Predicting Interest Rate Changes From Federal Reserve Transcripts

Indira Puri
Stanford University

MOTIVATION

US interest rates play a large role in the domestic and international economy. In this paper, we predict interest rate changes via processing of Federal Reserve transcripts.

The contribution of our research is twofold. First, we provide a means via which machine learning can be incorporated in macroeconomics to solve an important question. Second, our research can be used for prediction of interest rates far in advance. In particular, this work provides a means via which the public may anticipate changes in the interest rate, because transcripts are publicly accessible, as opposed to meetings with stockbrockers, access to former FOMC members, and so on, which much of the public does not have access to. Since we predict interest rate changes about six weeks in advance, when no new economic indicator data has been released, our research also provides a method for institutions and private bodies to anticipate interest rate changes well in advance of their occurrence.

DATASET

For target fed funds rate, we use the data series "target fed funds rate" provided by FRED, the Economic Research division of the Federal Reserve Bank of St. Louis. This series spans 1982-2008. Federal Open Market Committee transcripts are provided on the Federal Reserve website, federalreserve.gov. We are grateful to Miguel Acosta, who has previously examined transparency of the federal reserve by comparing their different types of documentation via latent semantic analysis [1], for providing us with transcripts already downloaded and processed from the Federal Reserve website, from 1976-2008. In total, we have 210 observations.

Fig. 1: Number of times the Fed chose to raise, hold, and lower rates, from 1982-2008.

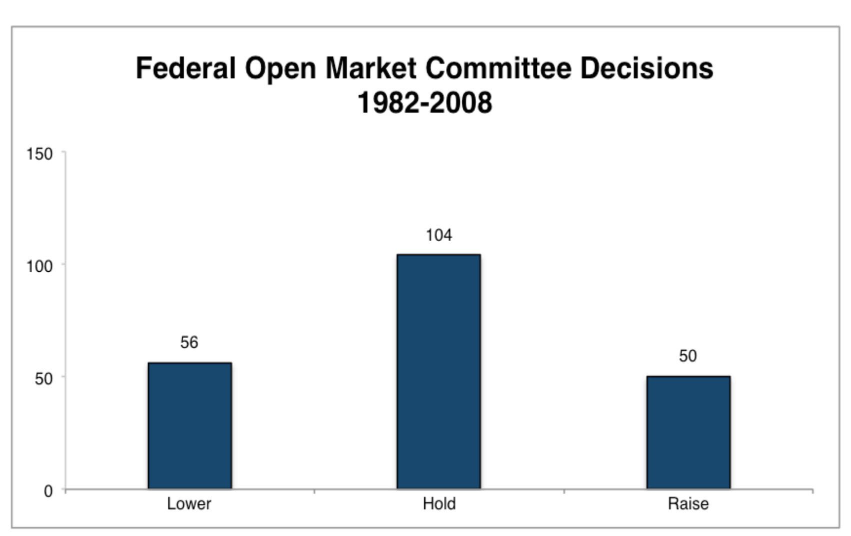

METHODS AND RESULTS

We considered Multinomial Naïve Bayes, SVMs, and Logistic Regression. While not discussed here, we make important decisions on how to represent transcripts and which evaluation metric to use. Standard NLP techniques are not suitable for this problem; nor is using accuracy as an evaluation metric.

Performance of Classifiers

We train models with varying numbers of previous transcripts. To account for the small size and skew of the dataset, we use leave-one-out cross validated F1 score as our evaluation metric.

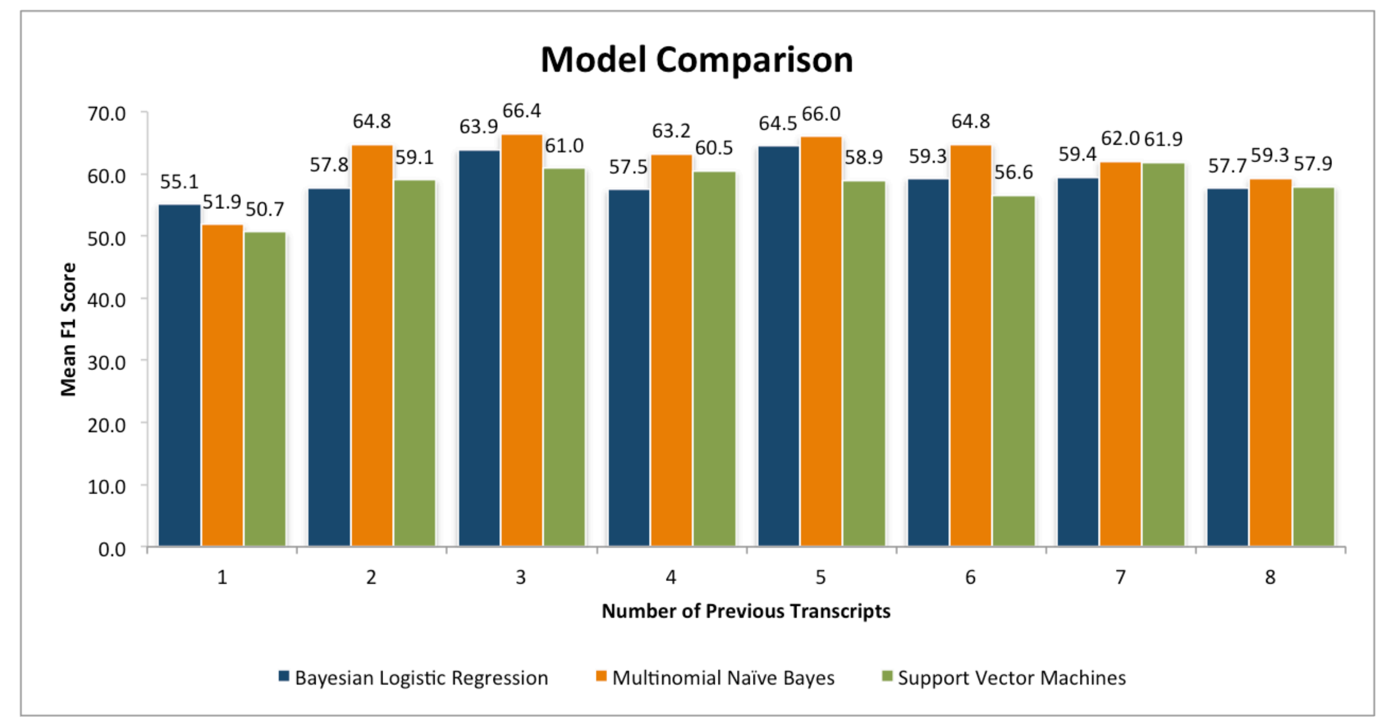

Fig. 2: Leave-One-Out Cross Validation Score Across Several Models.

Feature Selection

There were initially 30,000 features; we choose the top n features via the mutual information score. We then score each model on leave-one-out cross validated F1 score.

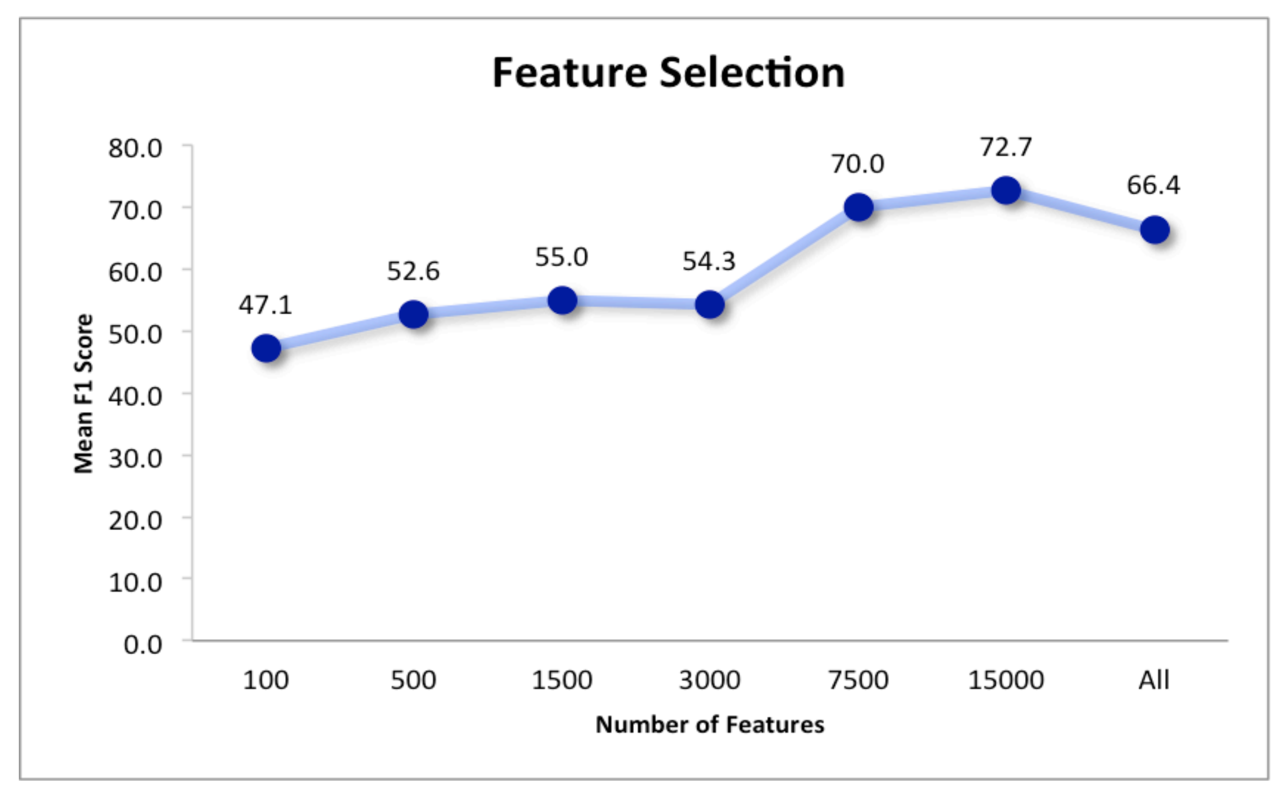

Fig. 3: Feature Selection for Multinomial Naïve Bayes with Three Previous Transcripts. N = 100, 500, 1400, 3000, 7500, 15000 tried.

Best Model

Our best model is MNB with 3 previous transcripts and 15,000 top features chosen via mutual information score, which has average out-sample accuracy 73%, in-sample accuracy 78%, and average out-sample F1 score 73%.

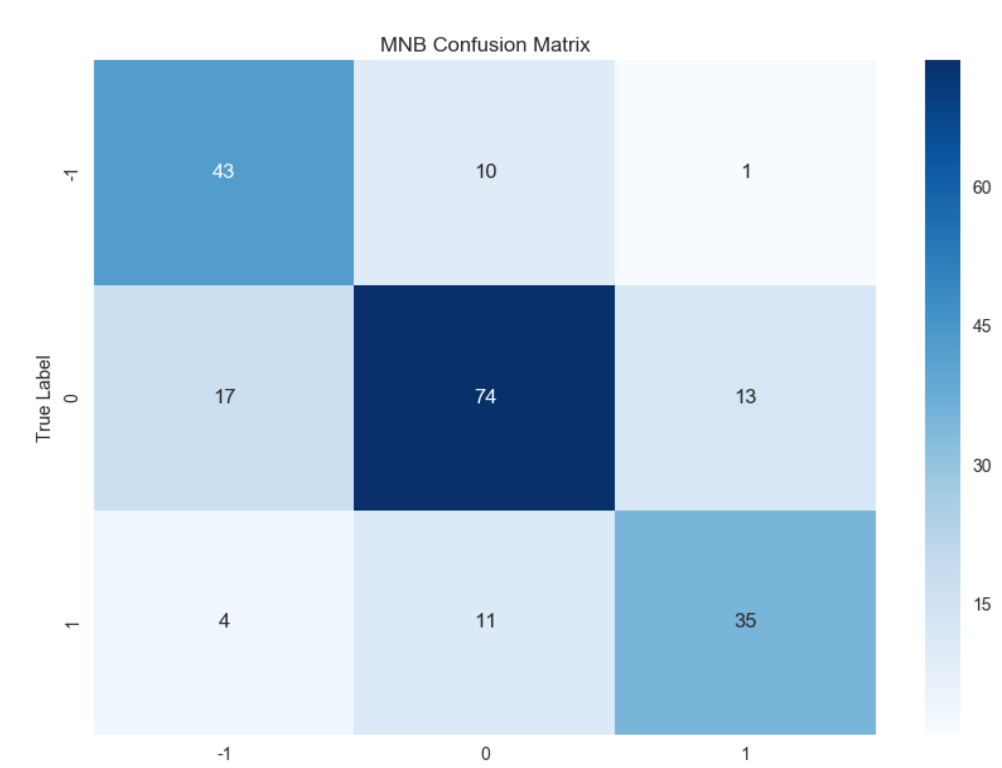

Fig. 4: Confusion matrix for the best model, MNB With Three Previous Transcripts and 15,000 Features.

DISCUSSION

Because this is the first paper to process Federal Reserve transcripts to predict interest rates, there is no analogous benchmark, but we can look at the success of similar problems:

1) Models that predict interest rate changes using publicly available data.
2) Models that predict interest rate changes using private data.

In the first category, [3] uses publicly available to predict the direction of Fed decisions. The major difference in their setup relative to this paper is that [3] tries to predict the direction of Fed decisions one day in advance, as opposed to six weeks in advance like this paper. This means that [3] allows access to new economic indicators, which this paper does not do. The highest in-sample accuracy on any of the models in [3] is 75% ([3], Table 3). In comparison, the in-sample accuracy of our best model, is 78%. The real test, of course, is out-sample accuracy, which [3] does not consider. Still, it is notable that our model outperforms one using explicit economic indicators in a shorter timeframe to announcement.

In the second category, [2] shows that from 1989-1993, the fed futures markets anticipated 41% of interest rate changes, and from 1994-2000, 76% of interest rate changes. In comparison, our out-sample accuracy is 73%. Their setup differs in that (1) they look at prediction accuracy about one month in advance (as they consider futures) whereas ours is one and a half months in advance, so the fed funds market has access to slightly more information and (2) they look at numerical accuracy instead of directional accuracy. Because the setups are not analogous, we say we ballpark private sector accuracy.

FUTURE WORK

Further research could include examining the accuracy of transcripts two meetings in advance (two months before decision) and looking at how well machine learning predicts numerical, as opposed to directional, interest rate changes.

REFERENCES

[1] Acosta, M. 2015. FOMC Response to Calls for Transparency, Finance and Economics Discussion Series 2015-060. Washington: Board of Governors of the Federal Reserve System.
[2] Lange, J., B. Sack, and W. Whitesell. 2003. Anticipations of Monetary Policy in Financial Markets, Journal of Money, Credit and Banking 35, 889–909.
[3] Lapp, J., D.K. Pearce, and S. Luksanasut. 2003. The Predictability of FOMC Decisions: Evidence from the Volcker and Greenspan Chairmanships, Southern Economic Journal 70, 312–327.