Transparency and Deliberation within the FOMC: a Computational Linguistics Approach

CFM-DP2014-11

Stephen Hansen[6,8], Michael McMahon[1,2,3,4,5,9] and Andrea Prat[3,7]
[1]Centre for Applied Macroeconomic Analysis; [2]Centre for Economic Performance; [3]Centre for Economic Policy Research; [4]Centre for Macroeconomics; [5]Centre for Competitive Advantage in the Global Economy; [6]Graduate School of Economics, Barcelona; [7]Columbia University; [8]Universitat Pompeu Fabra; [9]University of Warwick

Some policymakers, while accepting the argument that some degree of transparency is necessary for democratic accountability, argue that too much transparency would likely reduce the usefulness of monetary policy committee (MPC) meetings. Perhaps as a result, there is considerable heterogeneity in how much any given central bank reveals about internal deliberation. Understanding the extent of these costs of transparency is especially important now as some major central banks consider changing how transparent they are; the Bank of England and the ECB are both considering releasing more information about their MPC meetings.

Assuming that monetary policy makers are trying to protect their reputation as experts, economic theory suggests that transparency potentially has both positive and negative effects. As more transparency makes it easier to see a given members actions or contributions in the meeting, it may cause members to react by changing their behaviour in order to protect their reputation. The positive, or discipline, effect is that members will work harder in preparing for the meeting which likely makes them better participants. The negative, or conformity, result is that members may make special effort to say the same thing as everyone else which reduces the usefulness of the meeting (which is supposed to be free and frank exchange of expert views).

In order to explore which of these two potentially off-setting effects dominates, we conduct an empirical evaluation of the effects of greater transparency on central bank deliberation in the setting of the Federal Reserve's Federal Open Market Committee (FOMC). This analysis raises two main empirical challenges. Firstly, we use a difference-in-difference empirical strategy in order to try to uncover a causal relationship from changes in transparency to changes in deliberation. Second, since we want to measure deliberation, we need to make use of computational linguistic models (particularly Latent Dirichlet Allocation) to measure various aspects of the debate.

We find evidence that increased transparency both gave rise to discipline and conformity effects. In order to gauge the overall effect on the information content of the meeting, we examine whether those members who react most to the change in transparency become more or less influential (measured in terms of whether a member affects what the rest of the committee talks about after they speak). Given that we find that the new behaviour is more influential, we conclude that, overall, transparency appears to have been positive at the Fed since it brings more information to the meeting in spite of the fact that it creates a more sterile debate.

The challenge for central banks is to structure the deliberation process in order to maximize the discipline effect while minimizing the conformity effect.