

Changing Focus of the FOMC Through the Financial Crisis

John E. Miller

Computer & Information Sciences
University of Delaware
Newark, DE 19711
jmilller@udel.edu

Kathleen F. McCoy

Computer & Information Sciences
University of Delaware
Newark, DE 19711
mccoy@udel.edu

Abstract

The financial crisis of 2007-8 with effects continuing to this day is a complex beast. We use Latent Dirichlet Allocation (LDA) topic analysis and time series analysis to examine behavioral changes in the Federal Open Market Committee (FOMC) during 2005-08. We are the home fix-it guy with a new pair of pliers (topic analysis) and we will try it out on anything that needs fixing such as “Who and what was the financial crisis.” It’s not the only tool for the task, but it is a fun and informative tool to use.

1 Introduction

We investigate the financial crisis of 2007-8 using natural language processing tools to focus on “who and what was the financial crisis” (PoliInformatics, 2014). We are not very good at coloring within the lines and we have a new pair of pliers, LDA topic analysis, so we apply topic analysis to the Federal Open Market Committee (FOMC, 2014) transcripts of 2005-08 to see how FOMC behavior changed over time.

Our analysis process consists of:

- Preparation of FOMC transcripts,
- LDA topic analysis,
- Mean square successive differences to identify changing topic mixtures, and
- Sequence plots of topic mixture proportions.

Our work products include: FOMC topics, sequence plots of topic proportions, interpretation of findings, and critique of our procedure.

1.1 LDA Topic Analysis

LDA topic analysis uses a per document bag of words approach to determine topic compositions of words and document mixtures of topics. Figure 1 (Steyvers and Griffiths, 2007) shows a corpus explained as the product of topic compositions

(Φ) and document mixtures (Θ). Analysis constructs topics and mixtures of topics by assigning words to topics within documents. Topic compositions are interpreted as topics or themes of documents, conversations, or discussions. Document topic mixtures can be examined to see how document mixture proportions vary over time. Topic analysis reduces the dimensionality of a corpus by orders of magnitude from millions of words to frequency distributions of hundreds of topics.

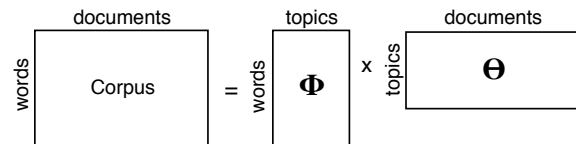


Figure 1: Topic Model

LDA topic analysis (Blei, 2012; Blei et al., 2003) is based on a generative probabilistic model where document mixtures and topic compositions are generated according to multinomial probability distributions Θ and Φ respectively. Analysis reverses the generative model, calculating weighted document topic mixtures Θ and topic word compositions Φ from the corpus. The topic analysis implementation used in this study borrows from UMass Mallet (McCallum, 2002).

1.2 Transcript Preparation and Analysis

Document preparation of the FOMC transcripts included:

- Clean the corpus and segment transcripts into ≈ 2000 word documents.
- Transform currency and numeric amounts, dates and durations, infrequent words, stop-words, and punctuation into reserved words.
- Form bigrams of word sequences with counts $> 1/3$ of the individual word counts; drop stop-word and punctuation reserve words.

FOMC transcripts covered 41 meetings and conference calls over 4 years. The original 41 transcripts of $\approx 1,700,000$ words were transformed to 708 documents of 686,742 words with 6,627 unique word types. We examined 50, 100, 200, and 400 topic models, choosing the 400 topic model for its slightly better held-out log likelihood and perplexity, and good quality topics as measured by coherence (Mimno et al., 2011), similarity, word distribution, and importance. Mixture proportions were initially estimated for the 708 documents and then aggregated into mixture proportions corresponding to the 41 original transcripts.

1.3 Topic Selection and Sequence Plots

We eliminated ≈ 200 of the 400 topics for having a low number of document occurrences. Of the remainder, we tested whether there were sequence effects, and constructed plots for topic mixture proportions showing sequence effects.

Our test proceeds as follows: mixture proportions are transformed to reals using the logit transform, $z = \log(p/1-p)$; estimates of the variance, $\hat{\sigma}^2$, and mean square of successive differences (MSSD), $\hat{\delta}^2$, are calculated on the logits; f ratios are formed, $f = \frac{2\hat{\sigma}^2}{\hat{\delta}^2}$, and compared to reference values. The MSSD (Montgomery, 1997) is calculated over the logit transformed mixture proportions as

$$\hat{\delta}^2 = \frac{\sum_{i=2}^n (z_i - z_{i-1})^2}{n-1}. \quad (1)$$

The variance estimate, $\hat{\sigma}^2$, includes variability over time, such as trends or shifts, while the MSSD estimate, $\hat{\delta}^2$, largely excludes such effects. If there is no sequence effect, $\hat{\sigma}^2$ and $\frac{\hat{\delta}^2}{2}$ estimate the population variance, σ^2 , and their f ratio should vary about 1.0. The two variance estimates are dependent and the variability is not stable over time, even with the logit transformation, so the resulting f ratio is not distributed as the F distribution and does not permit a valid statistical test. We use this f ratio as an indicator of a likely effect with reference values of 2, 2.5, 3, and 4. Mixture proportions of topics with f -values beyond the reference values are plotted over time.

2 Topics Showing Change

Table 1 shows topics ordered by document frequency along with f -value indicators of change

shown by the number of asterisks. A few high frequency topics for which there was no change are given for background. A subsequent one way analysis of variance (ANOVA) with group size of 2 was used to validate the MSSD method, and a few more topics are selected as shown by \dagger .¹

f	Description
*	Inflation - Price, Increase, Higher, Cost
*	Inflation, Policy, Economy, Market, Rate concerns
*	Interest rates
**	Credit, Financial, Bank, Capital, Loss concerns
****	Weakened financial market and economy
*	Strong - Cost, Side, Effect, Labor
***	Monetary policy, Financial market, Liquidity, Crisis
†	Move, Tighten, Point, Stop, Risk
*	Bank, Credit, Fund, Large, Liquidity, Borrowing, Cost
**	Productivity - Growth and Structural, Unit labor
*	Market participant, Firm, Policy, Action, FRB/US Model
*	Facility, System, Institution, Regulatory, Extend, Limit
*	Balance Sheet - Institution, Capital, Financial - Scenario
†	Savings rate, Low, Consumption, GDP
****	Commercial paper, Asset-backed
†	Money market, Mutual fund, Investor, Treasury
**	Hurricane Katrina impact
†	Financial market, Tail risk, Weak
****	New England, Region, Nation
* $f > 2$, ** $f > 2.5$, *** $f > 3$, **** $f > 4$, †other.	

Table 1: Topics with f -value Indicators

We have concisely identified meaningful themes from FOMC transcripts. Topics with non-empty f -value indicators show change over time and address “What was the financial crisis.”

3 Sequence Line Plots of FOMC Topics

Sequence plots of mixture proportions for the topics in table 1 offer more insight into the the financial crisis. We learned conference calls were different from meetings. They were shorter and showed more near zero mixture proportions for topics, likely a result of their focused nature.

The general topics *Inflation, Policy, Economy, Market, Rate concerns* and *Interest rates* (figure 2) show little change over time. The plot seems less than clear-cut, but removing conference call effects, removes any appearance of change as well.

The more sharply defined *Inflation - Price, Increase, Higher, Cost* topic (figure 3) shows a reduced concern with inflation starting late 2006 anticipating the mortgage crisis. Inflation concerns never dropped to zero and remained a topic of discussion even in conference calls.

¹The discrepancy between MSSD and ANOVA is small. Sequence plots reveal substantial variation between consecutive conference call and meeting proportions resulting in inflated MSSD estimates and reduced power for our test.

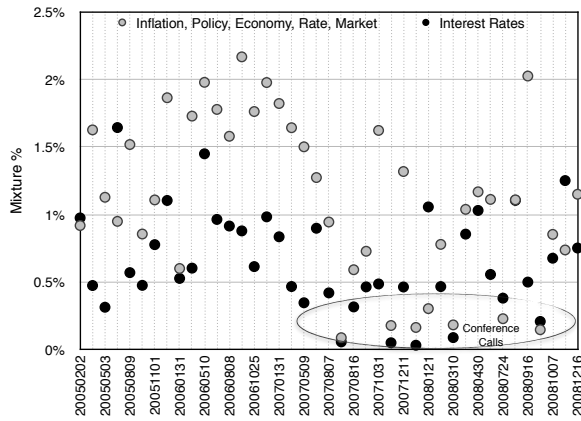


Figure 2: Important Topics No Change

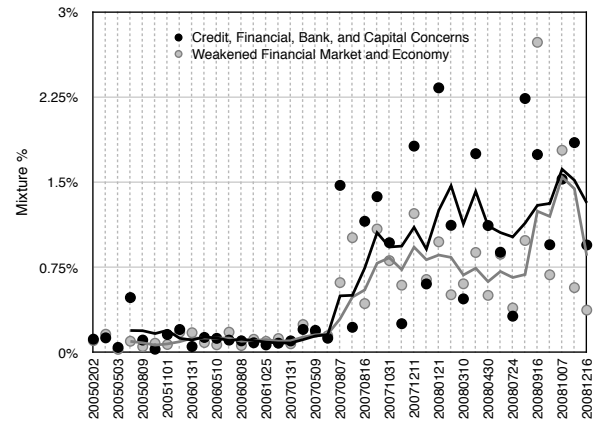


Figure 4: Credit Concerns; Weakened Economy

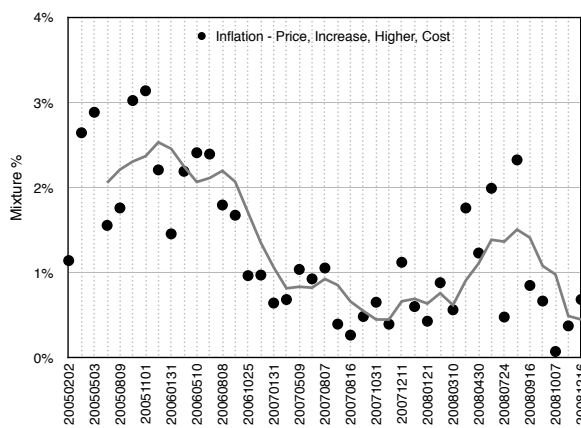


Figure 3: Specific Inflation Topic - Change

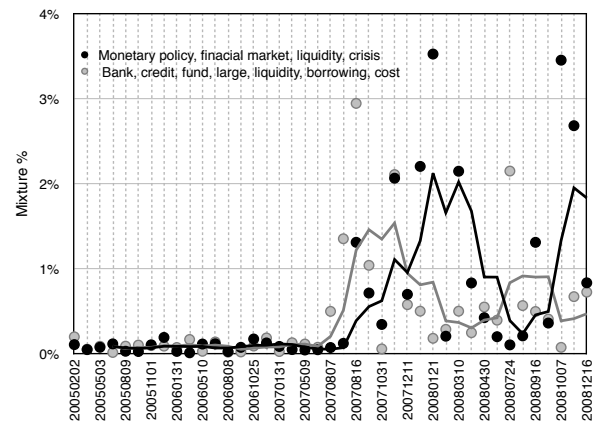


Figure 5: Monetary Policy; Banking

Credit, Financial, Bank, Capital, Loss concerns and *Weakened financial market and economy* topics (figure 4) became relevant with the August, 2007 meetings at the same time that Countrywide Financial's mortgage problems surfaced. *Monetary policy, Financial market, Liquidity, Crisis* and *Bank, Credit, Fund, Large, Liquidity, Borrowing, Cost* topics became relevant at the same time (figure 5). The financial crisis was then all of *Credit, Banking, Capital, Losses; Weakened financial market and economy; Liquidity and Credit Crisis*.

Productivity - Growth and Structural and *Strong - Cost, Side, Effect, Labor* topics both dropped to irrelevance at the same time that the Countrywide Financial and mortgage problems in general became evident (figure 6). The strategic *Productivity* topic remained irrelevant throughout 2008, while *Strong - Cost, Side, Effect, Labor* gained importance during the Bear Stearns takeover and terroristic September 2008. Similarly, *Market Participant, Policy, Action, FRB/US Model* and *New Eng-*

land, Region, Nation, while never a high percent of discussion, also dropped to irrelevance in August, 2007 (figure 7). The financial crisis was *ignored strategic productivity concerns and reduced attention to non-crisis factors in general*.

Commercial paper, Asset-backed and *Hurricane Katrina impact* were point effect topics during the 2005-08 period (figure 8). Hurricane Katrina, before the financial crisis, had a six percent portion of discussion in the September, 2005 meeting. *Commercial paper, Asset-backed* neared four percent portion of discussion during August, 2007 meetings including conference calls.

4 Summary and Discussion

Topic analysis effectively reduced the dimensionality of examining FOMC transcripts by orders of magnitude. Elimination of low importance topics and selection of topics by our MSSD test resulted in another order of magnitude reduction. With less than 30 remaining topics to investigate, it became practical to examine behavior for selected topics

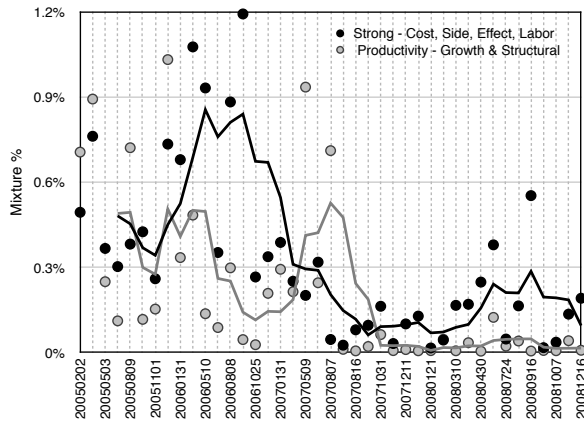


Figure 6: Strong - Cost, Effect; Productivity

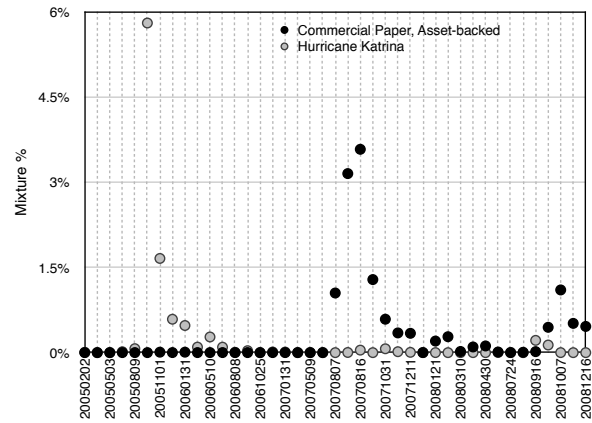


Figure 8: Commercial Paper; Hurricane Katrina

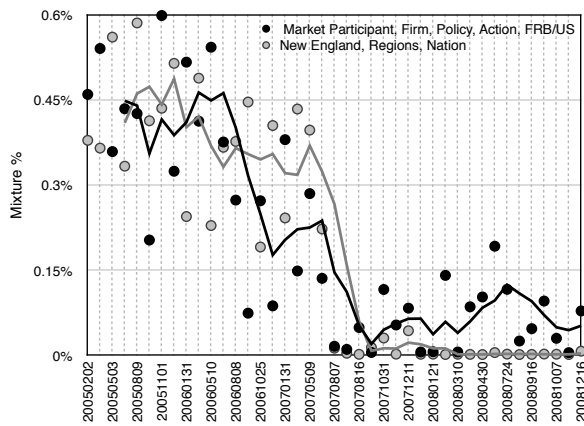


Figure 7: Market Participant; New England

using sequence plots. The review was highly informative thus showing the utility of topic analysis to identify overall themes and time series analysis to reveal changes in time with combined result addressing ‘what’ was the financial crisis.

The financial crisis was a horrible and complex beast. We looked at only the operation of the FOMC as recorded in its transcripts for 2005-08. Very evident were the changes that took place starting with the August, 2007 meetings, coincident with the exposure of the mortgage market crisis. Some previously important discussion topics neared zero percent of discussion and many other topics became vital at this same time.

We did not take into account differences between meetings and conference calls. Sequence plots showed the frequent conference calls during 2007-08 as having near zero mixture percent for many topics that were otherwise important in meetings during this period. The increased variability on topics for conference calls is a result of the small document size and short, focused

discussions. Such variability confuses the document mixture percents and reduces the power of our MSSD test for selecting topics showing change. Subsequent work should weight documents by size and note qualitative differences as meeting versus conference call.

Document preparation reduced the number of word-types and constructed bigrams that made interpreting topics easier. We will consider easier qualification of bigrams as well as trigrams in subsequent work. Named entity recognition might also help with the hope of identifying more than *banks* as the ‘who’ of the financial crisis.

We chose the 400 topic model as the slightly better quality model. However, the task of dealing with 400 topics was daunting even with censoring unimportant topics and our MSSD test for selecting topics. We will consider to using fewer topics in subsequent studies. We would also like to improve topic descriptions, perhaps through the improved use of bigrams and trigrams.

Interdisciplinary Work

The dimensionality reduction by orders of magnitude makes it easier for content experts to understand problems. Better topic descriptions and an improved process for selecting relevant topics would also make work easier. Going beyond our current state, a topic analysis delivering a *hierarchy of topics* and *mixture proportions* for such hierarchy would be even more propitious. There is mutual benefit from working together – natural language processing and machine learning experts developing finished tools for more focused and effective analysis in PoliInformatics.

References

- David M. Blei, Andrew Y. Ng, Michael I. Jordan. 2003. Latent Dirichlet Allocation. *The Journal of Machine Learning Research*, 3, p.993-1022.
- David M. Blei. 2012. Probabilistic Topic Models. *Communications of the ACM*, v.55 i.4, p.77-84.
- Johnathan Chang, Jordan Boyd-Graber, Sean Gerrish, Chong Wang, and David M. Blei. 2009. Reading Tea Leaves: How Humans Interpret Topic Models. *Neural Information Processing Systems*.
- Federal Open Market Committee. 2014. *Board of Governors of the Federal Reserve System*. http://www.federalreserve.gov/monetarypolicy/fomc_historical.htm.
- Andres K. McCallum. 2002. MALLET: A Machine Learning for Language Toolkit. <http://mallet.cs.umass.edu>.
- David Mimno, Hanna M. Wallach, Edmund Talley, Miriam Leenders, and Andrew McCallum. 2011. Optimizing Semantic Coherence in Topic Models. *2011 EMNLP Conference, Proceedings*:262-272.
- Douglas C. Montgomery. 1997. *Introduction to Statistical Quality Control*. John Wiley and Sons, New York, USA.
- J. Paisley, C. Wang, D. Blei and M. Jordan. 2014. Nested hierarchical Dirichlet processes. *IEEE Transactions on Pattern Recognition and Machine Intelligence*, (to appear).
- Mark Steyvers and Tom Griffiths. 2007. Probabilistic Topic Models. *Latent Semantic Analysis: A road to Meaning*. In T. Landauer, S. D. McNamara & W. Kintsch (eds.) Laurence Erlbaum.
- NLP Unshared Task on PoliInformatics 2014. 2014. <https://sites.google.com/site/unsharedtask2014>.
- John von Neuman. 1941. Distribution of the Ratio of the Mean Square Successive Difference to the Variance. *The Annals of Mathematical Statistics*, v.12, n.4, p.367-395.