, in contrast to the picture one gets from the so-called "misery index" which sums inflation and the unemployment rate. Our results in Section 6.2 suggest this difference can be explained by the late 1980s having a combination of a higher natural rate of unemployment, a higher inflation target, and lower sensitivity to inflation in the FOMC's loss function.

4. MAIN RESULTS

4.1. *Symmetric loss function*

We start with estimating versions of a U-shaped (i.e. quadratic) loss function, as shown in equation (5). The estimation uses 1,676 speaker-meeting observations over the period January 2000 through December 2011. We use January 2000 as the sample start date for our baseline analysis for two reasons. First, the Greenbook does not contain forecasts prior to 2000 for PCE

26. We similarly plot the net negativity for the minutes from each FOMC meetings and speeches of FOMC members in Supplementary Appendix Figure G.1. Net negativity in the minutes moves quite similarly to that in the transcripts through 2015. After 2015, net negativity in the minutes remained relatively low until 2020; it spiked around the time of the onset of the COVID-19 pandemic. Also, it is worth noting that the time series pattern of our FOMC negativity measure appears to be quite consistent with the patterns of the topic-specific time series of net positivity in FOMC minutes constructed by Jegadeesh and Wu (2017), plotted in their Figure 2.

Downloaded from https://academic.oup.com/restud/advance-article/doi/10.1093/restud/rdab094/6484651 by Federal Reserve Bank of San Francisco Research Library user on 25 Ja

inflation, which is the measure that the FOMC publicly stated in 2012 was its preferred measure of inflation (FOMC, 2012).[27] Second, it is widely believed that the inflation target was different (and likely time-varying) prior to the Great Moderation starting in the late 1990s (see Ireland, 2007; Cogley, Primiceri and Sargent, 2010; Coibion and Gorodnichenko, 2011). In Section 6.2, we expand the sample period to 1986-2011 (which entails splicing in a "PCE-adjusted" core CPI inflation Greenbook forecast) and estimate a version of the loss-function with a time-varying π_t^* and also allowing for the time of the regime change to a constant target period to float freely. The results from this exercise indicate that the constant-target period begins in 2000, consistent with our baseline sample period. We use December 2011 as the final month of the baseline sample period because the FOMC first officially announced an inflation target (of 2%) in January 2012. We use our sentiment approach to separately estimate the post-2011 target on the FOMC minutes in Section 6.1.

The results of estimating a U-shaped loss function are shown in Table 1. Note that standard errors (shown in parentheses) are robust to heteroskedasticity and two-way clustering by time (meeting) and speaker. The first column shows the results of a simple specification that includes only π_t and π_t^2. This estimated quadratic relationship also is plotted visually in Figure 3(a), overlaid on a bin-scatter plot of negativity against short-run inflation.[28] In the scatter plot, both inflation and negativity have been residualized with respect to speaker fixed effects. The negative sign on π_t and the positive sign on π_t^2 implies that loss has a convex relationship with inflation and that there exists a loss-minimizing inflation rate, π^*.

As indicated in column (1) of the table, the estimated quadratic loss function implies a $\hat{\pi}^*$ of approximately 1.25%, with a 95% confidence interval of 0.8 to 1.7. How we arrive at this value of $\hat{\pi}^*$ is clear from Figure 3(a). Negativity tends to be lowest when inflation is around 1.3–1.5% and rises as inflation moves outside that range in either direction.

This estimate is considerably lower than both the 2.0% inflation target officially announced by the FOMC in January 2012 and survey measures of longer-run inflation expectations, as seen in Figure 4.[29] Moreover, as noted in the introduction, the transcripts reveal that at least some committee members were well aware of the discrepancy between the FOMC's target and market participants' perception of the target.[30] The estimate is also below the 2000–11 average of actual year-over-year core PCE inflation, which is 1.8 (measured using either current vintage or real-time data).[31]

27. While it is clear from the Greenbook and from the FOMC transcripts that the FOMC shifted from focusing on the CPI to the PCEPI measure in the years around 2000, the exact timing and extent of the shift is somewhat ambiguous. As a robustness exercise, we estimate the loss function using Greenbook CPI forecasts and, as shown in Supplementary Appendix Table G.3, found very similar results but for a higher estimated inflation target. Specifically, comparing Table G.3 to the corresponding columns in Table 2, we see that the estimated inflation target based on the CPI is 20–50 basis points higher than that based on the PCE, consistent with the fact that CPI inflation averaged around 30 basis points higher than PCE inflation over 2000–11.

28. A bin-scatter plot divides observations into equal bins—100 in this case—and then plots the mean of x against the mean of y for each bin. Bin-scatter plots are frequently used to visualize data when there are a very large number of underlying data points.

29. Market-based measures of inflation expectations were also above this value. For example, the TIPS-based 5-year-ahead, 5-year-forward break-even inflation rate (based on CPI) averaged 2.4 between 2003 and 2012.

30. In addition to William Poole, Timothy Geithner and Eric Rosengren made similar comments at the September 2006 and January 2009 meetings, respectively. The result is also reminiscent of Afrouzi *et al.* (2015)'s finding that inflation expectations of firms in New Zealand have been persistently higher than the central bank's announced target over the past 25 years.

31. The average excluding the zero lower bound period (Dec. 2008 through the end of our sample period), which arguably hampered the Fed's ability to boost inflation, is even higher at 1.9.

Downloaded from https://academic.oup.com/restud/advance-article/doi/10.1093/restud/rdab094/6484651 by Federal Reserve Bank of San Francisco Research Library user on 25 Ja

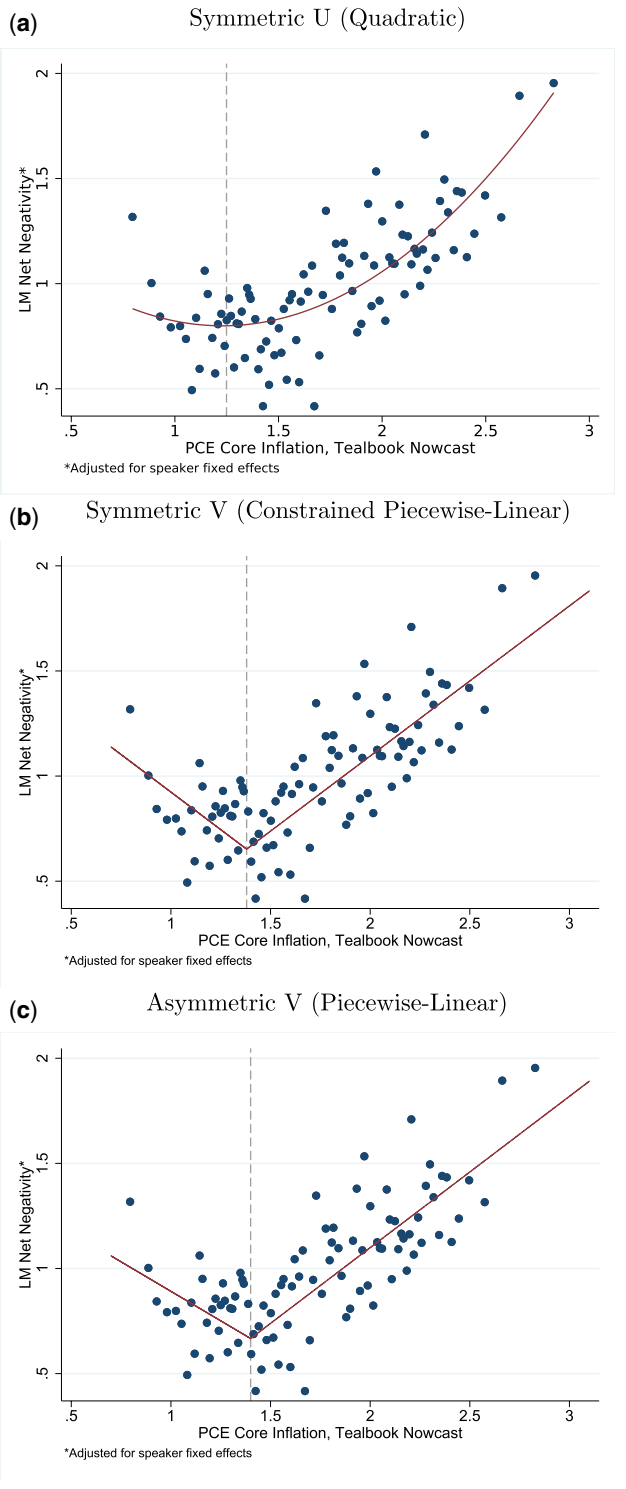

FIGURE 3
Estimated loss function for inflation—alternative specifications

Downloaded from https://academic.oup.com/restud/advance-article/doi/10.1093/restud/rdab094/6484651 by Federal Reserve Bank of San Francisco Research Library user on 25 Ja

TABLE 2
Estimated V-shaped loss function FOMC meetings: 2000–11

	(1) Symmetric V	(2) Symmetric V	(3) Symmetric V	(4) Asymmetric V	(5) Asymmetric V	(6) Asymmetric V
$\lvert \pi_t - \pi^* \rvert$, GB	0.714***	0.429***	0.490***			
	(0.148)	(0.126)	(0.132)			
$(\pi_t > \pi^*) = 0 \times (\pi_t - \pi^*)$, GB				−0.561	−0.838***	−1.037***
				(0.345)	(0.276)	(0.212)
$(\pi_t > \pi^*) = 1 \times (\pi_t - \pi^*)$, GB				0.720***	0.433***	0.330**
				(0.147)	(0.119)	(0.135)
$(u_t - u_t^*)$, GB			0.0234			−0.00507
			(0.0763)			(0.0846)
$(u_t - u_t^*)^2$, GB			0.00313			0.00402
			(0.0127)			(0.0128)
Δy_t , GB		−0.103***	−0.0322		−0.106***	−0.0383
		(0.0178)	(0.0265)		(0.0176)	(0.0252)
Δy_t^2 , GB		−0.0263***	−0.0331***		−0.0267***	−0.0332***
		(0.00577)	(0.00597)		(0.00579)	(0.00646)
S&P return, 3m			−2.101***			−2.122***
			(0.607)			(0.628)
Shiller PE ratio			0.0207*			0.0169
			(0.0106)			(0.0104)
VIX uncertainty			0.00770			0.00788
			(0.00617)			(0.00635)
Short-term yieldvolatility			0.249			0.226
			(0.240)			(0.246)
Observations	1,676	1,676	1,676	1,676	1,676	1,676
Adjusted R^2	0.049	0.175	0.229	0.049	0.177	0.230
Type	Meeting	Meeting	Meeting	Meeting	Meeting	Meeting
Sample	2000m1–2011m12	2000m1–2011m12	2000m1–2011m12	2000m1–2011m12	2000m1–2011m12	2000m1–2011m12
Symmetry p-value	0.	0.	0.	0.622	0.115	0.001
$\hat{\pi}^*$	1.38	1.38	1.55	1.4	1.35	1.35
SE	(0.11)	(0.11)	(0.13)	(0.22)	(0.18)	(0.19)
95% C.I.	[1.16–1.6]	[1.16–1.6]	[1.3–1.8]	[0.97–1.83]	[1–1.7]	[0.98–1.72]

Notes: The variables y and y* denote log GDP and log potential GDP, respectively. Standard errors (in parentheses) are robust to heteroskedasticity and two-way clustering by sample-month and speaker. Regressions include speaker fixed effects. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Downloaded from https://academic.oup.com/restud/advance-article/doi/10.1093/restud/rdab094/6484651 by Federal Reserve Bank of San Francisco Research Library user on 25 Ja

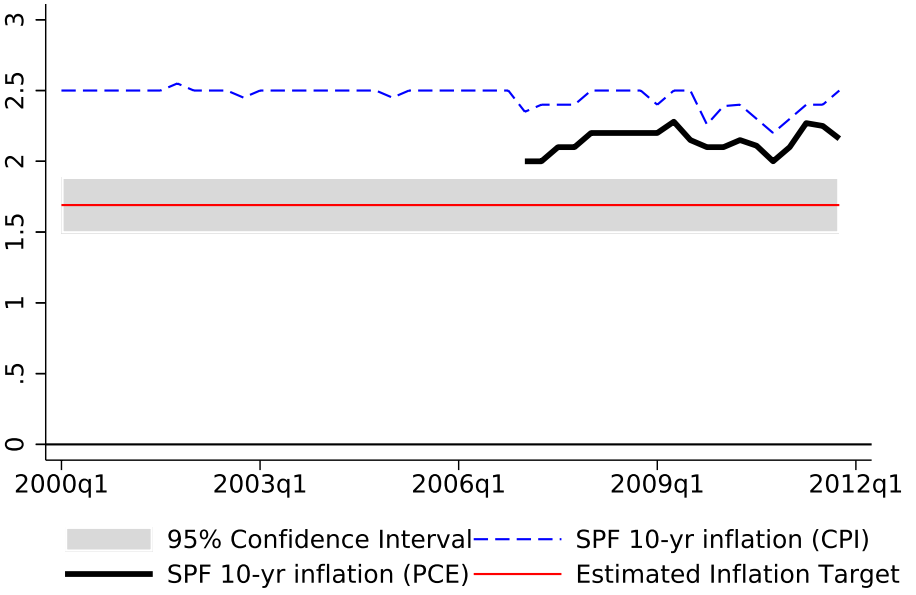

FIGURE 4
$\hat{\pi}^*$ Compared with SPF inflation expectations

We next add non-inflation variables, linear and squared, to in the remaining columns of Table 1. In column (2), we estimate the standard loss function including the unemployment gap, with both a linear and squared term. The unemployment gap is measured using the Greenbook's forecast of the unemployment rate minus its estimate of the natural rate of unemployment (as of that meeting). As with other Greenbook forecasts, we average the current-quarter and one-quarter ahead forecasts. The coefficients on both the linear and squared unemployment gaps are found to be statistically insignificant.

In column (3), we replace the unemployment gap with output growth, again measured by the Greenbook short-run forecast. As noted earlier, some prior studies have found a role for output growth in central bank preferences and interest rate policy (e.g. Walsh, 2003; Orphanides and Williams, 2006; Coibion and Gorodnichenko, 2011). We find that negativity falls significantly with both output growth and its squared term—implying loss falls monotonically with output growth. This result is consistent with the prior studies finding a role for output growth in interest rate rules. It is also consistent with the fact, documented in Thornton (2011), that the FOMC explicitly stated its policy objectives in terms of "economic growth" or "growth in output" from 1991 through 2008, when it changed its policy directive to refer to "maximum employment."[32]

Some FOMC participants and outside observers have suggested that financial variables may also enter the central bank's loss function (Bernanke and Gertler, 2001; Kohn, 2006; Peek *et al.*, 2015). We examine this hypothesis by adding a number of common financial measures to our specification. Specifically, in column (4), we include the trailing 90-day change (preceding the

32. In an additional regression, we interacted a post-2008 dummy with output growth to test the hypothesis that the FOMC's language in the transcripts coincided with the change in the language of its stated objectives. The coefficient on the interaction was positive and significant, indicating that movements in output growth changed the FOMC's loss by a smaller amount in the post 2008 period, corroborating the change in the stated objectives.

Downloaded from https://academic.oup.com/restud/advance-article/doi/10.1093/restud/rdab094/6484651 by Federal Reserve Bank of San Francisco Research Library user on 25 Ja

meeting) in the S&P 500 index, the previous month's Shiller P/E ratio (a measure of the recent *level* of stock market prices),[33] the stock market volatility index (VIX) and real-time short-term yield volatility taken from Christensen and Rudebusch (2012).[34] The short-term yield volatility measure is included to control for a possible interest-rate smoothing motive.[35] Negativity moves strongly in the opposite direction with recent stock market returns and in the same direction as financial volatility[36] but statistically insignificant.[37] These results are consistent with Peek *et al.* (2015), who argue that financial variables do enter the FOMC's loss function based on their evidence that references to financial instability in the FOMC transcripts significantly improve the fit of an estimated Taylor Rule. They are also consistent with Cieslak and Vissing-Jorgensen (2018) who find that negative mentions of the stock market in FOMC minutes and transcripts have high explanatory power for interest rate policy changes.

We next estimate a loss function specification that replaces the quadratic functional form for inflation with a constrained/symmetric piecewise-linear form. This is equivalent to estimating the piecewise-linear specification shown in equation (6) while constraining $\delta_1 = -\delta_2$. In other words, we estimate a V-shaped instead of a U-shaped loss function (in terms of inflation) while still maintaining symmetry. The results are shown in columns (1), (2), and (3) of Table 2. The specification underlying column (1) includes only the inflation gap terms (i.e. restricting ω and κ to equal 0). We obtain a $\hat{\pi}^*$ of 1.4, which is slightly higher than that obtained using the quadratic specification. The results in column (1) are shown visually in Figure 3(b). As with Figure 3(a), we overlay this fitted relationship on a bin-scatter plot of negativity against inflation. We see that the V-shape appears to fit the data slightly better than the U-shape (as verified by the higher R^2 in the regression). Column (2) adds output growth and output growth squared. Both are found to reduce the FOMC's net negativity and their inclusion has no effect on the estimated inflation target, which remains at 1.4. Finally, the specification underlying column (3) includes all additional factors which results in a modestly higher estimate of π^* of 1.55. The coefficients for the other variables are similar to the analogous quadratic specification (column (4) of Table 1), with a negative effect on loss from output growth squared and recent stock market returns and a positive effect on loss from the Shiller P/E ratio.

4.2. *Asymmetric loss function*

The results above are based on the assumption of a symmetric loss function. We now relax that assumption by estimating the piecewise-linear specification shown in equation (6).[38] The results

33. The Shiller PE ratio is available monthly, so we can only include the preceding month's value as opposed to the value immediately preceding the meeting.

34. Both volatility measures are measured as the average value over the 3 days preceding the meeting.

35. Results did not change when using the squared values of the first-difference in real-time short-rate projections (see Supplementary Appendix Tables G.4 and G.5). The short-rate projection is measured as the implied 3 months-ahead fed future rate from Christensen and Rudebusch (2012) at the time right before the meeting (i.e. the average value of the projection over the 3 days preceding the meeting). We prefer the implied-volatility measure because it directly captures the market's perceived risk about interest-rate movements. Interest-rate risk is something the FOMC may have preferences over and try to mitigate by interest-rate smoothing.

36. We also ran specifications that included forecast uncertainty as control (measured by the SPF dispersion variables) which we report in tables G.6 and G.7.

37. Note that when financial variables are added, the coefficient on linear output growth becomes approximately zero while squared output growth remains negative and significant. Taken literally, this implies that loss falls as growth deviates from zero, even in the negative direction. However, we note that the coefficient on the linear term is imprecisely estimated because negative growth is rarely observed in our sample (10% of the sample) and occurs contemporaneously with strong movements in the financial variables.

38. We continue to assume symmetry for output growth and the unemployment gap for two reasons. First, the inflation target is our primary focus and this exercise is meant primarily to assess whether our estimated inflation target is

Downloaded from https://academic.oup.com/restud/advance-article/doi/10.1093/restud/rdab094/6484651 by Federal Reserve Bank of San Francisco Research Library user on 25 Ja

when we allow for asymmetry are very similar and shown in column (4) with only inflation terms included, column (5) with output growth included, and column (6) with all additional factors included. The results in column (4) are shown visually in Figure 3(c). The slope by which loss increases away from target is about the same (in absolute value) above and below the target. The coefficients on the inflation gap above and below zero indicate that negativity rises faster as inflation moves further below target than it does when inflation moves further above target. That is, the absolute value of the slope coefficient for the inflation gap below zero is larger than the slope coefficient for the inflation gap above zero. The "Symmetry p-val" row at the bottom of Table 2 shows the p-value on the null hypothesis that $\hat{\delta}_1 = -\hat{\delta}_2$ based on a t-test using the two-way clustered variance-covariance matrix. In the full specification, the difference in slopes is found to be significant, with a p-value of 0.001. We performed a number of robustness checks in terms of including additional controls, which we report in Supplementary Appendix Tables G.6–G.9. This includes specifications with the squares of financial variables, a lagged dependent variable, lagged inflation, 10-year inflation expectations, and the dispersion in SPF nowcasts and forecasts of inflation, GDP growth, and unemployment. Overall, the estimates remain stable and the inflation target hovers around 1.5 in all of these additional specifications.

One caution in interpreting these results deserves mention. The asymmetry of the loss function in terms of inflation relative to target can, of course, only be estimated over the range of observed inflation rates during the sample period. Greenbook near-term forecasts of core PCE inflation between 2000 and 2011 rarely fell much below 1% and never got very close to zero. FOMC policymakers have frequently spoken about the particular dangers of deflation (especially given the zero lower bound on interest rate policy). It is entirely possible that FOMC preferences are much more asymmetric in the range of inflation close to zero. That is, the loss function slope could sharply steepen (in absolute value) as inflation rates get close to or below zero, which we have not observed in the data.

4.3. *Loss function implied by alternative forecast horizons*

As discussed in Section 2.2, the Greenbook forecasts of inflation, the unemployment gap, and real GDP growth are short-run forecasts and so, by omitting the FOMC's longer-run expectations of these variables, our results could potentially suffer from omitted variable bias. To assess the potential for such bias, we investigate how the results—especially the estimated inflation target—change as one includes farther-ahead forecasts in these macroeconomic variables. Table 3 shows the results of this exercise. The first column simply redisplays the results of estimating our baseline specification, in which inflation, the unemployment gap, and real GDP growth are measured as averages of the Greenbook's current-quarter and one-quarter-ahead forecasts. In column (2), these variables are measured as averages of current-quarter through two-quarter-ahead forecasts; column (3) uses averages of current-quarter through three-quarter ahead forecasts; and column (4) uses averages of current-quarter through four-quarter-ahead forecasts.[39]

We see that the estimated loss function is quite similar across these alternative specifications. In particular, the estimated inflation target is little affected by including farther-ahead forecasts in the measures of the macroeconomic variables. For example, both the baseline specification, using

robust to asymmetry in inflation preferences. Second, the methodology for allowing for asymmetry involves a grid search and bootstrapping over the preferred inflation target. Allowing for asymmetry in the other variables would be extremely computationally intensive.

39. This sensitivity check mimics that of Orphanides (2004), who estimated forward-looking interest rate rules using Greenbook inflation forecast averages encompassing horizons ranging from 1 to 4 quarters ahead. He found the results were insensitive to the alternative forecast horizons.

Downloaded from https://academic.oup.com/restud/advance-article/doi/10.1093/restud/rdab094/6484651 by Federal Reserve Bank of San Francisco Research Library user on 25 Ja

TABLE 3
Estimated V-shaped loss function. Alternative Greenbook Forecast horizons: FOMC meetings: 2000–11

	(1) 0-1 qtr ahead avg	(2) 0-2 qtr ahead avg	(3) 0-3 qtr ahead avg	(4) 0-4 qtr ahead avg
$\lvert \pi_t - \pi^* \rvert$, GB	0.490***	0.589***	0.651***	0.661***
	(0.132)	(0.153)	(0.164)	(0.172)
$(u_t - u_t^*)$, GB	0.0234	0.105	0.180**	0.249***
	(0.0763)	(0.0776)	(0.0821)	(0.0886)
$(u_t - u_t^*)^2$, GB	0.00313	−0.0143	−0.0298**	−0.0432***
	(0.0127)	(0.0123)	(0.0131)	(0.0144)
Δy_t , GB	−0.0322	0.0253	0.0893	0.169**
	(0.0265)	(0.0383)	(0.0594)	(0.0860)
Δy_t^2 , GB	−0.0331***	−0.0436***	−0.0562***	−0.0703***
	(0.00597)	(0.00761)	(0.0106)	(0.0145)
S&P return, 3m	−2.101***	−1.977***	−1.759***	−1.604***
	(0.607)	(0.590)	(0.594)	(0.609)
Shiller PE ratio	0.0207*	0.0253**	0.0293***	0.0314***
	(0.0106)	(0.0106)	(0.0108)	(0.0109)
VIX uncertainty	0.00770	0.0102*	0.0122**	0.0136**
	(0.00617)	(0.00609)	(0.00607)	(0.00630)
Short-term yield volatility	0.249	0.314	0.357	0.388
	(0.240)	(0.249)	(0.258)	(0.266)
Observations	1,676	1,676	1,676	1,676
Adjusted R^2	0.229	0.223	0.218	0.213
Type	Meeting	Meeting	Meeting	Meeting
Sample	2000m1–2011m12	2000m1–2011m12	2000m1–2011m12	2000m1–2011m12
$\hat{\pi}^*$	1.55	1.56	1.55	1.52
SE	(0.13)	(0.13)	(0.12)	(0.11)
95% C.I.	[1.3–1.8]	[1.31–1.81]	[1.31–1.79]	[1.3–1.74]

Notes: Columns differ in what horizons are included in the forecast average for each Greenbook (GB) variable. Column (1) is our baseline specification and includes current-quarter (0-qtr ahead) and 1-quarter ahead forecasts. Column (2) includes 0- through 2-quarter ahead forecasts; column (3) includes 0- through 3-quarter ahead forecast; and column (4) includes 0- through 4-quarter ahead forecasts. Standard errors (in parentheses) are robust to heteroskedasticity and two-way clustering by sample-month and speaker. Regressions include speaker fixed effects. * $p<0.10$, ** $p<0.05$, *** $p<0.01$.

averages of current-quarter and one-quarter-ahead forecasts, and the specification using averages of current-quarter through four-quarter-ahead forecasts, yield $\hat{\pi}^* \approx 1.5$.[40]

4.4. *Loss function implied by filtered transcripts*

We next investigate the loss function using the net negativity measures from particular subsets of the full corpus of meeting transcript text. As described in Section 3, our baseline negativity measure is constructed from transcripts text that filtered out speaker remarks that did not contain at least one word from the Oxford Dictionary of Economics (ODE). This filter was intended to remove remarks in the FOMC meetings that are unrelated to economic discussions. Using similar methodology, we also explored (1) not filtering at all (keeping all remarks), (2) keeping only remarks associated with inflation, and (3) keeping only remarks associated with slack. For

40. We also ran specifications including median inflation expectations from the SPF (10-year-ahead CPI) as an additional regressor (see Supplementary Appendix Tables G.8 and G.9). Of course, it is possible that the FOMC's loss function contains constant-interest-rate forecasts beyond four quarters ahead, which are unobserved to us, that could create an omitted variable bias. We show in Supplementary Appendix A.2 that under this assumption the estimated coefficients on the short-run forecasts reflect both their direct impact on short-run loss plus their effect on longer-run loss given their degree of persistence.

Downloaded from https://academic.oup.com/restud/advance-article/doi/10.1093/restud/rdab094/6484651 by Federal Reserve Bank of San Francisco Research Library user on 25 Ja

inflation and slack, we constructed our own lists as subsets of ODE words and phrases. From these alternative sets of text, we computed the LM net negativity score by speaker-meeting as before (fraction of negative words minus fraction of positive words). Time series plots of the monthly median values of these three alternative measures are provided in Supplementary Appendix Figure E2. Loss function estimates using these alternative negativity measures are provided in Table 4. The loss function estimates using negativity calculated from the unfiltered text (column (2)) are very similar to the baseline estimates. This is not surprising given that nearly 90% of the transcript remarks contain economic terms. When we filter to include only remarks related to inflation (column (3)) or slack terms (column (4)), the estimates are generally consistent with the baseline estimates but with less precise estimates of the inflation target. Interestingly, the estimated inflation target is somewhat higher when we focus on remarks related to slack, consistent with slack remarks being more negative in low-inflation/high-slack periods and more positive in high-inflation/low-slack periods.

Lastly, we consider two potential concerns one might have with our baseline measure of net negativity. The first concern is the lexical sentiment analysis could mistakenly categorize a purely factual statement as positive or negative if that statement happens to contain a word in the LM positive or negative dictionaries. We address this concern in two ways. First, we construct an alternative net negativity measure that excludes/ignores statements that do *not* include a personal pronoun (I, I'm, I'll, I've, I'd, W/we, W/we're, W/we've, W/'d, M/my, M/mine, O/our, us) based on the logic that utterances containing personal pronouns are more likely to express an opinion and not be a purely factual statement. The loss function regression using this measure are shown in column (5) of Table 4. The coefficients are very similar to the baseline results (in column (1)), and the estimated inflation target is unchanged at 1.55. Second, we construct a measure after manually removing factual phrases that would be scored positive or negative according to the LM lexicons. For example, the statement, "inflation declined last month," is scored as negative using the LM lexicons because "declined" is in the LM negative dictionary. To filter out such phrases, we first searched for directional words like "declined" in the LM positive and negative dictionaries. We identified 73 directional words (including variations such as "gain," "gained," "gaining"). We then identified all economic variables within the 1500 most frequently occurring words (unigrams) in the transcripts—words such as "inflation," "price," "output," etc.[41] Finally, to construct this alternative "no fact phrase" measure of net negatively, we filter out (i.e. treat as neutral) any bigram represented by the pairing of a directional word and an economic variable. The regression results based on this alternative measure are shown in column (6) of Table 4. The baseline results are virtually unchanged by filtering out such fact phrases.

The second concern is that some LM terms classified as positive or negative could plausibly be construed as neutral in central bank parlance. For example, the words "easing" and "stability," which are classified as positive by LM, might be mentioned in regard to neutral economic terms such as "quantitative easing" or "price stability."[42] More generally, the LM dictionary does not consider the word's context—the simplest example would be bigrams and trigams that contain positive or negative LM unigrams. We went through the 50 most common bigrams and trigrams that contain either a positive or negative LM word in the transcripts, and flagged those that have plausibly neutral sentiment (see Supplementary Appendix Table G.11). We then filtered out these flagged bigrams and trigrams.[43] The results using this filtered subset of text are shown in column

41. See Supplementary Appendix Table G.12 for the full list of directional words and economic variables.
42. Note, our baseline measure removes the words "unemployed" and "unemployment" for a similar reason.
43. We also removed the unigrams "volatile," "volatility," "ease," and "easing."

Downloaded from https://academic.oup.com/restud/advance-article/doi/10.1093/restud/rdab094/6484651 by Federal Reserve Bank of San Francisco Research Library user on 25 Ja

TABLE 4
Negativity using filtered subsets of transcript text FOMC meetings: 2000–11

	(1) Baseline	(2) All text	(3) Inflation text	(4) Slack text	(5) Opinion text	(6) No fact phrases	(7) Adjusted LM
$\lvert \pi_t - \pi^* \rvert$, GB	0.490***	0.494***	0.412***	0.479**	0.570***	0.490***	0.498***
	(0.132)	(0.133)	(0.136)	(0.187)	(0.131)	(0.132)	(0.126)
$(u_t - u_t^*)$, GB	0.0234	0.0184	−0.0232	−0.0471	0.0792	0.0234	0.00539
	(0.0763)	(0.0749)	(0.0899)	(0.107)	(0.0609)	(0.0763)	(0.0752)
$(u_t - u_t^*)^2$, GB	0.00313	0.00383	0.0137	0.0195	−0.00957	0.00313	0.00663
	(0.0127)	(0.0124)	(0.0152)	(0.0185)	(0.0113)	(0.0127)	(0.0127)
Δy_t , GB	−0.0322	−0.0330	−0.0527	−0.0929*	−0.00987	−0.0322	−0.0306
	(0.0265)	(0.0262)	(0.0343)	(0.0492)	(0.0287)	(0.0265)	(0.0258)
Δy_t^2 , GB	−0.0331***	−0.0333***	−0.0403***	−0.0352***	−0.0266***	−0.0331***	−0.0313***
	(0.00597)	(0.00590)	(0.00720)	(0.00983)	(0.00641)	(0.00597)	(0.00578)
S&P return, 3m	−2.101***	−2.087***	−2.272***	−3.037***	−2.031***	−2.101***	−2.171***
	(0.607)	(0.604)	(0.670)	(0.604)	(0.599)	(0.607)	(0.578)
Shiller PE ratio	0.0207*	0.0208**	0.0219*	0.0314**	0.0237***	0.0207*	0.0182*
	(0.0106)	(0.0104)	(0.0117)	(0.0144)	(0.00901)	(0.0106)	(0.0103)
VIX uncertainty	0.00770	0.00758	0.00689	0.00358	0.00891	0.00770	0.00642
	(0.00617)	(0.00623)	(0.00691)	(0.00686)	(0.00608)	(0.00617)	(0.00578)
Short-term yield volatility	0.249	0.254	0.468	0.699**	0.223	0.249	0.243
	(0.240)	(0.239)	(0.289)	(0.346)	(0.235)	(0.240)	(0.233)
Observations	1,676	1,676	1,609	1,118	1,676	1,676	1,676
Adjusted R^2	0.229	0.230	0.237	0.202	0.157	0.229	0.222
Type	Meeting	Meeting	Meeting	Meeting	Meeting	Meeting	Meeting
Sample	2000m1–2011m12	2000m1–2011m12	2000m1–2011m12	2000m1–2011m12	2000m1–2011m12	2000m1–2011m12	2000m1–2011m12
Speaker	ChairGovPres	ChairGovPres	ChairGovPres	ChairGovPres	ChairGovPres	ChairGovPres	ChairGovPres
$\hat{\pi^*}$	1.55	1.56	1.55	1.79	1.55	1.55	1.57
SE	(0.13)	(0.12)	(0.24)	(0.2)	(0.14)	(0.13)	(0.12)
95% C.I.	[1.3–1.8]	[1.32–1.8]	[1.08–2.02]	[1.4–2.18]	[1.28–1.82]	[1.3–1.8]	[1.33–1.81]

Notes: The variables y and y* denote log GDP and log potential GDP, respectively. Standard errors (in parentheses) are robust to heteroskedasticity and two-way clustering by sample-month and speaker. Regressions include speaker fixed effects. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Downloaded from https://academic.oup.com/restud/advance-article/doi/10.1093/restud/rdab094/6484651 by Federal Reserve Bank of San Francisco Research Library user on 25 Ja

(7) of Table 4 of the paper (labelled "Adjusted LM"). The results are almost identical to the baseline.[44]

4.5. *Alternative sentiment measures*

Our baseline results are based on measuring sentiment using the difference between the fraction of negative words and the fraction of positive words (net negativity). In this subsection, we provide results based on alternative, independently constructed sentiment measures. First, we estimate the FOMC loss function using the negative and positive fractions separately. Second, we generate a net negativity measure using a popular open-source python tool called Valence Aware Dictionary and Sentiment Reasoner (VADER) (Hutto and Gilbert, 2014). VADER is also a lexical sentiment classifier, based on a combination of available general-english and social media lexicons, but it also applies a number of heuristic rules to take account of certain contextual characteristics of how each word is used in a given sentence. The rules relate to negation (e.g. a word being preceded by "not" reverses its valence), punctuation (e.g. exclamation points), capitalization, degree modification (e.g. being preceded by "very" or "slightly"), and coming before or after the word "but." Such rules can improve the accuracy of a sentiment classifier. However, an important downside of VADER for our application is that the sentiment scores associated with its lexicon and rules are constructed from social media text/language which often will have very different word meanings than those in the economics and finance domain (as shown in Shapiro *et al.* (2018)).

The results of using these alternative measures are shown in Supplementary Appendix Table G.1 and are remarkably consistent across all of the sentiment measures. In particular, like the baseline measure, all three of the alternative measures point to an inflation target—that is, the inflation rate minimizing negativity and maximizing positivity—of approximately $1\frac{1}{2}$. With regards to the other potential factors in the loss function, the results are generally qualitatively similar across the measures, though the statistical significance varies. One notable difference is that the financial variables appear to affect negativity more strongly than positivity.

5. CORROBORATING EVIDENCE ON THE INFLATION TARGET

5.1. *Narrative analysis*

The analysis in the prior two sections estimated, among other things, the implicit inflation target most consistent with the FOMC's preferences, as revealed by the tone of their discussions. Here, we complement that analysis with a narrative approach, aiming to identify explicit statements of a preferred inflation target. Specifically, we perform an automated regular expression (regex) search over all FOMC meeting transcripts (including conference calls) from 1986 through 2013 for terms or phrases related to inflation objectives.[45] We then reviewed and analysed the sections of text corresponding to those matches.

Discussion of explicit inflation targeting did not begin to appear in the FOMC meetings until around 1994. Prior to then, FOMC participants occasionally mentioned an objective of "price

44. We report cross-correlations for these negativity measures in Supplementary Appendix Table G.13.

45. We performed two separate searches and collected the union of the two results. The first search looked for text containing any of the following strings: inflation target | inflation objective | long-run inflation | long run inflation | longer-run inflation | longer run inflation | inflation goal | objective for inflation | target for inflation | goal for inflation | PCE objective | PCE goal | PCE target | target for PCE inflation | goal for PCE inflation | objective for PCE inflation | objective for core | goal for core | target for core. The second was a regex search that looked for lines of text containing: (inflation | PCE) AND (percent) AND (objective|target|long-run|longer-run|long run|longer run), where "|" represents "or."

Downloaded from https://academic.oup.com/restud/advance-article/doi/10.1093/restud/rdab094/6484651 by Federal Reserve Bank of San Francisco Research Library user on 25 Ja

stability" but without relating that to an explicit inflation target. One illuminating exception we found was a statement made by Federal Reserve Board Governor David Mullins in the November 1993 meeting that suggests an implicit target moving below 3% by that time: "I think there's a real payoff not just from stabilizing inflation in the 3 to 4% range but in moving lower".

The first serious consideration of an explicit inflation target, that we could find, comes from St. Louis Fed President Thomas Meltzer in 1994. At the July meeting that year, he said "If we don't make an explicit statement in this FOMC testimony with respect to our long-run expectations on inflation that goes beyond 'we think price stability is good,' and get more specific in terms of a target range, then at the very least I think we have to make it clear that we consider 3% inflation to be unacceptable...." A few meetings later, in November 1994, he stated: "...I feel that it may be time for us to consider setting a specific inflation target that looks out into the future. I think, and this point was made as well, that it could make our job considerably easier in circumstances like the present—with upward cyclical inflationary pressures—if people were willing to look out to a longer-range target and that added to credibility."

There was much debate over the following couple years about whether the FOMC should have an inflation target, even internally. For instance, Atlanta Fed President Robert Forrestal during the Jan. 31/Feb. 1 1995 meeting said "...I would be against an inflation target and I would associate myself entirely with the views of Governor [Janet] Yellen" (who had noted some potential risks to having a target) while other participants expressed being in favour of a target. The question remained unsettled into the late 1990s, as exemplified by this statement from President Melzer at the November 1997 meeting: "What are the FOMC's intentions? Do we like seeing inflation below 2%? Does the public know it? I think, as I have said before, that we ought to be more explicit about our longer-term objective. In that event, it would be much less likely that our actions would be misinterpreted as being anti-jobs or anti-growth."

Starting in the early 2000s, however, explicit statements by FOMC participants of their inflation preferences became much more common. Indeed, while the automated regular-expression search described above found only 50 matches over 1986–99—and only 2 of those mentioned a specific inflation target—the search yielded several hundred matches over the 2000–13 period. Reading through the sections of text corresponding to those matches, we identified and tabulated each instance in which a participant expressed their own preference for a specific numerical inflation target (or range such as between 1.5 and 2). Figure 7 plots the results of this narrative analysis. Each circle represents a stated preference, from the indicated speaker (y-axis) at the indicated FOMC meeting (x-axis), for a specific inflation target. The value of the preferred target is shown within the circle and also indicated by the circle's colour. Values go from 1.0 (light blue) to 1.5 (dark blue) to 1.75 (purple) to 2.0 (red).

In total, we identified 58 statements of an inflation target preference. 20 of these expressed a preference for $1\frac{1}{2}$%; 9 for either exactly $1\frac{3}{4}$ or between $1\frac{1}{2}$ and 2%; 27 for 2%; and just two for 1%. However, these aggregate counts mask a clear and important pattern over time. Statements in favour of a 2% target were non-existent before 2007 and did not become widespread until 2009, late in the Great Recession. In fact, a single participant, Governor Frederic Mishkin accounted for all but one of the pre-2009 2% statements.[46] Prior to the Great Recession, $1\frac{1}{2}$% was the overwhelming favourite among participants.

Although the consensus target preference was $1\frac{1}{2}$% for most of our baseline sample period this analysis also reveals that some participants shifted their preference up to 2% during the Great Recession. Specifically. the transition of the consensus favourite over the course of the

46. Governor Mishkin alluded to the fact that most other participants preferred $1\frac{1}{2}$% at the October 2007 meeting: "People know that I am a 2% kind of guy, and I know good people here who are $1\frac{1}{2}$% kinds of guys."

Downloaded from https://academic.oup.com/restud/advance-article/doi/10.1093/restud/rdab094/6484651 by Federal Reserve Bank of San Francisco Research Library user on 25 Ja

TABLE 5
Sample splits based on narrative evidence FOMC meetings: 2000–11

	(1) Baseline	(2) Narrative 1.5	(3) Narrative 2.0	(4) Other
$\lvert\pi_t-\pi^*\rvert$, GB	0.490***	0.521**	0.753***	0.399***
	(0.132)	(0.219)	(0.208)	(0.154)
$(u_t-u_t^*)$, GB	0.0234	0.117	-0.326**	0.0462
	(0.0763)	(0.115)	(0.162)	(0.0959)
$(u_t-u_t^*)^2$, GB	0.00313	-0.0105	0.0522**	0.00169
	(0.0127)	(0.0172)	(0.0233)	(0.0201)
Δy_t , GB	-0.0322	-0.0561	-0.0322	-0.0467*
	(0.0265)	(0.0416)	(0.0466)	(0.0280)
Δy_t^2 , GB	-0.0331***	-0.0311***	-0.0669***	-0.0265***
	(0.00597)	(0.00755)	(0.0146)	(0.00666)
S&P Return, 3m	-2.101***	-1.173	-4.666***	-2.271***
	(0.607)	(0.830)	(1.038)	(0.681)
Shiller PE ratio	0.0207*	0.0241	0.0426	0.0244*
	(0.0106)	(0.0165)	(0.0475)	(0.0132)
VIX uncertainty	0.00770	0.0102	-0.00321	0.00546
	(0.00617)	(0.00883)	(0.0105)	(0.00743)
Short-term yield volatility	0.249	0.465	0.713*	-0.194
	(0.240)	(0.303)	(0.411)	(0.285)
Observations	1676	587	240	849
Adjusted R^2	0.229	0.286	0.310	0.175
Type	Meeting	Meeting	Meeting	Meeting
Sample	2000m1–2011m12	2000m1–2011m12	2000m1–2011m12	2000m1–2011m12
$\hat{\pi}^*$	1.55	1.37	2.25	1.55
SE	(0.13)	(0.2)	(0.37)	(0.13)
95% C.I.	[1.3–1.8]	[0.98–1.76]	[1.52–2.98]	[1.3–1.8]

Notes: The variables y and y* denote log GDP and log potential GDP, respectively. Standard errors (in parentheses) are robust to heteroskedasticity and two-way clustering by sample-month and speaker. Regressions include speaker fixed effects. * $p<0.10$, ** $p<0.05$, *** $p<0.01$.

Great Recession is book-ended by the 20–21 March 2007 regular meeting and a conference call meeting on 16 January 2009. In the March 2007 meeting, FOMC participants were explicitly asked by the Chair if they thought the Fed should have an explicit numerical inflation target and, if so, what that target should be. Four members expressed a preference for $1\frac{1}{2}\%$, three for between $1\frac{1}{2}$ and 2%, two for 1%, and just one (Mishkin) for 2%. In the January 2009 meeting, seven out of the ten members that voiced a specific preference desired 2%. In fact, most of those seven had themselves previously expressed a lower preferred target.

We next use this narrative evidence as a way to validate our sentiment-based approach of estimating the inflation target. Here we split our speaker-meeting observations into three subsamples: (1) observations in which the speaker's most recent explicit statement of an inflation target preference was for 2.0, (2) observations in which the speaker's most recent explicit statement of an inflation target was for below 2.0 (typically 1.5), and observations from speakers who do not explicitly state a preference.[47] The results are shown in Table 5. We estimate an inflation target of 2.25 for doves ("narrative 2.0"), 1.36 for hawks ("narrative 1.5"), and in between at 1.55 for others. These results not only provide some external validation of the sentiment-based approach of measuring the inflation target, but also suggest that the majority of participants preferred an

47. Specifically, the "narrative 2.0" subsample consists of Cumming, Evans, Lockhart, Mishkin, Rosengren, and Kohn for all meetings, and Yellen and Pianalto for January 2009 onwards. The "narrative 1.5" sub-sample consists of Bullard, Plosser, Fisher, Lacker, Geithner, Poole, Meyer, Moskow, Broaddus, and Hoenig for all meetings, and Yellen and Pianalto prior to January 2009.

Downloaded from https://academic.oup.com/restud/advance-article/doi/10.1093/restud/rdab094/6484651 by Federal Reserve Bank of San Francisco Research Library user on 25 Ja

inflation target of $1\frac{1}{2}\%$ over the 2000–11 period, despite some FOMC participants shifting to a preference for 2% towards the end of the sample period. Of course, in January 2012, the 2% explicit target was codified officially by the FOMC in their public announcement of the target. One plausible explanation for this shift was that, at this time, nominal interest rates had hit the zero lower bound, hampering the Fed's ability to raise inflation. In such a low interest rate environment, raising the inflation target can provide a hedge against hitting the zero lower bound in future episodes (see Ball, 2014; Andrade, Galí, Bihan and Matheron, 2018).

5.2. *Heterogeneous preferences*

An appealing feature of the sentiment analysis approach to studying central bank preferences is that it allows one to investigate heterogeneity in the loss function across different types of FOMC participants. In this subsection, we first investigate heterogeneity in preferences of Fed governors (including the chair) vs. regional presidents and of FOMC voters vs. non-voters. We then consider heterogeneity across hawks and doves, using several different hawk/dove measures. This latter analysis also serves to provide additional external validation of our key result that the FOMC had an implicit inflation target of 1.5% on average over 2000–11.

Columns (1)–(4) of Table 6 report results where observations are split across particular types of speakers (defined by their status at the time of their remarks): (1) FOMC voting participants, (2) FOMC non-voting participants, (3) Governors, and (4) Regional Fed Presidents. Comparing voting and non-voting participants (columns (1) and (2), respectively), we find that voters have a somewhat higher implicit inflation target than non-voters, 1.75 vs. 1.51, though the difference is not statistically significant. Comparing the results for governors (including the chair) vs. regional Fed presidents, in columns (3) and (4), we find governors had a target close to 2 while presidents had a target of about 1.5. Moreover, the inflation target estimated from presidents' negativity is more precisely estimated than that from governors.[48] It also appears that non-voters and regional presidents have stronger preferences in favour of higher output growth.

Next, we investigate how the implicit inflation target differs between hawks and doves on the FOMC based on external information signalling speakers' hawkishness/dovishness. We consider three different measures of hawkishness/dovishness. First, we split speaker-meeting observations into hawkish and dovish following the lexical hawkish/dovish sentiment scoring approach of Malmendier, Nagel and Yan (2020) and Apel and Blix Grimadi (2014) (see Supplementary Appendix F for details). We then code a speaker-meeting observation as hawkish (dovish) if its text contains more such hawkish (dovish) phrases than dovish (hawkish) phrases. Ties are considered neutral and dropped for the purposes of this exercise. Finally, we estimate our baseline loss function regression separately for hawkish and dovish observations. The results are shown in columns (5) and (6) of Table 6. Consistent with the notion that hawks have a lower inflation target (i.e. are less tolerant of higher inflation), we estimate an inflation target of 1.45 for the hawkish sub-sample and 1.86 for the dovish sub-sample.

Second, and closely related, we classify speakers as either hawkish or dovish—restricting these classifications to be fixed over time—based on whether they had more or less hawkish phrases than dovish phrases over all of their utterances in the transcripts. We then estimate the baseline loss function regressions separately for hawkish and dovish speakers. The results are shown in columns (7) and (8) of Table 6. We find an inflation target of 1.35 for hawkish speakers and 1.73 for dovish speakers, consistent with hawkish speakers having a lower target.

48. We also tried estimating a separate loss function for FOMC chairs, however we found there were too few observations to obtain precise estimates.

Downloaded from https://academic.oup.com/restud/advance-article/doi/10.1093/restud/rdab094/6484651 by Federal Reserve Bank of San Francisco Research Library user on 25 Ja

TABLE 6
Allowing for heterogeneity across speaker types FOMC meetings: 2000–11

	(1) Voters	(2) Non-Voters	(3) Governors	(4) Presidents	(5) Hawkish Observations	(6) Dovish Observations	(7) Hawkish Speakers	(8) Dovish Speakers	(9) Hawkish Dissenters	(10) Dovish Dissenters
$\lvert \pi_t - \pi^* \rvert$, GB	0.405***	0.714***	0.283*	0.652***	0.488***	0.656***	0.402**	0.634***	0.668**	0.666***
	(0.146)	(0.175)	(0.161)	(0.131)	(0.144)	(0.150)	(0.168)	(0.171)	(0.324)	(0.198)
$(u_t - u_t^*)$, GB	0.00552	0.0239	0.0725	−0.000618	0.442***	−0.332***	0.0960	−0.101	0.181	−0.304***
	(0.0996)	(0.0877)	(0.0773)	(0.0875)	(0.0987)	(0.108)	(0.0935)	(0.107)	(0.118)	(0.117)
$(u_t - u_t^*)^2$, GB	0.00365	0.00516	−0.00919	0.00809	−0.0792***	0.0484***	−0.00537	0.0207	−0.0129	0.0221
	(0.0193)	(0.0160)	(0.0156)	(0.0132)	(0.0203)	(0.0185)	(0.0169)	(0.0164)	(0.0137)	(0.0216)
Δy_t , GB	−0.00357	−0.0744**	0.0220	−0.0541**	−0.112***	0.0619**	−0.0620**	0.000516	−0.107***	−0.000340
	(0.0307)	(0.0329)	(0.0385)	(0.0257)	(0.0422)	(0.0292)	(0.0297)	(0.0301)	(0.0347)	(0.0457)
Δy_t^2 , GB	−0.0335***	−0.0375***	−0.0295***	−0.0366***	−0.0312***	−0.0401***	−0.0294***	−0.0394***	−0.0149*	−0.0526***
	(0.00626)	(0.00941)	(0.00901)	(0.00664)	(0.00893)	(0.00646)	(0.00689)	(0.00815)	(0.00801)	(0.00989)
S&P return, 3m	−2.760***	−1.424**	−3.277***	−1.668***	−1.622*	−2.835***	−1.695**	−2.669***	−1.569**	−2.400***
	(0.713)	(0.631)	(0.793)	(0.618)	(0.881)	(0.551)	(0.715)	(0.685)	(0.681)	(0.207)
Shiller PE ratio	0.0186	0.0253**	0.0271**	0.0199	0.0732***	−0.00682	0.0223*	0.0194	0.0309**	−0.0288*
	(0.0129)	(0.0111)	(0.0113)	(0.0122)	(0.0105)	(0.0141)	(0.0117)	(0.0160)	(0.0138)	(0.0164)
VIX uncertainty	0.00613	0.00963	0.000401	0.0108	0.00683	0.0147**	0.00437	0.0116	−0.00208	−0.00712
	(0.00755)	(0.00592)	(0.00716)	(0.00699)	(0.00916)	(0.00648)	(0.00854)	(0.00732)	(0.00691)	(0.00823)
Short-term yield volatility	0.0813	0.448**	−0.00268	0.322	−0.504	0.309	−0.127	0.548**	−0.0193	−0.380
	(0.276)	(0.227)	(0.348)	(0.249)	(0.329)	(0.355)	(0.286)	(0.276)	(0.274)	(0.435)
Observations	1015	661	542	1134	402	408	817	819	348	257
Adjusted R^2	0.164	0.320	0.107	0.304	0.242	0.202	0.221	0.246	0.254	0.284
Type	Meeting	Meeting	Meeting	Meeting	Meeting	Meeting	Meeting	Meeting	Meeting	Meeting
Sample	2000m1–2011m12	2000m1–2011m12	2000m1–2011m12	2000m1–2011m12	2000m1–2011m12	2000m1–2011m12	2000m1–2011m12	2000m1–2011m12	2000m1–2011m12	2000m1–2011m12
$\hat{\pi}^*$	1.75	1.51	2.05	1.54	1.45	1.86	1.35	1.73	1.3	1.52
SE	(0.26)	(0.14)	(0.42)	(0.16)	(0.25)	(0.24)	(0.18)	(0.17)	(0.26)	(0.25)
95% C.I.	[1.24–2.26]	[1.24–1.78]	[1.23–2.87]	[1.23–1.85]	[0.96–1.94]	[1.39–2.33]	[1–1.7]	[1.4–2.06]	[0.79–1.81]	[1.03–2.01]

Notes: The variables y and y* denote log GDP and log potential GDP, respectively. Standard errors (in parentheses) are robust to heteroskedasticity and two-way clustering by sample-month and speaker. Regressions include speaker fixed effects. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Downloaded from https://academic.oup.com/restud/advance-article/doi/10.1093/restud/rdab094/6484651 by Federal Reserve Bank of San Francisco Research Library user on 25 Ja

Third, we classify speakers as hawkish or dovish based on observed dissents to FOMC interest rate decisions (i.e. statements) by participants. We used the data set of FOMC dissents from Thornton, Wheelock, *et al.* (2014) which, in addition to listing which FOMC participants dissented in each meeting, notes whether the dissent was in a hawkish or dovish direction (i.e. favouring tighter or looser interest rate policy than the rest of the committee). Of course, many speakers/participants never dissented over our sample period; these speakers' observations are dropped for this exercise. We estimate our baseline regression separately for hawkish and dovish dissenters. The results are shown in columns (9) and (10) of Table 6. Based on dissents, we again find that hawkish participants preferred a lower inflation target than dovish participants, though the results here are less precise (likely due to the smaller number of observations because of the rarity of dissents) and the difference in the inflation target point estimates is smaller than in the prior analysis of hawks vs. doves. These three separate analyses, in addition to the narrative analysis, not only reveal interesting heterogeneity in preferences across speaker types, but also provide strong evidence of external validity for our general sentiment analysis approach.[49]

5.3. *Taylor rule estimations of* π^*

As another way to corroborate the text-based estimate of π^*, we assess whether it is consistent with the FOMC's interest rate setting behaviour. Specifically, we compare how well a simple Taylor rule fits with the 1.5 estimate vs. two alternative measures of π^*—the commonly held view of 2% and the perceived target variable (PTR) from the FRB/US model used in many macroeconomic analyses. We estimate the following Taylor rule :

$$ffr_t = \pi_t + r_t^* + \alpha(\pi_t - \pi_t^*) + \beta(u_t - u_t^*) + \varepsilon_t \tag{7}$$

under three different measures of π^*: a constant equal to 1.5, a constant equal to 2, and PTR (which is time varying). The variable ffr_t is the nominal federal funds rate, π_t is the average of the Greenbook nowcast and 1-quarter ahead measure of Core PCE inflation, and $(u_t - u_t^*)$ is the analogous measure of the Greenbook unemployment gap. Similar to Coibion and Gorodnichenko (2011) and Coibion *et al.* (2012), we estimate the specification using OLS.[50] We use two alternative measures of the natural rate of interest (r_t^*): Christensen and Rudebusch (2019) and Lubik and Matthes (2015).[51]

The results for the 2000–08 period are shown in columns (1) through (6) of Table 7.[52] We do not include the zero-lower-bound (ZLB) period because there is no time variation in the federal funds rate. The 1.5 measure of π^* performs better in terms of R^2 across all three measures of π^* in terms of R^2 as well as well the expected sign on the coefficients. Specifically, coefficients using $\pi^* = 1.5$ appear to resemble the "balanced-approach" rule of Taylor (1999) , where $\alpha = 0.5$ and

49. As far as the coefficients on the other variables are concerned, there is no clear consistent pattern regarding the coefficients on the unemployment gap and output growth. However, it does appear that doves have a stronger preference in favour of higher stock market returns while hawks have a stronger dislike of higher stock P/E ratios. These results may suggest that hawks have greater concern about potential overheating in financial markets.

50. There are studies showing that the endogeneity bias associated with estimating Taylor rules using OLS is likely small (see e.g. de Vries and Li, 2013; Carvalho, Nechio and Tristao, 2019).

51. Both measures are advantageous because they are "macroeconomic model-free"—meaning they are less prone to bias from any mis-specification of output and inflation dynamics. Christensen and Rudebusch (2019) is based on a dynamic term-structure model and Lubik and Matthes (2015) is based on time-varying parameter vector auto-regressive model.

52. In Section 6.2, we test the reaction function using a time-varying π_t^* generated from FOMC transcripts going back to 1986 shown in columns (7) and (8) of Table 7.

Downloaded from https://academic.oup.com/restud/advance-article/doi/10.1093/restud/rdab094/6484651 by Federal Reserve Bank of San Francisco Research Library user on 25 Ja

TABLE 7
Estimated Taylor rules under alternative measures of π^*

	(1) π^*=1.5	(2) π^*=1.5	(3) π^*=2	(4) π^*=2	(5) π^*=PTR	(6) π^*=PTR	(7) π^*=SW	(8) π^*=PTR
$(u_t - u_t^*)$, GB	−1.716***	−1.339***	−1.837***	−1.404***	−1.919***	−1.505***	−1.318***	−1.244***
	(0.0979)	(0.112)	(0.108)	(0.149)	(0.116)	(0.162)	(0.0668)	(0.0685)
$(\pi_t - \pi_t^*)$, GB	0.558***	0.798***	−0.166	0.110	−0.358**	−0.179	0.799***	−0.603***
	(0.145)	(0.118)	(0.157)	(0.177)	(0.153)	(0.184)	(0.0525)	(0.0767)
Observations	108	108	108	108	108	108	254	254
Adjusted R^2	0.786	0.738	0.754	0.647	0.760	0.649	0.708	0.471
r^*	CR	LM	CR	LM	CR	LM	LM	LM
Period	2000m1–2008m12	2000m1–2008m12	2000m1–2008m12	2000m1–2008m12	2000m1–2008m12	2000m1–2008m12	1986m1–2008m12	1986m1–2008m12

Notes: The dependent variable is ffr_t - r^* - π_t, where ffr_t is the nominal federal funds rate. Columns (1) and (2) are estimated with π^* equal to 1.5, columns (3) and (4) with π^* equal to 2, and columns (5) and (6) with π^* equal to the FRB/US model variable PTR. Columns (1), (3), and (5) are estimated with r* taken from Christensen and Rudebusch (2019) (CR), and columns (2), (4), and (6) are estimated with r* taken from Lubik and Matthes (2015) (LM) * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Downloaded from https://academic.oup.com/restud/advance-article/doi/10.1093/restud/rdab094/6484651 by Federal Reserve Bank of San Francisco Research Library user on 25 Ja

$\beta = -2$. Not only do the other two measures of π^* have a worse fit, but they produce measures of α (the coefficient on inflation relative to its target) that are either insignificant or the wrong sign. This analysis gives evidence that the FOMC tended to raise (lower) interest rates when inflation was above (below) 1.5%, as opposed to the commonly held view of doing so around 2% or the PTR variable used in many macroeconomic models.

5.4. *FOMC meeting minutes*

A fourth piece of corroborating evidence on our baseline loss function estimates comes from applying our sentiment analysis approach to the FOMC meeting minutes. Three weeks after each FOMC meeting, the Federal Reserve releases a summary, known as the minutes, of the discussion at the meeting. The minutes are a general overview of the topics discussed and opinions expressed at the meeting, with no attribution of any comments or opinions to individual FOMC members. While the tone of the minutes is much less emotive than that of the actual spoken remarks made at the meeting, it is nonetheless interesting to see whether the expressed negativity in the minutes relates to inflation, slack, and output growth in a similar fashion to that in the transcripts.

We therefore repeat the same loss function estimations done in subsections 4.1 and 4.2 above but using the net negativity from the minutes.[53] Though minutes are available for meetings up to the present, we restrict the sample period here to 2000–11 to facilitate comparison to the transcripts' results above. We discuss results based on post-2011 minutes in the next section.

The results are shown in columns (1) and (2) of Table 8.[54] Whether controls are omitted (column 1) or included (column 2), we obtain an estimated $\hat{\pi}^*$ close to 1.5. Indeed, in our preferred specification (column 2), we get exactly the same $\hat{\pi}^*$, at 1.55, using the minutes as we did using the transcripts, though the standard error is larger. Thus, the language used in the public minutes of the FOMC meetings appears to point to a similar inflation target as that implied by the verbatim language actually used the meetings. We also find output growth has an even stronger effect on negativity in the minutes than in the transcripts, while stock market returns have a smaller (and statistically insignificant) effect.

6. THE INFLATION TARGET IN OTHER PERIODS

Our baseline results in Section 4 are based on the January 2000–December 2011 sample period. January 2000 is when the Greenbook first began routinely providing PCE inflation nowcasts and forecasts. December 2011 is the last month before the FOMC officially announced an inflation target. In January 2012, the FOMC announced an intended target of 2.0% based on the PCE inflation measure. In this section, we apply our sentiment analysis approach to estimating the implicit inflation target before and after the baseline 2000-2011 period.

6.1. *The post-2011 inflation target*

We begin by considering the post-2011 period. As of the time of this writing, the latest available FOMC transcripts and Greenbook forecasts are for December 2015 (due to the Federal Reserve's

53. We also estimate the loss function using sentiment from the public speeches of FOMC participants. The results are shown in Supplementary Appendix Table G.2. The estimated inflation target based on the speeches is higher (a little above 2) than that based on the FOMC meetings (either transcripts or minutes). A potential explanation is that policymakers may express less concern about inflation in public than in private.

54. We follow the same non-linear gridsearch estimation procedure described and used in subsection 4.2. Recall the procedure searches over potential inflation targets from 0.00 to 3.00, with step sizes of 0.01, for the estimated loss function yielding the minimum sum of squared residuals.

Downloaded from https://academic.oup.com/restud/advance-article/doi/10.1093/restud/rdab094/6484651 by Federal Reserve Bank of San Francisco Research Library user on 25 Ja

TABLE 8
Estimated loss function based on FOMC minutes

	(1) 2000-2011	(2) 2000-2011	(3) 2012-2019	(4) 2012-2019
$\lvert\pi_t-\pi^*\rvert$	0.716***	0.428**	0.204	0.301*
	(0.240)	(0.180)	(0.165)	(0.174)
$(u_t-u_t^*)$		0.161		−0.485***
		(0.109)		(0.136)
$(u_t-u_t^*)^2$		−0.0228		0.0617
		(0.0186)		(0.0403)
Δy_t		−0.198***		−0.276
		(0.0317)		(0.478)
Δy_t^2		−0.0335***		0.0239
		(0.00594)		(0.0847)
S&P return, 3m		−1.043		−0.399
		(0.718)		(1.005)
Shiller PE ratio		0.0348**		−0.130***
		(0.0133)		(0.0327)
VIX uncertainty		0.0112		−0.00505
		(0.00760)		(0.0123)
Short-term yield volatility		0.0381		0.157
		(0.402)		(0.607)
Constant	0.629***	0.300	0.436***	4.576***
	(0.122)	(0.407)	(0.0796)	(0.788)
Observations	95	95	64	64
Adjusted R^2	0.088	0.765	0.005	0.354
Type	Minutes	Minutes	Minutes	Minutes
Sample	2000m1–2011m12	2000m1–2011m12	2012m1–2019m12	2012m1–2019m12
$\hat{\pi^*}$	1.43	1.55	1.97	2.01
SE	(0.27)	(0.31)	(0.38)	(0.27)
95% C.I.	[0.95–1.55]	[1.2–2.15]	[1.42–2.58]	[1.43–2.43]

Notes: Because Greenbook forecasts are unavailable past 2015m12 we use IHS Markit/Macroeconomic Advisers forecasts for 2016m1 onward. Standard errors (in parentheses) are robust to heteroskedasticity. * $p<0.10$, ** $p<0.05$, *** $p<0.01$.

5-year publication lag for these historical documents). In principal, one could use this short 2012–15 period to estimate the loss function and test whether the inflation target shifted up to 2.0. However, aside from concerns related to relying on such a short time series dimension, a major complication comes from the fact that both actual inflation and Greenbook forecasts of inflation were well below 2 throughout the 2012–15 period. In particular, our measure of Greenbook inflation (average of current quarter and one-quarter ahead forecasts) fluctuated within a narrow range from 1.05 to 1.75 (averaging a little under 1.5) during these 4 years.[55] Any revealed preference approach, such as our sentiment analysis approach or a Taylor Rule approach, is unable to estimate a preferred inflation rate if that rate is outside of the range of observed values in the data. Thus, to examine the post-2011 period, we turn to the FOMC minutes, which are available up to the present. We demonstrated above that the minutes yield very similar results to those based on the transcripts.

Using the minutes' net negativity, we estimate our baseline (symmetric V) loss function for the period from January 2012 through December 2019. We end the sample with 2019 to avoid the large volatility in economic and financial variables associated with the COVID-19 pandemic, though the results (including $\hat{\pi}^*$) are very similar if we include data through 2020. Because Greenbook forecasts are not available through 2015 (as of the time of this writing), we fill in the time series data for those variables (inflation, unemployment gap, and output growth) using

55. See Supplementary Appendix Figure G.3.

Downloaded from https://academic.oup.com/restud/advance-article/doi/10.1093/restud/rdab094/6484651 by Federal Reserve Bank of San Francisco Research Library user on 25 Ja

forecasts from IHS Markit/Macroeconomic Advisers (MA).[56] As we show in Supplementary Appendix Figure G.3, these near-term forecasts from MA track those from the Greenbook very closely.

The results are shown in columns (3) and (4) of Table 8. We find that the estimated inflation target did indeed shift up to 2.0 in the post-2011 period, consistent with the FOMC's establishment of an official 2% target in January 2012.[57] We also find some evidence that the Fed's preferences regarding slack and output growth changed somewhat in the later period. In the 2012–19 period, output growth no longer has a significant effect on net negativity, while the unemployment gap has a significant quadratic effect. These results based on minutes should be viewed with caution until corroborated in the future with transcripts data. That said, such a shift in preferences away from growth and toward slack would be consistent with Thornton (2011)'s observation, noted in the introduction, that the FOMC's policy directive since 1991 explicitly mentioned promoting sustainable "growth in output," but that this was dropped in 2009.

6.2. *The pre-2000 inflation target*

Our baseline analysis is run during a period when inflation was relatively low and steady. Prior to the 1990s, however, inflation was higher and less stable. Indeed, numerous studies have shown that both trend inflation and inflation persistence dramatically declined (see Stock and Watson, 2007, 2016) from the high inflation period of the 1970s and 80s to the late 1990s and 2000s. Other studies (e.g. Ireland, 2007) have more explicitly linked this decline in trend inflation to a decline in the Federal Reserve's target. Committee members also noted that the implicit inflation target was likely declining over this period. For example in the January 2003 FOMC meeting, Governor Donald Kohn stated "...I do believe that, from 1987 at least into the second part of the 1990s, the Committee surely did not have a constant inflation target." Indeed, a close reading of the transcripts from the mid-1980s and the 1990s suggests that the FOMC was far from agreement on whether they should even have an explicit inflation target, much less what that target should be. An analysis that includes this period should therefore allow for the possibility that the implicit target—that is, the target consistent with the central bank's preferences—was time-varying and different than it was over the 2000–11 period. Thus, we relax the assumption that π^* is a constant, and estimate the FOMC's symmetric loss function allowing π_t^* to be time varying.

For tractability, we make two assumptions about the path of π_t^*. First, π_t^* is assumed to fit a constant over a "constant-target period"—defined as a time period between s and 2011m12. Second, π_t^* is assumed to fit a quadratic polynomial function of time during a "time-varying-target period"—defined as a period between 1986m2 and s.[58] We imputed the core PCE forecasts back to 1986m1 by splicing adjusted Greenbook forecasts of core CPI—adjusted using the average difference between core CPI and core PCE over pre-2000 period. We let the regime change, s, float freely which allows for the possibility that π_t^* is time-varying or constant for the entire sample period, 1986m1-2011m12.

We perform a grid search over four parameters that govern the path of π_t^*: a constant in the constant-target period, a slope and curvature in the time-varying-target period, and a date, s, of the regime switch between the time-varying and constant-target period. For aid in identification, we estimate π_t^* assuming a symmetric V-shaped loss function, however, we allow the slope of the

56. As we did with the Greenbook data, we measure these variables using an average of current quarter and one-quarter ahead forecasts.

57. Results based on the asymmetric V specification are very similar and are provided in Supplementary Appendix Table **??**.

58. 1986m2 is the first month when core CPI inflation forecasts were published in the Greenbook.

Downloaded from https://academic.oup.com/restud/advance-article/doi/10.1093/restud/rdab094/6484651 by Federal Reserve Bank of San Francisco Research Library user on 25 Ja

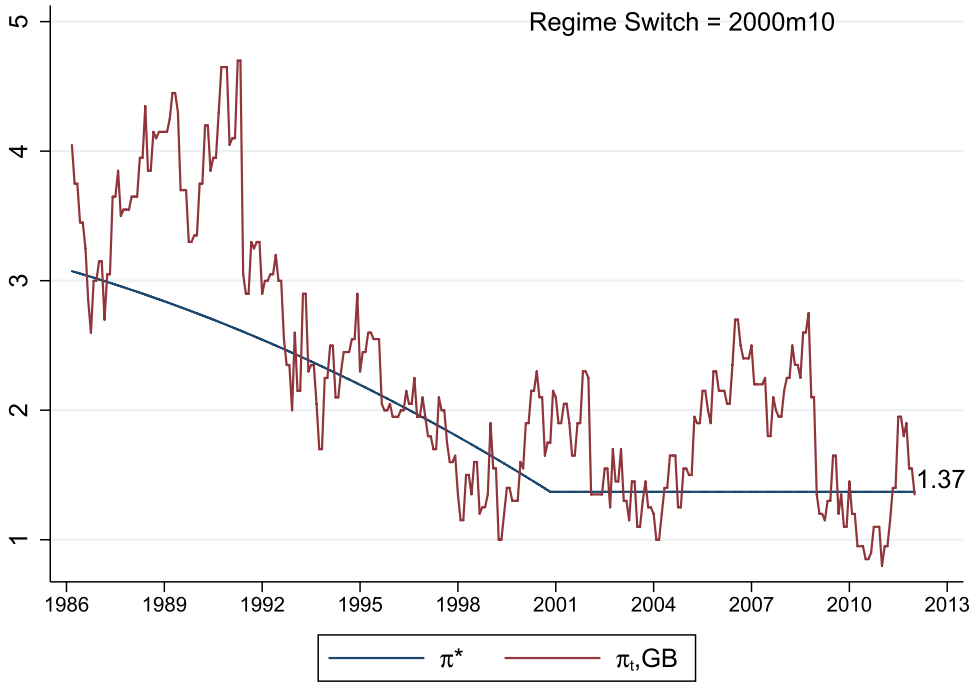

FIGURE 5
Estimated time-varying π^* FOMC meetings: 1986–2011 no additional factors

V-shape to differ in the time-varying- and constant-target periods. Technical details are available in Supplementary Appendix C.

The results in Figures 5 (without controls) and 6 (with controls)[59] show that the FOMC's implied inflation target did in fact decline over time. Specifically, π_t^* is around 2 % in 1986m1 and declines to approximately $1\frac{1}{2}$ beginning in 2000, where it remains until the end of the sample period (2011m12). This decline is consistent with Ireland (2007), Cogley *et al.* (2010), and Coibion and Gorodnichenko (2011), which also found declines in π^* over a similar time period, despite the very different approach.[60] While the level of π_t^* declined, the loss function became steeper (results not depicted in figure).[61] That is, the FOMC's language showed a stronger dislike of inflation deviations in the later period when inflation was in fact more stable and less persistent.

As an external validity check of this time-varying of measure of π_t^*, we use it in estimating the Federal Reserve's interest rate reaction function, equation (7) discussed in Section 5.3. The results of estimating this interest rate rule over the 1986–2008 period using this measure of π_t^* are shown in column (7) of Table 7.[62] The results are quite sensible, showing coefficients similar

59. Controls include the full set included in column (2) of Table 2. We also include the difference between GB headline and core inflation to control for energy price fluctuations, which were volatile during this period.

60. The finding that the FOMC's preferred inflation target fell from the mid-1980s to around 2000 contrasts somewhat with Ball and Mazumder (2018) who argue that the inflation target itself was flat during that period but that the public's inflation expectations shifted from being backward-looking to anchored at the Fed's target.

61. The coefficient increased from 0.2 in the time-varying period to 0.45 in the stable period.

62. We use the measure of π_t^* constructed using controls, shown in Figure 6. For this extended analysis of the reaction function, we cannot use the Christensen–Rudebusch measure of r^* since its only available back to 1999.

Downloaded from https://academic.oup.com/restud/advance-article/doi/10.1093/restud/rdab094/6484651 by Federal Reserve Bank of San Francisco Research Library user on 25 Ja

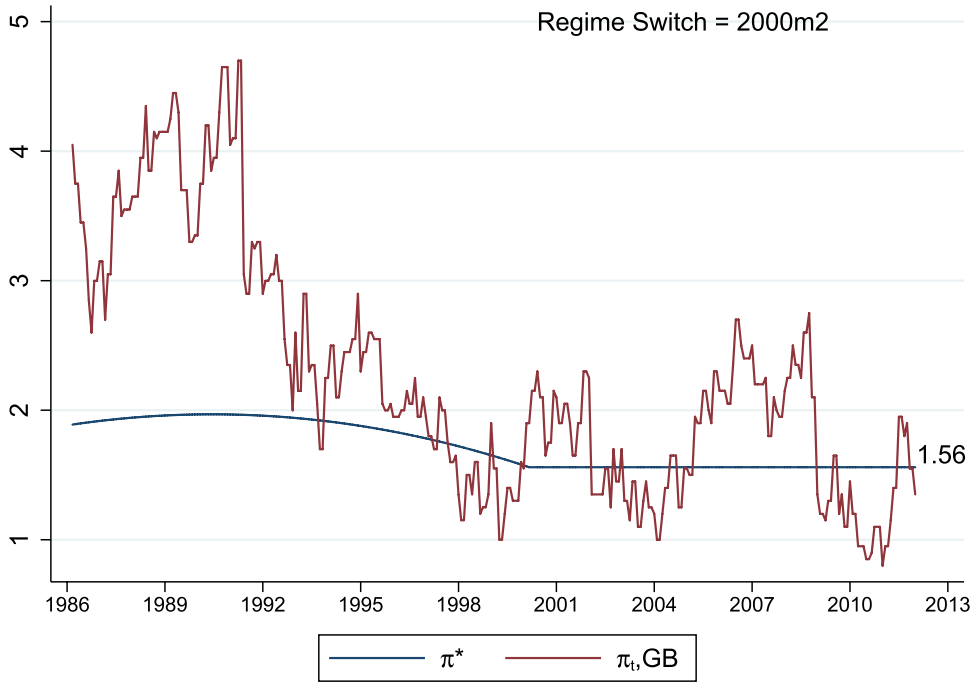

FIGURE 6
Estimated time-varying π^* FOMC meetings: 1986–2011 with additional factors

in magnitude and significance to those using the 1.5 measure over the 2000–08 period. For comparison purposes, we repeat this exercise using the inflation target measure from the FRB/US model (PTR). The results based on the PTR show a much lower R^2 and the wrong sign on the inflation gap term. In sum, the time-varying $\hat{\pi}_t^*$ obtained using our sentiment analysis approach yields a more reasonable and better-fitting interest rate rule than one based on conventional measures.

7. CONCLUSION

We proposed and applied an approach to directly estimating central bank preferences using sentiment text analysis of internal policy-making deliberations. This approach could be applied to the estimation of the objective function for any central bank that provides transcripts or detailed summaries of their internal meetings. In this article, we applied the approach to estimate the objectives of the U.S. Federal Open Market Committee, based on their internal meeting transcripts. Our baseline analysis, estimated on the period between 2000 and 2011, found that the FOMC's implicit inflation target was roughly $1\frac{1}{2}\%$. This result is robust to alternative functional forms for the loss function and is also obtained when we separately analyse the negativity expressed in the minutes of the FOMC meetings. This estimate of the inflation target is significantly below survey measures of longer-run inflation expectations over this period. However, we document that FOMC members often explicitly stated their preferred inflation target and $1\frac{1}{2}\%$ seemed to be the consensus choice at least up until 2009, when a shift up to 2% began to emerge. Indeed, we find corroborating evidence for this broader shift in preferences to 2% when we apply the sentiment analysis approach using data from FOMC minutes separately for the 2012–19 period.

Downloaded from https://academic.oup.com/restud/advance-article/doi/10.1093/restud/rdab094/6484651 by Federal Reserve Bank of San Francisco Research Library user on 25 Ja

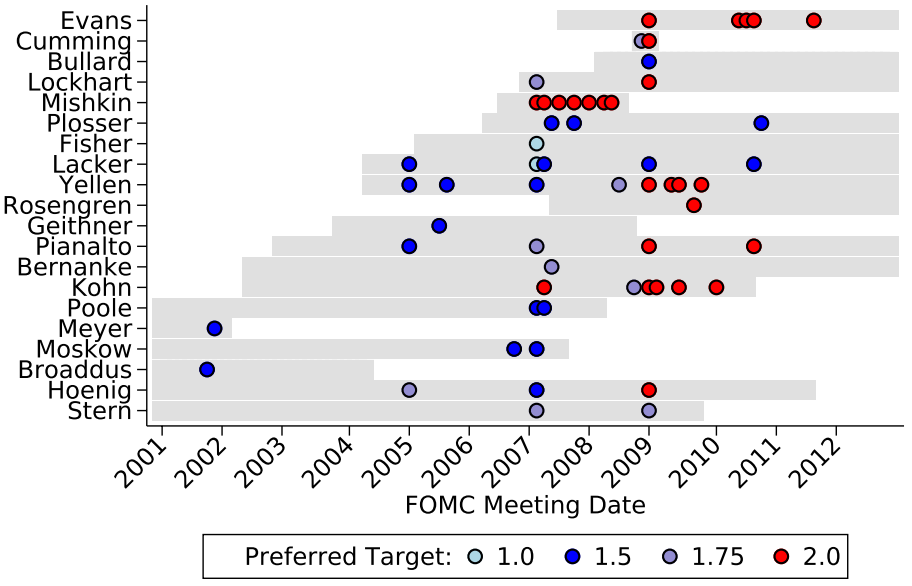

Note: Only speakers with at least one stated preference are shown. Each circle represents a stated inflation target preference from the indicated speaker (y-axis) at the indicated FOMC meeting (x-axis). Gray shading indicates speaker's FOMC tenure.

FIGURE 7

Explicit statements of inflation target preferences in FOMC meetings by speaker, over time

We also find that the FOMC's negativity was more strongly affected by output growth and financial conditions than by contemporaneous economic slack. The focus by central bankers' on output growth has been suggested in prior theoretical work (e.g. Walsh, 2003) and also found empirically in Taylor Rule estimations (e.g. Coibion and Gorodnichenko, 2011). It is also consistent with prior narrative analysis of FOMC communications by Thornton (2011), who showed that the FOMC routinely cited "growth in output," and not sustainable employment, the unemployment rate, or any concept of slack, as part of their policy directive from 1979 through 2008.

That financial variables may also enter the central bank's loss function is a much more contested issue. For example, Bernanke and Gertler (2001) argue that monetary policymakers *should not* respond to asset prices. Former vice chair Don Kohn (Kohn, 2006, 2008) went further, arguing that the Fed *does not* respond to asset prices. On the other side, Peek *et al.* (2015) augmented an empirical Taylor Rule with a variable measuring references to financial instability and found it to be significantly predictive of Fed interest rate changes. In a similar vein, Cieslak and Vissing-Jorgensen (2018) find that both actual negative stock market returns and references to stock market declines in the FOMC transcripts are predictive of target rate cuts. An interesting avenue of future research would be to better uncover why the FOMC appears to care about stock market returns, even after accounting for inflation, slack, and growth. One plausible reason is that FOMC members regard these variables as good indicators of the forward trajectory of the economy. That is, they serve as proxies for constant-interest-rate economic projections, measuring where the economy would be headed in the absence of any monetary policy actions.

Supplementary Data

Supplementary data are available at *Review of Economic Studies* online. And the replication packages are available at https://dx.doi.org/10.5281/zenodo.5598429.

Downloaded from https://academic.oup.com/restud/advance-article/doi/10.1093/restud/rdab094/6484651 by Federal Reserve Bank of San Francisco Research Library user on 25 Ja

Data Availability Statement

The data and code underlying this research is available on Zenodo at https://doi.org/10.5281/zenodo.5598429.

REFERENCES

AFROUZI, H., KUMAR, S., COIBION, O. and GORODNICHENKO, Y. (2015), "Inflation Targeting Does Not Anchor Inflation Expectations: Evidence from Firms in New Zealand", *Brookings Papers on Economic Activity* (FALL 2015), 151–208.
ANDRADE, P., GALÍ, J., BIHAN, H. L. and MATHERON, J. (2018), "The Optimal Inflation Target and the Natural Rate of Interest" (Technical Report, National Bureau of Economic Research).
APEL, M. and BLIX GRIMADI, M. (2014), "How Informative are Central Bank Minutes?", *Review of Economics/Jahrbuch für Wirtschaftswissenschaften*, **65**, 53–76.
APEL, M., GRIMALDI, M. and HULL, I. (2019), "How Much Information Do Monetary Policy Committees Disclose? Evidence from the FOMC's Minutes and Transcripts" (Technical Report, Sveriges Riksbank Working Paper Series).
BAKER, S. R., BLOOM, N. and DAVIS, S. J. (2016), "Measuring Economic Policy Uncertainty*", *The Quarterly Journal of Economics*, **131**, 1593–1636.
BALL, L. and MAZUMDER, S. (2018), "A Phillips Curve with Anchored Expectations and Short-term Unemployment", *Journal of Money, Credit and Banking*, **51**, 111–137.
BALL, L. M. (2014), "The Case for a Long-run Inflation Target of Four Percent" (Number 14-92. International Monetary Fund).
BARRO, R. J. and GORDON, D. B. (1983), "Rules, Discretion and Reputation in a Model of Monetary Policy", *Journal of Monetary Economics*, **12**, 101–121.
BAUER, M. D. and RUDEBUSCH, G. D. (2017), "Interest Rates under Falling Stars" (FRBSF Working Paper).
BERNANKE, B. S. and GERTLER, M. (2001), "Should Central Banks Respond to Movements in Asset Prices?" *American Economic Review*, **91**, 253–257.
BIANCHI, F., MELOSI, L. and ROTTNER, M. (2019), "Hitting the Elusive Inflation Target" (Technical Report, National Bureau of Economic Research).
BLANCHARD, O. (2018), "On the Future of Macroeconomic Models", *Oxford Review of Economic Policy*, **34**, 43–54.
BLINDER, A. S. (1997), "Distinguished Lecture on Economics in Government: What Central Bankers Could Learn from Academics—and Vice Versa", *The Journal of Economic Perspectives*, **11**, 3–19.
BLINDER, A. S. (2000), "Central-Bank Credibility: Why Do We Care? How Do We Build It?", *American Economic Review*, **90**, 1421–1431.
BOMFIM, A. N. and RUDEBUSCH, G. D. (2000), "Opportunistic and Deliberate Disinflation under Imperfect Credibility", *Journal of Money, Credit and Banking*, 707–721.
CAMERON, A. C. and MILLER, D. L. (2015), "A Practitioner's Guide to Cluster-Robust Inference", *Journal of Human Resources*, **50**, 317–372.
CARVALHO, C., NECHIO, F. and TRISTAO, T. (2019), "Taylor Rule Estimation by OLS" (Federal Reserve Bank of San Francisco).
CHAPPELL JR, H. W., HAVRILESKY, T. M. and MCGREGOR, R. R. (1997), "Monetary Policy Preferences of Individual FOMC Members: A Content Analysis of the Memoranda of Discussion", *Review of Economics and Statistics*, **79**, 454–460.
CHRISTENSEN, J. H. and RUDEBUSCH, G. D. (2012), "The Response of Interest Rates to US and UK Quantitative Easing", *The Economic Journal*, **122**, F385–F414.
CHRISTENSEN, J. H. and RUDEBUSCH, G. D. (2019), "A New Normal for Interest Rates? Evidence from Inflation-indexed Debt", *Review of Economics and Statistics*, **101**, 933–949.
CIESLAK, A. and VISSING-JORGENSEN, A. (2018), "The Economics of the Fed Put" (Working Paper).
CLARIDA, R., GALI, J. and GERTLER, M. (1999), "The Science of Monetary Policy: A New Keynesian Perspective", *Journal of Economic Literature*, **37**, 1661–1707.
COGLEY, T., PRIMICERI, G. E. and SARGENT, T. J. (2010), "Inflation-gap persistence in the US", *American Economic Journal: Macroeconomics*, **2**, 43–69.
COIBION, O. and GORODNICHENKO, Y. (2011), "Monetary Policy, Trend Inflation, and the Great Moderation: An Alternative Interpretation", *American Economic Review*, **101**, 341–370.
COIBION, O., GORODNICHENKO, Y. and WIELAND, J. (2012), "The Optimal Inflation Rate in New Keynesian Models: Should Central Banks Raise Their Inflation Targets in Light of the Zero Lower Bound?", *Review of Economic Studies*, **79**, 1371–1406.
DE VRIES, C. and LI, W. (2013), "Estimating Monetary Policy Rules with Serially Correlated Monetary Policy Shocks" (Unpublished Manuscript).
DEL NEGRO, M., EUSEPI, S., GIANNONI, M. P., SBORDONE, A. M., TAMBALOTTI, A., COCCI, M., HASEGAWA, R. and LINDER, M. (2013), "The FRBNY DSGE Model" (FRB of New York Staff Report, (647)).
DENNIS, R. (2006), "The Policy Preferences of the US Federal Reserve", *Journal of Applied Econometrics*, **21**, 55–77.
DI TELLA, R., MACCULLOCH, R. J. and OSWALD, A. J. (2001), "Preferences over Inflation and Unemployment: Evidence from Surveys of Happiness", *American Economic Review*, **91**, 335–341.

Downloaded from https://academic.oup.com/restud/advance-article/doi/10.1093/restud/rdab094/6484651 by Federal Reserve Bank of San Francisco Research Library user on 25 Ja

DOTSEY, M. and IRELAND, P. (1996), "The Welfare Cost of Inflation in General Equilibrium", *Journal of Monetary Economics*, **37**, 29–47.
ERCEG, C. J. and LEVIN, A. T. (2003), "Imperfect Credibility and Inflation Persistence", *Journal of Monetary Economics*, **50**, 915–944.
EROSA, A. and VENTURA, G. (2002), "On Inflation as a Regressive Consumption Tax", *Journal of Monetary Economics*, **49**, 761–795.
FAUST, J. and SVENSSON, L. E. (2001), "Transparency and Credibility: Monetary Policy with Unobservable Goals", *International Economic Review*, **42**, 369–397.
FAVERO, C. A. and ROVELLI, R. (2003), "Macroeconomic Stability and the Preferences of the Fed: A Formal Analysis, 1961-98", *Journal of Money, Credit and Banking*, **35**, 545–556.
FOMC (2012), "Statement on Longer-run Goals and Monetary Policy Strategy" [Online; posted 25-January-2012].
FUHRER, J. C., OLIVEI, G. P. and TOOTELL, G. M. (2012), "Inflation Dynamics When Inflation is Near Zero", *Journal of Money, Credit and Banking*, **44**, 83–122.
GIANNONI, M. P. and WOODFORD, M. (2003), "Optimal Interest-Rate Rules: I. General Theory" (NBER Working Papers 9419, National Bureau of Economic Research, Inc.).
GIVENS, G. E. (2012), "Estimating Central Bank Preferences under Commitment and Discretion", *Journal of Money, Credit and Banking*, **44**, 1033–1061.
HESTON, S. L. and SINHA, N. R. (2015), "News versus Sentiment: Predicting Stock Returns from News Stories" (Robert H. Smith School Research Paper).
HUTTO, C. and GILBERT, E. (2014), "Vader: A Parsimonious Rule-Based Model for Sentiment Analysis of Social Media Text", *Proceedings of the International AAAI Conference on Web and Social Media*, **8**, 216–225.
ILBAS, P. (2012), "Revealing the Preferences of the US Federal Reserve", *Journal of Applied Econometrics*, **27**, 440–473.
IRELAND, P. N. (2007), "Changes in the Federal Reserve's Inflation Target: Causes and Consequences", *Journal of Money, Credit and Banking*, **39**, 1851–1882.
JEGADEESH, N. and WU, D. A. (2017), "Deciphering Fedspeak: The Information Content of FOMC Meetings", *Monetary Economics: Central Banks–Policies & Impacts eJournal*.
KOHN, D. L. (2006), "Monetary Policy and Asset Prices" (European Central Bank).
KOHN, D. L. (2008), "Monetary Policy and Asset Prices Revisited" (Cato Institute).
LEDUC, S. and WILSON, D. J. (2017), "Has the Wage Phillips Curve Gone Dormant?", *FRBSF Economic Letter*, 2017–2030.
LIU, B. (2010), "Sentiment Analysis and Subjectivity", *Handbook of Natural Language Processing*, **2**, 627–666.
LOUGHRAN, T. and MCDONALD, B. (2011), "When is a Liability Not a Liability? Textual Analysis, Dictionaries, and 10-ks", *The Journal of Finance*, **66**, 35–65.
LOUGHRAN, T. and MCDONALD, B. (2016), "Textual Analysis in Accounting and Finance: A Survey", *Journal of Accounting Research*, **54**, 1187–1230.
LUBIK, T. A. and MATTHES, C. (2015), "Calculating the Natural Rate of Interest: A Comparison of Two Alternative Approaches", *Richmond Fed Economic Brief*, 15–10.
LUCAS, R. E. (2001), "Inflation and Welfare", in Leijonhufvud, A. (ed) *Monetary Theory as a Basis for Monetary Policy*, International Economic Association Series (Palgrave Macmillan) 96–142.
MALMENDIER, U., NAGEL, S. and YAN, Z. (2020), "The Making of Hawks and Doves", *Journal of Monetary Economics*, **117**, 19–42.
MEADE, E. E. and STASAVAGE, D. (2008), "Publicity of Debate and the Incentive to Dissent: Evidence from the US Federal Reserve", *The Economic Journal*, **118**, 695–717.
NAKAMURA, E. and STEINSSON, J. (2018), "High-frequency Identification of Monetary Non-neutrality: The Information Effect*", *The Quarterly Journal of Economics*, **133**, 283–1330.
ORPHANIDES, A. (2004), "Monetary Policy Rules, Macroeconomic Stability, and Inflation: A View from the Trenches", *Journal of Money, Credit and Banking*, **36**, 151–175.
ORPHANIDES, A. and WILLIAMS, J. C. (2006), "Monetary Policy with Imperfect Knowledge", *Journal of the European Economic Association*, **4**, 366–375.
PEEK, J., ROSENGREN, E. S. and TOOTELL, G. M. (2015), "Does Fed policy reveal a ternary mandate?" (Working Papers 16-11, FRBB).
PICAULT, M. and RENAULT, T. (2017), "Words Are Not All Created Equal: A New Measure of ECB Communication", *Journal of International Money and Finance*, **79**, 136–156.
ROMER, C. D. and ROMER, D. H. (2004), "A New Measure of Monetary Shocks: Derivation and Implications", *American Economic Review*, **94**, 1055–1084.
RUDEBUSCH, G. and SVENSSON, L. E. (1999), *Policy Rules for Inflation Targeting* (Chicago, IL, USA: University of Chicago Press) 203–262.
SHAPIRO, A. H., SUDHOF, M. and WILSON, D. (2018), "Measuring News Sentiment", (Working Paper, Federal Reserve Bank of San Francisco).
STOCK, J. H. and WATSON, M. W. (2007), "Why Has US Inflation Become Harder to Forecast?", *Journal of Money, Credit and Banking*, **39**, 3–33.
STOCK, J. H. and WATSON, M. W. (2016), "Core Inflation and Trend Inflation", *Review of Economics and Statistics*, **98**, 770–784.
SURICO, P. (2007), "The Fed's Monetary Policy Rule and U.S. Inflation: The Case of Asymmetric Preferences", *Journal of Economic Dynamics and Control*, **31**, 305–324.

Downloaded from https://academic.oup.com/restud/advance-article/doi/10.1093/restud/rdab094/6484651 by Federal Reserve Bank of San Francisco Research Library user on 25 Ja

SVENSSON, L. E. (2015), "The Possible Unemployment Cost of Average Inflation Below a Credible Target", *American Economic Journal: Macroeconomics*, **7**, 258–96.
TAYLOR, J. B. (1999), "A Historical Analysis of Monetary Policy Rules", in *Monetary Policy Rules* (NBER) 319–348.
THORNTON, D. L. and WHEELOCK, D. C. (2011), "What Does the Change in the FOMCs Statement of Objectives Mean?", *Economic Synopses*, 2011.
THORNTON, D. L., WHEELOCK, D. C., *ET AL.* (2014), "Making Sense of Dissents: A History of FOMC Dissents", *Federal Reserve Bank of St. Louis Review*, **96**, 213–227.
WALSH, C. (2003), "Speed Limit Policies: The Output Gap and Optimal Monetary Policy", *American Economic Review*, **93**, 265–278.
WALSH, C. E. (2004), "Robustly Optimal Instrument Rules and Robust Control: An Equivalence Result", *Journal of Money, Credit and Banking*, **36**, 1105–1113.
WALSH, C. E. (2017), *Monetary Theory and Policy* (Cambridge, MA: MIT Press).
WISCHNEWSKY, A., JANSEN, D.-J. and NEUENKIRCH, M. (2019), "Financial Stability and the Fed: Evidence from Congressional Hearings" (De Nederlandsche Bank Working Paper).
WOODFORD, M. (2003), "Optimal Interest-rate Smoothing", *The Review of Economic Studies*, **70**, 861–886.

Downloaded from https://academic.oup.com/restud/advance-article/doi/10.1093/restud/rdab094/6484651 by Federal Reserve Bank of San Francisco Research Library user on 25 Ja